2-DAY VIBE-CODING WORKSHOP PROPOSAL

PropSimplify Sdn Bhd

Prepared 18 Nov 2026

RECOMMENDED COHORT

22 attendees

A 22-person cohort matches the open-deal headcount and covers all current engineers within the ~35-seat cap. At RM 3,500/pax the base fee is RM 77,000, inside the RM 70–80K approved budget, though the 8% SST treatment should be confirmed with the cost-sensitive CEO. The audit shows Tooling 48/Critical and Adoption 45/Critical, so one full engineering cohort is the right move to replace three squads' fragmented tools with a single shared baseline rather than a partial rollout.

DAY 1 — TOOL STACK, PROMPT FLUENCY AND SHARED-RULES BASELINE FOR JAVA/KOTLIN/VUE/POSTGRES ENGINEERS

- **Cursor / Claude Code stack audit and hands-on baseline** — Tooling is 48/Critical with no evidenced Cursor, Claude Code, Copilot or MCP seats, so this lab confirms licences and gets every engineer shipping in one standard editor.
- **Prompt fluency clinic on real PropSimplify code: Java/Kotlin services, Vue.js UI, PostgreSQL SQL** — Adoption is 45/Critical and discovery reports three squads using three tools with wildly varying output, so this drills repeatable prompting and context-setting on their actual stack.
- **Shared CLAUDE.md / cursorrules and squad-level rule set** — The explicit pain point is no shared rules across squads; audit action P0 requires a committed shared rules artefact, so attendees author and merge it in-session for Java/Kotlin/Vue/Postgres repos.
- **Secrets handling, MCP wiring and blast-radius basics** — Tooling 48/Critical includes no secrets-handling posture or MCP server evidence, so this module sets minimum safe configuration before agents touch repos.
- **Adoption telemetry and licence utilisation setup** — Adoption 45/Critical is unmeasured with no licence utilisation or daily/weekly usage split, so the cohort defines what standardized Cursor + Claude Code workflow will be measured against.

DAY 2 — PR REVIEW, AGENTIC WORKFLOWS, GUARDRAILS AND AI-ASSISTED DEFINITION-OF-DONE

- **AI-assisted PR review on their repo host and CI** — Practices is 42/Critical with zero evidence AI touches PR review, and audit risk says Day 2 applicability is unproven, so this lab uses their actual PR tool and review checklist.
- **Agentic dev workflows with blast-radius guardrails** — Culture is 66 but CEO approval for autonomous agents is untested, so this module gives the COO/CEO a written blast-radius policy and escalation path before any agent rollout.
- **Code-gen discipline and definition-of-done updates for AI-assisted code** — Outcomes is 52/At risk with no defect or quality baseline, so teams update DoD and review criteria to make AI-generated Java/Kotlin/Vue/Postgres changes reviewable rather than unfalsifiable claims.

Throughput baseline and 30-day impact measurement — Velocity is 52/At risk with no cycle time, PRs/dev/week or deploy frequency, so this module defines the pre/post metrics the 30-day review will report to the cost-sensitive CEO.

PROPOSED DATES

- Week of Dec 14, 2026 (Mon–Tue Dec 14–15)
- Week of Jan 5, 2027 (Tue–Wed Jan 5–6)
- Week of Jan 12, 2027 (Mon–Tue Jan 12–13)

VENUE

On-site at client office

A single 22-engineer hands-on cohort benefits from on-site delivery: the audit flags no evidenced Cursor/Claude Code seats and no secrets/MCP posture, so facilitators can run the 30-minute stack audit and fix licence gaps in the room. On-site also protects the shared-rules workshop and PR-review labs from remote tooling friction, while giving the cost-sensitive CEO visible retention and a focused Friday decision.

INVESTMENT

22 pax × RM 4K/pax	RM 77,000
+ 8% SST	RM 6,160
Total incl. 8% SST	**RM 83,160**

SUGGESTED FOLLOW-UPS

- 30-day impact review against the pre-cohort baseline: cycle time, PRs/dev/week, deploy frequency and AI-assisted PR share, packaged as the CEO ROI one-pager.
- Stand up 5 named internal AI champions from this cohort to own shared CLAUDE.md/cursorrules updates and onboard later hires.
- Lead-engineer 1:1s at weeks 2 and 4 to troubleshoot real AI-assisted PRs and tighten the agentic guardrail checklist.
- 90-day AI readiness re-audit covering Tooling, Adoption, Practices and Outcomes to quantify progress and decide whether other functions need a separate cohort.

NEXT STEPS

- Fri 9:00am: COO and CEO confirm the 22-seat cohort, approve dates, and settle the 8% SST treatment against the RM 70–80K budget.
- Engineering lead completes the 30-minute stack audit: Cursor/Claude Code seats, MCP, repo host, PR/CI tool, and current secrets policy.
- Ops shares the pre-cohort baseline: cycle time, PRs/dev/week, deploy frequency and escaped bug count; calendar holds are placed for the chosen date.

- Issue the PO/order form for 22 pax at RM 3,500 + 8% SST and lock one of the proposed delivery weeks.