

# Linux Kernel Vulnerability CVE-2024-26582: Local Privilege Escalation and Root Shell Analysis

Author: Synthex

Site: denizhalil.com

CVSS 7.8 HIGH

Date: July 2026

## INTRODUCTION

The Linux kernel serves as the core of modern operating systems, handling resource management and enforcing fundamental security boundaries. When memory management flaws manifest within this privileged layer, they can render user-space security controls—such as firewalls, container isolation, and traditional access control lists—entirely ineffective. This article provides a rigorous technical analysis of **CVE-2024-26582**, a high-severity vulnerability discovered in the Linux kernel's native Transport Layer Security (TLS) parsing component ( `net/tls` ). While initially recognized for its potential to cause a systemic Denial of Service (DoS) via abrupt kernel panics, security researchers successfully demonstrated how local attackers can systematically weaponize this flaw. By manipulating the compromised subsystem, an unprivileged user can completely bypass modern operating system mitigations to achieve full administrative (root) code execution, effectively granting them unrestricted control over the host environment.

## LEARNING OBJECTIVES

Upon completing this article, you will be able to:

- Grasp the mechanics of Kernel-level Use-After-Free (UAF) vulnerabilities.
- Identify affected systems and apply appropriate vendor patches and mitigation workflows.
- Understand the underlying architectural root cause and tracking logic behind CVE-2024-26582.
- Learn how Local Privilege Escalation (LPE) exploits manipulate kernel structures and bypass Kernel Address Space Layout Randomization (KASLR).

## WHAT IS LINUX KERNEL VULNERABILITY CVE-2024-26582

**CVE-2024-26582** represents a critical structural flaw residing deep within the built-in cryptographic and networking subsystems of the Linux kernel. Specifically, the vulnerability manifests within the internal Transport Layer Security implementation, which is compiled via the `CONFIG_TLS` configuration flag. Assigned a Common Vulnerability Scoring System (CVSS) base score of **7.8 (High)**, this defect shifts the traditional understanding of network layer bugs. While many network-related vulnerabilities require remote traffic injection, this specific flaw exposes a fundamental breakdown in how the kernel tracks and references memory blocks internally during complex, multi-threaded cryptographic workflows.

The core danger of this vulnerability lies in its ability to facilitate Local Privilege Escalation (LPE). A local attacker possessing minimal, unprivileged access to a system—such as a standard user account or a restricted service account—can systematically trigger a reference counting error in kernel memory. Because the malicious manipulation occurs entirely within the highly privileged kernel space (Ring 0), it completely bypasses traditional application-level security controls, including user-space sandboxes, container runtimes, and standard Endpoint Detection and Response (EDR) agents. Once the kernel's memory management is compromised, the boundaries separating user privileges dissolve entirely.

From an operational standpoint, the weaponization of this vulnerability allows an attacker to transition from a restricted environment into an interactive root shell. By carefully structuring memory allocation patterns, exploit scripts turn what was originally classified as a standard kernel panic or Denial of Service (DoS) vulnerability into a precise write-primitive. This enables the unauthorized modification of process credentials directly in the system's memory. Consequently, this bug has become a focal point for security administrators, as it demonstrates that even optimized, native kernel-level acceleration features like kernel TLS (kTLS) can introduce severe architectural risks if memory states are not perfectly synchronized.

### Key Characteristics of the Threat:

- **Privileged Execution Context:** The flaw operates entirely within the kernel space, meaning any code executed via this exploit runs with the absolute highest system authority, rendering user-space access controls obsolete.
- **Local Exploitation Vector:** Attackers must already have local code execution capability on the target machine (e.g., via SSH, a web application exploit, or a compromised container) to interact with the vulnerable network sockets.
- **Bypass of Modern Isolation:** Standard defensive boundaries, such as Docker/Kubernetes container perimeters or systemd sandboxing, fail to contain the exploit because the attack directly subverts the underlying shared host kernel.
- **Transition to Arbitrary Read/Write:** The underlying Use-After-Free (UAF) mechanism provides a predictable primitive that allows attackers to map out the kernel memory space, defeat KASLR, and overwrite critical system pointers.

## EXAMPLE TERMINAL 1: IDENTIFYING A VULNERABLE KERNEL ENVIRONMENT

```
user@vulnerable-host:~$ uname -r
6.6.15-generic

user@vulnerable-host:~$ lsmod | grep tls
tls                114688  0

user@vulnerable-host:~$ zcat /proc/config.gz | grep CONFIG_TLS
```

```
CONFIG_TLS=y
CONFIG_TLS_DEVICE=y
```

## EXAMPLE TERMINAL 2: EXECUTING THE LOCAL PRIVILEGE ESCALATION POC

```
user@vulnerable-host:~$ id
uid=1001(user) gid=1001(user) groups=1001(user)

user@vulnerable-host:~$ ./cve_2024_26582_poc
[*] Linux Kernel CVE-2024-26582 LPE Exploit
[*] Grooming kernel heap via kTLS sockets...
[*] Asynchronous decryption triggered, forcing partial read state.
[*] Leaking kernel base address... KASLR bypass successful!
[*] Modifying target process credential structure in memory...
[*] Escalating privileges to root. Spawning shell...

root@vulnerable-host:~# id
uid=0(root) gid=0(root) groups=0(root)
```

## TECHNICAL DETAIL: HOW THE VULNERABILITY WORKS

The root cause of CVE-2024-26582 stems from a complex race-like condition yielding a Use-After-Free (UAF) flaw during asynchronous cryptographic decryption paired with partial read operations within the kernel's `net/tls` module. When applications use kernel-level TLS (kTLS) to offload cryptographic tasks, incoming records are split into scatterlists and dispatched to an asynchronous decryption pipeline. Under normal circumstances, the kernel meticulously increments and decrements a page's reference counter ( `page->_refcount` ) to track every thread or subsystem actively relying on that block of memory. However, when a user space application issues a partial read command while an asynchronous cryptographic operation is actively pending on the exact same packet buffer, a structural desynchronization occurs.

This desynchronization directly leads to a catastrophic reference counting failure during data ingestion. Specifically, when processing incoming ciphertext, the function `tls_decrypt_sg` allocates memory pages and prepares a separate buffer ( `clear_skb` ) to store the decrypted cleartext. Crucially, the function fails to increment the active reference counter on the underlying physical pages supporting this cleartext structure before handling the partial read state. As soon as the asynchronous hardware engine or cryptographic worker thread finishes processing the raw data block, execution automatically jumps to the decryption callback routine, `tls_decrypt_done` . This callback executes a standard cleanup protocol by calling `put_page()` , which explicitly tells the kernel that the cryptographic subsystem is done using that page. Because the initial reference counter was never incremented to account for the ongoing socket-read operation, this drop brings the page's global counter straight to zero, prompting the kernel page allocator to instantly mark that specific block of physical memory as free, unallocated, and open to redistribution.

With the page formally marked as free, a "dangling pointer" scenario is established within the active networking queue. The kernel's socket receive engine continues executing its regular loop, subsequently calling `process_rx_list` to fetch and return the decrypted bytes to the waiting user application. When `process_rx_list` attempts to parse and pull data from the memory address stored in the `clear_skb` pointer, it is accessing a region of the kernel heap that has already been recycled or left completely unmapped. Attackers can reliably weaponize this dangling reference by utilizing sophisticated heap grooming techniques—spraying the kernel heap (`kmalloc`) with customized objects immediately after the early free occurs. By substituting the legitimate cleartext page with a malicious payload containing forge-crafted kernel structures, the attacker forces the kernel to execute arbitrary memory operations, ultimately overriding security credentials and granting full privilege escalation.

### Anatomy of the Exploitation Path:

- **Reference Counting Failure:** The function `tls_decrypt_sg` fails to claim an active reference count on pages derived from `clear_skb` during async processing, setting the stage for a premature lifecycle termination.
- **Premature Memory Release:** The async callback routine `tls_decrypt_done` natively invokes `put_page()`, lowering the tracking counter to zero and causing the kernel allocator to immediately reclaim active pages while they are still queued for reading.
- **Dangling Pointer Exploitation:** The function `process_rx_list` processes the active socket receive queue, reading from the now-reclaimed addresses and allowing attackers to intercept or corrupt data via a standard Use-After-Free primitive.
- **Privilege Architecture Subversion:** By grooming the kernel heap to control the contents of the freed memory space, attackers leak critical kernel pointers to completely bypass KASLR and execute a targeted write primitive to rewrite process credentials to root status.

## EXAMPLE TERMINAL 3: MONITORING KERNEL HEAP STATE AND UAF EXPLOITATION VIA KASAN

```
[ 142.108432] =====  
[ 142.109121] BUG: KASAN: use-after-free in process_rx_list+0x1fc/0x430 [tls]  
[ 142.109843] Read of size 8 at addr ffff888104b2a000 by task poc_exploit/1245  
[ 142.110502]  
[ 142.111189] CPU: 2 PID: 1245 Comm: poc_exploit Not tainted 6.6.15 #1  
[ 142.111854] Call Trace:  
[ 142.112521] <TASK>  
[ 142.113192] dump_stack_lvl+0x45/0x60  
[ 142.113854] print_report+0xc8/0x600  
[ 142.114521] kasan_report+0xa3/0xe0  
[ 142.115201] process_rx_list+0x1fc/0x430 [tls]  
[ 142.115852] tls_sw_recvmsg+0x3b8/0x910 [tls]  
[ 142.116510] sock_recvmsg+0xbc/0xf0
```

```

[ 142.117189] __sys_recvfrom+0x120/0x1e0
[ 142.117854] __x64_sys_recvfrom+0x3a/0x50
[ 142.118521] do_syscall_64+0x3b/0x90
[ 142.119184] </TASK>
[ 142.119851]
[ 142.120512] Allocated by task 1245:
[ 142.121182] kmem_cache_alloc_node+0x12a/0x320
[ 142.121854] tls_decrypt_sg+0x184/0x520 [tls]
[ 142.122511] tls_sw_sendpage+0x210/0x6c0 [tls]
[ 142.123189]
[ 142.123854] Freed by async worker task 48:
[ 142.124510] kmem_cache_free+0x9a/0x2d0
[ 142.125191] tls_decrypt_done+0x8f/0x130 [tls]
[ 142.125854] crypto_finalize_request+0x52/0x90
[ 142.126521] =====
[ 142.127190] [*] Exploit trapped KASAN trace. Continuing heap payload spray...
[ 142.127854] [*] Successfully hijacked struct cred. New UID: 0 (root)

```

## AFFECTED SOFTWARE & PLUGINS

This security flaw represents an inherent logic error embedded directly within the core architectural subsystems of the Linux kernel itself. Because it is natively rooted inside the kernel's memory allocation and cryptographic tracking layers, it functions completely independently of any external enterprise applications, third-party user-space software, or optional server plugins. The entire vulnerability is self-contained within the native kernel code, meaning that any system running an unpatched version of the operating system's core remains structurally exposed. Consequently, typical surface-level defenses—such as application patching, disabling specific web server extensions, or modifying high-level application software configurations—provide absolutely zero protection against this specific vector.

To understand the exact surface area of exposure, administrators must look specifically at how the underlying operating system was compiled and deployed. The risk vectors are entirely defined by whether specific kernel features were built into the active image and whether the system architecture supports hardware or software-based Transport Layer Security acceleration. In enterprise data centers, virtualization clusters, and containerized cloud environments, the vulnerability presents a widespread risk because multiple distinct instances or container pods often share the exact same underlying host kernel. If that shared host kernel satisfies the vulnerable configuration parameters, every isolated user-space application running on top of it becomes susceptible to a complete breakout if a local asset is compromised.

### Scope of System Vulnerability:

- **Native Kernel Component:** The vulnerability resides permanently inside the `net/tls` subsystem, meaning it is tied completely to the operating system's low-level networking architecture rather than external modules.

- **Compilation Prerequisite:** Exposure is strictly dependent on the `CONFIG_TLS` flag being enabled during the kernel compilation phase, a setting that is standard across almost all major Linux enterprise and desktop distributions.
- **Affected Kernel Version Ranges:** A wide variety of mainline and stable kernel branches across the modern `5.x` and `6.x` series are vulnerable, specifically any unpatched upstream builds distributed prior to the mid-February 2024 security rollout.
- **Target Version Thresholds:** Specific upstream milestone versions that must be updated include any builds compiled prior to versions `6.7.6`, `6.6.18`, and `6.1.79` respectively, where the definitive reference counting patch was formally integrated.

## CONCLUSION

CVE-2024-26582 illustrates how a seemingly minor reference counting oversight deep within an isolated kernel subsystem can be chained together to achieve complete system compromise. The transition of this defect from a simple, localized Denial of Service (DoS) bug into a precise local privilege escalation vector underscores the volatile nature of kernel-level vulnerabilities. When low-privileged or local users can exploit structural memory flaws to consistently break out of unprivileged boundaries and establish full interactive root shells, the entire integrity of the operating system is undermined. This capability presents a direct and severe threat to multi-tenant servers, high-density cloud container environments, and critical enterprise infrastructures where logical isolation is the primary line of defense.

The structural breakdown of container boundaries during such an exploit highlights a broader architectural lesson for modern cybersecurity engineering: user-space isolation is only as secure as the underlying kernel that facilitates it. Technologies like Docker, Kubernetes, and specialized sandboxes rely fundamentally on the host kernel to enforce namespace separation and control groups (cgroups). When a vulnerability like this Use-After-Free bypasses those primitives at Ring 0, the entire logical defense matrix collapses simultaneously. Consequently, security teams must treat kernel-level memory management bugs not merely as stability flaws, but as high-priority risk vectors capable of granting adversaries unrestricted access to host resources, underlying virtual file systems, and sensitive cross-tenant data.

Ultimately, the definitive remediation for this vulnerability requires upgrading systems to a stable kernel version that natively integrates the official upstream fix: `net: tls: fix use-after-free with partial reads and async decrypt`. System administrators and DevSecOps engineers should immediately audit their active infrastructure deployments by running verification commands such as `uname -r`. To mitigate the risk completely, organizations must pull the latest security backports directly from their respective Linux distribution vendors (such as Red Hat, Ubuntu, or Debian), apply the updated kernel packages, and schedule a system reboot to guarantee that the unpatched memory management logic is fully purged and the new, secure images are live.