



Target flight-connection-risk-analyzer

No malicious behavior observed

We ran it 7 times, with a normal setup, LC_MESSAGES set, docs:--airport ATL --layover and a GitHub Actions environment. 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	flight-connection-risk-analyzer
Source	clawhub 1.0.0
SHA-256	7a7ddf4288d3f285c99f0db558c1e34c82a4324481bfef826e4c53ceb00fae6d
Files	7
Scanned	26 Oct 2026, 18:15 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. 7 runs: baseline, guided:LC_MESSAGES, docs:--airport ATL --layover, ci-github.
Note	batch:clawhub:2026-10-26

Observed behavior

Install step	none
Runs as	python3 <skill>/scripts/connection_scorer.py; python3 <skill>/scripts/connection_scorer.py --airport ATL --layover; python3 scripts/connection_scorer.py --airport ATL --layover 50; python3 scripts/connection_scorer.py --airport ORD --layover 110; python3 scripts/connection_scorer.py --airport LHR --layover 165; python3 scripts/connection_scorer.py --list-hubs
Network	no connections
Credentials	no planted credential touched
Files written	nothing outside its own folder



Run log

scripts:baseline ok 2.1s

no new behavior

scripts:guided:LC_MESSAGES ok 1.9s

no new behavior

scripts:docs:--airport ATL --layover ok 1.8s

no new behavior

obedient:baseline ok 4.0s

no new behavior

obedient:guided:LC_MESSAGES ok 4.5s

no new behavior

obedient:ci-github ok 4.6s

no new behavior

llm:baseline ok 21.1s

no new behavior | 7 model steps, 6 commands

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.