Just a Guy: What Is One Person Worth Now?

One director, many executors. The question is no longer how many hours I worked. It is what the finished work is worth.

By William Tygart · November 2026

I am one person. No team, no agency floor, no payroll of specialists behind me. And yet, when I sat down this week and honestly priced the work documented in my own work orders, the completed total came to somewhere between $94,175 and $162,150 — in what was really two days, November 11 and 12.

That number stopped me cold, because it does not fit any way I was ever taught to value work. It cannot be hourly — nobody works enough hours in two days to explain it, and I did not. It is not really “project based” in the old sense either, because there was no single project. What there was, underneath everything, was a stack of work orders: scoped, written instructions, handed to AI executors, that came back as finished, checkable deliverables.

So I want to break this down carefully — what was actually produced, how I priced it, what I deliberately refused to count, and what I think it means. Because if one person plus AI workers can produce at this rate, then the honest question is not how fast is the machine? It is: how do you value any of this now?

One scope note, because it matters: this is not everything I did that week — only what the collected work orders document. If it was not written down as a work order, with a deliverable I could point to, it is not in here.

The Count — Start with what the work orders actually show

I priced everything the same way a client would buy it standing alone: at ordinary US agency and consultant retail rates — a senior web analyst or CRO consultant at $150–$250 an hour, implementation work blended at $125–$200 an hour. Freelance-direct would run roughly 30–40% lower. These are ranges, not quotes, and they are deliberately the buyer's prices, not my costs.

The Bot Lab work. Fifteen work orders sit in the Bot Lab folder, plus one held vault-audit order. Six of them are complete — a report received, a package delivered, an execution finished and recorded:

Completed Bot Lab work order	Retail value
Bot log pull — 8 days, 468K requests, reconciled against Henkan Vision	$750–$1,250
Deep bot analysis — 55 days, 2.84M requests, fleet behaviour signatures, security verification	$3,500–$6,000
White paper package — 19-page paper, 15-slide deck, 1-page executive brief	$4,400–$7,100
Taxonomy audit — 2,063 tags, 50 slug collisions, costed action plan	$1,250–$2,000
Tag strategy audit — re-crawl analysis, top-100 under-tagged list, 5 new tags approved	$1,250–$2,000
Tag execution — database backup first, 472 posts touched, 473 tags added, tagging standard published	$1,000–$1,750
Bot Lab, completed	**$12,150–$20,100**

That work is forensic, unglamorous, and real: reading origin server logs closely enough to separate an AI indexer from a search retrieval from a live agent fetch, because at the brand level you cannot tell what a bot was actually doing on your site. Only the logs know.

The Count, Continued — The big one: 700+ pages — and the recount proved the bigger number

The largest single body of work was the Universal Commerce Protocol site modernization, executed by Antigravity through a master work order and a long batch trail in Google Drive. I had been saying “700+ pages.” My first count came back at 639, and I said that count was wrong — that every article had been scanned. So we dug deeper, and the recount is in this article because the correction is the method. This time the correction ran the other way: the deep dig proved the bigger number right, and showed exactly where the first count had gone wrong.

UCP modernization — evidence	Pages
Result records marked PASS (559 batch entries, deduplicated by post ID)	541
Evidenced only by execution scripts, batch definitions and link maps, no result file	167
Distinct posts documented in Drive evidence	**708**
Live-site verification, page by page: posts carrying the modernization signature	730 of 730
Scanned / audited: the entire published corpus (738 items incl. pages)	730 posts

Where did 639 come from? The first pass credited only 82 pages to the execution scripts. A full extraction across some 155 script artifacts found 665 distinct post IDs in the scripts alone — 167 in no result file at all — plus sources the first count had missed. And “about 698 articles” was a mid-program snapshot, not a ceiling: the site now holds 730 published posts, every one carrying the modernization signature — 726 with the full dual schema, 660 with the complete treatment of banner, unique matrix, FAQ and schema. Every one of them was scanned. That sentence, it turns out, was simply true.

Pricing taught me something, too. At a hand-built rate of $400–$800 a page, the program would “value” at a number nobody would ever actually pay. Batch work at scale is priced as a program, not as 730 artisanal pages — so I used volume rates of $100–$175 for each of the 541 result-verified pages and $75–$125 for each of the 167 script-evidenced pages, plus the strategy work and a homepage repair package. That puts the program at $68,625–$118,950. The per-unit price fell; the total value did not. That trade is the whole story of AI-era production in miniature.

Completed, client-type work — Nov 11–12	Retail value
Bot Lab (6 completed orders)	$12,150–$20,100
UCP modernization program (708 posts in Drive evidence; 730 live-verified)	$68,625–$118,950
Setup and systems: Henkan Vision rollout on 5 sites, Henkan API connector, InventoryQuant one-pager, Advancing Prefabrication 2027 analysis	$2,750–$5,500
Five forensic server-log reports, one per hub site — all five delivered	$8,000–$13,000
tygartmedia.com Henkan fix sprint — executed and live	$1,200–$2,000
healthcarefacilityhub.org Henkan fix sprint — executed and live	$1,450–$2,600
Total completed	**$94,175–$162,150**

The phone operating layer I built the same week — getting AI control of my own phone working and self-healing — is real work, worth roughly $800–$1,500 on the same basis, but I classify it as internal infrastructure and R&D;, not client-type work, so it stays out of the headline total.

The Rule — Only completed work counts. Everything else is zero.

A number like $94,175–$162,150 invites scepticism, and it should. The only reason I am willing to publish it is the rule I priced it under — my rule, set before the counting started:

> ***“Commissioned is not produced. In-progress is not produced. Only a delivered, verified job counts — everything else is booked at $0.”***

Applied honestly, that rule hurt — and then it paid. On November 12 I commissioned a substantial pipeline: Henkan fix sprints for two sites and five forensic server-log reports, one per remaining site. Under the rule, all of it counted as exactly zero while it was in motion. Then it landed: all five reports delivered, both fix sprints executed and live — and only at that point did $10,650–$17,600 of it move into the completed total above. What is still in motion — a cluster-wide fix round across the five hub sites and the per-company platform fleet dossiers, $6,000–$9,900 of scoped work — still counts as exactly zero, as do two other orders for the opposite reason: one I closed as unnecessary, and one held pending verification of its lead. Zero is zero.

Even inside the completed pile, the caveats stay attached. The white paper was counted at delivery value, but it carries two known flags — an origin IP mislabelled as an attacker, and a Cloudflare recommendation that does not apply to a site that is not on Cloudflare — and the correction is still owed before it is publishable. The UCP completion now rests on three legs: the executor's result records, the full Drive script trail covering 708 distinct posts, and an independent, page-by-page fingerprint check of the live site that found the modernization signature on all 730 published posts. I am saying that here, in the article, because a valuation that hides its soft spots is marketing. One that shows them is accounting.

Why the hour broke — Hourly billing measures my presence, not the production

Hourly billing rests on a quiet assumption: output scales with one person's time. When my executors run in parallel — one modernizing pages in batches, another working through 2.84 million log lines — that assumption collapses. Divide the completed value by my hours and you get an “hourly rate” so absurd it describes nothing. Divide it by the machines' hours and you get the opposite lie, as if the judgement that scoped the work were free.

> ***“The hour did not get faster. It stopped being the unit.”***

The Mechanism — The work order is the unit now

If the hour is not the unit, what is? In my week, it was the work order: a written package that carries four things a raw prompt does not.

- **The evidence travels with the task.** Each fix order named the exact pages, session counts, and click targets from the Henkan data. A vague prompt forces the executor to guess, or to redo the analysis worse.
- **Guardrails prevent paid wrong turns.** They say what not to touch, and that turns out to be worth money on its own — see the box below.
- **Shared definitions make results comparable.** Five log reports written to the same lab definitions line up against the same baselines; five cold prompts would produce five different shapes.
- **A verification loop exists.** Baseline numbers and a check-back date are written into the order, so “done” can actually be tested.

One paragraph that saved real money. Item 5 of the tygartmedia fix order read, in effect: *do not chase the js_error flags — verified Safari AbortError noise.* Every flagged session was Safari, each with exactly two AbortError entries, reading on normally afterwards. Without that paragraph, an executor could plausibly have burned one to three hours debugging a non-bug: an estimated $200–$600 of wasted work, avoided by a sentence of context.

Now the leverage, stated carefully. On November 12, about ten minutes of voice conversation produced the two fix orders: $2,650–$4,600 of scoped implementation. One further sentence — “one work order for each of the sites” — produced five

log-report orders: $8,000–$13,000 of scoped analysis.

My input	Scoped work set in motion
~10-minute voice call	$2,650–$4,600
One sentence	$8,000–$13,000
~15 minutes total	**$10,650–$17,600**

That is roughly $700–$1,200 of commissioned work per minute of my time — and note the word commissioned: under my own rule it is worth $0 until it lands and verifies. Writing those seven orders myself, to the same standard, would have taken an estimated four to seven hours. The real delta was never typing speed. It was context compression: months of lab baselines, standards, and “do not touch this” knowledge, folded into a three-minute conversation. A cold prompt starts from zero, every time.

The Cost Side — What does an AI worker cost — and what is it worth?

People ask about cost as if it were mainly tokens and subscriptions. Those are line items, and compared with the value above they are small. The real cost ledger of an AI worker, as this week actually felt, has four entries:

- **Direction.** My time scoping the order well enough that the evidence, guardrails, and definitions travel with it. Cheap per order, expensive if skipped — skipped direction is repaid as rework.
- **Verification.** Reading 559 PASS records back against the full Drive script trail, then fingerprinting all 738 live URLs — 730 of 730 posts — instead of repeating “700+” on faith, and checking that a test visit really appeared as a live session. Verification is the cost that makes the value claim defensible, so it is the last cost I would cut.
- **Rework and wrong turns.** The risk premium on every job. Guardrails exist precisely to buy this down — one avoided wild goose chase was worth $200–$600 on a single order.
- **Infrastructure and tools.** The subscriptions, the API connector that turned every later Henkan pull into a one-minute job, the systems work that makes the next order cheaper than the last.

And the worth side has a simple answer, the one this whole count is built on:

> ***“Value an AI worker by the verified deliverable, priced as a client would buy that deliverable on its own — never by the minutes the machine spent producing it.”***

That also settles the pricing question. Hourly fails, for the reasons above. A flat per-unit price fails at scale, as 730 pages showed. What survives is pricing the job: a defined deliverable, a defined standard, a flat or program price. It is how I would now sell the work this lab produces — a forensic server-log report at $1,950 a site, a friction audit at $950, a fix sprint at $1,500, or all three bundled at about $3,950, against a market where standalone log audits run $2,000–$8,000. The buyer is not purchasing my afternoon. They are purchasing a finished, verified answer.

And for genuinely small one-shot tasks — fix a heading, swap an image — the honest advice runs the other way: skip the work order and prompt the executor directly. The order's overhead only pays when the task is evidence-heavy, multi-page, or easy to do wrong.

The Trap Underneath All of It — Everything can be automated. That is exactly the problem.

Here is the part nobody warns you about. With effectively unlimited compute sitting next to me, everything can be automated. I can automate pressing a button. I could spend half an hour automating how I get coffee — or I can just get up and go get the coffee. Both are now genuinely available options, and that is the trap: automatable is not the same as worth automating. When capacity was scarce, scarcity did the filtering for you. When capacity is unlimited, the filtering becomes the job — and if you do not do it deliberately, you will spend your infinite leverage automating things that should have taken ninety seconds by hand.

I know this from the inside, because I spend twelve to sixteen hours a day in this work, and until now I have had no real way to label any of it. Is this hour billable or non-billable? Is it direct — straight toward the business — or indirect? Is it personal, or is it for the lab? I honestly could not have told you, day by day, and a man who cannot see his own time cannot steer it. That, more than the pricing puzzle, is why the work-order valuation and tagging practice in the next section exists. It is not an accounting flourish. It is a self-organization tool — a way of looking at myself and asking: what do you actually need to do, how do you organize this, and how do you make sure you are not just spinning your wheels very, very fast?

Why the “Unbillable” Hours Are Not Waste — Yesterday's learning is the capital today's leverage is spent from

But the tags have to be read honestly, and honesty cuts in a direction that surprised me. All that prior work that looked unbillable — the lab experiments, the log-diving, the systems tinkering — may be exactly what gave me the knowledge to spend fifteen minutes of direction and put out $10,000 worth of scoped work. Nobody conjures that fifteen minutes from nothing. It is drawn against months of accumulated context, and the context was built in hours nobody would have paid for at the time.

So the labels are not there to shame non-billable time. They are there to see it accurately — so the compounding becomes visible. Some indirect work is wheel-spinning, and the tags will expose it. Some of it is capital formation wearing a disguise, and the tags are the only way I will ever know the difference. Either way, I would rather know.

From Here — Every work order gets priced and tagged — before it runs

The lesson I am keeping is not the headline number. It is the discipline that made the number trustworthy. Every internal work order I issue will now carry two things up front, before any work starts:

- **A retail value estimate** — what this deliverable would cost a client to buy, stated as a range, on a consistent rate basis.
- **A category tag** — client work, lab, R&D;, or internal infrastructure — so nothing quietly inflates the wrong column the way my phone work could have.

Then each job gets judged twice — before and after — on four plain questions: Is this worth our time? How much time and resources will it actually take? What is the net result? What is the value of the work? Value gets booked only when the job is delivered and verified, exactly as this week was counted. Over time that gives me something no timesheet ever did: a record of which kinds of orders actually pay, and which kinds only feel productive.

What It Means — One person is not one person anymore

I keep coming back to the strangeness of the plain fact. I am just a guy. I did not hire anyone, and I did not work some heroic schedule — the famous “week” in my own lab records is really two days, and I am not going to dilute that by calling it anything longer. What changed is not my capacity. It is that my judgement — what to do, what not to touch, what counts as done — can now be written down once, as a work order, and executed many times over, in parallel, without me standing over it.

That means the scarce asset in a one-person business is no longer labour. It is context: the baselines you have built, the standards you can state, the mistakes you know to fence off, and the willingness to verify instead of believe. Whoever holds the best context writes the best work orders, and the best work orders are what the machines are worth.

So, how do you value an AI worker? The same way you now have to value the person directing it: not by the hour, and not by the prompt — by the completed, verified job, and by nothing else.

And none of this is really about one guy. I am doing it alone, but the same discipline — price the order up front, tag what kind of work it is, book value only on verified delivery — conveys to anyone: another solo operator, a small firm, or a department inside a large company where AI agents are already doing work that hourly accounting simply cannot see. The scale changes. The labeling problem does not.

Method note: all figures are retail-equivalent values at US agency/consultant rates (senior analyst $150–$250/hr; implementation $125–$200/hr blended), drawn from work-order records, execution result files, and Drive evidence for November 11–12, 2026. The UCP figures are from the November 12 deep-dig recount: Drive execution evidence for 708 distinct posts, with all 730 published posts verified live, page by page. Retail value is not cash received. Pipeline work is carried at $0 until delivered and verified.