



Vantalect

AI-NATIVE ENGINEERING · FIXED PRICE

Engineering Leverage

Measured 31× engineering leverage with Claude Code on our own build, installed in your team in six weeks.

Prepared by Charles Sieg · charles.sieg@vantalect.com · vantalect.com

September 29, 2026

The problem

If your board has asked what the company is getting from Claude Code and the honest answer is “faster autocomplete,” the tool is not the problem. The operating system around it is.

That operating system is what this engagement installs: nested guidance files that constrain every agent session, pre-commit gates that refuse untested code, review loops, and a leverage record kept at the time of the work, not reconstructed afterward. Both engagements start from your team’s real backlog and a measured baseline, and both are fixed in scope and price.

The Assessment

Two weeks

\$25,000

Up to 40 consultant hours

One team, one real backlog. We measure how your engineers and their agents actually work today, audit the repos for what agents need and do not find, and demonstrate the difference on one of your own tickets.

GUARANTEED

The Embed

Six weeks

\$35,000

Up to 110 consultant hours

One team of up to six engineers. Week one is the baseline. Weeks two through five install the operating system inside your repos, pairing on real tickets. Week six measures it and hands it over.

The second half of the fee is invoiced only if the measured lift at week six is at least 3x your week-one baseline.

Who it is for

- A VP of Engineering, CTO, or head of platform with 30 to 3,000 engineers.
- You already pay for Claude Code, Copilot, or Cursor, and you are not seeing it in throughput.
- One team is willing to work differently for six weeks and be measured doing it.
- You have budget authority for a fixed-price engagement without a procurement cycle.

It is not the right engagement for teams that want a tool recommendation (the tool is not the bottleneck), for organizations that cannot let an outside engineer work inside their repositories, or for anyone who needs a slide deck more than a measured number.

Week by week

The work has the same shape every time, which is why it is fixed price.

The Assessment: two weeks, \$25,000.

Week	What happens
1. Baseline and audit	Your team works its normal backlog while every task is logged: the senior hours it would take and the agent minutes it takes. In parallel, the agent-readiness audit of your repositories covers constraints, tests, gates, tooling, and what an agent needs and does not find.
2. Demonstrate and plan	One live ticket from your backlog, done the AI-native way in front of the team. Then the written installation plan, which you can run yourselves or with us, and a sponsor readout: the baseline number, the audit findings, and what six weeks would change.

The Embed: six weeks, \$35,000. The Assessment's first two weeks are included.

Week	What happens
1. Baseline	The Assessment's week one: the kickoff, the logging class, and office hours. No process changes yet. This is the number everything is measured against.
2. Constraints	The constraint graph goes in: nested guidance files that tell every agent session how the codebase is built, what it must never do, and what done means. One repository at a time, each proven on a real ticket with the engineer who owns it.
3. Gates	Pre-commit gates that refuse untested code, and review loops that catch what the gates cannot. Nothing is installed that a ticket did not justify.
4. Workflows	Agent workflows for the work your team actually does: the defect fix that ships with its regression test, the feature that ships with its tests, the review that reads the diff before you do. Paired on real tickets.
5. Tooling	Tool integrations where they pay for themselves, and nowhere else. By the end of the week, the team is logging its own leverage.
6. Measure and hand over	The same measurement as week one, on the same kind of work. You get the number, the playbook, and a team that already knows how to keep it moving.

What you receive

Every deliverable is an artifact in your repositories or your inbox, not a promise in a deck.

The Assessment

1. **Fulcrum, installed** in your environment under the one-year license, with a key for every engineer.
2. **The baseline record.** Two weeks of logged tasks and a written report: the aggregate execution factor, the median task factor, coverage, and the per-project breakdown, plus the raw export.

3. **The agent-readiness audit.** A per-repository scorecard of guidance files, tests and coverage, pre-commit gates, CI, tooling, secrets handling, and documentation, with findings ranked and each one stating what an agent needs and does not find.
4. **The live-ticket demonstration.** The merged pull request, the prompt transcript, and the recording if you allow one.
5. **The installation plan.** Week by week, for your repositories and your team, with one repository's guidance file drafted as the sample. It is what you run yourselves if you stop here.
6. **The sponsor readout.** A short written document, presented in the closing meeting.

The Embed

The Embed delivers everything in the Assessment, then eight more.

7. **The constraint graph** in every in-scope repository: nested guidance files, each proven on a merged ticket with the engineer who owns the repository.
8. **Gates and review loops** in those repositories: pre-commit gates that refuse untested code, a coverage floor, and a review step that reads the diff before a human does.
9. **Agent workflows**, installed and documented: the defect fix that ships with its regression test, the feature that ships with its tests, the review before merge, and the deploy verification.
10. **Tool integrations** where they paid for themselves and nowhere else, each documented.
11. **A team logging its own leverage** by week five, with the fatigue assessment on.
12. **The week-six measurement** against week one, same method, same kind of work. It decides the second invoice.
13. **The playbook.** An operating manual you own: how to write guidance, add a gate, log a task, read the number, and onboard a new engineer.
14. **The handover workshop** and its recording.

How the result is measured

Fulcrum takes both measurements, week one and week six, with the same method: senior-engineer-equivalent hours per task, divided by agent session minutes, summed over the period, on the same kind of work. The result is evaluated on the work that was actually paired.

The guarantee

The second half of the Embed fee, \$17,500, is invoiced only if the measured lift at week six is at least 3× your week-one baseline. If it is not, you keep everything that was installed.

Our own number, and its limits

The 31x is the aggregate execution factor: senior-engineer-equivalent hours per task, divided by agent session minutes, summed over 3,552 tasks logged since February 22, 2026. The median task is 22.5x. It comes from one engineer and Claude Code building 2,481,817 lines of code across 105 repositories, with 54,644 tests and 11,763 commits, measured August 26, 2026. Every row is published and recomputable at charlessieg.com/leverage. **Measured on our own build, not promised for yours. Yours gets measured.**

Terms at a glance

Standard terms. A signed statement of work repeats them for your team.

Term	The Assessment	The Embed
Fee	\$25,000, fixed	\$35,000, fixed
Length	Two weeks	Six weeks
Consultant hours included	Up to 40	Up to 110
Live with your team (estimate)	About 8 hours	About 38 hours
Beyond the cap	\$450 per hour, by written change order agreed before the hour is spent	Same
Invoicing	As its statement of work sets out	Half at signature; half only if week six measures at least 3x week one
Payment	Due within 15 days	Due within 15 days
Scope	One team of up to six engineers and up to four repositories, on your accounts and model subscriptions	Same

Both engagements run remotely by default, on any cloud, language, or repository host. Outside scope, and available only by written change order or a separate agreement: additional teams or repositories, an on-site day, product work beyond the paired tickets, infrastructure migration, and support after the closing readout.

You own outright everything delivered: the guidance files, gates, workflows, integrations, plan, playbook, recordings, and reports. Vantalect keeps the measurement and may cite it with your name removed unless you agree otherwise in writing. Either party may end an engagement with five business days' written notice; Vantalect then invoices consultant hours spent to date at \$450 per hour, capped at the engagement fee, and you keep everything delivered.

What we need from you

- Each paired engineer available at least four hours a week.
- The sponsor available for thirty minutes a week and for the readouts.
- Repository, continuous-integration, and container-host access within three business days of signature. Week one begins on the first business day after that access exists.
- A named contact for access, scheduling, and approvals.
- The model subscriptions your engineers already use.

If an obligation is not met, the schedule slides day for day.

What comes with every engagement

- **Charles, personally.** Every engagement is delivered by Charles Sieg, inside your systems, with your engineers. There is no hand-off to a junior team.
- **A Fulcrum license.** A one-year license to install and use Fulcrum in your environment is included with either engagement. After the first year, keep it for \$1,199 per year, or do not. The record you built is yours either way, and nothing it records is sent to Vantalect or to Renkara Media Group, Inc., Fulcrum's maker.
- **Pavilion, your workspace.** A private workspace set up in week one: the calendar, every recording with its transcript and summary, the documents your sponsor accepts, the decisions, the action items, and the measurement panel. At the close, one archive of all of it, yours.
- **Every session recorded** on Zoom with your permission, with transcripts and recaps.
- **A start within two weeks** of signature, once the access it needs exists. The calendar is set before week one.

Questions clients ask

Does our code leave our environment?

No. The work happens in your repositories, your accounts, and your model subscriptions. Nothing is copied out.

What does “31x” actually mean?

For every task, we estimate the hours a senior engineer familiar with the codebase would need and record the agent session minutes it actually took. Summed over 3,552 tasks, the ratio is 31. It is an upper-tail record of one operator with full product authority and no handoffs, which is why your number gets measured rather than assumed.

What if it does not work for our team?

For the Embed, you pay half. The second half of the fee is invoiced only if week six measures at least 3x week one, and you keep everything that was installed either way.

How much of Charles’s time do we actually get?

Up to 40 hours in the Assessment and up to 110 in the Embed, about 8 and 38 of them live with your team. The caps are in the statement of work; anything beyond them is a written change order at \$450 per hour, agreed before the hour is spent.

Start with an email

Email charles.sieg@vantalect.com with the team, the stack, and what your board asked; Charles answers within one business day. The first call is thirty minutes.

The full engagement page is vantalect.com/leverage; [the full terms at vantalect.com/terms](https://vantalect.com/terms) govern.

© 2026 Vantalect, LLC. Prices and terms as of September 29, 2026; the terms at vantalect.com/terms govern.

Claude and Claude Code are trademarks of Anthropic, PBC. GitHub Copilot is a trademark of GitHub, Inc. Cursor is a trademark of Anysphere, Inc. Zoom is a trademark of Zoom Video Communications, Inc. Vantalect is not affiliated with or endorsed by any of them.