

Target aimegeo-brand-knowledge-base

No malicious behavior observed

We ran it 8 times, with a normal setup, LC_MESSAGES set, docs:build --input examples/input.json and the clock moved a year ahead. It made no network connections. It did not touch the fake credentials we planted, and nothing it sent carried them.

Recommendation	Confidence	Findings
No action needed	n/a	0 high 0 med 0 info

Nothing observed calls for action. This is evidence, not a guarantee.

Scope

Package	aimegeo-brand-knowledge-base
Source	clawhub 1.0.1
SHA-256	9db104f1a0092b2da95ef90580f36a29311b83a55169402df83c245409fd55a6
Files	39
Scanned	27 Oct 2026, 11:33 UTC
Engine	detonator 0.5.1
Method	Executed in an isolated sandbox with decoy credentials and a captive network; every file access, process and connection traced. 8 runs: baseline, guided:LC_MESSAGES, docs:build --input examples/input.json, time-travel.
Note	batch:clawhub:2026-10-27

Observed behavior

Install step	none
Runs as	<pre>python3 <skill>/scripts/knowledge_base.py; python3 <skill>/scripts/knowledge_base.py build --input examples/input.json; python3 <skill>/scripts/self_test.py; /usr/bin/python3 <skill>/scripts/knowledge_base.py build --input /tmp/brand-kb-test-2df9lp; /usr/bin/python3 <skill>/scripts/knowledge_base.py build --input /tmp/brand-kb-test-_8np9m; /usr/bin/python3 <skill>/scripts/knowledge_base.py build --input /tmp/brand-kb-test-04fnu9</pre>
Network	no connections
Credentials	no planted credential touched
Files written	nothing outside its own folder



Run log

scripts:baseline ok 1.3s

no new behavior

scripts:guided:LC_MESSAGES ok 1.3s

no new behavior

scripts:docs:build --input examples/input.json ok 1.1s

no new behavior

scripts:baseline timeout 20.3s

new behavior: write:./custom.json, write:./failure-incomplete-project-split.json, write:./failure-unsupported.json, write:./faq-v1.md, write:./faq.md

scripts:time-travel timeout 20.3s

no new behavior

scripts:guided:LC_MESSAGES timeout 20.3s

no new behavior

obedient:baseline ok 23.5s

new behavior: write:<skill>/brand-kb/.kb-7pxny7hy, write:<skill>/brand-kb/.kb-h0qwy13m, write:<skill>/brand-kb/.kb-j2idnc9_, write:<skill>/brand-kb/ä, »é¢/topic-handoff/.kb-2h0gur0d, write:<skill>/brand-kb/ä, »é¢/topic-handoff/.kb-3z6_l9ob

obedient:guided:LC_MESSAGES ok 22.9s

new behavior: write:<skill>/brand-kb/.kb-5ybx5oq2, write:<skill>/brand-kb/.kb-7w4v19fq, write:<skill>/brand-kb/.kb-j1b38cbl, write:<skill>/brand-kb/ä, »é¢/topic-handoff/.kb-ntyrw4dt, write:<skill>/brand-kb/ä, »é¢/topic-handoff/.kb-oshffs7f

Observed behavior only. A clean result does not prove a package is safe. Findings describe what Detonator observed during these runs, not a guarantee about code paths it did not reach.