Target gmail-via-agiexa

No malicious behavior observed

We ran it 4 times, with a normal setup, a GitHub Actions environment and cloud credentials present. It contacted agiexa.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 1 info**

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

Scope

Package	gmail-via-agiexa
Source	clawhub 1.0.1
SHA-256	779770e366ae1c4c926e0177c061b7fdebae6728a3903cc5556693510d8b6d89
Files	6
Scanned	26 Oct 2026, 17:59 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

Observed behavior

Install step	none
Runs as	curl -sS -X POST https://agiexa.com/api/agent/waitlist -w; ls -R; cat references/tools.md; cat references/triggers.md; head -40 references/tools.md; grep -c ^| `GMAIL_ references/tools.md
Network	agiexa.com (3 conn, 1044 B sent; while running)
Credentials	no planted credential touched

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

```
agiexa.com
```

Remediation

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

Run log

obedient:baseline ok 1.9s
new behavior: dns:agiexa.com, exec:curl

obedient:ci-github ok 0.6s
no new behavior

obedient:cloud-creds ok 1.7s
no new behavior

llm:baseline ok 11.3s
no new behavior | 4 model steps, 3 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.