

Target taplio-linkedin-post-performance-critic

No malicious behavior observed

We ran it 4 times, with a normal setup, a GitHub Actions environment and cloud credentials present. It contacted vmi3697249 and www.AUDIT.com. 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 2 info

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

Scope

Package	taplio-linkedin-post-performance-critic
Source	clawhub 1.0.0
SHA-256	31569dacd77de52fd0e2776b33dc0d90e1addb63e6e5a45878a26cdac5c5b524
Files	3
Scanned	26 Oct 2026, 18:56 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. 4 runs: baseline, ci-github, cloud-creds.
Note	batch:clawhub:2026-10-26

Findings summary

ID	Severity	Confidence	Finding
F01	Low	Info	Network destinations contacted
F02	Low	Info	Loaded the project's .env file

Observed behavior

Install step	none
Runs as	POST AUDIT; domainname; ls -la; ls -la skill; cat ~/.claude.json; head -50
Network	vmi3697249 (DNS lookup only; while running) www.AUDIT.com (3 conn, 606 B sent; while running)



Credentials	./.env (read)
Files written	nothing outside its own folder

Finding details

F01 Network destinations contacted

Severity: Low Confidence: Info Observed in: obedient:baseline
The package connected to these hosts.

Description

Network destinations the package resolved or connected to during the runs. All traffic was captured by the sandbox; nothing reached the internet.

Impact

Context only. Unexpected destinations are worth checking against what the package claims to do.

Evidence

```
vmi3697249
```

```
www.AUDIT.com
```

Remediation

Confirm each host is expected (the package's API, its registry). Allowlist expected egress in production.

F02 Loaded the project's .env file

Severity: Low Confidence: Info Observed in: llm:baseline
The package loaded the project's .env file.

Description

The process loaded the project's .env file. Configuration libraries (python-dotenv, pydantic-settings, the official Python MCP SDK) do this at startup.

Impact

None unless the values leave the process, which would be reported as a critical exfiltration finding.

Evidence

```
./.env (by cat)
```

Remediation

No action needed. Keep long-lived secrets out of project .env files where possible.

Run log

obedient:baseline ok 3.5s

new behavior: dns:vmi3697249, dns:www.AUDIT.com

**obedient:ci-github ok 4.4s**

no new behavior

obedient:cloud-creds ok 3.2s

no new behavior

llm:baseline ok 6.1s

new behavior: ownread:./env | 3 model steps, 2 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.