Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

VULNERABILITY REPORT

January 2026 Linux Patch Roundup

Date of Publication	Admiralty Code	TA Number
January 27, 2026	A1	TA2026024

Summary

In January, more than **585** new vulnerabilities were discovered and addressed within the Linux ecosystem, impacting several major distributions such as Debian, Red Hat, OpenSUSE, and Ubuntu. During this period, over **1088** vulnerabilities were also highlighted, with corresponding hotfixes or patches released to resolve them. These vulnerabilities span from information disclosure to privilege escalation to code execution. HiveForce Labs has identified 10 severe vulnerabilities that are exploited or have a high potential of successful exploitation, necessitating immediate attention. To ensure protection, it is essential to upgrade systems to the latest version with the necessary security patches and appropriate security controls.

Threat Distribution

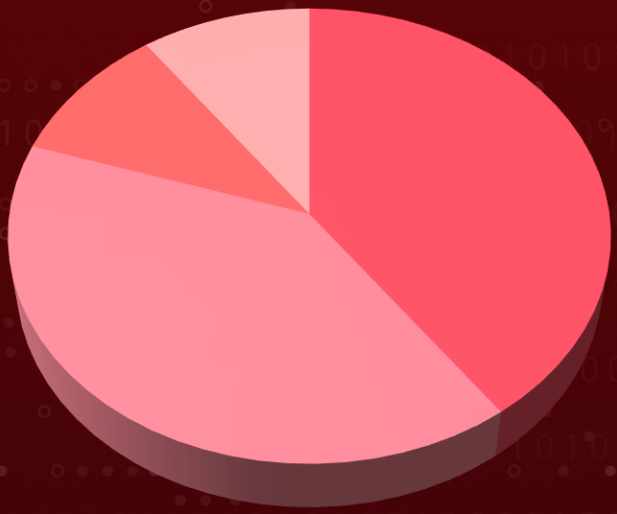

Adversary Tactics

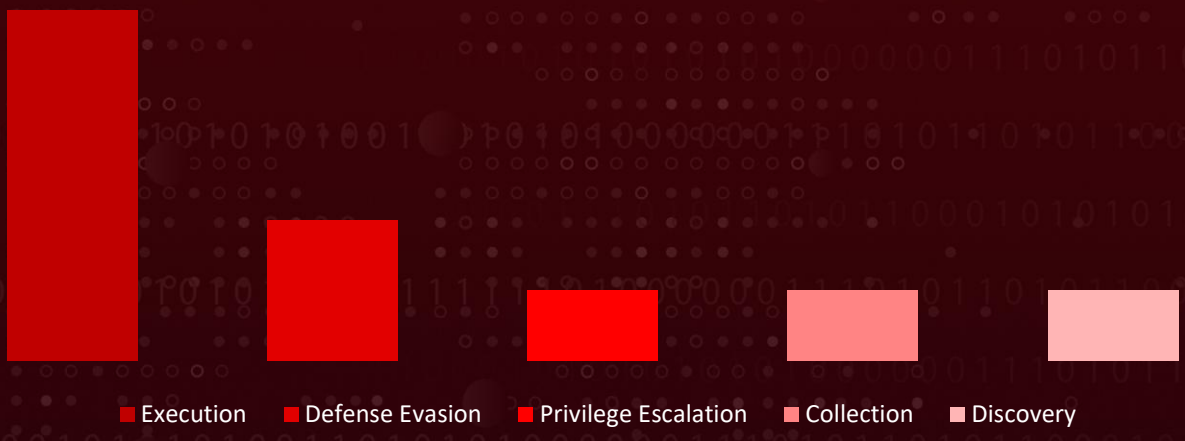

CVEs

CVE	NAME	AFFECTED PRODUCT	Impact	Attack Vector
CVE-2025-8110*	Gogs Symlink Bypass Remote Code Execution Vulnerability	Gogs, Alpine	Execution	Network
CVE-2026-0902	Chromium Out of Bounds Memory read Vulnerability	Debian, SUSE	Collection	Network
CVE-2026-0905	Chromium Information Disclosure Vulnerability	Debian, SUSE	Discovery	Network
CVE-2026-0906	Chromium Spoofing Vulnerability	Debian, SUSE	Defense Evasion	Network
CVE-2022-3564	Linux Kernel Use-After-Free Vulnerability	Linux, Ubuntu, RedHat, Debian	Execution	Local
CVE-2026-0628	Chromium Privilege Escalation Vulnerability	Debian, SUSE	Privilege Escalation	Network
CVE-2026-21876	OWASP CRS Charset Bypass Vulnerability	Debian, Ubuntu	Defense Evasion	Network

* Refers to **Notable CVEs**, vulnerabilities that are either exploited in zero-day attacks, included in the CISA KEV catalog, utilized in malware operations, or targeted by threat actors in their campaigns.

CVE	NAME	AFFECTED PRODUCT	Impact	Attack Vector
CVE-2025-21756	Linux Kernel Use-After-Free Vulnerability	Ubuntu, Rocky Linux, RedHat, Debian	Execution	Local
CVE-2025-68664*	LangGrinch (LangChain Serialization Injection Vulnerability)	LangChain	Execution	Network
CVE-2025-67268	gpsd Out-Of-Bounds Write Vulnerability	Alpine Linux, Debian, RedHat, Ubuntu	Execution	Network

* Refers to **Notable CVEs**, vulnerabilities that are either exploited in zero-day attacks, included in the CISA KEV catalog, utilized in malware operations, or targeted by threat actors in their campaigns.

Notable CVEs

Notable CVEs include vulnerabilities exploited in zero-day attacks, listed in the CISA KEV catalog, used in malware operations, or targeted by threat actors in their campaigns.

CVE ID	CELEBRITY VULNERABILITY	AFFECTED PRODUCTS	ASSOCIATED ACTOR
CVE-2025-8110	✗	Gogs (Before 0.13.4, all versions through 0.13.3), Alpine	-
	ZERO-DAY		
	✓	AFFECTED CPE	ASSOCIATED ATTACKS/RANSOMWARE
NAME	CISA KEV	cpe:2.3:a:gogs:gogs:*:*:*:*:*:*:*:* cpe:2.3:o:alpsalpine:*:*:*:*:*:*:*:*:*	-
Gogs Symlink Bypass Remote Code Execution Vulnerability	✓		
	CWE ID	ASSOCIATED TTPs	PATCH LINKS
	CWE-22	T1190: Exploit Public-Facing Application, T1059: Command and Scripting Interpreter, T1059.004: Unix Shell, T1083: File and Directory Discovery	Gogs, Alpine

CVE ID	CELEBRITY VULNERABILITY	AFFECTED PRODUCTS	ASSOCIATED ACTOR
CVE-2025-68664	LangGrinch	Langchain-core versions before: 0.3.81 and 1.2.5	-
	ZERO-DAY		
	✗	AFFECTED CPE	ASSOCIATED ATTACKS/RANSOMWARE
NAME	CISA KEV	cpe:2.3:a:langchain-ai:langchain:*:*:*:*:*:*:*:*	-
LangGrinch (LangChain Serialization Injection Vulnerability)	✗		
	CWE ID	ASSOCIATED TTPs	PATCH LINK
	CWE-502	T1059: Command and Scripting Interpreter, T1059.006: Python, T1552: Unsecured Credentials, T1552.001: Credentials In Files, T1195: Supply Chain Compromise, T1195.002: Compromise Software Supply Chain	LangChain

Vulnerability Details

#1

In January, the Linux ecosystem addressed over **1088** vulnerabilities across various distributions and products, covering critical issues such as information disclosure, privilege escalation, and code execution. Over **585** new vulnerabilities were discovered and patched. HiveForce lab has identified **10** critical vulnerabilities that are either currently being exploited or are highly likely to be exploited in the near future.

#2

These vulnerabilities could facilitate adversarial tactics such as Execution, Collection, Discovery, and Privilege Escalation. Notably, two of these vulnerabilities are under active exploitation, which requires urgent attention and remediation.

#3

The most urgent case involves Gogs, a popular self-hosted Git service. A newly identified zero-day flaw, CVE-2025-8110, is actively exploited in the wild and bypasses a previously patched remote code execution issue. The root cause lies in improper symbolic link handling in the PutContents API, which allows authenticated attackers to write files outside repository boundaries. By overwriting sensitive system files, attackers can achieve full code execution.

#4

More than 700 Gogs instances are already compromised, and over 1,400 exposed servers still allow open registration by default, creating a broad attack surface. Mitigations such as disabling open registration and isolating instances behind VPNs.

#5

At the kernel level, a long-standing vulnerability in the Linux Bluetooth stack remains a serious concern. CVE-2022-3564 is a use-after-free flaw in the l2cap_reassemble_sdu function of the L2CAP protocol. Because Bluetooth protocols in Linux are processed inside the kernel and exposed without authentication, the impact is severe.

#6

An attacker within Bluetooth range, armed only with the target's Bluetooth address, can trigger kernel crashes or potentially execute code with full kernel privileges by sending crafted L2CAP packets. Compromised or malicious Bluetooth hardware can exploit the same weakness.

#7

The AI ecosystem is not immune. LangChain Core contains a critical serialization injection vulnerability, CVE-2025-68664, known as "LangGrinch." The issue arises because the dumps() and dumpd() functions fail to properly escape user-controlled dictionaries containing the reserved lc key. During deserialization, especially when secrets_from_env was enabled by default, injected data is interpreted as legitimate LangChain objects.

#8

This allows attackers to extract environment variables holding API keys, database credentials, and other secrets, and in some cases reach code execution. The most common entry point is through LLM-controlled fields such as additional_kwargs or response_metadata, which can be influenced via prompt injection and later serialized during caching, logging, or streaming. A parallel flaw in LangChain.js highlights how this insecure pattern crosses language boundaries.

#9

Finally, critical infrastructure faces risk from a memory corruption bug in gpsd, the widely deployed GPS service daemon. CVE-2025-67268 is a heap-based out-of-bounds write in the hnd_129540 function of the NMEA2000 driver. The attack requires no authentication or user interaction and can be launched by injecting malicious packets on a CAN bus or NMEA2000 network. The result ranges from daemon crashes to potential code execution. Given gpsd's role in transportation systems, maritime navigation, and critical infrastructure monitoring, the consequences extend well beyond a single service failure.

Recommendations

Proactive Strategies:

- Enforce secure-by-default configurations across all services. Disable open registration, anonymous access, and unused APIs in platforms like Gogs and similar developer tooling. Treat exposed management interfaces as hostile by default.
- Harden trust boundaries in application logic. Validate file paths, symbolic links, array bounds, and serialization inputs rigorously. Assume all user-controlled data, including LLM outputs and metadata fields, is untrusted.
- Eliminate unsafe serialization patterns. Avoid deserializing user-influenced data into executable objects. In AI frameworks, remove implicit environment secret loading and restrict deserialization to strict, schema-validated formats.
- Apply defense-in-depth for memory safety. Use compiler hardening flags, memory-safe languages where possible, and runtime protections such as ASLR, DEP, and heap sanitization to limit exploit reliability.
- Implement continuous exposure monitoring. Actively scan for internet-facing services, misconfigurations, and known vulnerable versions to detect risk before exploitation occurs.

Reactive Strategies:

- Immediate Risk Containment and Hardening: Reduce exposure without delay by disabling open registration and limiting access strictly to trusted, authenticated users. Place all Gogs instances behind controlled network boundaries such as firewalls or VPNs. Conduct a thorough integrity review of repositories and underlying system files, rotate all credentials, and operate under the assumption of compromise until systems are fully validated.
- Secure Serialization and Secret Protection: Assume all serialized and deserialized data is untrusted. Remove or disable unsafe deserialization paths, explicitly disable automatic environment secret loading, and rotate all environment-based credentials. Review all workflows that serialize LLM-influenced fields, including metadata, caching, logging, and streaming, and enforce strict schema validation to prevent object injection and secret exfiltration.

Detect, Mitigate & Patch

CVE ID	TTPs	Detection	Mitigation	Patch
CVE-2025-8110*	T1190: Exploit Public-Facing Application, T1059: Command and Scripting Interpreter, T1059.004: Unix Shell, T1083: File and Directory Discovery	DET0080: Exploit Public-Facing Application - multi-signal correlation (request → error → post-exploit process/egress), DET0516: Behavioral Detection of Command and Scripting Interpreter Abuse, DET0215: Detection of Multi-Platform File Encryption for Impact	M1048: Application Isolation and Sandboxing, M1038: Execution Prevention, M1017: User Training	Gogs, Alpine
CVE-2026-0902	T1189: Drive-by Compromise, T1203: Exploitation for Client Execution	DET0189: Detection Strategy for Indicator Removal from Tools - Post-AV Evasion Modification, DET0203: Detection Strategy for Ptrace-Based Process Injection on Linux	M1051: Update Software, M1048: Application Isolation and Sandboxing	Debian, SUSE
CVE-2026-0905	T1071: Application Layer Protocol, T1071.001: Web Protocols, T1189: Drive-by Compromise	DET0071: Detection of Remote Data Staging Prior to Exfiltration, DET0189: Detection Strategy for Indicator Removal from Tools - Post-AV Evasion Modification	M1031: Network Intrusion Prevention, M1051: Update Software	Debian, SUSE

CVE ID	TTPs	Detection	Mitigation	Patch
CVE-2026-0906	T1204: User Execution, T1204.001: Malicious Link, T1566: Phishing, T1566.002: Spearphishing Link	DET0340: User Execution – Malicious Copy & Paste	M1017: User Training, M1054: Software Configuration	Debian, SUSE
CVE-2022-3564	T1068: Exploitation for Privilege Escalation, T1499: Endpoint Denial of Service, T1499.004: Application or System Exploitation, T1200: Hardware Additions	DET0068: Malicious IIS Components, DET0499: Behavioral Detection of Fallback or Alternate C2 Channels	M1019: Threat Intelligence Program, M1038: Execution Prevention, M1035: Limit Hardware Installation	Linux, Ubuntu, RedHat, Debian
CVE-2026-0628	T1059: Command and Scripting Interpreter, T1059.007: JavaScript, T1185: Browser Session Hijacking, T1176: Browser Extensions	DET0516: Behavioral Detection of Command and Scripting Interpreter Abuse. DET0185: Behavioral Detection Strategy for Use Alternate Authentication Material, DET0176: Drive-by Compromise — Behavior-based, Multi-platform Detection Strategy	M1038: Execution Prevention, M1017: User Training, M1047: Audit	Debian, SUSE

CVE ID	TTPs	Detection	Mitigation	Patch
CVE-2026-21876	T1190: Exploit Public-Facing Application, T1027: Obfuscated Files or Information	DET0080: Exploit Public-Facing Application – multi-signal correlation (request → error → post-exploit process/egress), DET0027: Detection of Web Protocol-Based C2 Over HTTP, HTTPS, or WebSockets	M1030: Network Segmentation, M1049: Antivirus/Antimalware	Debian, Ubuntu
CVE-2025-21756	T1068: Exploitation for Privilege Escalation, T1611: Escape to Host, T1499: Endpoint Denial of Service, T1499.004: Application or System Exploitation	DET0068: Detection Strategy for T1505.004 - Malicious IIS Components, DET0611: Detection of Access Notifications, DET0499: Behavioral Detection of Fallback or Alternate C2 Channels	M1026: Privileged Account Management, M1048: Application Isolation and Sandboxing, M1038: Execution Prevention	Ubuntu, Rocky Linux, RedHat, Debian

CVE ID	TTPs	Detection	Mitigation	Patch
CVE-2025-68664*	T1059: Command and Scripting Interpreter, T1059.006: Python, T1552: Unsecured Credentials, T1552.001: Credentials In Files, T1195: Supply Chain Compromise, T1195.002: Compromise Software Supply Chain	DET0516: Behavioral Detection of Command and Scripting Interpreter Abuse, DET0195: Behavioral Detection of System Network Configuration Discovery	M1038: Execution Prevention, M1022: Restrict File and Directory Permissions, M1016: Vulnerability Scanning	LangChain
CVE-2025-67268	T1190: Exploit Public-Facing Application, T1499: Endpoint Denial of Service, T1499.004: Application or System Exploitation, T1068: Exploitation for Privilege Escalation	DET0080: Exploit Public-Facing Application – multi-signal correlation (request → error → post-exploit process/egress), DET0499: Behavioral Detection of Fallback or Alternate C2 Channels, DET0068: Detection Strategy for T1505.004 - Malicious IIS Components	M1030: Network Segmentation, M1038: Execution Prevention, M1026: Privileged Account Management	Alpine Linux, Debian, RedHat, Ubuntu

References

https://lore.kernel.org/linux-cve-announce/

https://github.com/leonov-av/linux-patch-wednesday

https://www.debian.org/security/#DSAS

https://lists.ubuntu.com/archives/ubuntu-security-announce/

https://access.redhat.com/security/security-updates/

https://lists.opensuse.org/archives/list/security-announce@lists.opensuse.org/

https://hivepro.com/threat-advisory/the-gogs-blind-spot-a-zero-day-fueled-mass-compromise/

https://hivepro.com/threat-advisory/langgrinch-critical-langchain-serialization-flaws-enable-secret-exfiltration/

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a free demo with **HivePro Uni5**: Threat Exposure Management Platform.

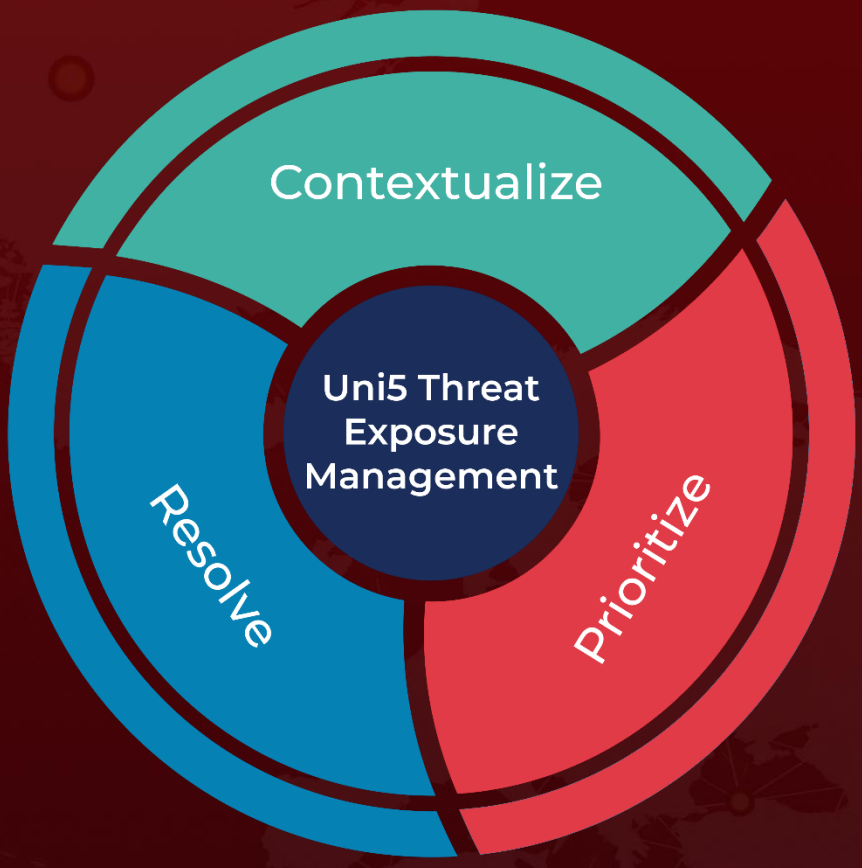

REPORT GENERATED ON

January 27, 2026 . 10:30 AM

More at www.hivepro.com