

EOTISEC

NATIONAL SECURITY INTELLIGENCE REPORT FOR BUSINESS

EOTISEC Analytical Division

SECTION 1: REPORT IDENTIFICATION

Report Title	Public Disclosure of Unpatched Linux Kernel Vulnerabilities: ZDI Reports to Public Mailing Lists, October 2026
Report Number	EOTISEC-2026-065
Date of Report	23 October 2026
Classification	BUSINESS SENSITIVE // EOTISEC SUBSCRIBER DISTRIBUTION
Originator	EOTISEC Analytical Division
Subject	Assesses the probability and severity of enterprise risk arising from six unpatched Linux vulnerability reports Trend Micro's Zero Day Initiative posted to public kernel mailing lists on 20 and 21 October 2026, and prioritizes subscriber response.
Customer Sector	Defense Industrial Base, Critical Infrastructure, Technology and Hosting Providers, Financial Services, Maritime Commerce
Prepared By	Sam Liles, PhD, EOTISEC Analytical Division
Reviewed By	Dr. Samuel Liles, CISO / Principal
Distribution	EOTISEC Subscriber Base

SECTION 2: SCOPE AND PURPOSE

This report assesses the enterprise risk arising from six unpatched vulnerability reports that Trend Micro's Zero Day Initiative sent to public Linux mailing lists on 20 and 21 October 2026, rather than through the private coordinated channel. It answers one question. What is the probability and severity that this disclosure produces material exposure in subscriber environments, and how should subscribers prioritize their response. The source base is the public mailing-list traffic, the originator's own advisory list, established trade and security press, one vendor technical analysis, and the Linux kernel project's published policy. The report covers the five kernel reports and the one userspace report in the batch. It does not cover PwnEscape (CVE-2026-99999) or the PHP2Shell flaws (CVE-2026-109618 and CVE-2026-109619), which are separate events from the same week and are addressed under separate cover.

SECTION 3: KEY JUDGMENTS

Each judgment below is a primary analytical conclusion drawn from the source material. ICD 203 probability language applies throughout. The confidence level and its rationale follow each judgment, then the business impact.

KEY JUDGMENT 1: The two KSMBD remote code execution flaws scored 9.0 (ZDI-CAN-33774 and ZDI-CAN-33549) are unlikely to produce material exposure across most subscriber environments, because the in-kernel SMB server is not enabled by default and most enterprises serve files through Samba or Windows rather than KSMBD.

Confidence: High. *The mechanism is well understood and corroborated across the mailing-list posts, the kernel documentation, and independent press and vendor reporting. KSMBD being off by default is a stable, verifiable fact, not an inference.*

BUSINESS IMPACT: *Defense Industrial Base, Critical Infrastructure, and Maritime Commerce subscribers should treat these as an inventory question, not a patch emergency. Confirm that no host has the ksmbd module*

loaded. If the answer is already in asset data, the cost is one query. If it is not, the disclosure has exposed a visibility gap that predates it.

KEY JUDGMENT 2: The `skb_zerocopy` local privilege escalation (ZDI-CAN-33367, scored 7.8) is the most consequential of the six for subscribers running multi-tenant or shared-compute Linux, and near-term exploitation attempts against exposed hosts are likely, because a working proof of concept is public and reportedly defeats the Dirty Frag mitigation from May 2026.

Confidence: Moderate. *A public proof of concept and the mechanism are well documented, but the claim of root on mainstream distributions rests substantially on a single vendor, TuxCare, which sells kernel patching and has not yet been independently corroborated.*

BUSINESS IMPACT: *Technology and Hosting Providers and any subscriber running CI runners, build hosts, or shared login nodes are the real exposure. On these systems a local user is a customer or a job by design. Disable unprivileged user namespaces where nothing needs them and treat this lower-scored flaw as higher priority than either 9.0.*

KEY JUDGMENT 3: The clustering of these disclosures in a single week very likely reflects heightened attention and an AI-assisted discovery cadence rather than a step change in the underlying rate of Linux kernel vulnerability discovery.

Confidence: Moderate. *The originator's own advisory list shows Linux reports arriving steadily through 2026, which supports the reading that this is salience rather than a spike. The judgment depends on an inference about attention that cannot be measured directly.*

BUSINESS IMPACT: *All sectors. Leadership will ask about this because they saw a headline. The correct response is to score each report as marginal load against existing posture, not against the noise floor and not against the headline. The reverse error is as costly as the alarmist one. Dismissing the local flaw because it is only a 7.8 is the mistake that gets an exposed host rooted.*

KEY JUDGMENT 4: The absence of CVE identifiers for all six reports almost certainly degrades automated detection and patch tracking for a period of weeks, shifting the near-term burden to manual verification.

Confidence: High. *The kernel project assigns CVEs only after a fix reaches a stable tree, which is documented policy. Scanners and patch dashboards that key on CVE identifiers have nothing to match against today. This is a direct consequence, not an estimate.*

BUSINESS IMPACT: *All sectors that rely on CVE-keyed vulnerability management. The genuine additive cost of this week is not any single flaw. It is that tooling goes blind at the moment attention peaks. Plan for manual checks using module and service state rather than CVE feeds until identifiers are assigned.*

SECTION 4: SITUATION AND BACKGROUND

What was disclosed and how

On the evening of 21 October 2026 UTC, Trend Micro's Zero Day Initiative sent five Linux kernel vulnerability reports to the kernel security address and copied them to public lists including the Linux kernel mailing list, `netdev`, `linux-cifs`, and `dri-devel`.¹ A sixth report, a ConnMan userspace flaw, had gone to the public ConnMan list a day earlier.² The kernel project normally handles security reports privately so a fix can ship before the details become public. Stable-kernel maintainer Greg Kroah-Hartman replied to several of the reports questioning why they had been sent to public lists without patches, telling one reporter to fix their submission process.³

The reports credit named outside researchers, and the originator's own advisory page lists each as reported on 21 October with a public disclosure deadline of 18 February 2027.⁴ The kernel's security policy, updated in May 2026, states that a bug found with AI assistance should be treated as already public and sent to subsystem lists, because such bugs tend to surface from several researchers at once. Whether this batch followed that policy or was a routing error is not resolved in the source material.³

The six reports

Two reports affect KSMDB, the SMB server built into the kernel, and the originator scored both 9.0. ZDI-CAN-33774 is a race in how KSMDB assigns its protocol handler. A burst of short SMB1 packets followed by an SMB2 packet in one TCP write causes a worker to handle SMB1 data with an SMB2 handler, reading past the end of the shorter buffer. It needs no authentication and reaches any host running KSMDB on port 445.⁵ ZDI-CAN-33549 is an integer underflow in the SMB Direct path. It needs an RDMA-capable interface, a writable share, and a client negotiating SMB Direct, which narrows its reach considerably.⁵

ZDI-CAN-33367, scored 7.8, is a local privilege escalation in the networking helper `skb_zerocopy`. It lets an unprivileged user write into the page cache of a file it can only read, the same class as Dirty Pipe. The public exploit needs unprivileged user namespaces and a set of auto-loading modules. The remaining reports are ZDI-CAN-34068, an out-of-bounds write in the VMware virtual GPU driver scored 8.8, and ZDI-CAN-34076, an uninitialized memory read in the NVMe target scored 8.6, both public as titles only. ZDI-CAN-32005 is a stack buffer overflow in ConnMan, a userspace network manager, not a kernel bug.⁶

Fix and reproduction status

No report has a CVE identifier, because the kernel project assigns those only after a fix reaches a stable tree.⁵ The KSMDB maintainer stated that ZDI-CAN-33774 is already fixed in his maintainer tree as commit 68f410faedaf, which has not reached mainline. The SMB Direct maintainer believes a branch aimed at Linux 7.5 covers ZDI-CAN-33549, unconfirmed. A fix for the `skb_zerocopy` flaw was posted to netdev for review on 22 October and has not merged.⁶ TuxCare reported reproducing the `skb_zerocopy` exploit to a root shell on AlmaLinux 8 through 10, CloudLinux 9 and 10, Debian 12 and 13, and Ubuntu 24.04 and 26.04, and getting past Ubuntu's AppArmor restriction. That reproduction has not been independently confirmed.⁷

SECTION 5: ANALYSIS

Assumptions

Assumption 1: KSMDB remains disabled by default across subscriber distributions and no subscriber has enabled it without inventory visibility. This is supported by distribution defaults and by the separation between KSMDB and Samba. It fails if a subscriber enabled KSMDB for high-performance file sharing and did not record it.

Assumption 2: TuxCare's root reproduction on mainstream distributions is accurate and representative of default configurations. It is supported by a detailed public write-up. It fails if independent reproduction cannot match it or shows the result depends on non-default conditions.

Assumption 3: Subscriber hosts that run untrusted local workloads also have the exploit prerequisites present, meaning unprivileged user namespaces enabled and the required modules loadable. It is supported by common defaults on Debian and Ubuntu. It fails where namespaces are disabled or the host runs no untrusted local code.

LINCHPIN ASSUMPTION: Assumption 3 is the linchpin. The entire severity case for Key Judgment 2 rests on the prerequisites being present on hosts that also run untrusted local code. Where user namespaces are disabled or the workload is single-purpose and trusted, the exposure collapses toward zero regardless of the public exploit. This assumption should be tested against the actual configuration of each exposed host and revisited at each production cycle.

Analysis of Alternatives

For Key Judgment 1, the primary alternative is that KSMDB exposure is wider than assumed because subscribers enabled it for performance without tracking it. We assess this alternative as unlikely, because KSMDB adoption remains low and the enterprise default is Samba or Windows file services.

For Key Judgment 2, the primary alternative is that the `skb_zerocopy` exploit does not generalize beyond TuxCare's test conditions and real-world exploitation stays rare. We assess this alternative as roughly even chance, because a single vendor's reproduction is the only evidence and independent confirmation is still absent.

For Key Judgment 3, the primary alternative is that the week marks a genuine acceleration in kernel vulnerability discovery driven by AI tooling rather than an attention artifact. We assess this alternative as unlikely in the near term, though the kernel team's own policy change acknowledges that AI-found bugs surface in clusters, so the rate question is worth monitoring.

For Key Judgment 4, the primary alternative is that CVE identifiers are assigned quickly enough that tooling blindness is negligible. We assess this alternative as unlikely, because CVE assignment follows a stable-tree fix and no fix has merged for the two most severe reports.

Implications by Customer Sector

Defense Industrial Base: The direct exposure is low because standard file services do not use `KSMDB`, but the inventory question carries weight under federal contracting obligations. Being able to show, from asset data, that no production host loads the `ksmbd` module is the deliverable here. If that cannot be answered quickly, the gap is a finding in its own right.

Critical Infrastructure: Long-lived Linux hosts and slow patch cycles mean the `skb_zerocopy` flaw and the title-only `vmwgfx` and `nvmet` reports should be tracked even where immediate exposure is unclear. Prioritize any host that runs untrusted or third-party local code.

Technology and Hosting Providers: This sector carries the real risk. Shared hosting, CI runners, and login nodes host untrusted local users by design, which is exactly the condition the `skb_zerocopy` exploit needs. Disable unprivileged user namespaces where nothing depends on them, after testing, and queue these hosts first for the kernel fix.

Financial Services: Exposure tracks the compute model. Dedicated application servers running trusted code are low priority. Shared Kubernetes nodes and developer platforms that permit user namespaces are higher. Use the module and namespace checks rather than CVE feeds, which are blind until identifiers exist.

Maritime Commerce: Shipboard and shoreside Linux systems rarely run `KSMDB` or multi-tenant workloads, so direct exposure is low. The practical action is to confirm that remote access to any Linux host does not expose port 445 to untrusted networks and to fold the kernel fix into the normal maintenance window.

Changes from Prior Assessment

No prior EOTISEC assessment exists on this topic. This report establishes the baseline for the October 2026 ZDI public-disclosure event and for the `skb_zerocopy` local privilege escalation. Future reporting should track the merge of the mainline fixes, the assignment of CVE identifiers, and any independent reproduction of the `skb_zerocopy` exploit.

SECTION 6: INFORMATION GAPS AND COLLECTION REQUIREMENTS

1. Gap: No CVE identifiers are assigned to any of the six reports.

Analytical significance: Key Judgment 4 is most sensitive. CVE assignment timing determines when automated tooling can track these flaws.

Collection requirement: Monitor the kernel CVE assignment list and distribution security trackers for identifiers tied to the `KSMDB`, `skb_zerocopy`, `vmwgfx`, and `nvmet` fixes.

2. Gap: ZDI-CAN-34068 (`vmwgfx`) and ZDI-CAN-34076 (`nvmet`) are public as titles only, with no technical detail or patch.

Analytical significance: These cannot yet be scored for subscriber exposure. Their severity could raise or lower the overall assessment.

Collection requirement: Watch the dri-devel and linux-nvme lists and the originator's advisory page for the full reports.

3. Gap: The TuxCare root reproduction of ZDI-CAN-33367 has not been independently confirmed.

Analytical significance: Key Judgment 2 is most sensitive. Independent confirmation would move its confidence from Moderate toward High.

Collection requirement: Watch for reproduction from a second party without a commercial patching interest, and for distribution advisories that confirm the affected version ranges.

4. Gap: Whether the public routing was an accident or an application of the kernel's AI-found-bug policy is unresolved.

Analytical significance: Key Judgment 3 is affected. The answer shapes whether this pattern recurs and how often subscribers should expect pre-fix public disclosure.

Collection requirement: Watch for a statement from the originator and for further maintainer commentary on the reporting process.

SECTION 7: SOURCE SUMMARY STATEMENT

This assessment draws on five source types. First, primary public mailing-list traffic from the Linux kernel lists, which carries the reports and the maintainer replies and is the authoritative record of what was disclosed. These posts were accessed through secondary reporting because the kernel archive blocks automated retrieval, so the maintainer quotations are corroborated across two independent outlets rather than read first-hand. Second, the originator's own advisory page, a first-party source for the report dates, scores, discoverers, and disclosure deadlines, which carries the originator's framing and is treated accordingly. Third, established security and trade press, used for the technical breakdown and the maintainer quotations. Fourth, one industry vendor analysis from TuxCare, which supplies the only root reproduction and distribution test matrix and which carries a clear commercial interest in kernel patching. The report rests a single-source dependency on TuxCare for Key Judgment 2, and that dependency is the reason its confidence is Moderate rather than High. Fifth, the Linux kernel project's published security policy, a primary source for the CVE assignment practice and the AI-found-bug guidance. No classified sources were used. EOTISEC products use open-source reporting only.

SECTION 8: ICD 503 CONSIDERATIONS AND ATT&CK / D3FEND MAPPING

ICD 503 considerations apply because this report concerns vulnerabilities in the Linux kernel and in a userspace network service, which are components of the information systems subscribers operate and must authorize. The disclosure affects the risk posture of those systems directly.

Risk to Information Systems and Data: The KSMBD flaws put network-reachable file servers at risk where the in-kernel server is running, with pre-authentication reach on port 445 for one of them. The skb_zerocopy flaw puts any multi-tenant or shared-compute Linux host at risk of local privilege escalation to root, which breaks the separation between an unprivileged workload and the host. The title-only vmwgfx and nvmet reports cannot yet be scoped.

NIST RMF Alignment: The Monitor phase is most directly implicated, since the near-term task is detection and tracking in the absence of CVE identifiers. The Assess and Authorize phases are implicated for any host where KSMBD exposure or shared-compute use changes the system's risk determination. Organizations should note in their system authorization documentation that CVE-keyed tooling is temporarily blind to these flaws and that compensating controls are in use pending merged fixes.

Supply Chain Risk: The exposure reaches subscribers through the distribution kernel supply chain. A fix in a maintainer tree or on a mailing list is not a fix on a running host until it moves through the distribution, the package, and a reboot or live patch. Subscribers with federal contracting obligations should track the gap between upstream fix and deployed kernel as a supply chain timing risk.

ATT&CK Technique Mapping

Technique ID	Technique Name	Tactic	Application to This Report	Status
T1190	Exploit Public-Facing Application	Initial Access	ZDI-CAN-33774 is reachable before authentication against any host running the in-kernel SMB server on TCP 445. An attacker who can reach that port can trigger the dispatch race without credentials.	ASSESSED
T1210	Exploitation of Remote Services	Lateral Movement	Both KSMBD flaws allow an attacker already inside the network to pivot to a file server exposing KSMBD. The SMB Direct path (ZDI-CAN-33549) additionally requires RDMA and a writable share.	ASSESSED
T1068	Exploitation for Privilege Escalation	Privilege Escalation	ZDI-CAN-33367 lets an unprivileged local user write into the page cache of a read-only file and escalate to root. A working proof of concept is public and reportedly defeats the Dirty Frag mitigation.	REPORTED
T1588.006	Vulnerabilities	Resource Development	Full technical detail for five of the six reports is public before fixes merged. Adversaries can acquire the vulnerability details directly from the public mailing list archives.	REPORTED
T1588.005	Exploit Capabilities	Resource Development	Three reports include complete proof-of-concept source. Adversaries can acquire working exploit code without developing it themselves, compressing the time from disclosure to use.	REPORTED

The two REPORTED resource-development entries reflect that full technical detail and, for three reports, working proof-of-concept source are already public. The initial-access and lateral-movement entries are ASSESSED from the network reachability described in the reports, not from observed intrusions. No in-the-wild exploitation of these specific flaws has been reported as of 23 October 2026.

D3FEND Countermeasure Mapping

D3FEND Technique	D3FEND ID	Addresses ATT&CK	Application and Priority
Network Traffic Filtering	D3-NTF	T1190, T1210	Block TCP 445 from untrusted networks on any host where the in-kernel SMB server is loaded. Priority HIGH on hosts confirmed to run KSMBD, LOW TO MODERATE elsewhere because most subscriber file services run Samba or Windows and do not expose KSMBD at all.
System Configuration Permissions	D3-SCP	T1068	Set user.max_user_namespaces to 0 on hosts that do not need unprivileged user namespaces, and block the nlmon and act_pedit modules. This removes prerequisites the public ZDI-CAN-33367 exploit depends on. Priority HIGH on multi-tenant and shared-compute hosts, LOW on single-purpose appliances where local code is already trusted.
Mandatory Access Control	D3-MAC	T1068	AppArmor and SELinux constrain the paths an escalation can take. Priority MODERATE, not HIGH, because the public exploit was reported to bypass Ubuntu's AppArmor restriction on unprivileged user namespaces by launching through an allowed profile. Treat it as defense in depth, not a control that closes the path.

The priority designations above are grounded in report-specific conditions, not generic best practice. Network filtering is HIGH only where a host actually runs KSMBD. System configuration hardening is HIGH only on multi-tenant hosts that run untrusted local code. Mandatory access control is held at MODERATE because the public exploit was reported to bypass the relevant AppArmor restriction.

SECTION 9: ANALYTIC TRADecraft SELF-CERTIFICATION

ICD 203 Requirement	Status	Notes
Objectivity: The analysis is free of advocacy, personal preference, or policy bias.	CONFIRMED	<i>The report separates the vendor marketing claim driving the headline from the verifiable facts, and it grades severity on exposure rather than on CVSS score alone.</i>
Independence of Political Consideration: No judgment was shaped to support a particular outcome or audience preference.	CONFIRMED	<i>No judgment was adjusted to favor any vendor, distribution, or subscriber. The single-vendor root reproduction is flagged as such.</i>
Timeliness: Source material is current and the report addresses matters actionable by the customer.	CONFIRMED	<i>All source material is dated 15 to 23 October 2026. The assessment is actionable now, before CVE assignment and before fixes reach stable kernels.</i>
Based on All Available Sources: All source material provided was considered. Gaps are documented in Section 6.	CONFIRMED	<i>Primary mailing-list posts, vendor analysis, the originator's advisory list, and the kernel project's own policy were all reviewed. Collection gaps are stated in Section 6.</i>
Source Credibility Described: Source quality is addressed in Section 7.	CONFIRMED	<i>Section 7 names each source type and flags the single-source dependency on TuxCare for the root reproduction claim.</i>
Uncertainty Expressed: Probability and confidence language conforms to ICD 203 standards.	CONFIRMED	<i>Each Key Judgment carries one probability term and a separate confidence level with rationale. The two are never combined in a sentence.</i>
Assumptions Distinguished from Facts: All assumptions are labeled. Linchpin assumptions are identified.	CONFIRMED	<i>Section 5 lists three numbered assumptions with failure conditions and designates the linchpin as the presence of the exploit prerequisites on hosts running untrusted local workloads.</i>
Alternatives Incorporated: At least one alternative hypothesis is addressed for each major judgment.	CONFIRMED	<i>Section 5 addresses one alternative hypothesis for each of the four Key Judgments using ICD 203 probability language.</i>
Customer Relevance Addressed: Implications for the relevant stakeholder are stated specifically.	CONFIRMED	<i>Section 5 states separate implications for each of the five customer sectors named in Section 1.</i>
Clear and Logical Argumentation: The main analytic message is stated up front in Section 3.	CONFIRMED	<i>The four Key Judgments precede the evidence and carry the analytic message, led by the finding that the lower-scored local flaw is the one that matters most to subscribers.</i>
ATT&CK and D3FEND Mapping Completed: REPORTED and ASSESSED attributions are correctly distinguished. Priority designations are grounded in the specific incident.	CONFIRMED	<i>Section 8 marks the public proof-of-concept items REPORTED and the network-reachability inferences ASSESSED. D3FEND priorities are tied to whether a host runs KSMDB or untrusted local code, not assigned generically.</i>

SECTION 10: ENDNOTES

¹ Michael Larabel. "Several Linux Kernel Vulnerabilities Accidentally Posted To Public Kernel Mailing List." Phoronix, 22 October 2026.

² "Linux kernel: unpatched ZDI bug reports with repro details leaked to public LKML (ksmbd RCE, CVSS 9.0)." pwn.watch, updated 22 October 2026. ConnMan report posted to the public ConnMan list 20 October 2026, 00:17 UTC.

³ Swati Khandelwal. "Two CVSS 9.0 Linux KSMDB RCE Flaws Disclosed on Public Mailing Lists Before Kernel Fixes." The Hacker News, 22 October 2026. Includes the Kroah-Hartman replies and the kernel policy on AI-assisted discovery.

⁴ "Upcoming Advisories." Trend Micro Zero Day Initiative. Lists ZDI-CAN-33549, 33774, 33367, 34068, and 34076 as Linux, reported 2026-10-21, deadline 2027-02-18, and ZDI-CAN-32005 as ConnMan, reported 2026-10-19, deadline 2027-02-16.

⁵ The Hacker News, 22 October 2026, op. cit. KSMDB mechanism, pre-authentication reach on port 445, SMB Direct conditions, and the kernel practice of assigning CVEs only after a stable-tree fix.

⁶ pwn.watch, op. cit. Report inventory, the skb_zerocopy prerequisites, the ConnMan classification, the ZDI-CAN-33774 maintainer-tree fix as commit 68f410faedaf, the SMB Direct for-7.5 branch, and the 22 October netdev patch posting for review.

⁷ Chris DeMars. "Leaked Linux Kernel Reports Include Two ksmbd Bugs and a Root Exploit That Beats the Dirty Frag Fix." TuxCare, 22 October 2026. Distribution test matrix, AppArmor bypass, and interim mitigations. Single-vendor source with a commercial kernel-patching interest, not independently confirmed.