

Target lfg-alarm

No malicious behavior observed

We ran it 1 time, with a normal setup. It contacted lfg.gm and mcp.lfg.gm. 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	lfg-alarm
Source	clawhub 1.0.1
SHA-256	00d8194baa2830d743f51036ecc1130d167847575d185438b8997f2ddcc3407e
Files	3
Scanned	26 Oct 2026, 19:35 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; head -5; cat skill/SKILL.md; ls skill; curl -s -m 15 -o /dev/null; curl -s -m 15 -X POST
Network	lfg.gm (2 conn, 144 B sent; while running) mcp.lfg.gm (10 conn, 1463 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

```
lfg.gm
```

```
mcp.lfg.gm
```

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

llm:baseline ok 16.9s

new behavior: dns:lfg.gm, dns:mcp.lfg.gm, exec:curl, ownread:./.env | 5 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.