

Target pto-maximizer

No malicious behavior observed

We ran it 7 times, with a normal setup, LC_MESSAGES set, docs:--country US --year and a GitHub Actions environment. It made no network connections. 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	pto-maximizer
Source	clawhub 1.0.0
SHA-256	59c4337b303a32b2337a473ffbbb9b1e96f40ea63ee6855a440b7cd2d06e1366
Files	7
Scanned	26 Oct 2026, 18:12 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. 7 runs: baseline, guided:LC_MESSAGES, docs:--country US --year, ci-github.
Note	batch:clawhub:2026-10-26

Findings summary

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

Observed behavior

Install step	none
Runs as	python3 <skill>/scripts/pto_planner.py; python3 <skill>/scripts/pto_planner.py --country US --year; python3 scripts/pto_planner.py --country US --year 2026; python3 scripts/pto_planner.py --country UK --year 2026; ls -la; ls -la skill
Network	no connections

Credentials	./.env (read)
Files written	./pto_2026.ics (written)

Finding details

F01 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.

F02 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

scripts:baseline ok 1.6s

no new behavior

scripts:guided:LC_MESSAGES ok 1.0s

no new behavior



scripts:docs:--country US --year ok 1.9s

no new behavior

obedient:baseline ok 9.4s

new behavior: write:./pto_2026.ics

obedient:guided:LC_MESSAGES ok 7.3s

no new behavior

obedient:ci-github ok 5.5s

no new behavior

llm:baseline ok 18.8s

new behavior: ownread:./env | 8 model steps, 11 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.