

Target magneta-contract-risk

No malicious behavior observed

We ran it 1 time, with a normal setup. It contacted vmi3697249 and x402.magneta.finance. 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	magneta-contract-risk
Source	clawhub 0.1.0
SHA-256	a369c90a1877c0821179157f56d3515a47d03a53a61bceddc78d2cf5d00327ff
Files	3
Scanned	26 Oct 2026, 19:30 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. 1 runs: baseline.
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	ls -la ./skill; ls -la; env; sed s/=.*/=<set>; grep -iE BUYER_PRIVATE_KEY CDP_; sed s/=.*/=<redacted>/ .env
Network	vmi3697249 (DNS lookup only; while running) x402.magneta.finance (8 conn, 976 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: 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
```

```
x402.magneta.finance
```

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 sed)
```

Remediation

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

Run log

llm:baseline ok 29.4s

new behavior: dns:vmi3697249, dns:x402.magneta.finance, exec:curl, ownread:./.env | 8 model steps, 7 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.