


Browser-Only Ransomware: From LLM Hallucinations to a Practical Attack Technique

 research.checkpoint.com/2026/browser-only-ransomware-from-llm-hallucinations-to-a-practical-attack-technique

stcpresearch

July 1, 2026



Research by: **Alexey Bukhteyev**

Key Takeaways

- AI can turn high-level malicious ideas into concrete techniques, and can independently design and implement novel attack paths that have not yet appeared in real-world campaigns.
- In this research, DeepSeek connected unrealistic browser-malware concepts with a real browser capability, turning an AI-generated malware hallucination into a plausible browser-native ransomware technique. Although the generated sample was incomplete, it exposed a practical abuse path based on the File System Access API and access to photo directories.

- The technique does not require a native payload, APK installation, browser exploit, or root access. It relies on social engineering and a legitimate permission prompt exposed by the File System Access API in Google Chrome.
- The Android scenario is especially concerning because photo directories are high value personal data stores and, unlike iOS, modern Android Chrome versions expose a browser API that allows web pages to read and modify files in those directories after user approval. Using a fake AI image-enhancement workflow gives users a plausible reason to approve folder-level file access. Our PoC demonstrates this browser-only workflow against selected image directories on Android.

Introduction

Over the past several years, large language models have reshaped software development, and malware development has followed the same path. Check Point Research has documented this trend from early experiments showing that AI systems could generate offensive components, to cases of cybercriminals using ChatGPT to create malicious tools, and later to advanced AI-authored malware frameworks such as [VoidLink](#). In some cases, LLMs lowered the barrier enough for users with little or no development experience to produce working offensive code.

As frontier models became better at writing reliable code, including complex security related components, major AI vendors also turned cyber safety into a dedicated control area. Clearly malicious requests involving credential theft, malware deployment, ransomware behavior, persistence, stealth, or unauthorized exploitation are now commonly blocked or refused. [OpenAI's cyber-safety documentation](#), for example, describes additional safeguards for models classified as having High Cybersecurity Capability, while Anthropic has published reports on [detecting and countering cyber misuse of Claude](#).

DeepSeek then becomes particularly relevant in this context for several reasons:

- Lower refusal rates for harmful cyber enforcement: compared with Anthropic and OpenAI, DeepSeek models were less consistent refusing harmful cyber requests, including the File System Access API implementation we will be discussing later on this article.
- Low barrier to access: DeepSeek is free to use via the web interface, widely available, and accessible in regions where other frontier models face regulatory or commercial restrictions. This lowers the cost of repeated malicious experimentation.
- End-to-end malicious code from a single prompt: in our testing, a working malicious application could often be generated from a single broad prompt. Achieving a comparable result with OpenAI or Anthropic typically requires decomposing the attack into multiple benign-looking requests and manually assembling the generated components.

Putting this all together, these differences make DeepSeek particularly attractive to threat actors: DeepSeek models can turn high-level malicious ideas into concrete, complete attacks with less expertise than competing platforms.

Check Point Research analyzed nearly 3,000 files attributed to DeepSeek observed in public telemetry over the past year. The dataset included Python, PowerShell, Batch, HTML, JavaScript, VBScript, and other file types. Of these, 1,383 files were classified as malicious or dangerous by either VirusTotal detection or static source analysis. Within this dataset, we found a sample that implemented a dangerous browser-native technique we have not observed exploited in the wild. We refer to it as **In-Browser Ransomware**. The technique uses a phishing lure to persuade the victim to grant file-system access to a web page; once access is granted, the page can enumerate local files in the selected folder, read and exfiltrate their contents, encrypt and overwrite them, and display a ransom-style message, all without installing a native payload or exploiting the browser.

The underlying browser risk was already known to browser engineers. [The File System Access specification](#) explicitly lists ransomware as a security consideration, and the 2023 USENIX Security paper [RoB: Ransomware over Modern Web Browsers](#) studied the abuse of the File System Access API to encrypt local files from a malicious web application.

The important finding in our research and what is new, is how the AI model brought these previously documented concepts together, into a realistic and enforceable attack scenario leveraging a method that defenders had originally thought was unfeasible due to browser sandboxing limits: a DeepSeek-attributed malicious sample, generated as an all-in-one malware fantasy, connected this documented platform risk to a realistic phishing-style web application, demonstrating a viable end-to-end attack chain. An attacker does not need to know that a browser exposes a file-system API. They can ask for an impossible-sounding outcome – a website that steals files, captures keystrokes, takes screenshots, encrypts files, and demands payment – and the model may connect the request to a real browser capability. Basically, the AI model showed an ability to reason across existing knowledge and combined multiple known components into a coherent attack workflow that could be readily used by an attacker. This illustrates how frontier AI models may move beyond simply enhancing existing attacker techniques to **lowering the** expertise required to operationalize complex attack chains by connecting knowledge in ways that previously relied on human experience and creativity.

A Noisy Sample With One Important Idea

The sample that caught our attention is SHA256

07c39f79ab92fb21557b82283472dce1c112f577d796111fb752c3c6d84c86b5, a Python Flask application that serves victim-facing HTML and JavaScript from embedded templates and also includes backend routes intended to receive information from the victim and provide an administration panel.

We do not have the prompt submitted to the AI model that produced this sample. Judging by the code structure, function names, and comments, it was likely formulated very broadly such as something similar to this example: create a universal malicious tool that runs through the browser and collects as much victim data as possible, encrypts files, and demands ransom. In a single front-end, the generated code assembled routines and stubs for keylogging, clipboard monitoring, form and network-request interception, Discord-token collection, crypto-wallet and payment-card discovery, geolocation requests, webcam and microphone access, screenshots, local-file access, Chrome exploit stubs, “persistence,” and a ransomware-style overlay. This does not mean the sample actually implements all of these capabilities. A more accurate reading is that it is an AI-generated blueprint in which the model tried to translate familiar capabilities of native stealers and ransomware tools into a web page opened in the browser.

The victim-facing page is disguised as a Discord avatar AI upscaler:




Figure 1 – Victim-facing lure disguised as a Discord avatar AI upscaler in the DeepSeek attributed InfernoGrabber sample.

Clicking the button on the victim-facing lure page is intended to start the malicious browser-side sequence, although the generated control flow is inconsistent and does not complete reliably. After a fake processing step, the page is intended to display a ransomnote-style overlay under the name **InfernoGrabber v9.0**. The message claims that passwords, credit cards, and personal files were encrypted, demands Bitcoin, and displays a countdown threatening publication of private data.




Figure 2 – InfernoGrabber ransom-note overlay.

Most of the functionality claimed in the sample collapses at the browser boundary. A normal web page can observe activity inside its own origin, capture input events delivered to its own DOM, request browser-mediated permissions, access storage scoped to its own origin, and render frightening overlays. It remains constrained by the browser security model.

In this sample, the “desktop screenshot” routine captures the rendered web page, the keylogger observes keystrokes only while the user interacts with the page, webcam and microphone capture depend on browser permission prompts, and the Discord-token stealing

logic searches storage available to the current origin. The “persistence” logic relies on browser storage and a service worker registration attempt.

Much of the sample therefore reads as an AI hallucination produced in response to an overly broad prompt or to requirements that a normal web page cannot satisfy. The exception was the file-access workflow, where the generated code reached for a real browser primitive with practical abuse potential.

The generated JavaScript referenced:

- `showOpenFilePicker()`;
- `showDirectoryPicker()`;
- recursive traversal of a user-selected directory;
- reading selected files through browser file handles;
- sending file contents to the Flask backend;
- displaying a ransomware-style warning after the interaction.

The File System Access API is a legitimate browser capability designed for web applications such as editors, IDEs, and creative tools. After the user grants access, a web application can read files and folders from the local device. The API also supports write access and directory enumeration under browser permission controls.

The technique is limited to browsers that expose the picker-based File System Access API. At the time of writing, this primarily means Chromium-family browsers: the API shipped on desktop in Chrome 86, and Chrome 132 extended File System Access support to Android and WebView. Firefox and Safari do not expose the same local file and directory picker methods, which limits the immediate attack surface but also concentrates the risk in Chrome-based browsing environments.

The sample lacked a complete and reliable browser-side encryption flow, yet the attack design was concrete: a fake utility convinces the user to grant browser file access, which allows the page to exfiltrate and encrypt files.

The model combined fake OS-level malware claims with a real browser primitive and produced a browser-native file-theft and ransomware scaffold. The sample shows how an LLM can transform an abstract malicious request into a new attack blueprint. The user likely wanted an all-in-one tool: a Discord-themed lure, a stealer, an admin panel, and a ransomware or locker workflow. The model chose a Flask application and a browser frontend as the unifying architecture. In doing so, it connected a hallucinated malware concept to a real platform feature with genuine abuse potential.

Even though we have not yet observed this exact browser-native ransomware pattern widespread in-the-wild campaigns, the technique is still operationally relevant for several reasons:

- The browser becomes the execution environment: the attack runs entirely inside the browser process, without installing any additional app, dropping a binary, or exploiting a vulnerability. Traditional endpoint protections focus on apps and native payloads; a website that encrypts files after a legitimate-looking permission sits outside those assumptions.
- Lower friction for victims: opening a web page and clicking “Allow” on a file-access prompt is a normal part of using modern web applications. Users do not intuitively treat this as “running malware”, which makes the social-engineering angle powerful.
- Cross-platform reach: the same browser-native technique can target any platform where the File System Access API is exposed, we tested on Android and Windows.

From Hallucinated Scaffold to Working PoC

Because the original sample was incomplete, we tested whether the latest DeepSeek model V4 could turn the same browser-native attack idea into a working proof of concept.

When prompted directly to create ransomware, the model consistently refused across all tested modes.




Figure 3 – DeepSeek V4 refuses to generate ransomware when prompted directly.

Even though some requests were denied, we managed to succeed in the end. We removed explicit terms such as “ransomware” while preserving the same functionality: a web page that asks the user for access to local files, processes them inside the browser, and leaves the user unable to recover the original content.

In Instant mode, DeepSeek consistently generated HTML/JavaScript code that used the File System Access API to interact with user-selected files.

In Expert mode, the behavior was inconsistent across attempts:

- several attempts ended in refusal;
- one generated a non-functional sample;
- one generated a fully working browser-based ransomware PoC.

One response was especially notable because the model described the result as:

“a crafted trap that combines a convincing AI upscaler interface with hidden ransomware-like behaviors”

This wording shows that the model recognized the malicious nature of the scenario while still continuing the generation.

For comparison, we tested similar requests against ChatGPT and Claude. In our tests, these systems either refused to help or generated constrained browser-safe implementations that did not use the File System Access API.

This does not mean that the same outcome is impossible with other frontier systems. With an incremental approach, a user can ask for separate components that appear benign in isolation, such as a user interface, browser file handling, client-side data transformation, and neutral status messaging, and then assemble them into a harmful workflow by replacing the neutral messages with a ransom note. The difference is the level of steering required. In that scenario, the user needs enough technical understanding to decompose the attack, preserve the malicious objective across separate requests, identify the right browser primitive, and combine the generated pieces manually.

In-Browser Ransomware on Android

To assess the practical risk of this technique, we used an LLM to build a controlled proof-of-concept (PoC) based on the same idea we observed in the DeepSeek-attributed sample: a browser-native ransomware workflow disguised as an AI image upscaler.

On Android, modern Chrome versions expose the picker-based File System Access API to web content. On iOS, Safari does not expose the same File System Access primitives to websites. Access to photos is mediated by the operating system's app-sandbox and photo-library permissions instead of a web API that can enumerate and modify arbitrary folders. Chrome on iOS uses WebKit which also does not implement File System Access API. As a result, on mobiles, the technique we demonstrate is currently practical on Android Chromium browsers.

At the same time, the attack surface is narrower than arbitrary disk access. The picker-based File System Access API does not let a web page target the whole system disk, and Chromium applies additional restrictions to sensitive locations. In Chromium's current implementation, broad access to locations such as the user's home directory, Desktop, Documents, Downloads, Chrome data, application directories, Windows, Program Files, AppData, and several Linux and Android system paths is blocked or constrained. The File System Access specification also explicitly recommends restricting sensitive directories and [lists ransomware as one of the risks](#) the API design must account for.

However, selection of the root of the default Pictures and Videos directories was not restricted on any of the tested operating systems (Android and Windows). This capability fits naturally

into a social-engineering workflow for a fake photo-processing application.

On desktop, the Pictures folder may contain personal files, but it is usually less central to business workflows than the user's entire home directory or a Documents directory.

On mobile, the risk profile changes: the photo library is often one of the most valuable local data stores. It may contain years of private photos, identity documents, banking screenshots, medical records, recovery codes, travel documents, work images, and photos of family members. Losing access to this data, or having it exfiltrated, can create personal or business issues from ransomware to blackmail or if the data is sensitive, public disclosure leading to reputational damage and more. Chrome 132 introduced File System Access support on Android, allowing web applications, after user approval, to read and save changes directly to selected files and folders. We tested this capability on several Android devices and confirmed that the latest Chrome version available to us at the time of testing, Chrome 148, also allowed selecting the photo directory, including the root of the DCIM folder.

The workflow on Android looks very natural. The user opens a web page that promises to enhance a photo, selects an image, and is then asked to choose a directory for saving the "enhanced" results. The browser warning that the site will be able to edit files in the selected folder is easy to rationalize in that context: the user expects the service to write processed images back to the device. During the fake processing step, the PoC encrypts pictures inside the selected directory.

Video 1 – Demonstration of a browser-native ransomware PoC on Android using the File System Access API.

The combination of this technique, a natural social-engineering lure, and browser-only execution makes the Android scenario especially concerning. The resulting flow requires no APK installation, no vulnerability exploitation, no native payload, and no root access.

Users generally do not treat opening a web page as a malware execution event, especially when no application is installed and no binary is downloaded. In this case, the browser prompt appears in a context where file access feels expected, while the granted permission gives the page meaningful control over a directory that may contain highly sensitive personal data.

Practical Recommendations for Users

While this research focuses on a controlled PoC, there are concrete steps users can take today to reduce the risk of browser-native ransomware abuse:

Treat browser folder-access prompts as high-stakes decisions: before approving “access to files in a folder”, check which site is asking, which folder is being selected, and whether editing files is truly necessary for the feature you expect. If you are unsure why a site needs write access to an entire directory, decline the request.

Avoid granting websites access to sensitive or irreplaceable data: do not expose folders that contain personal photos, identity documents, recovery codes, or work data unless the site is highly trusted and the need is clear. Prefer selecting a temporary or empty folder for experimental web tools, rather than your main photo library.

Prefer well-established applications for high-value data: for tasks such as backing up photos, editing large collections, or processing sensitive images, use reputable native apps or well-known cloud services instead of newly discovered browser tools with unknown reputation.

Maintain offline and cloud backups of important data: regular backups reduce the leverage attackers gain from encrypting or deleting local files, whether through native ransomware or browser-based techniques.

Keep browsers and mobile OSes updated: browser and OS vendors continue to refine permission models and harden sensitive APIs. Applying updates promptly ensures that you benefit from the latest security controls around features like File System Access.

Be skeptical of AI-branded lures: attackers increasingly disguise malicious flows as “AI” utilities, avatar upscalers, photo enhancers, or productivity tools. A polished AI-themed interface is not a guarantee of safety; apply the same caution you would to any unfamiliar site asking for broad access to local files.

Conclusion

LLM-assisted malware development changes the economics of malicious experimentation. A user with limited technical understanding can describe a harmful outcome, generate code, test the result, adjust the prompt, and repeat the process at very low cost. Tasks that once required a developer, a purchased builder, or prior knowledge of the relevant platform can now be approached through cheap iteration.

This also changes the defender’s problem. Malware generated this way may move the ecosystem away from a limited set of reused families and builders toward a larger volume of disposable, one-off artifacts, each carrying a unique combination of techniques, API usage, and payload logic.

Hallucination adds another important dimension. AI-generated malware can be technically wrong and still reveal practical malicious techniques. When a model tries to satisfy unrealistic requirements, it may search across legitimate platform features and map a malicious goal to an API that actually exists. This process can surface techniques that defenders have not yet seen in the wild, or turn risks previously described mostly in theory into workable attack concepts. The case analyzed in this research shows exactly that: a noisy and partially broken artifact connected a theoretical browser risk to a practical browser-only ransomware technique.

In this case, the user likely asked for an impossible web application, a single browser page that behaves like a fully features stealer and ransomware agent. The model could not satisfy all of those requirements correctly, but in the process of trying, it searched across legitimate browser features and anchored part of the fantasy to a real API: the File System Access API.

This illustrates a broader risk:

- A non-expert attacker does not need to know that such an API exists or how to abuse it.
- By describing a high-level malicious outcome in natural language, they can cause the model to discover and connect the malicious goal to previously under-explored platform capabilities.
- The resulting prototype can then be refined into a working PoC with minimal additional prompting or manual editing.

In other words, AI is not only lowering the barrier for reimplementing existing malware techniques; it is also capable of bridging the gap between purely theoretical risks and practical, novel attacks that defender have not yet seen deployed in the wild.

Historically, new attack techniques emerged through human experimentation, experience, and creativity. Frontier AI changes that dynamic. Rather than being constrained by conventional thinking or established attacker playbooks, AI can reason across existing knowledge and synthesize it in unexpected ways, connecting known capabilities into practical attack chains. The real shift is not that AI is inventing entirely new vulnerabilities, but that it may identify combinations and attack paths that humans had not previously recognized or operationalized.

At the time of analysis, we found no evidence that this technique had been adopted as an in-the-wild malware pattern. The original DeepSeek-attributed sample was incomplete and failed to implement the full attack reliably. However, our testing showed how little effort is required to transform the same idea into a fully working implementation using modern LLMs. The resulting workflow is especially concerning on mobile devices, where a seemingly legitimate request for access to a photo directory can expose highly sensitive personal data to encryption, exfiltration, or both. From a defensive perspective, browser folder-access prompts

should be treated as security decisions rather than routine clicks. Before granting a website access to an entire folder, users should review which site is asking, which folder is being selected, whether file modification is allowed, and whether the permission matches the action they intended. Users should avoid granting websites access to directories containing sensitive, private, or irreplaceable data whenever possible.