


Markdown Is the New Executable: Instruction-File Poisoning of AI Coding Agents

Security Implications and Guidance for AGENTS.md,
CLAUDE.md, Rules Files, and Agent Settings

2026-10-27

 AI-assisted Rapid Research



CSAI

cloud
security
alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

Instruction files such as `AGENTS.md`, `CLAUDE.md`, `.cursor/rules/*.mdc`, and `.github/copilot-instructions.md` are plain Markdown. AI coding agents load them automatically and treat them as trusted guidance from the developer. In practice, a file the agent loads at session start can function like executable configuration: it can direct the agent to install packages, change its own settings, or send credentials to an endpoint the attacker controls. The newest variants exploit no software bug, because the agent is doing what it was designed to do. Earlier cases relied on load-time flaws that vendors have since patched, and attackers have moved on to persuading the agent itself.

The most recent example surfaced on October 25–26, 2026, when Forever Security reported finding a malicious `AGENTS.md` in a customer's `~/Downloads` folder. The file tried to get OpenAI Codex to install the malicious npm package `function-color`. In the researchers' reconstruction, Codex ran a benign stand-in payload in more than 90% of attempts, and a companion `CLAUDE.md` got Claude Code to redirect its API traffic about half the time [1][2]. A related technique has already appeared in a supply-chain attack. In June 2026, the Miasma worm planted agent settings and rules files for Claude Code, Gemini CLI, Cursor, and VS Code in `Azure/durabletask`, which ran a credential stealer when the repository was opened in an AI tool; GitHub responded by disabling 73 repositories across four Microsoft organizations [3].

Vendors have patched the load-time flaws disclosed to and accepted by them so far, where repository settings ran before the user consented, although at least one reported workspace-load path in VS Code Copilot remains unfixed [15]. Those fixes appear necessary but not sufficient: Forever Security showed that an agent inside a trusted session can still be persuaded to write the dangerous configuration itself [1][2][4][5]. Written rules offer little backstop. One preprint study of 481 public `CLAUDE.md` files found that only 4–16% of the security rules developers wrote had a matching built-in control that could enforce them [7].

Organizations should therefore treat instruction and agent-settings files as code. That means code-owner review, scanning for hidden Unicode, not running agents from untrusted directories, and moving safety rules out of Markdown into enforced permissions, sandboxes, and managed policy.

Background

AI coding agents read persistent instruction files to learn a project's conventions, build commands, and constraints. `AGENTS.md` started as a cross-vendor convention in August 2025. OpenAI donated it to the Linux Foundation's Agentic AI Foundation in December 2025 and reported that more than 60,000 open-source projects and agent frameworks had adopted it, including Codex, Cursor, Devin, Gemini CLI, GitHub Copilot, Jules, and VS Code [8]. Claude Code reads `CLAUDE.md` and can also read `AGENTS.md`. Cursor uses `.cursor/rules/`, and GitHub Copilot reads `.github/copilot-instructions.md`.

How these files are discovered matters for security. Codex walks from the project root down to the working directory and concatenates every `AGENTS.md` (or `AGENTS.override.md`) it finds, up to a default limit of 32 KiB [9]. When there is no project root, it checks only the current directory, so an agent launched in `~/Downloads` reads whatever `AGENTS.md` is sitting there. Claude Code loads `CLAUDE.md` files from every directory above the working directory at launch, and loads files in subdirectories on demand when the agent touches a file there [10]. Anthropic's documentation states that these files are treated "as context, not enforced configuration," and that blocking an action reliably requires a hook or permission rule [10]. The result is a channel that accepts natural language from anyone who can place a file on disk or commit to a repository, and that channel has strong influence over an agent able to run shell commands.

Public disclosures specific to agent instruction files date to at least March 2025, when Pillar Security disclosed the "Rules File Backdoor." That technique hid instructions in Cursor and Copilot rules files using zero-width joiners and bidirectional text markers, which were invisible to reviewers, including in GitHub pull request views [11]. According to Pillar Security, Cursor responded that the risk was the user's responsibility, while GitHub added a hidden-Unicode warning on github.com in May 2025 [11]. In December 2025, the IDEsaster research by Ari Marzouk reported more than 30 vulnerabilities in AI IDEs and coding assistants, 24 of which received CVE identifiers. The research identified what the researcher described as "universal" attack chains that combined prompt injection with legitimate, auto-approved IDE features [12]. CSA's June 2026 analysis of coding agents as an attack surface brought these threads together and identified repository configuration poisoning as one of three converging threat vectors [13].

Security Analysis

Two Classes of Poisoned Files

This note separates two kinds of file because they need different defenses. *Instruction files* are Markdown that the model reads as guidance, such as `AGENTS.md`, `CLAUDE.md`, Cursor `.mdc` rules, Copilot instructions, and `SKILL.md` files. *Agent settings files* are structured configuration that the tool itself acts on without asking the model, such as `.claude/settings.json`, `.mcp.json`, `.gemini/settings.json`, `.vscode/settings.json`, and `.vscode/tasks.json`. Settings files can run code directly through hooks, MCP server definitions, folder-open tasks, or environment variables. Instruction files can only persuade, but they can persuade the agent to edit a settings file. Several of the cases below combine the two: an instruction file persuades the agent, and a settings file provides the execution or persistence.

The table below summarizes publicly reported incidents and disclosures that show this progression, from settings that executed at load time toward instructions that persuaded the agent to reconfigure itself.

Case	Date	File(s) abused	Mechanism	Status
Rules File Backdoor (Pillar Security) [11]	Mar 2025	Cursor rules, Copilot instructions	Hidden Unicode instructions cause the agent to inject a script into generated code	Vendors treated it as user responsibility; GitHub added a hidden-Unicode warning
Claude Code project hooks and MCP consent bypass (Check Point) [4]	Reported Jul–Oct 2025; disclosed Feb 2026	<code>.claude/settings.json</code> , <code>.mcp.json</code>	SessionStart hooks and auto-approved MCP servers ran commands; <code>ANTHROPIC_BASE_URL</code> sent the API key before the trust dialog	Patched (GHSA-ph6w-f82w-28w6, CVE-2025-59536, CVE-2026-21852)
GitHub Copilot auto-approve (CVE-2025-53773) [5]	Aug 2025	<code>.vscode/settings.json</code>	Injected instructions got the agent to set <code>chat.tools.autoApprove</code> , removing confirmation prompts	Patched
Cursor CLI project config (CVE-2025-61592) [14]	Oct 2025	<code>.cursor/cli.json</code> , <code>.cursor/rules/*.mdc</code>	Project config overrode global settings; together with rules-based injection, this led to code execution	Patched
Cursor cross-tool hooks (CVE-2026-48124) [6]	Jun 2026	<code>.claude/settings.local.json</code>	Cursor ran hooks from Claude Code's settings file without asking the user	Fixed in Cursor 3.0.0
Miasma worm, Azure/durabletask (StepSecurity) [3]	Jun 5, 2026	<code>.claude/settings.json</code> , <code>.gemini/settings.json</code> , <code>.cursor/rules/setup.mdc</code> , <code>.vscode/tasks.json</code>	A compromised contributor account committed files that run a 4.6 MB credential stealer when the repository is opened in an AI tool	73 repositories disabled; no CVE

Case	Date	File(s) abused	Mechanism	Status
Downloads-folder <code>AGENTS.md</code> (Forever Security) [1][2]	Oct 2026	<code>AGENTS.md</code> , <code>CLAUDE.md</code>	A file dropped by a website steers an agent launched in <code>~/Downloads</code> to install malware or rewrite its API endpoint	No patch issued; researchers state no vendor fix applies [1]

From Load-Time Bugs to Agent-Authored Configuration

The Check Point findings show what vendors can fix. A repository's `.claude/settings.json` could define hooks and environment variables that took effect before the user accepted the workspace trust dialog. In one case, setting `ANTHROPIC_BASE_URL` to an attacker's server caused Claude Code to send the user's API key in session-initialization requests before any consent [4]. Anthropic fixed each issue by moving the trust boundary earlier: MCP servers no longer start before approval, and network operations wait until the user confirms trust [4]. Microsoft addressed CVE-2025-53773 in a similar way, adding a confirmation gate when Copilot's own edit tool tries to modify the workspace settings file [5][15]. Not every load-time path has been closed, however. Repello AI reports that a workspace setting (`chat.permissions.default` set to `auto-approve`) still takes effect when a VS Code workspace is opened, verified on VS Code 1.123.0, and that Microsoft declined to treat this as a vulnerability [15].

The Forever Security research suggests that this kind of fix has a ceiling. The attacker did not rely on Claude Code reading a malicious setting before trust. Instead, the planted `CLAUDE.md` asked Claude to write the `ANTHROPIC_BASE_URL` change itself, inside a session the user had already started. Help Net Security reports that Claude did so in about half of the researchers' attempts. Because requests carry the user's login token, the researchers could then replay the captured token and take over the account without running any code [2]. Coverage of the code-execution rate through the attacker's server differs: Help Net reports about one in five tests, and BleepingComputer describes roughly half [1][2].

The Codex variant followed a similar path. Help Net reports that the 13 KB file, written in an ordinary engineering-guidelines tone, first got Codex to change its own machine-wide settings and then install the package. Because those settings apply globally, the compromise carries over to later sessions in unrelated projects [2]. In the researchers' reconstruction, Codex executed a benign stand-in payload in more than 90% of attempts, a figure Help Net gives as "more than 90%" and BleepingComputer as a 90% success rate [1][2]. The researchers concluded that "there is nothing here for OpenAI or Anthropic to patch" [1]. A more cautious reading is that patching load-time behavior is necessary but not sufficient. Once an attacker controls the instructions, configuration files the agent can write become an escalation path, and those files need protection from the agent as well as from the repository.

Why Review and Written Rules Fall Short

Defenders often rely on two compensating controls: human review of instruction files, and safety rules written into those same files. The evidence shows limits to both. Review fails against hidden-Unicode payloads that do not render [11]. It also fails against long, plausible files such as the 13 KB Forever sample, where one malicious directive sits among many legitimate ones [2]. The Miasma commit was disguised in a different way. Its message described a code change, but it touched no source code and only added the agent configuration files [3].

Written safety rules, meanwhile, give a false sense of enforcement. Ting Yan's August 2026 preprint, which has not been peer reviewed, examined 481 public `CLAUDE.md` files. It found that, depending on how strictly a match was defined, only about 4–16% of the security rules developers wrote, such as "do not read `.env`," corresponded to a Claude Code control capable of enforcing them; under the strictest criterion the figure was 4.4% [7]. The study measured whether a matching built-in control exists, not whether one was configured in those repositories, and its rates apply to the 66.3% of rules its extraction captured [7]. The author calls `CLAUDE.md` a "write-only channel," because developers get no feedback on whether a stated prohibition is enforced [7]. An attacker-supplied instruction file therefore competes with defender-supplied rules in the same unenforced channel; which one prevails depends on the model's judgment, not on any control the tool applies.

The kind of work the agent is doing also affects exposure. The CIPR benchmark introduced by Zhu et al., with 1,920 poisoned-repository instances across 20 repositories, found that task type alone produced up to a 4.5-fold difference in attack success rate. Test-execution tasks formed a "silent attack surface" with high success and few agent alerts [16]. This suggests that the routine tasks organizations most want to automate, such as running and fixing tests, may be where poisoned instructions succeed most quietly.

The Agent as Delivery Mechanism

Instruction poisoning also connects to earlier supply-chain attacks that used locally installed agents as tools. The August 2025 Nx "singularity" compromise shipped a payload that checked for the Claude, Gemini, and Amazon Q command-line tools and prompted them to find credentials on the host [17]. In July 2025, AWS disclosed that a malicious commit made with an inappropriately scoped GitHub token was included in version 1.84.0 of the Amazon Q Developer extension. The injected code invoked the Q Developer CLI but failed to execute because of a syntax error (CVE-2025-8217) [18]. These incidents suggest that any agent with shell access is a capability an attacker can borrow. The Forever Security case suggests that a Markdown file may be one of the lowest-cost ways to do so, since it requires no code execution to deliver.

Forever Security's telemetry suggests the exposure is broad. The company reports that about half of its customers run a coding agent inside `~/Downloads` and that about a quarter of agent sessions open a browser, suggesting a plausible path by which the file could arrive [2]. These figures are self-reported by a single vendor and have not been independently replicated.

Recommendations

Immediate Actions

Organizations should first stop the most direct exposure: agents running in directories that untrusted parties can write to. Developers should not launch coding agents in `~/Downloads`, home directories, or broad parent folders. Agents should instead be scoped to the specific project directory, as OpenAI and Anthropic told BleepingComputer [1]. Security teams should search developer endpoints and repositories for unexpected `AGENTS.md`, `CLAUDE.md`, `.cursor/rules/`, `.claude/`, `.gemini/`, and `.vscode/tasks.json` files. They should also check agent settings for endpoint overrides such as `ANTHROPIC_BASE_URL` that no one on the team added [2][3].

If a planted file is found, the affected login tokens and API keys should be treated as exposed and rotated. Any machine that installed `function-color`, or that opened `Azure/durabletask` in an AI tool after the June 5, 2026 commit, should be treated as compromised [2][3]. Organizations should also confirm that Claude Code, Cursor, Copilot, and Codex are at current versions so that the load-time fixes described above are in place.

Short-Term Mitigations

Over the next quarter, instruction and agent-settings files should go through the same change control as build scripts. The controls in the table below address the mechanisms seen in the cases above, and each maps to a specific weakness this note has described.

Control	Implementation	Mechanism addressed
Code-owner review	Add <code>CODEOWNERS</code> entries that require security review for <code>AGENTS.md</code> , <code>CLAUDE.md</code> , <code>.cursor/</code> , <code>.claude/</code> , <code>.gemini/</code> , <code>.mcp.json</code> , <code>.vscode/</code> , and <code>.github/copilot-instructions.md</code> [13]	Repository-borne poisoning such as Miasma
Hidden-Unicode scanning	Scan for zero-width and bidirectional Unicode characters in pre-commit hooks and CI, and block merges when they appear in instruction files [11]	Payloads invisible to reviewers, as in the Rules File Backdoor
Settings write protection	Deny agents write access to their own settings files, enforced at the file-system or sandbox level rather than only through the agent's edit-tool permissions [15]	Agent-authored configuration changes, as in the Forever Security case
Instruction-loading telemetry	Use telemetry where the tool provides it; Claude Code's <code>InstructionsLoaded</code> hook can log which instruction files loaded and when [10]	Undetected loading of planted files
Approval and sandbox integrity	Keep approval prompts and sandboxing enabled, and treat any instruction to bypass them as an indicator of compromise [1][2]	Auto-approve abuse, as in CVE-2025-53773

Two of these controls carry caveats. Tool-level permission rules on the edit tool may not stop an agent from writing a settings file through a shell command, which is why file-system or sandbox enforcement is needed so that an instruction file cannot easily persuade the agent to raise its own privileges. Similarly, the `InstructionsLoaded` hook does not fire for an `AGENTS.md` read through Claude Code's project-instructions setting, which limits its value in the Downloads-folder scenario [10].

For enterprises using Claude Code, managed settings provide organization-level instructions that users cannot override. A `managed-only` project-instructions value drops project, user, and local `CLAUDE.md` files and all `AGENTS.md` files at launch, the `claudeMdExcludes` setting can suppress specific ancestor files in monorepos, and external `@path` imports require one-time approval for each project [10]. These controls do not prevent subdirectory `CLAUDE.md` files, `.claude/rules/`, or path-scoped rules from loading on demand when the agent touches files in those directories [10]. Comparable controls for other agents should be inventoried and turned on where they exist.

Strategic Considerations

We recommend that organizations stop writing security policy as prose. The Yan study shows that rules written in Markdown mostly lack a matching enforcement control, and a poisoned file faces no obstacle that a legitimate one would not [7]. Organizations should express prohibitions as permission rules, sandbox profiles, and egress allow-lists that the agent cannot override, and reserve instruction files for conventions that are safe to ignore. Agent vendors would make this practical by separating the two channels more clearly. Useful steps include provenance or signing for instruction files, a clear distinction between instructions from the user, the repository, and the working directory, and requiring confirmation from outside the session before an agent can change its own endpoint, approval mode, or hook configuration.

Organizations should also include coding agents in their software supply-chain risk models, a position CSA has argued in its earlier analysis of coding agents as an unaudited supply-chain node [19]. The Miasma worm showed that a stolen contributor credential plus a set of agent configuration files is enough to spread through repositories without touching a package registry [3]. Tabletop exercises should therefore cover a chain in which one compromised agent leaks a token that is then used to poison widely cloned repositories.

CSA Resource Alignment

CSA's *AI Coding Agents as Attack Surface: MCP, Poisoning, and Miasma* (June 2026) is the most directly relevant prior work [13]. It documents repository configuration poisoning and the Miasma campaign as converging threat vectors and recommends treating agent configuration directories as security-relevant files. This note extends that analysis in two ways. It covers the October 2026 Downloads-folder vector, which needs no repository at all, and it shows that vendor fixes have moved much of the risk from load-time execution to configuration written by the agent itself.

CSA's research note *AI Coding Agents: An Unaudited Supply Chain Node* (July 2026) treats coding agents as an unreviewed dependency in the software supply chain [19]. Instruction-file poisoning is a concrete mechanism by which that dependency can be subverted, and the supply-chain risk-modeling recommendation above builds directly on that note's thesis.

CSA's research note *Sandboxing Agentic AI: Least-Privilege Patterns for MCP and Coding Agents* (May 2026) argues that agent incidents keep following the same pattern: an agent with more authority than its task needs is turned against its user by attacker-controlled input [20]. Instruction-file poisoning is a clear example of that pattern. The isolation, capability scoping, and tamper-resistant audit architecture in that note is the structural control that the recommendations above depend on, especially denying agents write access to their own settings.

CSA's *Agent Context Poisoning: SKILL.md and the New AI Supply Chain Attack Surface* (May 2026) examines the same class of attack through agent skills [21]. `SKILL.md` and `AGENTS.md` share the property that matters here: natural-language files loaded into an agent's trusted context with little provenance. Defenses for one should be designed to cover the other.

Finally, the *CSA AI Controls Matrix (AICM) v1.1* provides the control framework for operationalizing this guidance [22]. Its Application and Interface Security, Threat and Vulnerability Management, Change Control and Configuration Management, and Identity and Access Management domains map onto the controls recommended here: code-owner review of instruction files, configuration integrity for agent settings, least-privilege agent credentials, and monitoring of agent behavior. For threat modeling, CSA's *MAESTRO framework* offers a layered method for placing instruction-channel compromise among the other risks to an agent ecosystem [23].

References

- [1] Toulas, B. "[AI coding agents hijacked using AGENTS.md file in Downloads folder](#)." BleepingComputer, October 25, 2026.
- [2] Markovic, S. "[AI coding agents fall for malicious instructions planted in the Downloads folder](#)." Help Net Security, October 26, 2026.
- [3] Kurmi, A. "[Miasma Worm Hits Microsoft Again: Azure Functions Action and 72 Other Repositories Disabled After Supply Chain Attack Targeting AI Coding Agents](#)." StepSecurity, June 5, 2026.
- [4] Donenfeld, A. and Vanunu, O. "[Caught in the Hook: RCE and API Token Exfiltration Through Claude Code Project Files](#)." Check Point Research, February 25, 2026.
- [5] NIST. "[CVE-2025-53773 Detail](#)." National Vulnerability Database, August 2025.
- [6] SentinelOne. "[CVE-2026-48124: Cursor Desktop Hook Execution Vulnerability](#)." SentinelOne Vulnerability Database, June 2026.
- [7] Yan, T. "[When 'Do Not' Is Not Deny: Security Rules in CLAUDE.md vs Built-In Controls](#)." arXiv:2608.23550 (preprint), August 2026.
- [8] OpenAI. "[OpenAI co-founds the Agentic AI Foundation under the Linux Foundation](#)." OpenAI, December 9, 2025.
- [9] OpenAI. "[Custom instructions with AGENTS.md](#)." OpenAI Codex Documentation, accessed October 2026.
- [10] Anthropic. "[How Claude remembers your project](#)." Claude Code Documentation, accessed October 2026.
- [11] Karliner, Z. et al. "[New Vulnerability in GitHub Copilot and Cursor: How Hackers Can Weaponize Code Agents](#)." Pillar Security, March 18, 2025.
- [12] Lakshmanan, R. "[Researchers Uncover 30+ Flaws in AI Coding Tools Enabling Data Theft and RCE Attacks](#)." The Hacker News, December 6, 2025.
- [13] Cloud Security Alliance. "[AI Coding Agents as Attack Surface: MCP, Poisoning, and Miasma](#)." CSA AI Safety Initiative, June 28, 2026.
- [14] NIST. "[CVE-2025-61592 Detail](#)." National Vulnerability Database, October 2025.
- [15] Borpe, S. "[VS Code Copilot Workspace Trust Bypass](#)." Repello AI, June 4, 2026.
- [16] Zhu, F. et al. "[Beyond the Payload: How User Invocation Shapes Coding Agent Vulnerability to Repository Poisoning](#)." arXiv:2608.30686 (EMNLP 2026), August 2026.
- [17] Okta Threat Intelligence. "[The singularity attack: When attackers prompt your AI agents to do their bidding](#)." Okta, 2025.
- [18] Amazon Web Services. "[Security Bulletin AWS-2025-015: Amazon Q Developer for VS Code Extension \(CVE-2025-8217\)](#)." AWS, July 23, 2025.
- [19] Cloud Security Alliance. "[AI Coding Agents: An Unaudited Supply Chain Node](#)." CSA AI Safety Initiative, July 8, 2026.
- [20] Cloud Security Alliance. "[Sandboxing Agentic AI: Least-Privilege Patterns for MCP and Coding Agents](#)." CSA AI Safety Initiative, May 30, 2026.
- [21] Cloud Security Alliance. "[Agent Context Poisoning: SKILL.md and the New AI Supply Chain Attack Surface](#)." CSA AI Safety Initiative, May 6, 2026.
- [22] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." Cloud Security Alliance, 2025.
- [23] Huang, K. "[Agentic AI Threat Modeling Framework: MAESTRO](#)." Cloud Security Alliance, February 6, 2025.