INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-20 18:52 UTC

Critical Anthropic MCP Vulnerability Enables Remote Code Execution Attacks

CVE VULNERABILITY | CRITICAL | CVSS 9.8

SCC Item ID	SCC-CVE-2026-0058
Type	CVE Vulnerability
CVE ID	CVE-2025-49596
Severity	CRITICAL
CVSS Base Score	9.8
EPSS Score	0.0288 (86th percentile)
Affected Products	Anthropic MCP Inspector (mcp-inspector); mcp-remote package; multiple MCP SDK implementations across programming environments
Published	8 hours ago
Discovery Source	Serper

Executive Summary

A critical remote code execution vulnerability (CVE-2025-49596, CVSS 9.8) was disclosed in Anthropic's Model Context Protocol Inspector tool and the mcp-remote package, affecting developers and organizations using MCP-based tooling for AI agent orchestration. An attacker who can deliver malicious tool definitions or interact with an exposed MCP server can execute arbitrary code on the host system without authentication. Organizations integrating MCP tooling into their AI orchestration pipelines face direct risk of system compromise, supply chain poisoning, and lateral movement into broader development or production environments.

Technical Analysis

CVE-2025-49596 carries a CVSS base score of 9.8 (published by NVD; vendor CVSS vector was not published at time of this writing) and is rooted in three weakness classes: CWE-94 (Code Injection), CWE-74 (Injection), and CWE-20 (Improper Input Validation). The vulnerability exists in the Anthropic MCP Inspector tool and the mcp-remote package, with OX Security documenting a broader architectural concern across multiple MCP SDK implementations. The core flaw lies in how MCP processes tool definitions and server-side interactions - malicious content injected into these pathways results in arbitrary code execution on the affected host. Relevant MITRE ATT&CK techniques include T1195.002 (Compromise Software Supply Chain), T1059 (Command and Scripting Interpreter), and T1190 (Exploit Public-Facing Application). Multiple independent security researchers

disclosed variants of this issue: Oligo Security identified the MCP Inspector RCE; Tenable Research independently disclosed a critical RCE on Anthropic's MCP infrastructure; OX Security characterized the issue as systemic across MCP implementations. EPSS score is 0.02881 (86th percentile), indicating elevated exploitation probability relative to the broader CVE population. The item is not currently listed in CISA KEV. Patch and version-specific remediation details should be confirmed directly against the NVD entry (https://nvd.nist.gov/vuln/detail/CVE-2025-49596) and Anthropic's official advisory, as disclosure was recent and details may be updated.

Action Checklist

1. Step 1: Containment. Identify all internal systems running mcp-inspector or the mcp-remote package. Immediately restrict network access to MCP Inspector instances; remove internet-facing exposure. Audit CI/CD pipelines, developer workstations, and AI orchestration infrastructure for active MCP tool usage. Note: While MCP is marketed as a developer tool, it is frequently integrated into production AI orchestration pipelines. Audit both development and production infrastructure.
2. Step 2: Detection. Search developer workstation and server logs for unexpected process spawning from MCP-related processes (mcp-inspector, mcp-remote). Review npm audit output and package-lock.json or equivalent dependency manifests for mcp-inspector and mcp-remote entries. Check for anomalous outbound connections or command execution events originating from AI orchestration services. No confirmed public IOC hashes or IPs are available at this time as of publication date; behavioral detection is the primary signal. Check CISA, VulnCheck, and threat intel feeds for emerging IOCs.
3. Step 3: Eradication. Verify that Anthropic has released patched versions of mcp-inspector and mcp-remote through their official security advisory. If patches are available, apply immediately and confirm patched versions against the NVD entry (https://nvd.nist.gov/vuln/detail/CVE-2025-49596) and Anthropic's official security advisory before deployment. If no patch is yet available, uninstall or disable the affected packages and suspend MCP-dependent workflows until vendor remediation is complete.
4. Step 4: Recovery. After patching, re-run dependency scans across all affected repositories and pipelines. Validate that no unauthorized processes or persistence mechanisms were established on systems where MCP tooling was running. Monitor process execution and network telemetry on previously exposed hosts for a minimum of 72 hours post-remediation. Confirm no malicious tool definitions were injected into shared MCP server configurations.
5. Step 5: Post-Incident. This vulnerability exposes a control gap in AI tooling supply chain governance. Review your organization's inventory process for developer libraries and AI orchestration dependencies - MCP packages were likely not tracked as production-risk assets. Establish a policy requiring security review before adopting new AI SDK or agent framework components. Evaluate whether MCP server interactions in your environment are subject to input validation controls at the application layer.

Detection Guidance

As of publication date, no confirmed public IOCs (hashes, IPs, domains) have been attributed to active exploitation of CVE-2025-49596. Detection should focus on behavioral indicators: (1) Unexpected child process spawning from mcp-inspector or mcp-remote processes - look for shells (sh, bash, cmd, powershell) or interpreters launched as children of Node.js processes associated with MCP tooling. (2) Anomalous outbound network connections from developer workstations or AI pipeline hosts, particularly from Node.js processes. (3)

npm audit findings flagging mcp-inspector or mcp-remote in dependency trees across repositories. (4) Unexpected modifications to MCP server configuration files or tool definition files. EDR telemetry is the highest-value detection source. SIEM rules should target process lineage anomalies in environments where MCP tooling is known to run. This is a developer-tooling vector; extend detection scope beyond production servers to include developer endpoints and CI/CD runners. Check CISA, VulnCheck, and threat intel feeds for updates if IOCs emerge post-publication.

Framework Mappings

MITRE-ATTACK

- **T1195.002** — Compromise Software Supply Chain
- **T1059** — Command and Scripting Interpreter
- **T1190** — Exploit Public-Facing Application

NIST-800-53R5

- **CM-7** — Least Functionality
- **SA-9** — External System Services
- **SR-3** — Supply Chain Controls and Processes
- **SI-7** — Software, Firmware, and Information Integrity
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-10** — Information Input Validation
- **SR-2** — Supply Chain Risk Management Plan

OWASP-TOP10-2021

- **A03:2021** — Injection

CIS-V8

- **16.10** — Apply Secure Design Principles in Application Architectures
- **15.1** — Establish and Maintain an Inventory of Service Providers

ISO-27001-2022

- **A.8.26** — Application security requirements
- **A.8.8** — Management of technical vulnerabilities
- **A.5.21** — Managing information security in the ICT supply chain

NIST-CSF-2

- **GV.SC-01** — Cybersecurity supply chain risk management program

SOC2-TSC

- **CC9.2** — Manages risks associated with vendors and business partners

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1195.002	Compromise Software Supply Chain	Initial-Access
T1059	Command and Scripting Interpreter	Execution
T1190	Exploit Public-Facing Application	Initial-Access

Sources

Source	URL	Tier
	https://cyberpress.org/critical-anthropic-mcp-vulnerability/	T3
The Architectural Flaw at the Core of Anthropic's MCP - OX Security	https://www.ox.security/blog/the-mother-of-all-ai-supply-chains-cri...	T3
How Tenable Research Discovered a Critical Remote Code ...	https://www.tenable.com/blog/how-tenable-research-discovered-a-crit...	T3
Critical mcp-remote Vulnerability Enables Remote Code Execution ...	https://thehackernews.com/2025/07/critical-mcp-remote-vulnerability...	T3
Critical RCE Vulnerability in Anthropic MCP Inspector - Oligo Security	https://www.oligo.security/blog/critical-rce-vulnerability-in-anthr...	T3
NVD	https://nvd.nist.gov/vuln/detail/CVE-2025-49596	T1

Generated 2026-04-20 18:52 UTC by TJS Security Command Center