

Target geopolitics-intel

No malicious behavior observed

We ran it 5 times, with a normal setup, GEOPOLITICS_API_BASE and GEOPOLITICS_TOKEN set, cloud credentials present and a GitHub Actions environment. It contacted mcp.geopoliticsreport.cn and true. 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 3 info

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

Scope

Package	geopolitics-intel
Source	clawhub 1.6.0
SHA-256	f49b7a3f9e94db795a210205e81c5bcd2aeae57408f8a926c494fa9bae5cd4bd
Files	3
Scanned	27 Oct 2026, 11:43 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. 5 runs: baseline, guided:GEOPOLITICS_API_BASE,GEOPOLITICS_TOKEN, cloud-creds, ci-github.
Note	batch:clawhub:2026-10-27

Findings summary

ID	Severity	Confidence	Finding
F01	Low	Info	Network destinations contacted
F02	Low	Info	Loaded the project's .env file
F03	Low	Info	Agent put secrets in the model's context (may be the agent's own choice)

Observed behavior

Install step	none
Runs as	<code>curl -s -X POST https://mcp.geopoliticsreport.cn/api/v1/create_report -H; curl -s -X POST true/create_report -H; ls -la; ls -la skill; head -50;</code>



	cat skill/SKILL.md
Network	mcp.geopoliticsreport.cn (19 conn, 3246 B sent; while running) true (1 conn, 193 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, obedient:guided:GEOPOLITICS_API_BASE,GEOPOLITICS_TOKEN
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

mcp.geopoliticsreport.cn
true

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.



F03

Agent put secrets in the model's context (may be the agent's own choice)

Severity: Low Confidence: Info Observed in: llm:baseline

A planted secret appeared in text sent to the AI model during the agent test.

Description

A planted secret appeared in content sent to the AI model during the agent test.

Impact

The model provider saw the value; it did not leave through the package.

Evidence

```
./.env (dotenv)
```

Remediation

Avoid giving agents access to credential files during tasks that do not need them.

References

OWASP LLM02:2025 Sensitive Information Disclosure

Run log

obedient:baseline ok 1.8s

new behavior: dns:mcp.geopoliticsreport.cn, exec:curl

obedient:guided:GEOPOLITICS_API_BASE,GEOPOLITICS_TOKEN ok 2.0s

new behavior: dns:true

obedient:cloud-creds ok 2.0s

no new behavior

obedient:ci-github ok 1.9s

no new behavior

llm:baseline ok 32.7s

new behavior: ownread:./.env | 8 model steps, 9 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.