HOW WE BUILD

We build AI that does real work in your business. Then we prove it.

- Approved Anthropic partner

AI projects get killed for three reasons: nobody can say what they're worth, the cost runs away, or nobody trusts them. We built our method around all three. Five steps, and every one ends with something on paper.

THE FIVE STEPS

01 Understand

Where the hours go, what it costs today, and what AI must never touch in your business.

You get: The baseline, your AI inventory, the never-list.
Moves on when: You agree the baseline is right.

02 Assess

What to automate first, and what could go wrong.

You get: A risk level for each agent, and a readiness and gap review against NIST and AWS.
Moves on when: Every serious finding has a fix with a date, or a reason you accepted in writing.

03 Design

The number it has to move, signed by you. Controls sized to what it can break.

You get: The outcome charter, the control plan, a data-flow map, the system card.
Moves on when: You sign the charter. Your IT person signs the control plan.

04 Deploy

Tested on your real data, attacked on purpose, stopped on purpose. Then live.

You get: The test report, the kill-switch drill, the incident runbook, the go-live sign-off.
Moves on when: It hits the bar, the drill passed, and you, your IT person and we all sign.

05 Measure

Every month: the number, the cost, the risks. Keep it, fix it, grow it or shut it off.

You get: A monthly report and a fresh proof pack.
Moves on when: One call in writing every month: keep, fix, grow or shut it off.

Light rules for light jobs. Heavy rules for heavy ones.

LEVEL 1

Helps a person

Drafts, summarizes, looks things up. A person decides and sends.

Logs, data limits, a spend cap, the core attack tests.

LEVEL 2

Works inside

Updates your own systems: CRM notes, ticket routing, invoice matching.

Plus its own login, spot checks, and a kill switch tested for real.

LEVEL 3

Talks to customers

Answers questions, books meetings, sends anything outside the company.

Plus a person on the edge cases, attack tests on every outside channel, and a way to reach a human.

LEVEL 4

Moves money or decides about people

Pays, credits, prices, hires or denies.

Plus a person approves every action, a reason for every no, and the full proof pack before it goes live.

Most first builds land at level 2 or 3.

THE PROOF PACK

Every agent ships with its paperwork.

WHO ASKS FOR IT

Your IT person or IT provider
What it can touch, how it was tested, who can shut it off.

Your insurer
Insurers have started writing AI exclusions into policies. The pack answers their questions.

Your bigger customers
When their security questionnaire asks about AI, you answer with documents, not promises.

Your bank, investors or board
What you run, what it's worth, and who owns it.

FOR YOUR IT AND SECURITY TEAM

The part they'll ask about.

Every build is aligned to the NIST AI Risk Management Framework and its Generative AI Profile, reviewed against the AWS Well-Architected Agentic AI Lens, and tested against the OWASP Top 10 for LLM and agentic apps. It runs on Claude from Anthropic, whose AI management system is certified to ISO/IEC 42001. On Amazon Bedrock, it sits inside AWS's ISO/IEC 42001 scope.

Nobody is "NIST certified" - NIST doesn't certify anyone. We say aligned, and we hand you the map that shows it.

WHAT'S IN IT

- ✓ System card: what it does and what it won't
- ✓ Risk level and risk register
- ✓ Outcome charter with the signed target
- ✓ Control plan and data-flow map
- ✓ Test report, attack tests included
- ✓ Kill-switch drill record
- ✓ Incident runbook
- ✓ Go-live sign-off
- ✓ Latest monthly report
- ✓ Control map to NIST, AWS and OWASP

WHAT WE BUILD AGAINST

NIST AI Risk Management Framework 1.0
Every document maps to its Govern, Map, Measure and Manage subcategories.

NIST AI 600-1, the Generative AI Profile
The risks and actions specific to generative AI.

AWS Well-Architected Agentic AI Lens
Agent identity and least privilege, guardrails, human approval, tracing.

OWASP Top 10 for LLM Applications and for Agentic Applications
The attack tests in every test report.

ISO/IEC 42001
Anthropic's AI management system is certified. Amazon Bedrock is in AWS's certified scope.

WHAT THEY GET

- A control map from every document to NIST, AWS and OWASP
- An agent with its own login and only the access it needs
- A log of everything it does
- A kill switch they've watched work
- A proof pack for your insurer and your customers' security questionnaires

WHAT WE WON'T TELL YOU

- — That we're "NIST certified." NIST doesn't certify anyone.
- — That your agent is ISO 42001 certified. That takes an accredited audit of your own company. The proof pack gets you closer.
- — That following a framework is a legal safe harbor. Ask your lawyer.
- — That your insurance premium will drop. Ask your broker.

Bring us the job.

Work with us: graniteai.co/work-with-us