CSAI Foundation | Cloud Security Alliance

CISA's 72-Hour Window for Three Linux Kernel Flaws

2026-09-20

AI-assisted Rapid Research

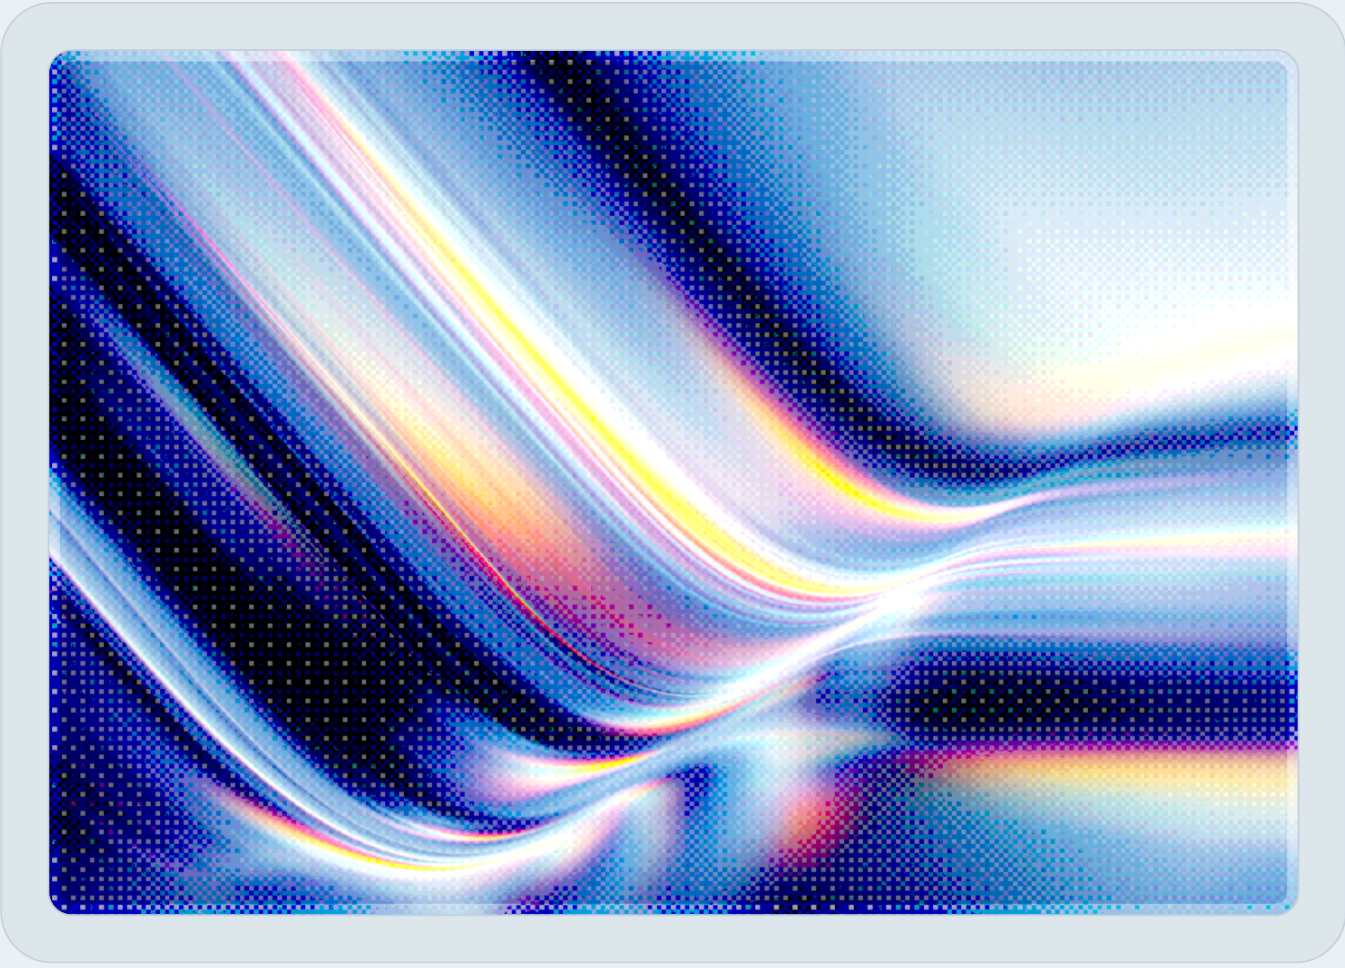

CSAI

CSA cloud security alliance®

Key Takeaways

On September 18, 2026, CISA added three Linux kernel vulnerabilities to its Known Exploited Vulnerabilities (KEV) catalog and, invoking Binding Operational Directive 26-04, gave Federal Civilian Executive Branch agencies until September 21, 2026 to remediate them, a window of roughly 72 hours [1][2][4][5]. The three flaws sit in unrelated kernel subsystems: CVE-2025-39682 is an improper-condition-check defect in the kernel TLS (kTLS) receive path, CVE-2026-53266 is an out-of-bounds write in the netfilter bridge ebtables SNAT target, and CVE-2025-39964 is a race condition in the AF_ALG cryptographic socket interface [1][2][9][10].

CISA has not disclosed who is exploiting these vulnerabilities, which organizations have been affected, or what specific techniques attackers are using; the KEV listing itself, rather than any published forensic detail, is what confirms exploitation has occurred in the wild [2][4]. Notably, CISA again flagged all three entries as requiring forensic triage, meaning covered agencies must investigate potentially exposed assets for evidence of compromise rather than treat a completed patch as the end of the response [2][5]. That requirement, layered onto an already compressed remediation window, signals that CISA regards these as vulnerabilities that may already have been used against specific targets rather than opportunistic, broad-scale scanning activity.

Severity scoring for two of the three CVEs varies across sources: CVE-2025-39682 carries a 7.1 rating from NVD against a 9.8 rating from the CVE.org CNA record, and CVE-2025-39964 carries a National Vulnerability Database and Red Hat score of 5.5 against a 7.8 rating from the CVE.org CNA record, illustrating how base CVSS metrics can undersell or oversell real-world risk for a given deployment context [1][9][13][14]. This is the third Linux kernel local-privilege-escalation or memory-corruption cluster that CSA's AI Safety Initiative has tracked since June 2026, following the pedit COW and DirtyClone exploits and the GhostLock container-escape flaw, and it reinforces a pattern of accelerating kernel-disclosure cadence testing organizations' patch discipline, a cadence CISA has separately attributed, at the policy level, to AI-accelerated exploit development [6][7].

For federal agencies and the enterprises that increasingly treat KEV additions and BOD 26-04 timelines as a de facto baseline [7], this event is also a live test of the four-variable prioritization model BOD 26-04 introduced in June 2026, since it replaces CVSS-centric triage with a framework built around asset exposure, KEV status, exploit automation, and technical impact [8]. Organizations that have not yet

operationalized a rapid, non-CVSS-dependent way to identify which systems actually use kTLS, ebtables SNAT rules, or AF_ALG sockets will find this particular KEV addition difficult to act on within a three-day window regardless of their general patch cadence.

Background

The Known Exploited Vulnerabilities catalog and the compressed patch mandate underlying this event both originate in a broader shift in CISA's approach to vulnerability prioritization. On June 10, 2026, CISA issued Binding Operational Directive 26-04, "Prioritizing Security Updates Based on Risk," which superseded the agency's prior directives on vulnerability remediation and shortened the federal patch timeline to as little as three calendar days for vulnerabilities meeting the highest tier of a four-variable risk matrix: confirmed exposure of the affected asset to the internet, presence in the KEV catalog, availability of automated exploitation tooling, and a technical impact rated as complete compromise [7] [8]. CSA's AI Safety Initiative has previously analyzed this directive at length, noting that CISA's stated rationale rests on evidence of AI-accelerated exploit development compressing the gap between disclosure and weaponization to hours rather than weeks [7].

The September 18 KEV additions arrived through two separate CISA alerts published the same day: one covering CVE-2025-39682 alone, and a second covering both CVE-2026-53266 and CVE-2025-39964 together [4][5]. All three affect the mainline Linux kernel rather than a specific distribution's packaging, which means the exposure spans the entire ecosystem of RHEL, Ubuntu, Debian, and other kernel-derived distributions to varying degrees depending on which kernel version and configuration each ships. CVE-2025-39682 concerns kernel TLS, a feature that offloads TLS record processing into kernel space for efficiency; the flaw allows a zero-length record retrieved from an internal receive queue to bypass the kernel's intended record-type handling, corrupting the assumptions later code makes about whether data has already been decrypted and reassembled [1][2][13]. Because kTLS is invoked from the socket receive path, Red Hat's advisory notes that the affected logic can be triggered by data arriving over the network on systems that have kTLS enabled and attached to TCP sockets, making internet-facing services that rely on kTLS acceleration a priority for review [13].

The second and third CVEs sit deeper in more specialized kernel subsystems. CVE-2026-53266 affects the ebtables SNAT target in the netfilter bridge code, which is used to rewrite network addresses on bridged traffic; when an optional feature rewrites the ARP sender hardware address, the kernel can call a memory-write function without first confirming that the destination memory range is actually writable, and if that memory happens to be backed by a special class of socket-buffer fragment tied to a file imported via a splice operation, the write can land directly in kernel-managed page memory rather than the intended buffer [1][10]. Exploiting this path requires a system configured with specific bridge netfilter

rules and, per some technical analyses, capabilities equivalent to CAP_NET_ADMIN, narrowing but not eliminating the population of affected systems, since containerized and network-appliance workloads frequently grant that capability to service processes [1][2]. CVE-2025-39964, by contrast, is a purely local, low-privilege issue: the AF_ALG interface, which lets user-space programs use kernel cryptographic primitives directly, did not previously enforce exclusive ownership of the write path on a given socket, so two threads or processes writing concurrently could interleave their operations and leave the kernel's internal cryptographic state inconsistent, corrupting output, crashing the system, or degrading the integrity of cryptographic results depending on the specific race outcome [2][9].

Security Analysis

The most operationally significant feature of this KEV addition is not any single CVE's technical severity but the combination of three unrelated attack surfaces, a compressed remediation deadline, and a mandatory forensic-triage requirement applied uniformly across all three. Unlike a single high-profile kernel flaw such as GhostLock, which concentrated remediation effort on one well-understood exploitation path, this addition forces security teams to simultaneously assess exposure across kTLS usage, bridge netfilter configuration, and AF_ALG consumption, three subsystems that are unlikely to be inventoried together in most organizations' existing asset-management tooling. In CSA's assessment, many organizations do not track which of their Linux hosts have kTLS enabled at the socket level, which containerized network functions apply ebtables SNAT rules with ARP rewriting, or which applications invoke AF_ALG directly, since these are lower-level kernel features typically exercised transparently by libraries or specialized workloads rather than tracked as discrete software inventory items. That visibility gap, rather than patch availability itself, is likely to be the binding constraint on federal agencies' ability to meet the September 21 deadline, and it illustrates a broader challenge with BOD 26-04's exposure-based prioritization model: identifying the asset-exposure variable requires configuration-level visibility that many vulnerability management programs were not built to provide.

The scoring divergence across CVE-2025-39682 and CVE-2025-39964 is itself a useful case study for the shift BOD 26-04 represents. CVE-2025-39682 carries a 7.1 rating from NVD against a 9.8 rating from the CVE.org CNA record, and CVE-2025-39964 carries a 7.8 rating from the CVE.org CNA record against a lower 5.5 rating from both NVD and Red Hat, differences that likely stem from disagreement over attack vector, required privileges, and the practical reachability of the vulnerable code path rather than any factual dispute about the underlying bug [1][9][13][14]. A CVSS-only triage process applied naively to this KEV addition could plausibly deprioritize CVE-2025-39964 as a secondary concern relative to a 9.8-rated flaw, even though CISA has assigned it the identical 72-hour remediation window and identical forensic-triage requirement as the other two. This is precisely the scenario BOD 26-04's

four-variable model was designed to correct: KEV status and confirmed exploitation, not the base CVSS score, are what drove the compressed timeline here, and organizations still running CVSS-anchored triage risk misjudging relative urgency within this specific KEV addition if they rely on score alone.

The mandatory forensic-triage requirement also deserves attention distinct from the patch action itself. CISA's implementation guidance for BOD 26-04 establishes forensic triage as a discrete, time-bound activity: agencies must look for indicators that an affected asset was actually compromised before or during the exploitation window, not simply confirm a patch has been applied [8]. For the three vulnerabilities in this KEV addition, meaningful forensic triage is complicated by the fact that CISA has not published exploitation indicators, attacker infrastructure, or a description of observed attack techniques, leaving organizations to rely on generic kernel-level compromise indicators such as unexpected privilege escalation events, anomalous process ancestry, or evidence of kernel memory corruption rather than any vulnerability-specific signature. This gap between the forensic-triage mandate and the available threat intelligence to act on it can recur in KEV additions that lack public exploitation detail, and it places a premium on organizations already having baseline behavioral monitoring in place rather than trying to build vulnerability-specific detection after the fact.

Finally, this event extends a pattern CSA's AI Safety Initiative has now documented three times in four months: a serious, actively or plausibly exploited Linux kernel vulnerability cluster surfacing on a cadence that, in CSA's assessment, outpaces many organizations' kernel patch review cycles. The pedit COW and DirtyClone exploits in June and the GhostLock container-escape flaw in July both involved local privilege escalation through page-cache or memory-corruption primitives that had gone unnoticed in mature, widely audited kernel code for years, in GhostLock's case for roughly fifteen years [6][11]. The recurrence of this pattern, three distinct kernel vulnerability clusters warranting rapid-response guidance over the same four-month span, suggests that kernel-level vulnerability management deserves the same continuous, tracked discipline that, in CSA's assessment, most organizations already apply to internet-facing application vulnerabilities, rather than being treated as a lower-frequency, lower-priority category because exploitation has historically required local access.

Recommendations

Immediate Actions

Federal agencies and organizations that treat CISA KEV timelines as a baseline should immediately inventory Linux and Linux-derived container hosts for kernel versions vulnerable to CVE-2025-39682, CVE-2026-53266, and CVE-2025-39964, prioritizing systems where the underlying feature is actually enabled: kTLS attached to TCP sockets, bridge netfilter configurations using ebtables SNAT with ARP

rewriting, and any workload invoking AF_ALG directly. Distribution-specific patches are already available; Red Hat has published RHSA-2025:16880 and RHSA-2025:16904 for CVE-2025-39682 and RHSA-2026:36645, RHSA-2026:39082, and RHSA-2026:39083 for CVE-2026-53266, and organizations running Red Hat Enterprise Linux or RHEL-derived platforms should apply these directly rather than waiting for a bundled kernel update cycle [3]. Because a package update alone does not remediate a running kernel, teams should verify actual kernel version post-patch and confirm affected systems have been rebooted before closing out remediation. Consistent with CISA's forensic-triage requirement, security teams should also begin reviewing available logs and endpoint telemetry on internet-facing systems using kTLS and on hosts with elevated netfilter capabilities for indicators of prior compromise, even in the absence of vulnerability-specific threat intelligence from CISA.

Short-Term Mitigations

Where immediate patching of the full fleet is not feasible within the compressed window, interim mitigations can reduce exposure on a per-subsystem basis. Organizations can disable kTLS on systems or services that do not require the performance benefit it provides, particularly on internet-facing endpoints, and can review and remove bridge netfilter ebtables SNAT rules that perform ARP sender hardware-address rewriting where that specific feature is not operationally required, since the vulnerable code path is only reachable through that configuration [2]. Where AF_ALG is not needed by any installed application, preventing the af_alg kernel module from loading closes off that attack surface without requiring a kernel patch, though this should be validated against any cryptographic libraries or container runtimes that may depend on it transparently [2]. None of these measures address the underlying kernel defects, and all three should be treated as temporary risk-reduction steps pending a full kernel update rather than substitutes for patching.

Strategic Considerations

Security teams should use this KEV addition as a forcing function to close the asset-visibility gap this analysis identified: knowing which hosts have kTLS enabled, which have specific netfilter rule configurations, and which processes use AF_ALG is a prerequisite for acting on the asset-exposure variable in BOD 26-04's four-variable model, and most vulnerability management tooling does not currently capture this level of kernel-feature configuration detail. Organizations should also build a standing, pre-authorized forensic-triage runbook for kernel-level KEV additions rather than improvising one under each 72-hour deadline, given that CISA has now applied the forensic-triage requirement across multiple unrelated kernel vulnerabilities in 2026. More broadly, given the recurrence of three serious kernel-level vulnerability clusters over roughly fourteen weeks this year, about two weeks between the pedit COW/DirtyClone and GhostLock disclosures, followed by roughly nine weeks before

this KEV addition, organizations should evaluate whether their current patch cadence and change-management processes for kernel updates specifically, as distinct from application and container-image patching, can realistically sustain three-day remediation for the highest-risk tier without requiring exceptional, ad hoc effort each time.

CSA Resource Alignment

This KEV addition is best understood in the context of CSA's existing analysis of Binding Operational Directive 26-04 itself. CSA's research note "CISA BOD 26-04: AI Threat Forces 3-Day Critical Patch Mandate" examines the four-variable risk matrix underlying the directive and argues that federal patch timelines are becoming a de facto enterprise baseline through insurance underwriting, contractual flow-downs, and audit frameworks [7]. This event is a direct, real-world application of that framework: CISA applied the identical 72-hour tier to three vulnerabilities with materially different CVSS scores, illustrating in practice the scoring divergence the CSA note anticipated as a driver of enterprise adoption of exposure-based, rather than CVSS-only, prioritization.

CSA's "Linux Kernel Root Exploits: pedit COW and DirtyClone" and "GhostLock: 15-Year-Old Linux Kernel Flaw Grants Root, Escapes Containers" research notes are the most directly relevant prior CSA work on the recurring pattern this analysis describes [6][11]. Both documents argue that kernel and driver components should be treated as first-class, urgently patched assets rather than deprioritized because exploitation nominally requires local access, and both provide distribution-specific remediation guidance and interim mitigation frameworks that remain applicable to organizations building a standing kernel-patch response process rather than reacting individually to each new KEV addition. Organizations that used either prior note to establish kernel patch SLAs or asset inventories should extend that same process to cover kTLS, netfilter, and AF_ALG configuration data specifically, since this KEV addition demonstrates that kernel exposure now spans networking and cryptographic subsystems beyond the privilege-escalation-focused scope of the earlier two events.

As a topical fallback for the broader vulnerability and patch management control expectations this event raises, the AI Controls Matrix (AICM) v1.1 provides applicable control language addressing the timely identification, tracking, and remediation of infrastructure vulnerabilities across hybrid and multi-cloud environments [12]. Organizations pursuing STAR assurance can document their response to this KEV addition, including asset inventory methodology, patch timeline, forensic-triage findings, and interim mitigations, as supporting evidence against AICM's threat and vulnerability management control expectations.

References

[1] The Hacker News. "CISA Flags Three Linux Kernel Vulnerabilities Exploited in the Wild." The Hacker News, September 2026.

[2] Cybersecurity News. "CISA Warns of Linux Kernel Vulnerabilities Actively Exploited in Attacks." Cybersecurity News, September 2026.

[3] Tech Times. "CISA Flags Three Actively Exploited Linux Kernel Flaws, Orders Federal Patch by Sunday." Tech Times, September 19, 2026.

[4] Cybersecurity and Infrastructure Security Agency. "CISA Adds One Known Exploited Vulnerability to Catalog." CISA, September 18, 2026.

[5] Cybersecurity and Infrastructure Security Agency. "CISA Adds Two Known Exploited Vulnerabilities to Catalog." CISA, September 18, 2026.

[6] Cloud Security Alliance. "Linux Kernel Root Exploits: pedit COW and DirtyClone." CSA AI Safety Initiative, June 29, 2026.

[7] Cloud Security Alliance. "CISA BOD 26-04: AI Threat Forces 3-Day Critical Patch Mandate." CSA AI Safety Initiative, June 13, 2026.

[8] Cybersecurity and Infrastructure Security Agency. "BOD 26-04: Implementation Guidance for Prioritizing Security Updates Based on Risk." CISA, June 2026.

[9] CVE Security. "CVE-2025-39964." CVE Security, September 2026.

[10] Recent Breaches. "CVE-2026-53266: Linux Kernel Out-of-Bounds Write Vulnerability." Recent Breaches, September 2026.

[11] Cloud Security Alliance. "GhostLock: 15-Year-Old Linux Kernel Flaw Grants Root, Escapes Containers." CSA AI Safety Initiative, July 14, 2026.

[12] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2026.

[13] Red Hat. "CVE-2025-39682." Red Hat Customer Portal, 2026.

[14] Red Hat. "CVE-2025-39964." Red Hat Customer Portal, 2026.

