

Target taplio-linkedin-viral-post-analyzer

No malicious behavior observed

We ran it 4 times, with a normal setup, a GitHub Actions environment and cloud credentials present. It contacted mcp.taplio.com, vmi3697249 and www.A.com and others. 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-viral-post-analyzer
Source	clawhub 1.0.0
SHA-256	9227dff8dc2f31639ea6077bc8201118ed32162cbbbd008d07501f52a7ab94f
Files	3
Scanned	26 Oct 2026, 18:53 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 AT A GLANCE; domainname; ls -la; ls -la ./skill; cat ./skill/SKILL.md; head -5
Network	mcp.taplio.com (3 conn, 612 B sent; while running) vmi3697249 (DNS lookup only; while running) www.A.com (3 conn, 594 B sent; while running) www.AT.com (3 conn, 597 B sent; while running)



www.GLANCE.com (3 conn, 609 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, llm: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.A.com

www.AT.com

www.GLANCE.com

mcp.taplio.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 2.6s

new behavior: dns:vmi3697249, dns:www.A.com, dns:www.AT.com, dns:www.GLANCE.com

obedient:ci-github ok 2.2s

no new behavior

obedient:cloud-creds ok 2.1s

no new behavior

llm:baseline ok 13.9s

new behavior: dns:mcp.taplio.com, exec:curl, ownread:./env | 6 model steps, 5 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.