The final business declaration

TITANOS is a one-person AI engineering practice in Brisbane, run with a team of AI agents and a strict test gate. It builds software, security checks and AI systems for small and mid-sized businesses.

This page is everything in the open: what has been built, what it cost, how fast it ran, how it is tested, how it is kept safe, and how you can do the same. It is free to read, copy and share.

Online version: https://titanos.tech/declaration

Key achievements

- 7,236 commits in the engine repository in 43 days (first commit 29 September 2026).
- About 168 commits a day on average across those 43 days.
- 527 automated test suites in the engine's main gate.
- 6,314 automated test cases passing across the estate on 9 October 2026.
- 256,217 lines of production code across 12 repositories (counted 13 October 2026).
- Cost per AI job fell 8.4 times, from US$0.97 to US$0.12, while quality on the matched set rose from 23 of 25 to 25 of 25.
- 1,549 agent jobs in one 24 hour window, with a 7.3% failure rate (failed or stuck) and a median job time of 12 minutes.
- 1,041 commits in this website's repository since 4 June 2026.

AI engineering capability

Cost

- About US$30 a month of AI tools while the main engine was built (to 29 September 2026).
- About AU$300 (estimate, source pending) a month of AI tool subscriptions since then.
- Cost per AI job fell 8.4 times, from US$0.97 to US$0.12, while quality on the matched set rose from 23 of 25 to 25 of 25.
- Receipted subscription costs are not measured yet; the two monthly figures above are the operator's own numbers and are labelled that way in the appendix.

Speed

- About 168 commits a day on average across those 43 days.
- 1,549 agent jobs in one 24 hour window, with a 7.3% failure rate (failed or stuck) and a median job time of 12 minutes.
- Case studies below give start and finish times for whole builds, taken from the git log.

Quality

- Nothing counts as done until the full gate is green and the result has been read, not assumed.
- A second agent reviews each change; no agent signs off its own work.
- Cost per AI job fell 8.4 times, from US$0.97 to US$0.12, while quality on the matched set rose from 23 of 25 to 25 of 25.

Tests and automated testing

- 527 automated test suites in the engine's main gate.
- 149 test commands in this website's test manifest, plus built-page checks that run after every build.

- 114 test files and 1,054 test cases in the customer backend (sign-in, billing, security).
- 21 mutation-test scripts that break the code on purpose to prove the tests notice.
- Every failure that reaches a person becomes a regression test the same day.
- Copy is linted too: no dashes, no filler words, and no addresses or credentials in anything public.

Security

- Outside parties are only ever checked passively, from public DNS records. Nothing of theirs is touched.
- Money, credentials, legal commitments, messages to third parties and anything irreversible need a human approval tap first.
- Secrets live outside the code and are never pasted into chat; a secret scan is part of the engine's test gate.
- The website sets a content security policy and publishes a security contact file (security.txt).
- The customer backend verifies signed payment events, rate limits sensitive routes per visitor, and sends no-store and anti-framing headers on every response.

Hardware

- The workstation is a laptop: Intel Core i7-7700HQ (8 logical CPUs), 23 GB of RAM, running Linux under Windows.
- Full test suites run on a cloud test runner: AMD EPYC 7B12, 4 cores, 15 GB of RAM.

Subscriptions and prices

- **Free email security scan**: AU$0 https://titanos.tech/scan
- **TITANOS Monitor**: AU$149 a month or AU$1,490 a year https://titanos.tech/monitor

- **Job Flow for tradies**: AU$249 a month https://titanos.tech/tradies
- **AI Growth Partner**: AU$1,500 a month, 3-month minimum https://titanos.tech/ai-delivery
- **AI Ops Partner**: AU$2,500 a month, 3-month minimum https://titanos.tech/ai-delivery
- **Embedded AI Partner**: from AU$5,000 a month, 3-month minimum https://titanos.tech/ai-delivery
- **Lead lists**: AU$497 to AU$3,497 one-off, or AU$990 a month https://titanos.tech/leads
- **Compliance pack**: AU$5,997 one-off https://titanos.tech/compliance
- **Enterprise programmes**: AU$49,997 to AU$149,997, scoped on a call https://titanos.tech/enterprise

Prices are read from the site's own price file and link to the page that sells each one. See also /pricing and /costs.

Case studies

Used vehicle dealership · REAL

A dealership command centre, built and independently reviewed in about 70 minutes

- **About 70 min** from the first job starting (15:36) to the reviewed build (16:46), 14 October 2026
- **23** jobs run in parallel: 4 discovery, 10 builders, 8 reviewers and 1 lead
- **598 / 598** automated tests passing
- **12 ms** to list 10,000 vehicles (slowest 5% of runs, one database measurement)

https://titanos.tech/case-studies/dealer-command-centre

Sales coaching · REAL

A 1,900-lesson sales course with its own coach, built in under 90 minutes

- **83 minutes** from the first commit (10:44) to the full course (12:07), 9 October 2026; 409 of its 412 commits
- **1,900** lessons in 190 modules, about 2.7 million words, each lesson with a five-question quiz
- **412** commits in the project history
- **35 / 35** automated tests passing

https://titanos.tech/case-studies/sales-coaching-course

Education and practice trading · REAL

A full trading-education app, idea to working product in one morning

- **3.5 hours** from first idea (05:10) to the finished app (08:40), 7 October 2026
- **56** commits in that window, all on the record
- **250 / 250** automated tests passing
- **About 6%** of one week of a flat-rate AI subscription

https://titanos.tech/case-studies/trading-app

Each case links to its own page with dates, costs and what was tested. Client names are never published, and every figure traces to the build's git log or a dated usage reading.

How we do it

1. **Write down what done means first.** Before anything is built, list the acceptance checks: what must work, what must not break, what the limits are. Everything later is judged against that list, not against a feeling.

2. **Run one orchestrator and many small agents.** One session holds the plan and the context. It hands self-contained briefs to agents, each with one job, its own branch and the exact files it may touch. Small, bounded jobs fail cheaply.
3. **Tests before claims.** A change is not described as working until a test shows it. Each bug that gets through becomes a regression test, and mutation tests confirm the tests really fail when the code is wrong.
4. **Run the full suite on a cloud test runner.** The workstation stays free for building. The whole suite runs on a separate machine from a clean copy, and the pass and fail counts are read, not assumed.
5. **Use independent review.** A different agent, with no stake in the change, reviews it. It only blocks on a demonstrated fault. Everything else goes on a list for later, so work keeps moving.
6. **Put approval gates around the irreversible.** Spending money, handling credentials, signing, messaging a third party and deleting things go to a person as a one-tap card. Everything reversible and in scope just gets done.
7. **Keep receipts.** Every change leaves a commit, a test result and a note of what is still unknown. A number without a source and a date does not get published.
8. **Turn repeats into code.** Anything done twice becomes a script with a test. That is how a one-person practice keeps pace: the machine does the known, the person decides the unknown.

This is not magic and it is not hands-off. A person sets the direction, approves the gated steps and owns every decision that matters.

Courses

Trading platform (includes the AI course)

A phone-first trading education platform: a 200-lesson course, live charts, an AI tutor on every screen and a practice-money journal. The course on how to use AI sits inside it.

Sign up: https://learn-trading.titanos.tech. Case study: https://titanos.tech/case-studies/trading-app

Sales course

A sales course of 190 modules and 1,900 lessons with quizzes and an AI coach for questions, role-play and speaking practice, built for a sales coach.

Join the waitlist at https://titanos.tech/declaration#courses. Case study: https://titanos.tech/case-studies/sales-coaching-course

Daily artifacts

Join the mailing list at https://titanos.tech/declaration#daily-artifacts

Daily artifacts from TITANOS. Unsubscribe any time.

Appendix: sources for every number

Every figure in this declaration, with where it came from and when. Measured on 10 November 2026 unless a different date is shown. MEASURED means counted from a real source. COMPUTED means arithmetic on measured figures. PUBLISHED means read from the live price file. OPERATOR_STATED means the operator's own figure with no receipt on file. Figures dated 9 October 2026 or 13 October 2026 come from the site ledger and were not re-measured for this document.

Figure	Status	Source and method	Date
7,236	MEASURED	git rev-list --count origin/main, engine repository. Commit count on	2026-11-10

Figure	Status	Source and method	Date
		the main branch; first commit date from git log.	
168	COMPUTED	engine_commits divided by 43 calendar days (29 September to 10 November inclusive). 7,236 / 43 = 168.3.	2026-11-10
1,041	MEASURED	git rev-list --count origin/main, website repository. Commit count on the main branch; first commit date from git log.	2026-11-10
527	MEASURED	distinct test_*.sh entries in the SUITES list of the engine's verify script on origin/main. Entries counted by script, duplicates removed (528 listed, 527 distinct).	2026-11-10
149	MEASURED	non-comment lines of scripts/test_manifest.txt, website repository, origin/main. Line count of the manifest the test runner reads.	2026-11-10
1,054	MEASURED	portal-api/test/*.test.ts in the website repository. Files counted with ls; test() and it() declarations counted with grep (declarations, not executions).	2026-11-10
21	MEASURED	scripts/test_*mutants.sh, website repository. File count.	2026-11-10
6,314	MEASURED	site ledger (lib/ledger.json, tests_main), from the engine's recorded verify runs. Recorded gate results. Not re-run for this document.	2026-10-09
256,217	MEASURED	site ledger (lib/ledger.json, estate_sloc); method on titanos.tech/engineering. Non-blank, non-comment lines; mirrors, clones and third-party code excluded. Not re-counted for this document.	2026-10-13

Figure	Status	Source and method	Date
8.4x	MEASURED	site ledger (lib/ledger.json, cost_job_x), from the kernel compression gate result. 25 matched jobs run before and after the compressed instruction set.	2026-10-09
1,549	MEASURED	site ledger (lib/ledger.json, fleet_jobs_24h, fleet_fail_rate_24h, fleet_median_job_min), engine job ledger. Job ledger totals for the 24 hours to the stated time. Not re-measured for this document.	2026-10-13
8 cores, 23 GB	MEASURED	nproc, free -g and lscpu on the workstation. Commands run on the day.	2026-11-10
4 cores, 15 GB	MEASURED	nproc, free -g and lscpu on the cloud test runner. Commands run on the day. The runner's address is deliberately not published.	2026-11-10
about US$30 a month	OPERATOR_STATED	titanos.tech/engineering, operator figure; no receipt on file. Stated by the operator. Not receipted.	2026-10-09
about AU$300 (estimate, source pending) a month	OPERATOR_STATED	titanos.tech/engineering and /costs, operator figure; no receipt on file. Stated by the operator. Not receipted. A receipted subscription ledger is not measured yet.	2026-10-09
AU$0 to AU$149,997	PUBLISHED	lib/pricing.ts in the website repository. Read from the file the live site renders its prices from.	2026-11-10

Case study figures: taken from lib/caseStudies.ts as at the date above. Each is traceable to that build's git log or a dated usage reading, and the case page shows the detail.

Not measured yet: a receipted subscription ledger, lines of code re-counted after 13 October 2026, and any revenue figure. They are left out rather than guessed.