



PROFILE

CVE-2022-22947

ZEROFOX INTELLIGENCE

P-2022-12-06a

Classification: TLP:CLEAR

Criticality: High

Intelligence Requirements: Common Vulnerabilities & Exposures

December 6, 2022



PROFILE: CVE-2022-22947

P-2022-12-06a

TLP:CLEAR

Created on: November 28, 2022

Intelligence Cut-off: 12:00 PM (EST) on December 1, 2022

First Observed: October 20, 2022

CVE Dictionary Entry: CVE-2022-22947

NVD Published Date: March 3, 2022

NVD Last Modified: October 28, 2022

Severity: CRITICAL

Base Score: 10.0

Vector: CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H

Metasploit Available: Yes

Current Description

In Spring Cloud Gateway versions prior to 3.1.1+ and 3.0.7+, applications are vulnerable to a code injection attack when the Gateway Actuator endpoint is enabled, exposed, and unsecured. A remote attacker could craft a malicious request that could allow arbitrary remote execution on the remote host.

Risk

CVE-2022-22947 is an RCE vulnerability that allows threat actors to gain an initial foothold within a corporate network. When successfully exploited, this vulnerability provides the running malware with the same permissions as the user running the Spring Cloud Gateway application. This vulnerability allows a threat actor the ability to easily run arbitrary shell commands on a front-facing application. Accordingly, ZeroFox rates its impact as high because a successful exploitation attempt can lead to severe economic and reputational damage for an organization.

Additionally, this vulnerability has a widely-available Metasploit module, and a multitude of proof-of-concept exploits are readily available. The existence of a Metasploit module allows for easy weaponization of this vulnerability, making the overall likelihood of an exploitation attempt high. Cybersecurity researchers report that threat actors have leveraged this vulnerability to scan for and exploit vulnerable endpoints, further supporting this rating.



Known Affected Software Configurations

Configuration	Common Platform Enumeration
1	cpe:2.3:a:vmware:spring_cloud_gateway:*.~*~*~*~*~*~*
1	cpe:2.3:a:vmware:spring_cloud_gateway:3.1.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:commerce_guided_search:11.3.2:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_bindin g_support_function:1.11.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_bindin g_support_function:22.1.3:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_consol e:22.2.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_exposure_function:22.1.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_function_cloud_native_environment:1.10.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_repository_function:1.15.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_repository_function:1.15.1:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_repository_function:22.1.2:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_repository_function:22.2.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_slice_selection_function:1.8.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_networ k_slice_selection_function:22.1.0:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_securit y_edge_protection_proxy:22.1.1:~*~*~*~*~*~*
2	cpe:2.3:a:oracle:communications_cloud_native_core_service _communication_proxy:1.15.0:~*~*~*~*~*~*



Weakness Enumeration

CWE-ID	CWE Name	Source
CWE-94	Improper Control of Generation of Code ('Code Injection')	NIST, VMware

Vulnerability Analysis

CVE-2022-22947 is a Remote Code Execution (RCE) within Spring Cloud Gateway Versions 3.10 and 3.06-3.0. This vulnerability was discovered by Wyatt Dahlenburg, a security researcher, and disclosed to VMware on January 15, 2022; the researcher posted about his discovery and his reporting of it to VMware on his blog on February 26, 2022.¹ On October 6, 2022, a Metasploit module was made publicly available via a GitHub repository. Throughout November, GreyNoise[.]io has detected 14 unique IP addresses attempting to exploit CVE-2022-22947 (Figure 1).




Figure 1. Unique IP scanning for CVE-2022-22947 over time

Source: [hXXps://viz.greynoise\[.\]io/tag/spring-cloud-gateway-code-injection-attempt?days=30](https://viz.greynoise.io/tag/spring-cloud-gateway-code-injection-attempt?days=30)

On his blog, Dahlenburg described how he discovered CVE-2022-22947 by realizing that the Java application Spring Cloud Gateway used Spring Expression Language (SpEL). In his investigation, he found that submitting a SpEL string as an argument to RewritePath to the `/actuator/gateway/routes/new_route` endpoint will be interpreted without any input sanitization.

¹ [hXXps://wya\[.\]pl/2022/02/26/cve-2022-22947-spel-casting-and-evil-beans/](https://wya[.]pl/2022/02/26/cve-2022-22947-spel-casting-and-evil-beans/)



Dahlenburg was able to take advantage of this vulnerability by submitting a SpEL string that contained the following Java Bean:

T{(java.lang.Runtime).getRuntime().exec(\"touch /tmp/x\")}. In this proof of concept, submitting an SpEL string provides access to a *getRuntime* function call, which eventually leads to running arbitrary code execution on the host system.

Recommendations

Since the disclosure by a security researcher on January 15, 2022, a patch has been released for Spring Cloud Gateway. ZeroFox Intelligence recommends updating all Spring Cloud Gateway applications to versions 3.1.1+ and 3.0.7+. If updating Spring Cloud Gateway is not immediately possible, put the Spring Cloud Gateway application behind a firewall to limit access. ZeroFox also recommends installing auditing software, such as filebeat and auditd, along with a SIEM to detect if the server has been compromised.



Appendix 1: IOCs (Sigma Rules)

```
title: CVE-2022-22947-Metasploit-Attempt
id: 8b24092d-3b9c-490e-8430-26cc1bc3f3ed
description: Rule to detect metasploit payloads once they have successfully exploited
CVE-2022-22947. This sigma rule does not trigger on CVE-2022-22947 exploit attempt but on
metasploit payloads that work in a standard linux environment. Sigma rule will trigger on a
successful execution attempt.
status: experimental
date: 2022/12/01
modified: 2022/12/01
author: ZeroFox
tags:
  - mitre.T1190
  - mitre.T1059
  - attack.InitialAccess
  - attack.Execution
references:
  - https://nvd.nist.gov/vuln/detail/CVE-2022-22947
  - https://wya.pl/2022/02/26/cve-2022-22947-spel-casting-and-evil-beans/
logsource:
  product: linux
  service: auditd
detection:
  selection:
    log.file.path: '*audit.log'
  message:
    - '*6563686f*657865635c285f5f696d706f72745f5f*'
    - '*303C26*3c3E*2F6465762F7463702F*'
  condition: selection
falsepositives:
  - Legitimate use of invoking python that requires executing __import__ directly from command
line
level: high
```

Figure 2. Sigma Rule to detect successful exploit attempts leveraging Metasploit framework

Source: ZeroFox Intelligence Collections



Appendix 2: Traffic Light Protocol for Information Dissemination

	Red	Amber
WHEN SHOULD IT BE USED?	Sources may use TLP:RED when information cannot be effectively acted upon by additional parties and could lead to impacts on a party's privacy, reputation, or operations if misused.	Sources may use TLP:AMBER when information requires support to be effectively acted upon but carries risks to privacy, reputation, or operations if shared outside of the organizations involved.
HOW MAY IT BE SHARED?	Recipients may NOT share TLP:RED with any parties outside of the specific exchange, meeting, or conversation in which it is originally disclosed.	Recipients may ONLY share TLP:AMBER information with members of their own organization and its clients, but only on a need-to-know basis to protect their organization and its clients and prevent further harm. Note that TLP:AMBER+STRICT restricts sharing to the organization only.
	Green	Clear
WHEN SHOULD IT BE USED?	Sources may use TLP:GREEN when information is useful for the awareness of all participating organizations, as well as with peers within the broader community or sector.	Sources may use TLP:CLEAR when information carries minimal or no risk of misuse in accordance with applicable rules and procedures for public release.
HOW MAY IT BE SHARED?	Recipients may share TLP:GREEN information with peers and partner organizations within their sector or community but not via publicly accessible channels.	Recipients may share TLP:CLEAR information without restriction, subject to copyright controls.