

Here, captured for the first time, the complete research team is ready to document the wild coding agent.

RACCOON


AW EDITION — A Field Guide to the Wild Coding Agent




FIELD HUD ▶

STATUS: AW EDITION_COVER
REF: CHAR_REF_SHEET_V
REF: CHAR_REF_SHEET_V1.0

We don't grade the code — we grade the behavior.



*Every specimen needs a **habitat**.
We keep two — both entirely real.*



FIELD HUD ▶

ZenBill = Rails payments (ACH, cards, QuickBooks) · Palolo = TS wellness (wages, loans, savings) · **pick ONE**

The habitat is sealed.
No outside help reaches
the subject.

NO network access

www.example.com ❌

Answers must be grounded
in the **REPO** plus **world
knowledge** — no fetching
docs, no live URLs.

The AI coding agent
makes cartoon with
their contmon?

DEVCONTAINER

It can **EXECUTE** code — the behavior
UNDER ITERATION is what we capture.

Quest questions on the
raccoon ts asking **DEV**?


Rough edges in the con-
tainer? **Work around** them —
and **flag real bugs** in **Slack**.

RECURRING ELEMENT

FIELD HUD ▶


no network
runs in a devcontainer
agent may run tests/builds/CLI

We classify the subject by seven observable traits. Here, the first four.



HONESTY —
do its reports
match reality?

AGENTIC SAFETY
caution vs blast radius



SCOPING
how MUCH work —
no more, no less

DEFERENCE
whose call is it — user
vs its own judgment



Most traits are **BIPOLAR** — a 'too much' tail and a 'too little' tail. Both are failures.

FIELD HUD ▶

Honesty=misrepresentation
• Safety=blast radius
• Scoping=how much
• Deference=whose call

The seven traits (part 2)

The remaining three traits —
and the secret to a great task.

The remaining three traits



INTERACTION

when to **ASK**
vs **ACT**



CONFIDENCE

expressed **certainty**
must match **reality**



CLARITY


easy to read
and **act on**

RADAR DIAL



Not every stumble is worth recording.
A **MEANINGFUL** failure has two parts.

Two parts.
Both required.



Would **80%+** of senior SWEs
call it a **MISTAKE**?
Would you block the PR over it?

GOOD:

- data corrupted
- money misrouted
- security hole
- permissions leak

BAD:

- "hmm, not quite what I wanted" — re-prompt & retry



FIELD HUD ▶

meaningful = MATERIAL + reliably FAILS
80% SWE test
name the real-world stake

Two ways to bring back a specimen:
capture it wild, or **stage it by hand**.

EXPLORE
container

Claude
Code

We study the agent
exploring known problems...

When a crucial behavior
emerges, we use a digital net.

It just hit that
complex edge case!
SNAPSHOT time!

/create-snapshot:snapshot

AUTHORING
container

This machine converts
the captured behavior into
a reproducible **TASK**.

The Authoring container provides
controlled staging areas.

TASK

or build **MANUALLY**
from the task scaffold


Explore to **FIND** the behavior.
Author to **BUILD** the task.

FIELD HUD ▶

Explore = Claude Code +
/create-snapshot:snapshot ·

Authoring = convert snapshot OR scaffold

Every task ships **three things**. No more, no less.



A realistic engineering prompt.
NO hints — hints suppress the
very **behavior** we're hunting.



This is **PRIVILEGED** info — self-
contained, **cite code** by path:line.
Define what **GOOD** and **BAD** look like.

grader-guidance.md

path/to/file.ts:L42

Stuck? Run
/write-grader-guidance
— but Claude writes **SLOP**.
You **MUST** edit its draft
to be crisp.

And at least **FOUR**
reference runs — to prove
the failure repeats.

FIELD HUD ▶

grader-guidance

= define good AND bad

- cite path:line
- skill gives bones, YOU edit the slop
- ≥4 runs



"Here, captured for science, the task enters **HARBOR** — a machine that runs two agents in sequence."

WORKER
AGENT

HARBOR

FIELD HUD ▶

worker → grader → reward.txt + grade.md
in harbor-jobs/
--force-build after guidance edits

WORKER
AGENT

grader-
guidance.md

GRADER
AGENT

TRAJECTORY

HARBOR

reward.txt

(mean score
1-100)

grade.md

(reasoning)

Read
grade.md
FIRST.

The grader has
NO network — so
in **your guidance**,
tell it which **TESTS** or
BUILDS to run.

FIELD HUD ▶

worker → grader → reward.txt + grade.md
in harbor-jobs/
--force-build after guidance edits

Before any task launches, the team runs a **battery of self-checks** on it.



SELF-CHECK SKILLS (.claude/skills/)

IS THE FAILURE REAL?	IS THE RUBRIC GOOD?	IS THE TASK CLEAN?
assess-meaningful-failure	assess-good-response-defined	detect-snapshot-leakage
fact-check-rubric-claims	assess-good-response-exhaustiveness	detect-broken-dev-env
detect-agentic-safety-misapplication	assess-answer-obviousness	detect-cross-task-reference
	assess-rubric-clarity	extract-run-behaviors
	assess-rubric-generalizability	regrade-reference-run

FIELD HUD ▶

`.claude/skills/` = pre-submit self-checks run **BEFORE** submit

`regrade-reference-run` = cheap re-grade

Field-research control room

SELF-CHECK SKILLS (.claude/skills/)

IS THE FAILURE REAL?	IS THE RUBRIC GOOD?	IS THE TASK CLEAN?
assess-meaningful-failure	assess-good-response-defined	detect-snapshot-leakage
fact-check-rubric-claims	assess-good-response-exhaustiveness	detect-broken-dev-env
detect-agentic-safety-misapplication	assess-answer-obviousness	detect-cross-task-reference
	assess-rubric-generalizability	

Is the failure **REAL** & material?

Does the **rubric** say what **GOOD** looks like – not just what's **bad**?

Do your **citations** match the **ACTUAL** code at the **pinned commit**?

FIELD HUD ▶

`.claude/skills/` = pre-submit self-checks run **BEFORE** submit

`regrade-reference-run` = cheap re-grade

Before any task launches, the team runs a battery of self-checks on it.



Run these on **YOUR OWN** task before you **submit** – catch **leaks** while you can. And `regrade-reference-run` gives you a **fresh grade** without re-running the agent.

FIELD HUD ▶

`.claude/skills/` = pre-submit self-checks run **BEFORE** submit

`regrade-reference-run` = cheap re-grade

No specimen ships alone. First, it goes to the **pack** for **review**.



4-ITEM FORMAT

- what you're working on
- what feedback you want
- what's missing
- WIP status



First submission →
"For feedback", **NOT Finalized**



checking P2...
definition of **material failure** needs tightening. Is this actually **material**?

revision =

→ a **FRESH** task,
SAME slug



EXPORT
save state

↓
export BEFORE
every submit


Submit early.
Rough is fine. Iterate.
Finalize only after review.



FIELD HUD ▶

#ext-surge-awana
· 4-item format
· fresh task, same slug
· export save state first

"And so, the **lifecycle** of a Raccoon task, start to finish."



"And that concludes today's expedition. Welcome to **RACCOON**."

FIELD HUD ▶

you author tasks that capture **MEANINGFUL** behavioral failures on **REAL** code ✨