CrewAI contains multiple vulnerabilities including SSRF, RCE and local file read

Vulnerability Note VU#221883

Original Release Date: 2026-03-30 | Last Revised: 2026-03-30

Overview

Four vulnerabilities have been identified in CrewAI, including remote code execution (RCE), arbitrary local file read, and server-side request forgery (SSRF). CVE-2026-2275 is directly caused by the Code Interpreter Tool. The other three vulnerabilities result from improper default configuration settings within the main CrewAI agent and associated Docker images. An attacker who can interact with a CrewAI agent that has the Code Interpreter Tool enabled may exploit these issues through prompt injection, ultimately chaining the vulnerabilities together. The vendor has provided a statement addressing some, but not all, of the reported vulnerabilities.

ABOUT VULNERABILITY NOTES

CONTACT US ABOUT THIS VULNERABILITY

PROVIDE A VENDOR STATEMENT

Description

CrewAI is a tool for building and orchestrating multi-agent AI systems. These agents are intended to work together to complete tasks, and developers define those tasks and workflows. CrewAI supports various tools, including one called the "Code Interpreter Tool", intended for execution of Python code within a secure Docker container.

CVE-2026-2275 origintate from the Code Interpreter tool itself. The remaining vulnerabilities stem from insecure fallback behaviors and configuration issues in the CrewAI agent and Docker environment. Exploitation of CVE-2026-2275 may enable attackers to trigger the additional vulnerabilities.

The vulnerabilities are listed below:

CVE-2026-2275 The CrewAI CodeInterpreter tool falls back to SandboxPython when it cannot reach Docker, which can enable code execution through arbitrary C function calls. This vulnerability can be triggered if: `allow_code_execution=True` is enabled in the agent configuration, or if the Code Interpreter Tool is manually added to the agent by the developer.

CVE-2026-2286 CrewAI contains a server-side request forgery (SSRF) vulnerability that enables content acquisition from internal and cloud services, facilitated by the RAG search tools not properly validating URLs provided at runtime.

CVE-2026-2287 CrewAI does not properly check that Docker is still running during runtime, and will fall back to a sandbox setting that allows for RCE exploitation.

CVE-2026-2285 CrewAI contains a arbitrary local file read vulnerability in the JSON loader tool that reads files without path validation, enabling access to files on the server.

CVE-2026-2275 can be triggered if 'allow_code_execution=True' is enabled in the agent settings or the tool is manually added to the agent by the creator.

Impact

An attacker with the ability to influence a CrewAI agent using the Code Interpreter Tool through either direct or indirect prompt injection can use the four vulnerabilities discovered to perform arbitrary file read, RCE, and server side request forgery. The results of the attacks can vary, as the attacker will achieve sandbox bypass and RCE/file read if the host machine is using Docker, or full RCE if the host machine is in configuration mode or unsafe mode. An attacker can use the arbitrary file read and SSRF vulnerabilities to perform credential theft, or the RCE vulnerabilities to perform further leveraging of the compromised device.

Solution

During coordinated disclosure, the vendor provided a statement addressing CVE-2026-2275 and CVE-2026-2287.

The vendor has indicated plans to take the following actions to improve security of CrewAI framework:

- Add `ctypes` and related modules to BLOCKED_MODULES in an upcoming release
- Evaluate configuration changes to fail closed rather than fall back to sandbox mode
- Provide clearer runtime warnings when sandbox mode is active
- Improve security-related documentation

At the time of writing, no complete patch is available for all disclosed vulnerabilities. Until fixes are released, users should:

- Remove or restrict or disable the Code Interpreter Tool wherever possible
- Remove (or avoid) enabling `allow_code_execution=True` setting unless absolutely necessary
- Limit the agent exposure to untrusted input or santiize input as appropriate
- Monitor Docker availability and prevent fallback to insecure sandbox modes

Acknowledgements

Thanks to the reporter, Yarden Porat of Cyata. This document was written by Christopher Cullen.

Vendor Information

Filter by status: All

Filter by content:

Additional information available

Sort by: Status

Expand all

CrewAI	Affected

References

- **https://docs.crewai.com/en/tools/ai-ml/codeinterpretertool**

Other Information

CVE IDs:	CVE-2026-2275 CVE-2026-2285 CVE-2026-2286 CVE-2026-2287
API URL:	VINCE JSON | CSAF
Date Public:	2026-03-26
Date First Published:	2026-03-30
Date Last Updated:	2026-03-30 15:50 UTC
Document Revision:	1

Download PGP Key

Read CERT/CC Blog

Learn about Vulnerability Analysis

Carnegie Mellon University
Software Engineering
Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-2612
412-268-5800

Contact SEI

Contact CERT/CC

412-268-5800
cert@cert.org