SingCERT
Singapore Cyber Emergency Response Team

Security Bulletin 22 July 2026

Generated on 22 July 2026

SingCERT's Security Bulletin summarises the list of vulnerabilities collated from the National Institute of Standards and Technology (NIST)'s National Vulnerability Database (NVD) in the past week.

The vulnerabilities are tabled based on severity, in accordance to their CVSSv3 base scores:

Critical	vulnerabilities with a base score of 9.0 to 10.0
High	vulnerabilities with a base score of 7.0 to 8.9
Medium	vulnerabilities with a base score of 4.0 to 6.9
Low	vulnerabilities with a base score of 0.1 to 3.9
None	vulnerabilities with a base score of 0.0

For those vulnerabilities without assigned CVSS scores, please visit NVD for the updated CVSS vulnerability entries.

CRITICAL VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2026-56699	Wazuh Manager before 5.0.0-beta3 fails to escape the DataValue.index field when constructing OpenSearch bulk requests, allowing enrolled agents to inject arbitrary NDJSON operations. Attackers can smuggle delete, index, or update operations into bulk requests executed under the manager's admin credentials, enabling document deletion, alert tampering, and cross-agent SIEM state manipulation.	10.0	More Details
CVE-2026-16117	Impact: @fastify/http-proxy versions up to and including 11.5.0 fail to rewrite the request prefix when the prefix segment is URL-encoded. Fastify's router URL-decodes paths for route matching, but request.url retains the original encoded form, and the prefix-rewrite step uses a literal string replace against the decoded prefix. A request that encodes one or more characters of the configured prefix therefore matches the route but skips the rewrite, so the raw encoded path is forwarded to the upstream unchanged. The upstream then decodes the path and serves it, letting an attacker reach upstream paths that the proxy was configured to hide via rewritePrefix, including internal or administrative endpoints. Patches: upgrade to @fastify/http-proxy 11.6.0. Workarounds: none.	10.0	More Details
CVE-2026-44359	Meshtastic is an open source mesh networking solution. Prior to version 2.7.21.1370b23, the Meshtastic GitHub repository's main_matrix.yml workflow is triggered by pull_request_target and multiple jobs check out the attacker's fork code and execute it with access to repository secrets and elevated GITHUB_TOKEN permissions. No approval gate exists. Pull requests from external users with author_association: "NONE" triggered the CI workflow automatically. The workflow directly executes attacker-controlled files from the fork checkout. This issue could have resulted in supply chain compromise, self-hosted runner compromise, and/or repository takeover for the repo. This issue is separate from GHSA-6mwm-v2vv-pp96, which addressed a command injection via github.head_ref in the setup job of the same workflow. That fix correctly moved to environment variables. However, the more critical fork checkout vulnerability across the check, build, and build-debian-src jobs was not addressed. Version 2.7.21.1370b23 contains a patch for thie issue.	10.0	More Details
CVE-2026-45336	HireFlow is a web-based interview management system for managing candidates, scheduling interviews, and tracking hiring progress. In 1.2 and earlier, app.py assigns a hard-coded Flask secret_key used to sign session cookies, allowing unauthenticated attackers who know the public source value to forge cookies containing role=admin and user_id values and bypass authentication. The advisory lists version 1.3 as fixed.	10.0	More Details
CVE-2026-	PrestaShop ps_facetedsearch is a module that adds layered navigation filters. From 3.0.0 until 4.0.4, the ps_facetedsearch module rebuilds selected search filters from the request URL, and the value of a slider filter, price or weight, is taken from the URL without sufficient validation and stored in an internal filter-block cache where it is serialized and later read back with a raw native unserialize() in		More

54159	src/Filters/Block.php. By crafting that value, an unauthenticated attacker can smuggle a malicious serialized PHP object into the cache, and when it is deserialized, a gadget chain writes an arbitrary PHP file inside the modules/ps_facetedsearch/ directory, which is then used as a webshell to run commands on the server. This issue is fixed in version 4.0.4.	10.0	Details
CVE-2026-46339	9Router is an AI router & token saver. From 0.4.30 until 0.4.37, 9Router's src/proxy.js middleware did not protect /api/cli-tools/* and /api/mcp/*, allowing unauthenticated registration of customPlugins through src/app/api/cli-tools/cowork-settings/route.js and command execution through the MCP bridge. This vulnerability is fixed in 0.4.37.	10.0	More Details
CVE-2026-46412	@beproduct/nestjs-auth is a NestJS authentication module for BeProduct IDS (Identity Server) with OpenID Connect support. Between 2026-05-11 20:19 UTC and 22:56 UTC, an attacker used a compromised npm publish token to publish 18 malicious versions of `@beproduct/nestjs-auth` (0.1.2 through 0.1.19). The postinstall payload attempted to harvest npm tokens (from `~/.npmrc`); GitHub personal access tokens, OAuth tokens (`gho_*`), and Actions OIDC tokens; AWS credentials (from environment variables and `~/.aws/credentials`); HashiCorp Vault tokens; and other secrets present in environment variables. Version `0.1.20` is a clean republish from the original `0.1.1` source tree. Anyone who installed any version in the range `>=0.1.2 <=0.1.19` should remove the package and clean the npm cache; install the clean version; rotate every credential present in the install environment, including all npm publish tokens, all GitHub PATs and OAuth tokens, AWS access keys, HashiCorp Vault tokens, and any other secret that was in env vars or config files at install time; scan affected hosts for indicators of compromise and, if any are found, treat the host as compromised and reimage; and check committed repository history for unexpected additions in `.claude/` or `.vscode/` directories. The worm is known to commit `setup.mjs` + hook configs to PR branches via automated agent runtimes.	10.0	More Details
CVE-2026-52887	NocoBase is an AI-powered no-code/low-code platform for building business applications and enterprise solutions. Prior to 2.0.61, NocoBase @nocobase/plugin-notification-in-app-message exposed GET /api/myInAppChannels:list, where the filter[latestMsgReceiveTimestamp][$lt] value was inserted into a Sequelize.literal() template string without escaping or parameter binding, allowing a signed-up authenticated user to run stacked PostgreSQL statements and potentially execute commands with COPY ... TO PROGRAM. This vulnerability is fixed in 2.0.61.	10.0	More Details
CVE-2026-63795	In the Linux kernel, the following vulnerability has been resolved: 9p: avoid putting oldfid in p9_client_walk() error path When p9_client_walk() is called with clone set to false, fid aliases oldfid. If the walk subsequently fails after the request has been sent, the error path jumps to clunk_fid, which currently calls p9_fid_put(fid) unconditionally. This drops a reference to oldfid even though ownership of oldfid remains with the caller. If this is the last reference, oldfid can be clunked and destroyed while the caller still expects it to be valid. A later use or put of oldfid can then trigger a use-after-free or refcount underflow. Fix this by only putting fid in the clunk_fid error path when it does not alias oldfid, matching the existing guard in the error path below. This can be triggered when a multi-component walk is split into multiple p9_client_walk() calls and a later non-cloning walk fails. A reproducer and refcount warning logs are available on request.	10.0	More Details
CVE-2026-50148	Metabase is an open-source business intelligence and embedded analytics tool. From 1.54.0 until 1.54.24, 1.55.24, 1.56.25, 1.57.19, 1.58.14, 1.59.10, and 1.60.4, a Metabase user with permission to add or edit a database connection can achieve remote code execution on the Metabase server by configuring a Snowflake connection to an attacker-controlled server, because a flaw in the Snowflake JDBC driver can write arbitrary files anywhere on the Metabase host, including replacing one of Metabase's own database driver files that later executes inside the Metabase process. This issue is fixed in versions 1.54.24, 1.55.24, 1.56.25, 1.57.19, 1.58.14, 1.59.10, and 1.60.4.	10.0	More Details
CVE-2026-54051	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.9.1, the agent sandbox gates shell commands behind an allowlist (`SandboxPolicy.isCommandAllowed`), which THREAT_MODEL.md calls the main control against a compromised agent (Adversary 3.2). The allowlist glob-matches the whole command string, but `ShellExecutor` runs that string through `/bin/sh -c`. So any wildcard allow such as `git *`, `npm *` or `node *` also matches `git status; <anything>`, and a scoped command becomes arbitrary execution. The issue is fixed in v5.9.1. `ShellExecutor` now executes via `spawn(file, args, { shell: false })` using a quote-aware parsed argv, so no shell is invoked. `SandboxPolicy.isCommandAllowed` and the new `SandboxPolicy.tokenizeCommand` reject any unquoted shell metacharacter (`; & | $ ` ` ` ( ) < > { }` newline) or unterminated quote before the allowlist glob match; quoted metacharacters are preserved as literal argument data. Users should upgrade to `network-ai@5.9.1` or later. As defense in depth, avoid broad wildcard allowlist entries such as `node *` / `npm *` which are direct code execution by design.	9.9	More Details
CVE-2026-46512	Frogman provides headless PBX control through MCP and HTTP API. Prior to 1.6.2, fm_dialplan_apply accepted template parameters including greeting, dest, url, extension, code, and file, and Tools/DialplanApply.php wrote Dialplan/Templates.php output to extensions_custom.conf while only Dialplan/TemplateBase.php:38-42 sanitized contextName(), allowing a PERM_WRITE caller using confirm:true to inject arbitrary Asterisk directives such as System(), Set(SHELL(...)), Goto, or Macro. This issue is fixed in version 1.6.2.	9.9	More Details

CVE-2026-64878	Unvalidated input in asset filter parameters allows shell metacharacters to escape command argument handling, resulting in remote code execution as a low-privileged OS user via the Analysis REST endpoint.	9.9	More Details
CVE-2026-8476	IBM Langflow OSS 1.0.0 through 1.10.0 contain a critical remote code execution vulnerability in the disk-based caching mechanism. The AsyncDiskCache class uses Python's unsafe pickle.loads() function to deserialize cached objects from disk without validation, integrity verification, or authentication, enabling arbitrary code execution when malicious pickle payloads are processed. Attackers who can influence cached data through file system access, malicious workflow inputs, custom components, or API manipulation can achieve complete system compromise with the privileges of the Langflow server process.	9.9	More Details
CVE-2026-8481	IBM Langflow OSS 1.0.0 through 1.10.0 contain a critical remote code execution vulnerability in the code validation API endpoint. The POST /api/v1/validate/code endpoint accepts user-supplied Python code and executes it directly using Python's built-in exec() function without sandboxing, input validation, or privilege restrictions, enabling any authenticated user to execute arbitrary system commands with the full privileges of the Langflow server process.	9.9	More Details
CVE-2026-8635	IBM Langflow OSS 1.0.0 through 1.10.0 allows authenticated users to escalate privileges to superuser by directly manipulating the database, execute arbitrary system commands, and achieve full system compromise with Langflow service permissions.	9.9	More Details
CVE-2026-52891	Wekan is open source kanban built with Meteor. Prior to 9.07, Wekan avatar upload functionality embeds user-supplied filenames into paths later passed to child_process.exec() for MIME-type detection. Because models/avatars.js and models/fileValidation.js used a shell command with the avatar filename, shell metacharacters such as backticks and $() in the filename could execute commands on the server. This issue is fixed in version 9.07.	9.9	More Details
CVE-2026-8859	IBM Langflow OSS 1.0.0 through 1.10.0 Langflow could allow an attacker to write arbitrary files to unintended locations due to improper input validation in the APIRequest component. A path traversal vulnerability exists when the "Save to File" feature is enabled, where filenames extracted from HTTP response Content-Disposition headers are not sanitized before being joined to the temporary directory path. An attacker controlling an external HTTP server can supply crafted filename values containing path traversal sequences (e.g., ../), enabling arbitrary file writes to locations accessible by the Langflow process.	9.9	More Details
CVE-2026-64879	A filename supplied during file upload is not properly sanitized before being used in system command execution, allowing an attacker to inject shell metacharacters and achieve command injection via the audit file upload functionality.	9.9	More Details
CVE-2026-9135	IBM Langflow OSS 1.0.0 through 1.10.0 Langflow versions up to 1.9.2 (commit 94981c443d4918517b9e8163d70fc598dc33a32d) contain a code injection vulnerability in the Policies component's ToolGuard integration that bypasses the allow_custom_components=false security control. The vulnerability exists because the validation mechanism only checks the main component source code in node_template["code"]["value"] but fails to validate dynamic CodeInput fields that store generated ToolGuard Python files. Attackers can embed malicious Python code in these unvalidated dynamic fields, which are persisted in Flow.data and later executed server-side when a guarded tool is invoked through the ToolGuard runtime. This allows authenticated users with flow creation privileges to achieve arbitrary Python code execution on the backend despite custom component restrictions. The vulnerability can be escalated through cross-tenant flow manipulation via the agentic MCP update_flow_component_field tool, which accepts attacker-controlled user_id parameters, enabling attackers to inject malicious code into victim users' flows. When combined with publicly accessible flows and specific misconfigurations (AUTO_LOGIN=true, NEW_USER_IS_ACTIVE=true), the attack can be conducted with reduced authentication requirements.	9.9	More Details
CVE-2026-54052	n8n-MCP is an MCP server that provides AI assistants access to n8n node documentation, properties, and operations. Prior to 2.56.1, in HTTP mode with multi-tenancy enabled through ENABLE_MULTI_TENANT=true, n8n-mcp's local workflow version history backups were not isolated per tenant, allowing an authenticated tenant to read workflow version snapshots belonging to other tenants and delete or destroy other tenants' stored backups, including full node definitions, credential references, and authorization headers. This issue is fixed in version 2.56.1.	9.9	More Details
CVE-2026-44986	Penpot is an open-source design tool for design and code collaboration. Prior to 2.14.5, Penpot exposed teams_invitations.clj invitation tokens from create-team-invitations, embedded an existing profile id in auth.clj prepare-register-profile, and had auth.clj register-profile issue a session based on the invitation email match without password verification, allowing a registered user to take over any non-blocked profile. This issue is fixed in version 2.14.5.	9.9	More Details
CVE-2026-47392	PraisonAI is a multi-agent teams system. Prior to version 4.6.40 of PraisonAI, corresponding to version 1.6.40 of praisonaiagents, `execute_code()` in `praisonaiagents/tools/python_tools.py` (v1.6.37, subprocess sandbox mode) can be fully bypassed using `print.__self__` to retrieve the real Python `builtins` module, from which `__import__` can be extracted via `vars()` and runtime string construction. This achieves arbitrary OS command execution on the host, completely defeating the sandbox. This is a	9.9	More Details

	novel bypass that survives all patches for CVE-2026-39888 (frame traversal), CVE-2026-34938 (str subclass), and CVE-2026-40158 (`type.__getattribute__` trampoline). PraisonAI version 4.6.40 and praisonaiagents version 1.6.40 contain an updated fix.		
CVE-2026-51027	An issue in FileThingie v.2.5.7 allows a remote attacker to obtain sensitive information via the ft2.php component.	9.9	More Details
CVE-2026-64046	In the Linux kernel, the following vulnerability has been resolved: net: tls: prevent chain-after-chain in plain text SG Sashiko points out that if end = 0 (start != 0) the current code will create a chain link to content type right after the wrap link: This would create a chain where the wrap link points directly to another chain link. The scatterlist API sg_next iterator does not recursively resolve consecutive chain links. meaning this is illegal input to crypto. The wrapping link is unnecessary if end = 0. end is the entry after the last one used so end = 0 means there's nothing pushed after the wrap: end start i v v v [ ]...[ ][d][d][d][d][rsv for wrap] Skip the wrapping in this case. TLS 1.3 can use the "wrapping slot" for it's chaining if end = 0. This avoids the chain-after-chain. Move the wrap chaining before marking END and chaining off content type, that feels like more logical ordering to me, but should not matter from functional perspective.	9.8	More Details
CVE-2026-63994	In the Linux kernel, the following vulnerability has been resolved: tunnels: load network headers after skb_cow() in iptunnel_pmtud_build_icmp[v6]() Sashiko found that iptunnel_pmtud_build_icmp() and iptunnel_pmtud_build_icmpv6() were caching ip_hdr() and ipv6_hdr() before an skb_cow() call which can reallocate skb->head. Fix this possible UAF by initializing the local variables after the skb_cow() call. Remove skb_reset_network_header() calls which were not needed.	9.8	More Details
CVE-2026-64608	Heap type confusion and out-of-bounds read/write in the Apache Fory C++ implementation. When deserializing data in compatible mode, the field-skip paths do not correctly validate the declared field types against the actual data, so input with an inconsistent schema can cause type confusion and out-of-bounds memory access. Only the C++ implementation is affected; other language implementations of Apache Fory are not. This issue affects Apache Fory C++: from 0.14.0 before 1.4.0. Users are recommended to upgrade to version 1.4.0, which fixes the issue.	9.8	More Details
CVE-2026-64033	In the Linux kernel, the following vulnerability has been resolved: RDMA/rtrs: Fix use-after-free in path file creation cleanup In the error path of rtrs_srv_create_path_files(), the sysfs root folders may already have been created and srv_path->kobj may already have been initialized. If a later step fails, the cleanup currently calls kobject_put(&srv_path->kobj) before rtrs_srv_destroy_once_sysfs_root_folders(srv_path). kobject_put() may drop the last reference to srv_path->kobj and invoke the release callback, rtrs_srv_release(), which frees srv_path. The following call to rtrs_srv_destroy_once_sysfs_root_folders(srv_path) then dereferences srv_path internally to access srv_path->srv, resulting in a use-after-free. This failure path is reached before rtrs_srv_create_path_files() returns success, so the successful-path lifetime handling is not involved. Fix this by destroying the sysfs root folders before calling kobject_put(&srv_path->kobj), so srv_path is still valid while the helper accesses it. This issue was found by a static analysis tool I am developing.	9.8	More Details
CVE-2026-64035	In the Linux kernel, the following vulnerability has been resolved: igc: set tx buffer type for SMD frames Sashiko pointed out that igc_fpe_init_smd_frame() initializes igc_tx_buffer fields for an SMD skb, but does not set the buffer type: https://sashiko.dev/#/patchset/20260415025226.114115-1-kohei%40enjuk.jp Since igc_tx_buffer entries are reused, a stale XDP or XSK type can remain and make TX completion use the wrong cleanup path. Set the buffer type to IGC_TX_BUFFER_TYPE_SKB.	9.8	More Details
CVE-2026-64025	In the Linux kernel, the following vulnerability has been resolved: bpf, skmsg: fix verdict sk_data_ready racing with ktls rx sk_psock_strp_data_ready() already checks tls_sw_has_ctx_rx() and defers to psock->saved_data_ready when a TLS RX context is present, avoiding a conflict with the TLS strparser's ownership of the receive queue (commit e91de6afa81c, "bpf: Fix running sk_skb program types with ktls"). sk_psock_verdict_data_ready() has no equivalent guard. When a socket is inserted into a sockmap (BPF_SK_SKB_VERDICT) before TLS RX is configured, tls_sw_strparser_arm() saves sk_psock_verdict_data_ready as rx_ctx->saved_data_ready. On data arrival: tls_data_ready -> tls_strp_data_ready -> tls_rx_msg_ready -> saved_data_ready() = sk_psock_verdict_data_ready() -> tcp_read_skb() drains sk_receive_queue via __skb_unlink() without calling tcp_eat_skb(), so copied_seq is not advanced. tls_strp_msg_load() then finds tcp_inq() >= full_len (stale), calls tcp_recv_skb() on the now-empty queue, hits WARN_ON_ONCE(!first), and returns with rx_ctx->strp.anchor.frag_list pointing at a psock-owned (potentially freed) skb. tls_decrypt_sg() subsequently walks that frag_list: use-after-free. Apply the same fix as sk_psock_strp_data_ready(): if a TLS RX context is present, call psock->saved_data_ready (sock_def_readable) to wake recv() waiters and return immediately, leaving the receive queue untouched. TLS retains sole ownership of the queue and decrypts the record normally through tls_sw_recvmsg().	9.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix durable reconnect error path file lifetime After a durable reconnect succeeds, ksmbd_reopen_durable_fd() republishes the same ksmbd_file into the session volatile-id table. If smb2_open() then takes a later error path, cleanup first calls ksmbd_fd_put(work, fp) and then unconditionally calls ksmbd_put_durable_fd(dh_info.fp). In this case fp and dh_info.fp are the same object. The first put drops the reconnect lookup reference, but the	9.8	More

64016	final durable put can run __ksmbd_close_fd(NULL, fp). Because the final close is not session-aware, it can free the file object without removing the volatile-id entry that was just published into the session table. Use the session-aware put for the final reconnect drop when the reconnect had already succeeded and the error path is cleaning up the republished file. Earlier reconnect failures, before fp is assigned to dh_info.fp, keep using the durable-only put path.		Details
CVE-2026-64037	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mld: fix TSO segmentation explosion when AMSDU is disabled When the TLC notification disables AMSDU for a TID, the MLD driver sets max_tid_amsdu_len to the sentinel value 1. The TSO segmentation path in iwl_mld_tx_tso_segment() checks for zero but not for this sentinel, allowing it to reach the num_subframes calculation: num_subframes = (max_tid_amsdu_len + pad) / (subf_len + pad) = (1 + 2) / (1534 + 2) = 0 This zero propagates to iwl_tx_tso_segment() which sets: gso_size = num_subframes * mss = 0 Calling skb_gso_segment() with gso_size=0 creates over 32000 tiny segments from a single GSO skb. This floods the TX ring with ~1024 micro-frames (the rest are purged), creating a massive burst of TX completion events that can lead to memory corruption and a subsequent use-after-free in TCP's retransmit queue (refcount underflow in tcp_shifted_skb, NULL deref in tcp_rack_detect_loss). The MVM driver is immune because it checks mvmsta->amsdu_enabled before reaching the num_subframes calculation. The MLD driver has no equivalent bitmap check and relies solely on max_tid_amsdu_len, which does not catch the sentinel value. Fix this by detecting the sentinel value (max_tid_amsdu_len == 1) at the existing check and falling back to non-AMSDU TSO segmentation. Also add a WARN_ON_ONCE guard after the num_subframes division as defense-in-depth to catch any future code paths that produce zero through a different mechanism.	9.8	More Details
CVE-2026-64007	In the Linux kernel, the following vulnerability has been resolved: netfilter: synproxy: refresh tcphdr after skb_ensure_writable synproxy_tstamp_adjust() rewrites the TCP timestamp option in place and then patches the TCP checksum via inet_proto_csum_replace4() on the caller-supplied tcphdr pointer. Both ipv4_synproxy_hook() and ipv6_synproxy_hook() obtain that pointer with skb_header_pointer() before calling in, so it may either alias skb->head directly or point at the caller's on-stack _tcph buffer. Between obtaining the pointer and using it, the function calls skb_ensure_writable(skb, optend), which on a cloned or non-linear skb invokes pskb_expand_head() and frees the old skb->head. After that point the cached th is stale: caller (ipv[46]_synproxy_hook) th = skb_header_pointer(skb, ..., &_tcph) synproxy_tstamp_adjust(skb, protoff, th, ...) skb_ensure_writable(skb, optend) pskb_expand_head() /* kfree(old skb->head) */ ... inet_proto_csum_replace4(&th->check, ...) /* writes into freed head, or into the caller's stack copy leaving the on-wire checksum stale */ The option bytes are written through skb->data and are fine; only the checksum update goes through th and so lands in the wrong place. The result is either a write into freed slab memory or a packet leaving with a checksum that does not match its payload. Fix by re-deriving th from skb->data + protoff immediately after skb_ensure_writable() succeeds, so the subsequent checksum update targets the linear, writable header.	9.8	More Details
CVE-2026-64000	In the Linux kernel, the following vulnerability has been resolved: net: hsr: fix potential OOB access in supervision frame handling Ensure the entire TLV header is linearized before access by adding sizeof(struct hsr_sup_tlv) to the pskb_may_pull() calls. Without this, a truncated frame could cause an out-of-bounds access.	9.8	More Details
CVE-2026-64047	In the Linux kernel, the following vulnerability has been resolved: net: tls: fix off-by-one in sg_chain entry count for wrapped sk_msg ring When an sk_msg scatterlist ring wraps (sg.end < sg.start), tls_push_record() chains the tail portion of the ring to the head using sg_chain(). An extra entry in the sg array is reserved for this: struct sk_msg_sg { [...] /* The extra two elements: * 1) used for chaining the front and sections when the list becomes * partitioned (e.g. end < start). The crypto APIs require the * chaining; * 2) to chain tailer SG entries after the message. */ struct scatterlist data[MAX_MSG_FRAGS + 2]; The current code uses MAX_SKB_FRAGS + 1 as the ring size: sg_chain(&msg_pl->sg.data[msg_pl->sg.start], MAX_SKB_FRAGS - msg_pl->sg.start + 1, msg_pl->sg.data); This places the chain pointer at sg_chain(data[start], (MAX_SKB_FRAGS - msg_start + 1) .. = &data[start] + (MAX_SKB_FRAGS - msg_start + 1) - 1 = data[start + (MAX_SKB_FRAGS - start + 1) - 1] = data[MAX_SKB_FRAGS] instead of the true last entry. This is likely due to a "race" of the commit under Fixes landing close to commit 031097d9e079 ("bpf: sk_msg, zap ingress queue on psock down") Convert to ARRAY_SIZE and drop the data[start] / - start (as suggested by Sabrina).	9.8	More Details
CVE-2026-1617	Improper neutralization of special elements used in an SQL command ('SQL injection') vulnerability in Turkmesh Communication Services Inc. Turkhotspot 5651 Loglama allows SQL Injection. This issue affects Turkhotspot 5651 Loglama: from 5.1.2 before 5.1.3.	9.8	More Details
CVE-2026-63993	In the Linux kernel, the following vulnerability has been resolved: vxlan: do not reuse cached ip_hdr() value after skb_tunnel_check_pmtu() skb_tunnel_check_pmtu() can change skb->head. Reusing old_iph afer skb_tunnel_check_pmtu() can cause an UAF. Use instead ip_hdr(skb) as done in drivers/net/bareudp.c and drivers/net/geneve.c. Found by Sashiko.	9.8	More Details
CVE-2026-64606	Deserialization of untrusted data vulnerability that may allow class-registration checks to be bypassed during Java lambda deserialization. Only lambda capture class is affected This issue affects Apache Fory: from before 1.4.0. Users are recommended to upgrade to version 1.4.0, which fixes the issue.	9.8	More Details

CVE-2026-63984	In the Linux kernel, the following vulnerability has been resolved: ipv6: rpl: fix hdrlen overflow in ipv6_rpl_srh_decompress() ipv6_rpl_srh_decompress() computes: outhdr->hdrlen = (((n + 1) * sizeof(struct in6_addr)) >> 3); hdrlen is __u8. For n >= 127 the result exceeds 255 and silently truncates. With n=127 (cmpri=15, cmpre=15, pad=0, hdrlen=16): (128 * 16) >> 3 = 256, truncated to 0 as __u8 The caller in ipv6_rpl_srh_rcv() then places the compressed header at buf + ((ohdr->hdrlen + 1) << 3). With hdrlen=0 this is buf + 8, but the decompressed region occupies buf[0..2055] (8-byte header plus 128 full addresses). The compressed header overlaps the decompressed data, and ipv6_rpl_srh_compress() writes into this overlap, corrupting the routing header of the forwarded packet. The existing guard at exthdrs.c:546 checks (n + 1) > 255, which prevents n+1 from overflowing unsigned char (the segments_left field), but does not prevent the computed hdrlen from overflowing __u8. n=127 passes because 128 <= 255, yet hdrlen=256 does not fit. Tighten the bound to (n + 1) > 127. This caps n at 126, giving hdrlen = (127 * 16) >> 3 = 254, which fits in __u8. The compressed header then lands at buf + ((254 + 1) << 3) = buf + 2040, exactly past the decompressed region (buf[0..2039]). No overlap. 127 segments is well beyond any realistic RPL deployment.	9.8	More Details
CVE-2026-63979	In the Linux kernel, the following vulnerability has been resolved: net/handshake: hand off the pinned file reference to accept_doit handshake_req_next() removes the request from the per-net pending list and drops hn_lock before handshake_nl_accept_doit() reads req->hr_sk->sk_socket and dereferences sock->file (once in FD_PREPARE() and again in get_file()). In that window a consumer running tls_handshake_cancel() followed by sockfd_put() (svc_sock_free) or __fput_sync() (xs_reset_transport) releases sock->file. sock_release() then runs sock_orphan(), zeroing sk_socket, and frees the struct socket. The accept-side code either reads NULL through sk_socket or chases freed memory. The submit-side sock_hold() does not prevent this. sk_refcnt protects struct sock, but struct socket and sock->file are independently refcounted via the file descriptor the consumer owns. Pinning sk leaves sock and sock->file unprotected. Retarget the accept-side dereferences at req->hr_file, which was pinned at submit time, instead of req->hr_sk->sk_socket->file. Pinning on its own is not sufficient: a consumer that cancels between handshake_req_next() returning and accept_doit reaching FD_PREPARE() takes the !remove_pending() branch in handshake_req_cancel() and drops hr_file before the accept side takes its own reference. Hand off an additional file reference inside handshake_req_next(), under hn_lock, so the accept side operates on a reference that no concurrent handshake_req_cancel() can revoke. FD_PREPARE() consumes that handed-off reference, either by transferring it to the new fd in fd_publish() or by dropping it in the cleanup destructor on error; the explicit get_file() that previously balanced FD_PREPARE() is therefore redundant and goes away. Update handshake_req_cancel_test2 and _test3 to simulate the FD_PREPARE() consumption with an fput() so the kunit file-count assertions stay balanced.	9.8	More Details
CVE-2026-63978	In the Linux kernel, the following vulnerability has been resolved: net/handshake: Drain pending requests at net namespace exit The arguments to list_splice_init() in handshake_net_exit() are reversed. The call moves the local empty "requests" list onto hn->hn_requests, leaving the local list empty, so the subsequent drain loop runs zero iterations. Pending handshake requests that had not yet been accepted are not torn down when the net namespace is destroyed; each one keeps a reference on a socket file and on the handshake_req allocation. Pass the source and destination in the documented order (list_splice_init(list, head) moves list onto head) so the pending list is transferred to the local scratch list and drained through handshake_complete(). Fixing the splice direction exposes a list-corruption race. After the splice each req->hr_list still has non-empty link pointers, threading the stack-local scratch list rather than hn_requests. A concurrent handshake_req_cancel() -- for example, from sunrpc's TLS timeout on a kernel socket whose netns reference was not taken -- finds the request through the rhashtable, calls remove_pending(), and sees !list_empty(&req->hr_list). __remove_pending_locked() then list_del_init()s an entry off the scratch list while the drain iterates, corrupting it. The same call arriving after the drain loop has run list_del() on an entry hits LIST_POISON instead. Have remove_pending() check HANDSHAKE_F_NET_DRAINING under hn_lock and report not-found when drain is in progress. The drain has already taken ownership; handshake_complete()'s existing test_and_set on HANDSHAKE_F_REQ_COMPLETED still arbitrates between drain and cancel for who calls the consumer's hp_done. Use list_del_init() rather than list_del() in the drain so req->hr_list does not carry LIST_POISON after drain releases the entry. The DRAINING guard in remove_pending() makes cancel return false, but cancel still falls through to test_and_set_bit on HANDSHAKE_F_REQ_COMPLETED and drops the request's hr_file reference. Without another pin, if that is the last reference, sk_destruct frees the request while it is still linked on the drain loop's local list. Pin each request's hr_file under hn_lock before releasing the list, and drop that drain pin after the loop finishes with the request.	9.8	More Details
CVE-2026-64056	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: cortina: Make RX SKB per-port The SKB used to assemble packets from fragments in gmac_rx() is static local, but the Gemini has two ethernet ports, meaning there can be races between the ports on a bad day if a device is using both. Make the RX SKB a per-port variable and carry it over between invocations in the port struct instead. Zero the pointer once we call napi_gro_frags(), on error (after calling napi_free_frags()) or if the port is stopped. Zero it in some place where not strictly necessary just to emphasize what is going on. This was found by Sashiko during normal patch review.	9.8	More Details
CVE-2026-	Grav 2.0.4 (fixed in 2.0.7) contains a remote code execution vulnerability in Blueprint::dynamicData() (system/src/Grav/Common/Data/Blueprint.php), which passes a Class::method callable string and its arguments directly to call_user_func_array() without any allowlist. Because the form plugin routes page	9.8	More

65008	frontmatter through this path, an authenticated account with the admin.pages (or api.pages.write) permission can plant a malicious callable directive in a page. The command then executes as the web-server user whenever anyone — including an unauthenticated visitor — accesses the page.		Details
CVE-2026-63924	In the Linux kernel, the following vulnerability has been resolved: ipv6: exthdrs: refresh nh pointer after ipv6_hop_jumbo() ipv6_hop_jumbo() calls pskb_trim_rcsum(), which can change skb pointers. Let's recompute nh pointer to make sure any change won't mess things up.	9.8	More Details
CVE-2026-63922	In the Linux kernel, the following vulnerability has been resolved: ipv6: exthdrs: refresh nh after handling HAO option ip6_parse_tlv() caches skb_network_header(skb) in nh while walking IPv6 TLVs. ipv6_dest_hao() may call pskb_expand_head() for a cloned skb, which can move the skb head and invalidate the cached network header pointer. Refresh nh after ipv6_dest_hao() returns so any trailing padding or TLVs are parsed from the current skb head. This matches the existing pattern used in ip6_parse_tlv() after helpers that can modify skb header storage.	9.8	More Details
CVE-2026-63912	In the Linux kernel, the following vulnerability has been resolved: xfrm: esp: restore combined single-frag length gate The ESP out-of-place fast path appends the trailer in esp_output_head() before esp_output_tail() allocates the destination page frag. The head-side gate currently checks skb->data_len and tailen separately, but the tail code allocates a single destination frag from the combined post-trailer skb->data_len. Reject the page-frag fast path when the combined aligned length exceeds a page. Otherwise skb_page_frag_refill() may fall back to a single page while the destination sg still spans the combined skb->data_len. Restore this combined-length page gate for both IPv4 and IPv6.	9.8	More Details
CVE-2026-63888	In the Linux kernel, the following vulnerability has been resolved: scsi: target: iscsi: Fix CRC overread and double-free in iscsit_handle_text_cmd() Two latent bugs in the Text-phase handler, both present since the original LIO integration in commit e48354ce078c ("iscsi-target: Add iSCSI fabric support for target v4.1"): 1) DataDigest CRC buffer overread (4 bytes past text_in). text_in is kzalloc()'d at ALIGN(payload_length, 4). rx_size is then incremented by ISCSI_CRC_LEN to make room for the received DataDigest in the iovec, but the same (now-bumped) rx_size is passed as the buffer length to iscsit_crc_buf(): if (conn->conn_ops->DataDigest) { ... rx_size += ISCSI_CRC_LEN; } ... if (conn->conn_ops->DataDigest) { data_crc = iscsit_crc_buf(text_in, rx_size, 0, NULL); iscsit_crc_buf() walks rx_size bytes of text_in with crc32c(), so when DataDigest is negotiated it reads 4 bytes past the end of the text_in allocation. KASAN reproduces this directly on the unpatched mainline tree as slab-out-of-bounds in crc32c() called from the Text PDU path. The OOB bytes feed crc32c() and are then compared against the initiator-supplied checksum, so the value does not flow back to the attacker, but the kernel does read past the buffer on every Text PDU with DataDigest=CRC32C. Fix by passing the actual padded payload length (ALIGN(payload_length, 4)) that was used for the kzalloc(). 2) Stale cmd->text_in_ptr re-free (double-free) on ERL>0 bad DataDigest drop. On DataDigest mismatch with ErrorRecoveryLevel > 0 the handler silently drops the PDU and lets the initiator plug the CmdSN gap: kfree(text_in); return 0; cmd->text_in_ptr still points at the freed buffer. The next Text Request on the same ITT re-enters iscsit_setup_text_cmd(), which unconditionally does kfree(cmd->text_in_ptr); cmd->text_in_ptr = NULL; freeing the same pointer a second time. Session teardown via iscsit_release_cmd() has the same shape and hits the same double-free if the connection is dropped before a second Text Request arrives. On an unmodified mainline tree the bug-1 CRC overread fires first on the initial valid Text Request and perturbs the subsequent state, so #4 was isolated by building a kernel with only the bug-1 hunk of this patch applied plus temporary printk() observability around the three relevant kfree() sites. The observability prints are not part of this patch. On that build, a three-PDU Text Request sequence after login produces two back-to-back splats: BUG: KASAN: double-free in iscsit_setup_text_cmd+0x?? BUG: KASAN: double-free in iscsit_release_cmd+0x?? showing the same pointer freed in the ERL>0 drop path and again in iscsit_setup_text_cmd() (next Text Request on the same ITT) and once more in iscsit_release_cmd() (session teardown). On distro kernels with CONFIG_SLAB_FREELIST_HARDENED=y (default) the double-free becomes a remote kernel BUG(); on non-hardened kernels it corrupts the slab freelist. Fix by clearing cmd->text_in_ptr after the kfree() in the ERL>0 drop path. With both hunks applied #4 is directly observable on the stock tree without observability printks; fixing bug-1 alone would mask #4 less, not more, so the hunks are submitted together. Both fixes are one-liners. The Text PDU state machine is unchanged and the wire protocol is unaffected.	9.8	More Details
CVE-2026-63887	In the Linux kernel, the following vulnerability has been resolved: scsi: target: iscsi: Bound iscsi_encode_text_output() appends to rsp_buf iscsi_encode_text_output() concatenates "key=value\0" records into login->rsp_buf, an 8192-byte kzalloc(MAX_KEY_VALUE_PAIRS) buffer allocated in iscsit_alloc_login_setup_buffer(). The three sprintf() call sites in this function (lines 1398, 1411, 1424 in v7.1-rc2) never check the remaining buffer capacity: *length += sprintf(output_buf, "%s=%s", er->key, er->value); *length += 1; output_buf = textbuf + *length; The 8192-byte ceiling at iscsi_target_check_login_request() bounds the *input* Login PDU payload, but a single PDU can carry up to 2048 minimal four-byte "a=b\0" pairs, each unknown key expanding to a 16-byte "a=NotUnderstood\0" output record via iscsi_add_notunderstood_response(). 2048 * 16 = 32 KiB of output into an 8 KiB buffer, producing a ~24 KiB heap overrun in the kmalloc-8k slab. The fix introduces a static iscsi_encode_text_record() helper that uses snprintf() with a per-call bounds check against the remaining buffer, and threads a u32 textbuf_size parameter through iscsi_encode_text_output(). Both call sites in iscsi_target_handle_csg_zero() (PHASE_SECURITY) and iscsi_target_handle_csg_one() (PHASE_OPERATIONAL) pass MAX_KEY_VALUE_PAIRS. On overflow the encoder logs the condition, calls	9.8	More Details

	iscsi_release_extra_responses() to drop queued records, and returns -1; both caller sites now emit ISCSI_STATUS_CLS_INITIATOR_ERR / ISCSI_LOGIN_STATUS_INIT_ERR via iscsit_tx_login_rsp() before returning, so the initiator sees an explicit failed-login response rather than a silent connection drop. (Prior to this patch only the PHASE_OPERATIONAL caller did that; the PHASE_SECURITY caller is converted to the same shape.)		
CVE-2026-63886	In the Linux kernel, the following vulnerability has been resolved: scsi: target: iscsi: Validate CHAP_R length before base64 decode chap_server_compute_hash() allocates client_digest as kzalloc(chap->digest_size) and then, for BASE64-encoded responses, passes chap_r directly to chap_base64_decode() without checking whether the input length could produce more than digest_size bytes of output. chap_base64_decode() writes to the destination unconditionally as long as there is input to consume. With MAX_RESPONSE_LENGTH set to 128 and the "0b" prefix stripped by extract_param(), up to 127 base64 characters can reach the decoder. 127 characters decode to 95 bytes. For SHA-256 (digest_size=32) this overflows client_digest by 63 bytes; for MD5 (digest_size=16) the overflow is 79 bytes. The length check at line 344 fires after the write has already happened. The HEX branch in the same switch statement already validates the length up front. Apply the same approach to the BASE64 branch: strip trailing base64 padding characters, then reject any input whose data length exceeds DIV_ROUND_UP(digest_size * 4, 3) before calling the decoder. Stripping trailing '=' before the comparison handles both padded and unpadded encodings. chap_base64_decode() already returns early on '=', so the full original string is still passed to the decoder unchanged. The mutual CHAP path decodes CHAP_C into initiatorchg_binhex, which is kzalloc(CHAP_CHALLENGE_STR_LEN). extract_param() caps initiatorchg at CHAP_CHALLENGE_STR_LEN characters, so at most CHAP_CHALLENGE_STR_LEN-1 base64 characters reach the decoder. The maximum decoded size, DIV_ROUND_UP((CHAP_CHALLENGE_STR_LEN-1) * 3, 4), is less than CHAP_CHALLENGE_STR_LEN, so no overflow is possible there. A comment is added at the call site to document this.	9.8	More Details
CVE-2026-63857	In the Linux kernel, the following vulnerability has been resolved: net: airoha: Do not read uninitialized fragment address in airoha_dev_xmit() The transmit loop in airoha_dev_xmit() reads fragment address and length during its final iteration, when the loop index equals skb_shinfo(skb)->nr_frags, at which point the fragment data is uninitialized. While these values are never consumed, the read itself is unsafe and may trigger a page fault. Fix this by avoiding the fragment read on the last iteration. Additionally, move the skb pointer from the first to the last used packet descriptor, so that airoha_qdma_tx_napi_poll() defers freeing the skb until the final descriptor is processed.	9.8	More Details
CVE-2026-64055	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: cortina: Carry over frag counter The gmac_rx() NAPI poll function assembles packets in an SKB from a ring buffer. If the ring buffer gets completely emptied during a poll cycle, we exit gmac_rx(), but the packet is not yet completely assembled in the SKB, yet the fragment counter frag_nr is reset to zero on the next invocation. Solve this by making the RX fragment counter a part of the port struct, and carry it over between invocations. Reset the fragment counter only right after calling napi_gro_frags(), on error (after calling napi_free_frags()) or if stopping the port. Reset it in some place where not strictly necessary just to emphasize what is going on. This was found by Sashiko during normal patch review.	9.8	More Details
CVE-2026-64089	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: fix negative last_changeset_len batadv_piv_tt::last_changeset_len len was declared as s16, but the field is never intended to hold a negative value. When a value greater than 32767 is assigned, it wraps to a negative signed integer. In batadv_send_my_tt_response(), last_changeset_len is temporarily widened to s32. The incorrectly negative s16 value propagates into the s32, causing batadv_tt_prepare_tvlv_local_data() to allocate a full sized buffer but populates only a small portion of it with the collected changeset. All remaining bits are kept uninitialized. Using an u16 avoids this type confusion and ensures that no (negative) sign extension is performed in batadv_send_my_tt_response().	9.8	More Details
CVE-2026-64061	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix early put of sink folio in netfs_read_gaps() Fix netfs_read_gaps() to release the sink page it uses after waiting for the request to complete. The way the sink page is used is that an ITER_BVEC-class iterator is created that has the gaps from the target folio at either end, but has the sink page tiled over the middle so that a single read op can fill in both gaps. The bug was found by KASAN detecting a UAF on the generic/075 xfstest in the cifsd kernel thread that handles reception of data from the TCP socket: BUG: KASAN: use-after-free in _copy_to_iter+0x48a/0xa20 Write of size 885 at addr ffff888107f92000 by task cifsd/1285 CPU: 2 UID: 0 PID: 1285 Comm: cifsd Not tainted 7.0.0 #6 PREEMPT(lazy) Call Trace: dump_stack_lvl+0x5d/0x80 print_report+0x17f/0x4f1 kasan_report+0x100/0x1e0 kasan_check_range+0x10f/0x1e0 __asan_memcpy+0x3c/0x60 _copy_to_iter+0x48a/0xa20 __skb_datagram_iter+0x2c9/0x430 skb_copy_datagram_iter+0x6e/0x160 tcp_recvmsg_locked+0xce0/0x1130 tcp_recvmsg+0xeb/0x300 inet_recvmsg+0xcf/0x3a0 sock_recvmsg+0xea/0x100 cifs_readv_from_socket+0x3a6/0x4d0 [cifs] cifs_read_iter_from_socket+0xdd/0x130 [cifs] cifs_readv_receive+0xaad/0xb10 [cifs] cifs_demultiplex_thread+0x1148/0x1740 [cifs] kthread+0x1cf/0x210	9.8	More Details
CVE-2026-64066	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix netfs_read_to_pagecache() to pause on subreq failure Fix netfs_read_to_pagecache() so that it pauses the generation of new subrequests if an already-issued subrequest fails.	9.8	More Details

CVE-2026-16235	Crypt::Password versions through 0.28 for Perl generate insecure random values for salts. These versions use the built-in rand function, which is predictable and unsuitable for cryptography.	9.8	More Details
CVE-2026-64620	FreeRDP before 3.28.0 (affected <=3.27.1) contains a heap-based buffer overflow in crypto_rsa_common() (libfreerdp/crypto/crypto.c). The function writes the modular-exponentiation result into the caller's output buffer via BN_bn2bin() and only afterward checks output_length > out_length, so out-of-bounds bytes are written before the bounds check. On the server side, when a client selects RDP Standard Security, the encrypted client random is decrypted into a fixed 32-byte buffer. Because the server publishes its RSA public key, an unauthenticated attacker can forge a ciphertext whose decrypted value is up to the full modulus length (e.g. 256 bytes for RSA-2048), overflowing the 32-byte heap buffer by up to ~224 attacker-controlled bytes pre-authentication, resulting in denial of service.	9.8	More Details
CVE-2024-51315	The Tenda TX9 V22.03.02.20 firmware has a stack overflow vulnerability in the sub_425964 function of the file /goform/SetOnlineDevName	9.8	More Details
CVE-2026-53405	Improper Isolation or Compartmentalization vulnerability in Apache Syncope. An administrator with adequate entitlements can import arbitrary BPMN process definitions via the REST API and then start the process. When a BPMN process containing a Groovy scriptTask is imported and started, the Groovy script is executed directly on the server, with no sandbox. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 Through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue by wrapping Flowable's Groovy scriptTasks with security sandbox.	9.8	More Details
CVE-2026-53421	Improper Isolation or Compartmentalization vulnerability in Apache Syncope. An administrator with adequate entitlements can achieve remote code execution through the connector subsystem by relying on scripted connectors' (REST and SQL) capability to run Groovy scripts. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 Through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue by hardening the Groovy security sandbox.	9.8	More Details
CVE-2026-57308	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Apache Syncope. An administrator with adequate entitlements can achieve execution of arbitrary SQL via stacked queries, leveraging unsanitized sort parameters. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 Through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue.	9.8	More Details
CVE-2026-62183	Improper Privilege Management vulnerability in Apache Syncope. When: * the all-Java user workflow adapter is configured, or * the Flowable user workflow adapter is configured, bearing a BPMN definition not requiring admin approval for user self registration of self update requests the following scenario could happen. A REST API call can allow the user to grant themselves one or more of defined Roles, thus gaining their Entitlements and becoming in fact an administrator; the actual Entitlements gained depend on the Roles that are effectively defined on the specific Syncope deployment. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 Through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue.	9.8	More Details
CVE-2026-63071	Improper Isolation or Compartmentalization vulnerability in Apache Syncope. An administrator with adequate entitlements for Implementations can create a malicious Groovy class containing untrusted code bypassing the Groovy security sandbox. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue by tightening the Groovy security sandbox.	9.8	More Details
CVE-2024-51314	The Tenda TX9 V22.03.02.20 firmware has a stack overflow vulnerability in the sub_424CE0 function of the file /goform/setMacFilterCfg.	9.8	More Details
CVE-2026-35048	The Piwigo installer in versions 16.3.0 and earlier accepts POST parameters for database configuration and writes them directly into a PHP configuration file without proper sanitization. On PHP 8+, the `addslashes()` protection is bypassed because it checks for `get_magic_quotes_gpc()`, a function removed in PHP 8.0. This allows raw user input to be interpolated directly into PHP source code. An unauthenticated attacker can inject arbitrary PHP code through POST parameters (prefix, dbpasswd, dbhost, dbname, or dbuser), which gets written to `local/config/database.inc.php` and executed on every page load.	9.8	More Details
CVE-2026-41252	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a missing bounds check in xrdp, which allows a heap-based buffer overflow when operating in vnc-any mode. The issue occurs during the handling of RFB protocol color map messages from a VNC server, where incoming color indices are not properly validated. A malicious VNC server can exploit this flaw by sending crafted messages with out-of-range values, leading to an out-of-bounds write on the heap. This memory corruption can result in a denial of service (DoS) or potentially allow remote code execution (RCE) prior to authentication. This issue has been fixed in version 0.10.6.1.	9.8	More Details
CVE-2024-	The Tenda TX9 V22.03.02.20 firmware has a stack overflow vulnerability in the sub_42EEE0 function of	9.8	More

51312	the file /goform/SetStaticRouteCfg.		Details
CVE-2026-64193	Net::DNS versions through 1.55 for Perl allow remote execution injection via EDNS EXTENDED ERROR. Net::DNS::RR::OPT::EXTENDED_ERROR::_decompose parses the EXTRA-TEXT field of an EDNS EXTENDED-ERROR option (RFC 8914) by tokenising the raw bytes and passing the result to Perl's eval. There is some escaping done for $ and @, but not for backticks. This can be exploited for command execution if $pkt->edns->option('EXTENDED-ERROR') is called in array context, for example with a payload of {0:`"<command>"`} in EXTRA-TEXT.	9.8	More Details
CVE-2026-63766	GPT-SoVITS through 20250606v2pro contains an OS command injection vulnerability in webui.py where ASR, slice, denoise, and uvr5 functions interpolate unsanitized Gradio textbox values directly into shell commands executed with shell=True. Attackers can inject shell metacharacters through path parameters to execute arbitrary OS commands as the server process user without authentication.	9.8	More Details
CVE-2026-63767	ktransformers through 0.6.3, fixed in commit def0f93, contains an unauthenticated pickle deserialization vulnerability that allows remote attackers to execute arbitrary commands by sending crafted pickle payloads to the SchedulerServer ZMQ ROUTER socket bound to all interfaces. Attackers can exploit malicious __reduce__ methods embedded in crafted pickle payloads to execute arbitrary shell commands as the server process.	9.8	More Details
CVE-2026-52656	An issue in SJCAM AllWinner Tech products SJ4000-Air V1.4C and before and Whitelabel based v.1.4C and before allows an attacker to execute arbitrary code via a crafted FEX file	9.8	More Details
CVE-2026-64162	In the Linux kernel, the following vulnerability has been resolved: idpf: fix read_dev_clk_lock spinlock init in idpf_ptp_init() In idpf_ptp_init(), read_dev_clk_lock is initialized after ptp_schedule_worker() had already been called (and after idpf_ptp_settime64() could reach the lock). The PTP aux worker fires immediately upon scheduling and can call into idpf_ptp_read_src_clk_reg_direct(), which takes spin_lock(&ptp->read_dev_clk_lock) on an uninitialized lock, triggering the lockdep "non-static key" warning: [12973.796587] idpf 0000:83:00.0: Device HW Reset initiated [12974.094507] INFO: trying to register non-static key. ... [12974.097208] Call Trace: [12974.097213] <TASK> [12974.097218] dump_stack_lvl+0x93/0xe0 [12974.097234] register_lock_class+0x4c4/0x4e0 [12974.097249] ? __lock_acquire+0x427/0x2290 [12974.097259] __lock_acquire+0x98/0x2290 [12974.097272] lock_acquire+0xc6/0x310 [12974.097281] ? idpf_ptp_read_src_clk_reg+0xb7/0x150 [idpf] [12974.097311] ? lockdep_hardirqs_on_prepare+0xde/0x190 [12974.097318] ? finish_task_switch.isra.0+0xd2/0x350 [12974.097330] ? __pfx_ptp_aux_kworker+0x10/0x10 [ptp] [12974.097343] _raw_spin_lock+0x30/0x40 [12974.097353] ? idpf_ptp_read_src_clk_reg+0xb7/0x150 [idpf] [12974.097373] idpf_ptp_read_src_clk_reg+0xb7/0x150 [idpf] [12974.097391] ? kthread_worker_fn+0x88/0x3d0 [12974.097404] ? kthread_worker_fn+0x4e/0x3d0 [12974.097411] idpf_ptp_update_cached_phctime+0x26/0x120 [idpf] [12974.097428] ? _raw_spin_unlock_irq+0x28/0x50 [12974.097436] idpf_ptp_do_aux_work+0x15/0x20 [idpf] [12974.097454] ptp_aux_kworker+0x20/0x40 [ptp] [12974.097464] kthread_worker_fn+0xd5/0x3d0 [12974.097474] ? __pfx_kthread_worker_fn+0x10/0x10 [12974.097482] kthread+0xf4/0x130 [12974.097489] ? __pfx_kthread+0x10/0x10 [12974.097498] ret_from_fork+0x32c/0x410 [12974.097512] ? __pfx_kthread+0x10/0x10 [12974.097519] ret_from_fork_asm+0x1a/0x30 [12974.097540] </TASK> Move the call to spin_lock_init() up a bit to make sure read_dev_clk_lock is not touched before it's been initialized.	9.8	More Details
CVE-2026-64160	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix potential for tearing in ->remote_i_size and ->zero_point Fix potential tearing in using ->remote_i_size and ->zero_point by copying i_size_read() and i_size_write() and using the same seqcount as for i_size. We need to make sure that netfslib and the filesystems that use it always hold i_lock whilst updating any of the sizes to prevent i_size_seqcount from getting corrupted.	9.8	More Details
CVE-2026-64102	In the Linux kernel, the following vulnerability has been resolved: RDMA/siw: Reject MPA FPDU length underflow before signed receive math A malicious connected siw peer can send an iWARP FPDU whose MPA length field (c_hdr->mpa_len, 16 bit big-endian, peer-controlled) is smaller than the fixed DDP/RDMAP header for the announced opcode. Soft-iWARP parses the full header in siw_get_hdr() based on iwarp_pktinfo[opcode] .hdr_len, but never compares mpa_len against that header length. siw_tcp_rx_data() then derives srx->fpdu_part_rem = be16_to_cpu(mpa_len) - fpdu_part_rcvd + MPA_HDR_SIZE; where fpdu_part_rcvd equals iwarp_pktinfo[opcode].hdr_len at this point. For a tagged WRITE (hdr_len 16, MPA_HDR_SIZE 2) the smallest on-wire mpa_len of 0 yields fpdu_part_rem = -14, and any mpa_len below hdr_len - MPA_HDR_SIZE underflows to a negative int. The signed value then flows into siw_proc_write()/siw_proc_rresp() as bytes = min(srx->fpdu_part_rem, srx->skb_new); is handed to siw_check_mem() as an int len (whose interval check addr + len > mem->va + mem->len is satisfied for a valid base when len is negative), and reaches siw_rx_data() -> siw_rx_kva() / siw_rx_umem() -> skb_copy_bits() as a signed copy length. The header copy branch in skb_copy_bits() promotes that to size_t, producing a multi-gigabyte read. KASAN under a KUnit harness that drives the real kernel TCP receive path -- a loopback AF_INET socketpair, the malformed FPDU written via kernel_sendmsg, sk_data_ready firing in softirq, tcp_read_sock dispatching to siw_tcp_rx_data -- reports: BUG: KASAN: use-after-free in skb_copy_bits+0x284/0x480 Read of size 4294967295 at addr ffff888... Call Trace: skb_copy_bits siw_rx_kva siw_rx_data siw_check_mem siw_proc_write siw_tcp_rx_data __tcp_read_sock	9.8	More Details

	siw_qp_llp_data_ready tcp_data_ready tcp_data_queue Add the missing invariant at the earliest point where the peer header is fully assembled. iwarp_pktinfo[*].hdr_len - MPA_HDR_SIZE is exactly the value the siw transmitter uses as the minimum mpa_len for each opcode (drivers/infiniband/sw/siw/siw_qp.c:33), so this matches the protocol contract. Out-of-range FPDUs terminate the connection with TERM_ERROR_LAYER_LLP / LLP_ETYPE_MPA / LLP_ECODE_FPDU_START -- which is RFC 5044 Section 8 error code 3 ("Marker and ULPDU Length fields do not agree on the start of an FPDU"), the correct framing-error class for this inconsistency.		
CVE-2026-64067	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix missing barriers when accessing stream->subrequests locklessly The list of subrequests attached to stream->subrequests is accessed without locks by netfs_collect_read_results() and netfs_collect_write_results(), and then they access subreq->flags without taking a barrier after getting the subreq pointer from the list. Relatedly, the functions that build the list don't use any sort of write barrier when constructing the list to make sure that the NETFS_SREQ_IN_PROGRESS flag is perceived to be set first if no lock is taken. Fix this by: (1) Add a new list_add_tail_release() function that uses a release barrier to set the pointer to the new member of the list. (2) Add a new list_first_entry_or_null_acquire() function that uses an acquire barrier to read the pointer to the first member in a list (or return NULL). (3) Use list_add_tail_release() when adding a subreq to ->subrequests. (4) Use list_first_entry_or_null_acquire() when initially accessing the front of the list (when an item is removed, the pointer to the new front iterm is obtained under the same lock).	9.8	More Details
CVE-2026-64068	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix missing locking around retry adding new subreqs Fix netfs_retry_read_subrequests() and netfs_retry_write_stream() to take the appropriate lock when adding extra subrequests into stream->subrequests.	9.8	More Details
CVE-2026-64069	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix cancellation of a DIO and single read subrequests When the preparation of a new subrequest for a read fails, if the subrequest has already been added to the stream->subrequests list, it can't simply be put and abandoned as the collector may see it. Also, if it hasn't been queued yet, it has two outstanding refs that both need to be put. Both DIO read and single-read dispatch fail at this; further, both differ in the order they do things to the way buffered read works. Fix cancellation of both DIO-read and single-read subrequests that failed preparation by the following steps: (1) Harmonise all three reads (buffered, dio, single) to queue the subreq before prepping it. (2) Make all three call netfs_queue_read() to do the queuing. (3) Set NETFS_RREQ_ALL_QUEUED independently of the queuing as we don't know the length of the subreq at this point. (4) In all cases, set the error and NETFS_SREQ_FAILED flag on the subreq and then call netfs_read_subreq_terminated() to deal with it. This will pass responsibility off to the collector for dealing with it.	9.8	More Details
CVE-2026-13439	The Easy Form Builder by WhiteStudio plugin for WordPress is vulnerable to Unauthenticated Privilege Escalation to Administrator in versions up to, and including, 4.0.11 This is due to the password recovery flow using the publicly-visible session identifier ('sid') as the password reset token stored in wp_emsfb_temp_links, combined with a publicly-accessible nonce refresh endpoint (Emsfb/v1/nonce/refresh) that issues valid WordPress REST nonces to unauthenticated visitors. This makes it possible for unauthenticated attackers to reset the password of any WordPress user — including administrators — by scraping the public sid from a published login form page, submitting a recovery request for any known user email via Emsfb/v1/forms/message/add, and then calling Emsfb/v1/forms/recovery/efb_set_password with the known sid to set an arbitrary new password and gain full administrator access.	9.8	More Details
CVE-2026-63808	In the Linux kernel, the following vulnerability has been resolved: exfat: fix potential use-after-free in exfat_find_dir_entry() In exfat_find_dir_entry(), the buffer_head obtained from exfat_get_dentry() is released with brelse(bh) before the fall-through TYPE_EXTEND branch reads the directory entry through ep (which points into bh->b_data): brelse(bh); if (entry_type == TYPE_EXTEND) { ... len = exfat_extract_uni_name(ep, entry_uniname); ... } After brelse() drops our reference, nothing guarantees that the underlying page backing bh->b_data remains valid for the subsequent exfat_extract_uni_name() read. This is the same pattern fixed in commit fc961522ddbd ("exfat: Fix potential use after free in exfat_load_upcase_table()"). Move brelse(bh) so it runs after ep is no longer dereferenced on each branch. Confirmed on QEMU x86_64 with CONFIG_KASAN=y + CONFIG_DEBUG_PAGEALLOC=y + CONFIG_PAGE_POISONING=y on linux-next, using a crafted exFAT image (long filename with same-hash collisions forcing the TYPE_EXTEND path). With a debug-only invalidate_bdev() inserted between brelse(bh) and the ep read to make the stale-deref window deterministic, the unpatched kernel faults: BUG: KASAN: use-after-free in exfat_find_dir_entry+0x133b/0x15a0 BUG: unable to handle page fault for address: ffff88801a5fa0c2 Oops: 0000 [#1] SMP DEBUG_PAGEALLOC KASAN NOPTI RIP: 0010:exfat_find_dir_entry+0x1188/0x15a0 With this patch applied, the same instrumented harness completes cleanly under the same sanitizer stack. I have not reproduced a crash on an uninstrumented kernel under ordinary reclaim; the instrumented A/B establishes the lifetime violation and that the patch closes it, not an unaided triggerability claim.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: fix TOCTOU race for reported vlans The local TT based TVLV is generated by first checking the number of VLANs which have at least one TT entry. A new buffer with the correct size for the VLANs is then allocated. Only then, the		

CVE-2026-64091	list of VLANs s used to fill the VLAN entries in the buffer. During this time, the meshif_vlan_list_lock is held. But the actual number of TT entries of each VLAN can still increase during this time - just not the number of VLANs in the list. But the prefilter used in the buffer size calculation might still cause an increase of the number of VLANs which need to be stored. Simply because a VLAN might now suddenly have at least one entry when it had none in the pre-alloc check - and then needs to occupy space which was not allocated. It is better to overestimate the buffer size at the beginning and then fill the buffer only with the VLANs which are not empty.	9.8	More Details
CVE-2026-64625	AVideo before 29.0 contains an incomplete fix for CVE-2026-45578 where execAsync() re-wraps escaped commands in double-quoted sh -c, allowing command substitution via $() and backticks. Attackers can inject arbitrary OS commands through the Live plugin on_publish.php endpoint despite escapeshellarg() protection.	9.8	More Details
CVE-2026-64150	In the Linux kernel, the following vulnerability has been resolved: netfilter: nft_inner: release local_lock before re-enabling softirqs Quoting sashiko: In the error path, local_bh_enable() is called before local_unlock_nested_bh().	9.8	More Details
CVE-2026-64113	In the Linux kernel, the following vulnerability has been resolved: ixgbevf: fix use-after-free in VEPA multicast source pruning ixgbevf_clean_rx_irq() prunes frames whose source MAC matches the VF's own address (VEPA multicast workaround) by freeing the skb and continuing to the next descriptor: dev_kfree_skb_irq(skb); continue; The skb pointer is declared outside the while loop and persists across iterations. Because the continue skips the "skb = NULL" reset at the bottom of the loop, the next iteration enters the "else if (skb)" path and calls ixgbevf_add_rx_frag() on the freed skb, dereferencing skb_shinfo(skb)->nr_frags - a use-after-free in NAPI softirq context. The sibling driver iavf already handles this correctly by nulling the pointer before continuing. Apply the same pattern here. I do not have ixgbevf hardware; the bug was found by static analysis (scan_drop_continue_loops.py + semgrep drop_continue_in_loop, multi-tool corroboration with the highest score in the scan). The UAF was confirmed under KASAN by loading a test module that reproduces the exact code pattern (alloc skb, kfree_skb, then read skb_shinfo(skb)->nr_frags): BUG: KASAN: slab-use-after-free in ixgbevf_uaf_test_init+0x100/0x1000 Read of size 8 at addr 000000006163ae78 by task insmod/30 freed 208-byte region [000000006163adc0, 000000006163ae90) QEMU emulates igb (82576) but not ixgbe (82599), and the igbvf VF driver does not include the VEPA source pruning path, so a full end-to-end reproduction with emulated hardware was not possible.	9.8	More Details
CVE-2026-64122	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Fix use-after-free in mlx5e_tx_reporter_timeout_recover mlx5e_tx_reporter_timeout_recover() accesses sq->netdev after mlx5e_safe_reopen_channels() has torn down and freed the channel (and its embedded SQs). Replace the three sq->netdev references with priv->netdev which is safe because priv outlives channel teardown. The netdev_err() call already used priv->netdev for this reason; make the trylock/unlock and health_channel_eq_recover calls consistent. This fixes the following KASAN splat: BUG: KASAN: use-after-free in mlx5e_tx_reporter_timeout_recover+0x1dd/0x360 [mlx5_core] Read of size 8 at addr ffff889860ed0b28 by task kworker/u113:2/5277 Call Trace: mlx5e_tx_reporter_timeout_recover+0x1dd/0x360 [mlx5_core] devlink_health_reporter_recover+0xa2/0x150 devlink_health_report+0x254/0x7c0 mlx5e_reporter_tx_timeout+0x297/0x380 [mlx5_core] mlx5e_tx_timeout_work+0x109/0x170 [mlx5_core] process_one_work+0x677/0xf20 worker_thread+0x51f/0xd90 kthread+0x3a5/0x810 ret_from_fork+0x208/0x400 ret_from_fork_asm+0x1a/0x30	9.8	More Details
CVE-2026-64125	In the Linux kernel, the following vulnerability has been resolved: net: bcmgenet: keep RBUF EEE/PM disabled Setting RBUF_EEE_EN | RBUF_PM_EN in RBUF_ENERGY_CTRL breaks the RX path on GENET hardware once MAC EEE becomes active. RX traffic stops flowing while the link stays up and the usual descriptor/RX error counters remain quiet. In that state the MAC still accepts frames (rbuf_ovflow_cnt keeps climbing) but RBUF no longer forwards them to DMA, so rx_packets is no longer incremented at the netdev level. On some boards the corruption ends up as a paging fault in skb_release_data via bcmgenet_rx_poll on an LPI exit. Reproduced on Pi 4B (BCM2711 + BCM54213PE) and confirmed by Florian Fainelli on an internal Broadcom 4908-family board with the same crash signature. RBUF_PM_EN is not publicly documented. This shows up more often now that phy_support_eee() enables EEE by default, but it also affects older kernels as soon as TX LPI is turned on via ethtool, so it is not specific to recent changes. Always clear RBUF_EEE_EN | RBUF_PM_EN in bcmgenet_eee_enable_set so the bits stay off across resets. UMAC and TBUF setup is left alone so TX-side EEE keeps working.	9.8	More Details
CVE-2026-64132	In the Linux kernel, the following vulnerability has been resolved: ipv6: ioam: refresh hdr pointer before ioam6_event() Reported by Sashiko: In ipv6_hop_ioam(), the hdr pointer is initialized to point into the skb's linear data buffer. Later, the code calls skb_ensure_writable(), which might reallocate the buffer: if (skb_ensure_writable(skb, optoff + 2 + hdr->opt_len)) goto drop; /* Trace pointer may have changed */ trace = (struct ioam6_trace_hdr *)(skb_network_header(skb) + optoff + sizeof(*hdr)); ioam6_fill_trace_data(skb, ns, trace, true); ioam6_event(IOAM6_EVENT_TRACE, dev_net(skb->dev), GFP_ATOMIC, (void *)trace, hdr->opt_len - 2); If the skb is cloned or lacks sufficient linear headroom, skb_ensure_writable() will invoke pskb_expand_head(), which reallocates the skb's data buffer and frees the old one, invalidating pointers to it. While the code recalculates the trace pointer immediately after the call to skb_ensure_writable(), it fails to recalculate the hdr pointer. This patch fixes the above by	9.8	More Details

	recalculating the hdr pointer before passing hdr->opt_len to ioam6_event(), so that we avoid any UaF.		
CVE-2026-64136	In the Linux kernel, the following vulnerability has been resolved: smb: client: protect tc_count increment in smb2_find_smb_sess_tcon_unlocked() Commit 96c4af418586 ("cifs: Fix locking usage for tcon fields") refactored cifs code to change cifs_tcp_ses_lock for tc_lock around tc_count changes. There was missing lock around tc_count increment inside smb2_find_smb_sess_tcon_unlocked().	9.8	More Details
CVE-2026-64142	In the Linux kernel, the following vulnerability has been resolved: ksmbd: close durable scavenger races against m_fp_list lookups ksmbd_durable_scavenger() has two related races against any walker that iterates f_ci->m_fp_list, including ksmbd_lookup_fd_inode() (used by ksmbd_vfs_rename) and the share-mode checks in fs/smb/server/smb_common.c. (1) fp->node list-head reuse. Durable-preserved handles can remain linked on f_ci->m_fp_list after session teardown so share-mode checks still see them while the handle is reconnectable. The scavenger collected expired handles by adding fp->node to a local scavenger_list after removing them from the global durable idr. Because fp->node is the same list_head used by m_fp_list, list_add(&fp->node, &scavenger_list) overwrites the m_fp_list links and corrupts both lists. CONFIG_DEBUG_LIST can report this on the share-mode walk path. (2) Refcount race against m_fp_list walkers. The scavenger qualifies an expired durable handle with atomic_read(&fp->refcount) > 1 and fp->conn under global_ft.lock, removes fp from global_ft, then drops global_ft.lock before unlinking fp from m_fp_list and freeing it. During that gap fp is still linked on m_fp_list with f_state == FP_INITED. ksmbd_lookup_fd_inode() under m_lock read calls ksmbd_fp_get() (atomic_inc_not_zero on refcount that is still 1) and takes a live reference; the scavenger then unlinks and frees fp while the holder owns a reference, leading to UAF on the holder's subsequent ksmbd_fd_put() and on any field reads performed by a concurrent share-mode walker that iterates m_fp_list without taking ksmbd_fp_get() (smb_check_perm_dleases-like paths). Fix both: * Stop reusing fp->node as a scavenger-private list node. Remove one expired handle from global_ft under global_ft.lock, take an explicit transient reference, drop the lock, unlink fp->node from m_fp_list under f_ci->m_lock, then drop both the durable lifetime and transient references with atomic_sub_and_test(2, &fp->refcount). If the scavenger is the last putter the close runs there; otherwise an in-flight holder that already raced through the m_fp_list lookup owns the final close via its ksmbd_fd_put() path. The one-at-a-time disposal can rescan the durable idr when multiple handles expire in the same pass, but durable scavenging is a background expiration path and the final full scan recomputes min_timeout before the next wait. * Clear fp->persistent_id inside __ksmbd_remove_durable_fd() right after idr_remove(), so a delayed final close from a holder that snatched fp does not re-issue idr_remove() on a persistent id that idr_alloc_cyclic() in ksmbd_open_durable_fd() may have already handed out to a brand-new durable handle. * Bypass the per-conn open_files_count decrement in __put_fd_final() when fp is detached from any session table (fp->conn cleared by session_fd_check() at durable preserve -- paired with the volatile_id clear at unpublish, so checking fp->conn alone is sufficient). The walker that owns the final close runs from an unrelated work->conn whose stats.open_files_count never tracked this durable fp; without this guard the holder would underflow that unrelated counter. The two races are folded into one patch because patch (1) alone cleans up the corrupted list but leaves a deterministic UAF window for m_fp_list walkers that the transient-reference and persistent_id discipline in (2) close; bisecting onto an intermediate state would land on a UAF that pre-patch chaos merely made less reproducible. Validation: * CONFIG_DEBUG_LIST coverage for the list_head reuse path. * KASAN-enabled direct SMB2 durable-handle coverage that exercised ksmbd_durable_scavenger() and non-NULL ksmbd_lookup_fd_inode() returns while durable handles expired under concurrent rename lookups, with no KASAN, UAF, list-corruption, ODEBUG, or WARNING reports. ---truncated---	9.8	More Details
CVE-2026-63825	In the Linux kernel, the following vulnerability has been resolved: gcov: use atomic counter updates to fix concurrent access crashes GCC's GCOV instrumentation can merge global branch counters with loop induction variables as an optimization. In inflate_fast(), the inner copy loops get transformed so that the GCOV counter value is loaded multiple times to compute the loop base address, start index, and end bound. Since GCOV counters are global (not per-CPU), concurrent execution on different CPUs causes the counter to change between loads, producing inconsistent values and out-of-bounds memory writes. The crash manifests during IPComp (IP Payload Compression) processing when inflate_fast() runs concurrently on multiple CPUs: BUG: unable to handle page fault for address: ffffd0a3c0902ffa RIP: inflate_fast+1431 Call Trace: zlib_inflate __deflate_decompress crypto_comp_decompress ipcomp_decompress [xfrm_ipcomp] ipcomp_input [xfrm_ipcomp] xfrm_input At the crash point, the compiler generated three loads from the same global GCOV counter (__gcov0.inflate_fast+216) to compute base, start, and end for an indexed loop. Another CPU modified the counter between loads, making the values inconsistent - the write went 3.4 MB past a 65 KB buffer. Add -fprofile-update=prefer-atomic to CFLAGS_GCOV at the global level in the top-level Makefile, guarded by a try-run compile test. The test compiles a minimal program with and without -fprofile-update=prefer-atomic using the full KBUILD_CFLAGS, then compares undefined symbols in the resulting object files. If prefer-atomic introduces new undefined references (such as __atomic_fetch_add_8 on i386 or __aarch64_ldadd8_relax on arm64 with outline-atomics), the flag is not added -- the kernel does not link against libatomic. On architectures where GCC inlines 64-bit atomic counter updates (x86_64, s390, ...) the test passes and the flag is enabled, preventing the compiler from merging counters with loop induction variables and fixing the observed concurrent-access crash. On architectures where the flag would introduce libatomic dependencies, it is silently omitted and behaviour is no worse than before this patch. Move the CFLAGS_GCOV block from its original position (before the arch Makefile include) to after the core	9.8	More Details

	KBUILD_CFLAGS assignments but before the scripts/Makefile.gcc-plugins include. This placement ensures the try-run test sees arch-specific flags (-m32, -march=, -mno-outline-atomics) while avoiding GCC plugin flags (-fplugin=) that would break the test on clean builds when plugin shared objects do not yet exist.		
CVE-2026-47865	VMware Avi Load Balancer contains an authentication bypass vulnerability. A malicious user with network access may be able to access the Avi Control plane by bypassing the authentication mechanism. Affected versions: 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	9.8	More Details
CVE-2026-63800	In the Linux kernel, the following vulnerability has been resolved: pNFS: Fix use-after-free in pnfs_update_layout() When hitting the NFS_LAYOUT_RETURN branch in pnfs_update_layout(), the code calls pnfs_prepare_to_retry_layoutget(lo). If it succeeds, pnfs_put_layout_hdr(lo) is called before trace_pnfs_update_layout(), which still references 'lo'. This results in a use-after-free when the tracepoint accesses lo's fields. Fix this by moving the tracepoint call before pnfs_put_layout_hdr(lo).	9.8	More Details
CVE-2023-49899	An unauthenticated remote attacker can execute any command on the affected device due to not correctly verifying the origin of a communication channel.	9.8	More Details
CVE-2026-47391	PraisonAI is a multi-agent teams system. Prior to version 4.6.40, PraisonAI's first-party A2A server example exposes an unauthenticated A2A JSON-RPC endpoint and registers a `calculate(expression)` tool implemented with Python `eval()`. The example also binds to `0.0.0.0`. A remote unauthenticated attacker can send `message/send` to `/a2a`; the request reaches `agent.chat()`, and a real LLM can invoke the registered `calculate` tool. In testing with `gemini/gemini-2.5-flash-lite`, this resulted in arbitrary Python execution in the server process, confirmed by creation of a marker file from an unauthenticated HTTP request. The issue affects deployments following the official A2A example or similar unauthenticated public A2A deployments with unsafe tools. The default unauthenticated A2A surface also exposes task history and task cancellation APIs, increasing confidentiality and integrity impact. Version 4.6.40 patches the issue.	9.8	More Details
CVE-2026-14956	The Bricksforge plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 3.1.8.6. This is due to improper validation of the fieldIds parameter in the Pro Forms registration action, which allows attacker-supplied field IDs to be added to the trusted form-field whitelist. This makes it possible for unauthenticated attackers to register a new administrator account by submitting a crafted request to a publicly accessible Bricksforge Pro Forms registration form. Successful exploitation requires that the site has a public Bricksforge Pro Forms element configured with the User Registration action.	9.8	More Details
CVE-2026-53412	Improper Input Validation in Zoom Desktop Client for Windows, Zoom VDI Client for Windows, and Zoom Meeting SDK for Windows may allow an unauthenticated user to conduct an account takeover via network access.	9.8	More Details
CVE-2026-44180	Jupyter Enterprise Gateway launches remote Jupyter Notebook kernels across distributed clusters like Apache Spark, Kubernetes, and Docker Swarm. Versions 2.0.0rc1 and above prior to 3.3.0 have a prohibited UID and GID feature that by default prevents launching kernels with UID or GID 0 (root), and this restriction can be bypassed using a specially crafted KERNEL_UID or KERNEL_GID value. This input validation vulnerability allows running Jupyter kernels as root, which can be dangerous as it allows more attack surface, and may lead to container escapes, compromising the worker node and all workloads running on it. Repeated exploitation can compromise all worker nodes, and thus the entire Kubernetes cluster. It is possible to specify volume mounts, so one vector for a container escape is to use a hostPath R/W volume mount, use this UID/GID bypass to run as root, and then gain code execution in the underlying worker node by creating a crontab entry in the mounted host file system. This issue has been fixed in version 3.0.0.	9.8	More Details
CVE-2026-47393	PraisonAI is a multi-agent teams system. CVE-2026-44338 (GHSA-6rmh-7xcm-cpxj) documents that PraisonAI ships a code-generator (`praisonai.deploy.api.generate_api_server_code`) that emits a Flask API server with authentication disabled by default. Users who follow the documented quickstart (`praisonai deploy --type api`) get a server that binds to `0.0.0.0` per the recommended sample YAML, exposes `/chat` and `/agents` endpoints, runs `praisonai.run()` on user-supplied JSON input — LLM orchestration with the API key materials present in the process environment, and does not require any authentication. Versions prior to 4.6.40 still ship the generator with `auth_enabled` defaulting to `False`. The fix shape is opt-in via `APIConfig(auth_enabled=True, auth_token=...)`. Version 4.6.40 fixes the issue.	9.8	More Details
CVE-2026-63087	Grafana OnCall through 1.16.11 contains an unauthenticated access vulnerability that allows remote attackers to obtain a valid PluginAuthToken by sending a POST request to the internal plugin install endpoint using hardcoded default stack_id and org_id values present in the public source tree. Attackers can leverage the acquired token to authenticate against all internal API endpoints, create arbitrary Admin users via the user-context header bootstrap path, revoke the legitimate plugin token, and redirect OnCall-to-Grafana API calls to an attacker-controlled host by overwriting the organization's grafana_url and api_token.	9.8	More Details

CVE-2026-47396	PraisonAI is a multi-agent teams system. Prior to version 4.6.40, PraisonAI's call server exposes a network-facing agent control API without authentication when `CALL_SERVER_TOKEN` is not configured. The affected component is the `praisonai.api.agent_invoke` router as mounted by `praisonai.api.call`. The authentication helper `verify_token()` fails open when `CALL_SERVER_TOKEN` is unset. Since every sensitive agent-control endpoint depends on this helper, starting the call server without a token allows any reachable client to list agents, inspect agent metadata and instructions, invoke agents, and unregister agents. This is security-relevant because the bundled call server includes the vulnerable router and binds to `0.0.0.0`. As a result, operators who launch the call server without explicitly setting `CALL_SERVER_TOKEN` may unintentionally expose an unauthenticated remote agent control plane. Version 4.6.40 fixes the issue.	9.8	More Details
CVE-2026-47410	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an insecure default cryptographic key. The JWT signing secret defaults to the hardcoded literal `"dev-secret-change-me"` when `PLATFORM_JWT_SECRET` is unset. A safety check exists but only fires when `PLATFORM_ENV != "dev"`; the default value of `PLATFORM_ENV` is `"dev"`, so the check is silently bypassed in any deployment that does not explicitly opt out. The attacker reads the literal from this public source file, mints a JWT with arbitrary `sub` and `email` claims, and authenticates as any existing user (including workspace owners and admins). PraisonAI Platform version 0.1.4 patches the issue.	9.8	More Details
CVE-2026-46562	Yamcs is a mission control framework. Prior to 5.12.7, the Nashorn ScriptEngine used to evaluate user-supplied JavaScript algorithm text in yamcs-core/src/main/java/org/yamcs/algorithms/ScriptAlgorithmExecutorFactory.java was constructed without a ClassFilter, so a user with the ChangeMissionDatabase privilege could override an algorithm through the MdbOverrideApi.updateAlgorithm endpoint and supply JavaScript that reaches arbitrary Java classes (for example Java.type("java.lang.Runtime").getRuntime().exec(...)) to execute arbitrary OS commands as the Yamcs process; in the default configuration with no security.yaml the built-in guest user has superuser=true, making the issue reachable without authentication. This issue is fixed in versions 5.12.7 and 5.13.0, which disable algorithm editing by default.	9.8	More Details
CVE-2026-3031	Image::EPEG versions through 0.15 for Perl embeds an unsupported version of the Epeg library. Image::EPEG includes Epeg 0.9.0 that was last updated in 2004. Epeg is a fast JPEG thumbnail library that was once part of the Englightenment Project.	9.8	More Details
CVE-2026-45695	Kopia is a cross-platform backup tool for Windows, macOS, and Linux with fast incremental backups, client-side end-to-end encryption, compression, and data deduplication. Prior to 0.23.0, Kopia's HTTP server started with --without-password accepts unauthenticated requests to /api/v1/repo/exists and forwards attacker-supplied SFTP storage configuration to blob.NewStorage, where externalSSH: true and sshArguments containing -oProxyCommand=<cmd> can cause exec.CommandContext("ssh") to invoke the command through OpenSSH. This issue is fixed in version 0.23.0.	9.8	More Details
CVE-2016-20096	Linknat VOS3000 and VOS2009 through version 2.1.2.0 contain an unauthenticated SQL injection vulnerability that allows remote attackers to execute arbitrary SQL commands by manipulating the name parameter in a POST request to the login endpoint. Attackers can inject malicious SQL through the login form and retrieve injected query results from a subsequent session request, enabling extraction of plaintext credentials and other database content with DBA-level privileges.	9.8	More Details
CVE-2023-49900	An unauthenticated remote attacker is able to perform remote code execution due to incorrectly sanitized user input in the SetParameter command.	9.8	More Details
CVE-2026-12492	The Happy Coders OTP Login for WooCommerce WordPress plugin before 2.8 does not verify that a one-time password was actually validated before authenticating a user based on a supplied identifier, allowing unauthenticated attackers to log in as any existing user, including administrators, as well as to create new accounts.	9.8	More Details
CVE-2026-9810	The AI Copilot WordPress plugin before 1.5.4 does not bind OAuth access tokens to a WordPress user, and accepts any valid token as an administrator session, allowing unauthenticated attackers who complete the public OAuth flow to execute privileged MCP tools as an administrator, including arbitrary user creation and role escalation.	9.8	More Details
CVE-2026-15013	The SAML Single Sign On – SSO Login plugin for WordPress is vulnerable to Authentication Bypass via SAML Signature Algorithm Confusion in all versions up to, and including, 5.4.3. The vulnerability exists because `Mo_SAML_Utilities::mo_saml_cast_key()` reads the `SignatureMethod` Algorithm attribute directly from the attacker-controlled `SAMLResponse` parameter rather than enforcing the locally configured algorithm, causing the plugin to recast the IdP's RSA public key as an HMAC-SHA1 shared secret and validate the forged signature against it. This makes it possible for unauthenticated attackers to forge a SAML assertion targeting any WordPress account — including administrators — obtain valid WordPress authentication cookies, and achieve full administrator-level account takeover.	9.8	More Details
CVE-2026-	Wekan is open source kanban built with Meteor. Prior to 9.46, header-login with HEADER_LOGIN_TRUSTED_IPS uses getRequestIp() in server/lib/headerLoginAuth.js to trust the client-		More

55652	supplied X-Forwarded-For header before the real socket address, allowing an unauthenticated attacker to send HEADER_LOGIN_ID for any username and receive a meteor_login_token session, including for admin. This issue is fixed in version 9.46.	9.8	Details
CVE-2026-30623	LiteLLM 1.18.10 contains a remote code execution vulnerability in its MCP server creation functionality. The application allows users to add MCP servers via a JSON configuration specifying arbitrary command and args values. LiteLLM executes these values on the host without validation, enabling attackers to run arbitrary operating system commands. Successful exploitation may result in remote code execution with the privileges of the LiteLLM process.	9.8	More Details
CVE-2026-30618	xszyou Fay 4.3.1 contains a remote code execution vulnerability in its MCP STDIO server management and command execution handling. A remote attacker can access the publicly exposed MCP management interface and configure an MCP STDIO server with attacker-controlled commands and parameters, resulting in execution of arbitrary commands on the server. Successful exploitation allows arbitrary command execution within the context of the Fay service.	9.8	More Details
CVE-2025-65720	An issue in Open Source GPT Researcher v3.3.7 allows attackers to execute arbitrary commands on a victim system via user interaction with a crafted HTML page.	9.8	More Details
CVE-2026-51380	Buffer Overflow vulnerability in Tenda AC10 v3 (firmware V03.03.16.09) allows attackers to cause a permanent Denial of Service (DoS) or potentially execute remote code via the /cgi-bin/UploadCfg endpoint	9.8	More Details
CVE-2026-49352	9Router is an AI router & token saver. From 0.2.21 until 0.4.44, 9Router used the hardcoded fallback JWT secret 9router-default-secret-change-me in src/app/api/auth/login/route.js, src/middleware.js, and later src/lib/auth/dashboardSession.js, allowing attackers to forge an auth_token cookie when JWT_SECRET was unset. This issue is fixed in version 0.4.44	9.8	More Details
CVE-2026-14960	Pegatron `Tdelo64.sys` improperly exposes privileged hardware access functionality through the `\\.\Tdelo` device interface. IOCTL handlers including `TDE_IOCTL_INDEXIO_READ` and `TDE_IOCTL_INDEXIO_WRITE` permit unprivileged user-mode callers to perform arbitrary hardware I/O port reads and writes without authorization checks. A local attacker can abuse this functionality to manipulate hardware registers, tamper with firmware-related interfaces, cause system instability, or establish persistent low-level compromise.	9.8	More Details
CVE-2026-35290	Vulnerability in Oracle Application Testing Suite. The supported version that is affected is 13.3.0.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via TCP to compromise Oracle Application Testing Suite. Successful attacks of this vulnerability can result in takeover of Oracle Application Testing Suite. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-46876	Vulnerability in Oracle Application Testing Suite. The supported version that is affected is 13.3.0.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via Oracle Net to compromise Oracle Application Testing Suite. Successful attacks of this vulnerability can result in takeover of Oracle Application Testing Suite. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-46924	Vulnerability in Oracle Application Testing Suite. The supported version that is affected is 13.3.0.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via TCP to compromise Oracle Application Testing Suite. Successful attacks of this vulnerability can result in takeover of Oracle Application Testing Suite. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-46982	Vulnerability in the Oracle Retail Integration Bus product of Oracle Retail Applications (component: RIB Kernal). The supported version that is affected is 14.1.3.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Retail Integration Bus. Successful attacks of this vulnerability can result in takeover of Oracle Retail Integration Bus. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-46983	Vulnerability in the Oracle Retail Integration Bus product of Oracle Retail Applications (component: RIB Kernal). The supported version that is affected is 16.0.3. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Retail Integration Bus. Successful attacks of this vulnerability can result in takeover of Oracle Retail Integration Bus. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-15982	The Aimogen Pro - All-in-One AI Content Writer, Editor, ChatBot & Automation Toolkit plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 2.8.4. This is due to due to a missing capability check on the 'aiomatic_call_google_ai_function' function. This makes it possible for unauthenticated attackers to leverage the 'aimogen_wp_god_mode' tool to clear function blacklists and execute arbitrary PHP functions, such as creating administrator accounts.	9.8	More Details

CVE-2026-38158	A SQL injection vulnerability in the /ureport/datasource/previewData component of ureport v2.2.9 allows attackers to access sensitive database information via crafted SQL statements.	9.8	More Details
CVE-2026-16395	Integer overflow in the Audio/Video component. This vulnerability was fixed in Firefox 153.	9.8	More Details
CVE-2026-9103	IBM Langflow OSS 1.0.0 through 1.10.0 could allow a remote attacker to gain unauthorized access due to improper authentication in the /api/v1/login/auto_login endpoint. The endpoint issues long-lived superuser bearer tokens without requiring authentication when the AUTO_LOGIN configuration is enabled (enabled by default), which may allow an unauthenticated network attacker to obtain full administrative access. Additionally, permissive cross-origin resource sharing (CORS) settings may allow tokens to be exposed to unintended origins, increasing the risk of unauthorized access.	9.8	More Details
CVE-2026-53399	In the Linux kernel, the following vulnerability has been resolved: nfsd: release layout stid on setlease failure nfs4_alloc_stid() publishes the new stid into cl->cl_stateids via idr_alloc_cyclic() under cl_lock before returning to nfsd4_alloc_layout_stateid(). When nfsd4_layout_setlease() then fails, the error path frees the layout stateid directly with kmem_cache_free() without ever calling idr_remove(), leaving the IDR slot pointing at freed slab memory. Any subsequent IDR walker (states_show, client teardown) dereferences the dangling pointer. The correct teardown for an IDR-published stid is nfs4_put_stid(), which removes the IDR slot under cl_lock, dispatches sc_free (nfsd4_free_layout_stateid) to release ls->ls_file via nfsd4_close_layout(), and drops the nfs4_file reference in its tail. A second issue blocks that switch: nfsd4_free_layout_stateid() unconditionally inspects ls->ls_fence_work via delayed_work_pending() under ls_lock, but INIT_DELAYED_WORK(&ls->ls_fence_work, ...) currently runs only after the setlease call. On the setlease-failure path the destructor would touch an uninitialized delayed_work. nfsd4_alloc_layout_stateid() nfs4_alloc_stid() /* idr_alloc_cyclic under cl_lock */ nfsd4_layout_setlease() /* fails */ nfs4_put_stid() nfsd4_free_layout_stateid() delayed_work_pending(&ls->ls_fence_work) /* needs INIT */ nfsd4_close_layout() /* nfsd_file_put(ls->ls_file) */ put_nfs4_file() Fix by hoisting the ls_fenced / ls_fence_delay / INIT_DELAYED_WORK initialization above the nfsd4_layout_setlease() call, and replace the manual nfsd_file_put + put_nfs4_file + kmem_cache_free cleanup with a single nfs4_put_stid(stp).	9.8	More Details
CVE-2026-53398	In the Linux kernel, the following vulnerability has been resolved: NFSD: Fix SECINFO_NO_NAME decode error cleanup nfsd4_decode_secinfo_no_name() currently initializes sin_exp after decoding sin_style. If the XDR stream is truncated, the decoder returns nfserr_bad_xdr before sin_exp is initialized. Since commit 3fdc54646234 ("NFSD: Reduce amount of struct nfsd4_compoundargs that needs clearing"), the inline iops array is not cleared between RPC calls. A failed SECINFO_NO_NAME decode can therefore leave sin_exp holding stale union contents from a previous operation. The error response path still invokes nfsd4_secinfo_no_name_release(), which calls exp_put() on a non-NULL sin_exp. Initialize sin_exp before the first failable decode step, matching nfsd4_decode_secinfo().	9.8	More Details
CVE-2026-53384	In the Linux kernel, the following vulnerability has been resolved: serial: 8250_dw: unregister 8250 port if clk_notifier_register() fails dw8250_probe() registers the 8250 port via serial8250_register_8250_port() and then, if the device has a clock, registers a clock notifier. If clk_notifier_register() fails, probe returns the error but leaves the 8250 port registered. The matching serial8250_unregister_port() lives in dw8250_remove(), which is not called when probe fails, so the port slot stays occupied until the device is rebound or the system is rebooted. The devm-allocated driver data is freed while the port still references it (via the saved private_data and serial_in/serial_out callbacks), so any access to that port slot before a rebind is a use-after-free hazard. Unregister the port on the clk_notifier_register() error path.	9.8	More Details
CVE-2024-51313	The Tenda TX9 V22.03.02.20 firmware has a stack overflow vulnerability in the sub_42EA38 function of the file /goform/SetVirtualServerCfg.	9.8	More Details
CVE-2026-52348	cool-admin-java 8.0.0 has a SQL injection vulnerability in the order() method of CrudOption.java.	9.8	More Details
CVE-2026-48062	CodeIgniter is a PHP full-stack web framework. Prior to 4.7.3, the ext_in upload validation rule in system/Validation/StrictRules/FileRules.php checked the MIME-derived guessed extension instead of the client-provided filename extension. As a result, an uploaded file named shell.php containing GIF-like content could pass validation such as uploaded[avatar]|is_image[avatar]|mime_in[avatar,image/gif]|ext_in[avatar,gif] because the detected MIME type maps to gif, even though the uploaded filename extension is php. Applications are impacted if they accept user-controlled uploads, rely on ext_in to validate the uploaded filename extension, save uploaded files using the original client filename with $file->move($path), store uploads in a web-accessible directory, and allow PHP or other executable files to run from that directory. In those conditions, this may lead to arbitrary code execution. This issue is fixed in version 4.7.3.	9.8	More Details
CVE-2026-13446	IBM Langflow OSS 1.0.0 through 1.10.1 contains hard-coded credentials, such as a password or cryptographic key, which it uses for its own inbound authentication, outbound communication to external components, or encryption of internal data.	9.8	More Details
	IBM Langflow OSS 1.0.0 through 1.10.0 has a vulnerability in Langflow's webhook authentication logic allows unauthenticated users to trigger the execution of any flow. The system incorrectly bypasses API		

CVE-2026-8505	key validation when the WEBHOOK_AUTH_ENABLE configuration is set to False (which is the default setting). This allows a remote attacker who knows a flow's UUID to execute it as if they were the owner, potentially leading to Remote Code Execution (RCE).	9.8	More Details
CVE-2026-63030	WordPress 6.9.x before 6.9.5 and 7.0.x before 7.0.2 is affected by a REST API batch endpoint route confusion issue which, combined with the author__not_in WP_Query SQL Injection (CVE-2026-60137), could allow an attacker to perform SQL Injection and achieve Remote Code Execution.	9.8	More Details
CVE-2026-16361	Memory safety bugs present in Firefox ESR 115.37 and Firefox ESR 140.12. Some of these bugs showed evidence of memory corruption and we presume that with enough effort some of these could have been exploited to run arbitrary code. This vulnerability was fixed in Firefox ESR 115.38 and Firefox ESR 140.13.	9.8	More Details
CVE-2026-51080	libpvestorage-perl v9.1.1 and libpve-storage-perl v8.3.7 were discovered to contain an XML External Entity (XXE) vulnerability.	9.8	More Details
CVE-2026-36669	An unauthenticated arbitrary file upload vulnerability in ck_upload_handler.php in Feng Office 3.11.13.11 allows remote attackers to upload malicious files (such as .html) to the web-accessible /tmp/ directory.	9.8	More Details
CVE-2026-16408	Integer overflow in the Audio/Video: Playback component. This vulnerability was fixed in Firefox 153.	9.8	More Details
CVE-2026-16411	Memory safety bugs present in Firefox 152. Some of these bugs showed evidence of memory corruption and we presume that with enough effort some of these could have been exploited to run arbitrary code. This vulnerability was fixed in Firefox 153.	9.8	More Details
CVE-2024-51311	The Tenda TX9 V22.03.02.05 firmware has a stack overflow vulnerability in the sub_4418CC function of the file /goform/SetNetControlList.	9.8	More Details
CVE-2026-9202	IBM Langflow OSS 1.0.0 through 1.10.0 allows unauthenticated attackers to create unlimited user accounts on any Langflow instance; when NEW_USER_IS_ACTIVE=true (documented deployment option), newly created accounts are immediately active and can authenticate to reach RCE endpoints, bypassing the need for AUTO_LOGIN.	9.8	More Details
CVE-2026-8297	Improper neutralization of special elements used in an SQL command ('SQL injection') vulnerability in Gis Informatics Engineering Consulting Laboratory R&D and Software Services Inc. GisLab Laboratory Management System allows SQL Injection. This issue affects GisLab Laboratory Management System: from 1.4.03 through 08072026.	9.8	More Details
CVE-2026-60024	The Joomla extension Events Booking prior version 5.8.0 did by default allow unauthenticated users to upload media assets.	9.8	More Details
CVE-2026-12692	Unverified password change vulnerability in Vimesoft Inc. Enterprise Video Platform allows Authentication Bypass. This issue affects Enterprise Video Platform: from 3.11.0.0 before 3.25.0.	9.8	More Details
CVE-2026-9198	IBM Langflow OSS 1.0.0 through 1.10.0 allows unauthenticated attackers to chain /api/v1/auto_login (mints SUPERUSER tokens to any network caller) with /api/v1/validate/code (executes user code via exec()) to achieve full RCE on default Langflow deployments	9.8	More Details
CVE-2026-16412	Memory safety bugs present in Firefox ESR 140.12 and Firefox 152. Some of these bugs showed evidence of memory corruption and we presume that with enough effort some of these could have been exploited to run arbitrary code. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	9.8	More Details
CVE-2026-61451	The Grav API plugin (grav-plugin-api) before 1.0.4 does not validate the origin of the client-supplied admin_base_url field in the POST /api/v1/auth/forgot-password endpoint. The sanitizeHttpUrl() function only checks that the URL scheme is http/https and never verifies the host against the server's own origin, so an attacker can supply an arbitrary host. As a result, an unauthenticated attacker can cause the password reset email sent to a victim to contain a reset link pointing at an attacker-controlled server; when the victim follows the link, the valid reset token is disclosed to the attacker, enabling full account takeover. The vulnerable base URL can also be influenced via the Referer or Origin headers.	9.6	More Details
CVE-2026-62948	OpenWrt is a Linux operating system targeting embedded devices. Prior to 25.12.5, odhcpd writes a DHCPv6 client FQDN option 39 hostname into /tmp/odhcpd.leases through src/statefiles.c statefiles_write_state6() and statefiles_write_state4() without escaping, allowing newline injection of forged lease lines that LuCI rpcd-mod-luci getDHCPLeases displays through htdocs/luci-static/resources/view/status/include/40_dhcp.js and htdocs/luci-static/resources/luci.js dom.append as live HTML in the Active DHCPv6 Leases admin page. This vulnerability is fixed in 25.12.5.	9.6	More Details
CVE-2026-22752	Authentication bypass by primary weakness vulnerability in Spring Security Spring Authorization Server. This issue affects Spring Authorization Server: from 7.0.0 through 7.0.4, from 1.5.0 through 1.5.6, from 1.4.0 through 1.4.9, from 1.3.0 through 1.3.10.	9.6	More Details
	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to		

CVE-2026-47416	0.1.4 are vulnerable to vertical privilege escalation. The `PATCH /workspaces/{workspace_id}/members/{user_id}` endpoint is gated by `require_workspace_member(workspace_id)`, which defaults to `min_role="member"` and is never overridden by the route. The handler then calls `MemberService.update_role(workspace_id, user_id, body.role)` which sets the target member's role to whatever the request body specifies, with no check that the caller has owner-or-admin privilege, no check that the new role is not higher than the caller's own, and no check that the caller is not silently promoting themselves. PraisonAI Platform version 0.1.4 patches the issue.	9.6	More Details
CVE-2026-55518	Avo is a framework to create admin panels for Ruby on Rails apps. Prior to 3.32.1 and 4.0.0.beta.51, Avo's association attach workflow checks attach_<association>? in the UI and GET /resources/:resource/:id/:related/new path, but the actual write endpoint, POST /resources/:resource/:id/:related, does not run the same authorization check before mutating the association through Avo::AssociationsController#create. An authenticated low-privileged Avo user can bypass hidden or disabled attach controls and directly attach related records to a parent record by sending a crafted POST request, which can lead to privilege escalation and cross-tenant data exposure where associations represent authorization-bearing relationships. This issue is fixed in versions 3.32.1 and 4.0.0.beta.51.	9.6	More Details
CVE-2026-54458	WWBN AVideo is an open source video platform. Versions prior to 29.0 contain a stored DOM Cross-Site Scripting vulnerability in the YPTSocket plugin. Any unauthenticated remote attacker can execute arbitrary JavaScript in the authenticated origin of every administrator currently viewing a page that renders the YPTSocket online-users debug panel. plugin/YPTSocket/getWebSocket.json.php issues a signed WebSocket token to any anonymous caller, and MessageSQLiteV2::onOpen at plugin/YPTSocket/MessageSQLiteV2.php lines 91 and 110 reads the attacker-controlled webSocketSelfURI and page_title query parameters from the WebSocket connection URL with no validation. Both values persist into the in-memory SQLite connections table and broadcast inside the users_id_online array sent to every connected client; on the client, plugin/YPTSocket/script.js::updateSocketUserCard interpolates the broadcast page_title into an HTML template literal that is passed to jQuery $.append(html), which parses attacker bytes into live DOM nodes including <img> with inline event handlers. Successful attackers can can read non-HttpOnly cookies and the CSRF token rendered into the admin dashboard, issue authenticated requests to any admin-only endpoint, exfiltrate the admin dashboard DOM, and chain into any admin-context mutation. When the victim is an AVideo administrator, the attacker turns a single anonymous WebSocket connection into full administrative takeover via the admin's own session. This issue has been patched by https://github.com/WWBN/AVideo/commit/8be71e53ccbe9b84b30870db386fb4d2b11e1c16.	9.6	More Details
CVE-2026-53513	Better Auth is an authentication and authorization library for TypeScript. Prior to 1.6.11, the @better-auth/sso plugin's POST /sso/register and POST /sso/update-provider endpoints accept attacker-controlled oidcConfig.userInfoEndpoint, tokenEndpoint, and jwksEndpoint URLs when skipDiscovery: true is set, store them on the ssoProvider row without origin validation, and fetch them during OIDC callback, allowing non-blind server-side request forgery and possible account linking when trustEmailVerified: true is configured. This issue is fixed in version 1.6.11.	9.6	More Details
CVE-2026-47413	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have aprivilege escalation / cross-tenant member injection. The `POST /workspaces/{workspace_id}/members` endpoint is gated only by `require_workspace_member(workspace_id)` (default `min_role="member"`) and forwards the request body's `user_id` and `role` straight into `MemberService.add(workspace_id, user_id, role)`, which has no caller-permission check. A user with the lowest workspace privilege can add any user (including a new attacker-controlled second account, or an existing account they want to grief) as owner of the workspace. PraisonAI Platform version 0.1.4 patches the issue.	9.6	More Details
CVE-2026-65007	The Grav api plugin (grav-plugin-api) before 1.0.8 fails to properly authorize API key generation and revocation: the plugin intercepts the apiKeyGenerate/apiKeyRevoke admin tasks before the account-management ACL runs and authorizes the caller on only the admin.login permission (the baseline permission held by every panel user). This allows any user with admin.login to mint a persistent API key bound to any account, and the forged key inherits the target account's API permissions. On installs where an API-enabled account holds broader permissions, this enables account impersonation and privilege escalation up to account takeover.	9.6	More Details
CVE-2026-16242	A flaw was found in the Konnectivity proxy-server configuration for hosted control planes. The agent-facing listener was started without --cluster-ca-cert (and without token-based agent authentication), so client certificates were not validated. A remote attacker who can reach the Konnectivity cluster endpoint could connect as an unauthenticated agent, join the routing pool, and potentially proxy, inspect, modify, or drop control-plane-to-node traffic.	9.4	More Details
	In the Linux kernel, the following vulnerability has been resolved: net: skmsg: preserve sg.copy across SG transforms The sk_msg sg.copy bitmap is part of the scatterlist entry ownership state. A set bit tells sk_msg_compute_data_pointers() not to expose the entry through writable BPF ctx->data. This protects entries backed by pages that are not private to the sk_msg, such as splice-backed file page-cache pages.		

CVE-2026-63830	Several sk_msg transform paths move, copy, split, or compact msg->sg.data[] entries without moving the matching sg.copy bit. This can make an externally backed entry arrive at a new slot with a clear copy bit. A later SK_MSG verdict can then expose sg_virt(sge) as writable ctx->data and BPF stores can modify the original page cache. Keep sg.copy synchronized with sg.data[] whenever entries are transferred, shifted, split, or copied into a new sk_msg. Clear the bit when an entry is replaced by a newly allocated private page or freed. This covers the BPF pull/push/pop helpers, sk_msg_shift_left/right(), sk_msg_xfer(), and tls_split_open_record(), including the partial tail entry created during TLS open-record splitting.	9.4	More Details
CVE-2026-12693	Authorization bypass through User-Controlled key vulnerability in Vimesoft Inc. Enterprise Video Platform allows Accessing Functionality Not Properly Constrained by ACLs. This issue affects Enterprise Video Platform: from 3.11.0.0 before 3.25.0.	9.4	More Details
CVE-2026-53595	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.224, the public endpoint `POST /user-setup/{hash}/{invite_sent_at}` (`OpenController@userSetupSave`) selects the target account solely by its `invite_hash` column, then overwrites that account's email and password and logs in as it. No authentication, cookie, or prior session is required. After a user activates, FreeScout sets `invite_hash` to the empty string. On MySQL and MariaDB, `VARCHAR` equality ignores trailing spaces, so a single URL-encoded space (`%20`) matches the stored empty string and selects the lowest-id activated user. The expiry guard decrypts `invite_sent_at` with the target's password hash, but `Helper::decrypt` returns its raw input unchanged when decryption fails. A plaintext numeric value such as `9999999999` therefore passes the time-to-live check without any secret. The result is that an anonymous attacker sets the email and password of the lowest-id activated FreeScout account (a support agent, or an administrator if one was added by invitation) and authenticates as that account. Version 1.8.224 contains a fix.	9.4	More Details
CVE-2026-64024	In the Linux kernel, the following vulnerability has been resolved: tcp: fix stale per-CPU tcp_tw_isn leak enabling ISN prediction Blamed commit moved the TIME_WAIT-derived ISN from the skb control block to a per-CPU variable, assuming the value would always be consumed by tcp_conn_request() for the same packet that wrote it. That assumption is violated by multiple drop paths between the producer (__this_cpu_write(tcp_tw_isn, isn) in tcp_v{4,6}_rcv()) and the consumer (tcp_conn_request()): - min_ttl / min_hopcount check - xfrm policy check - tcp_inbound_hash() MD5/AO mismatch - tcp_filter() eBPF/SO_ATTACH_FILTER drop - th->syn && th->fin discard in tcp_rcv_state_process() TCP_LISTEN - psp_sk_rx_policy_check() in tcp_v{4,6}_do_rcv() - tcp_checksum_complete() in tcp_v{4,6}_do_rcv() - tcp_v{4,6}_cookie_check() returning NULL When a packet is dropped on any of these paths, tcp_tw_isn is left set. The next SYN processed on the same CPU then consumes the non zero value in tcp_conn_request(), receiving a potentially predictable ISN. This patch moves back tcp_tw_isn to skb->cb[], getting rid of the per-cpu variable. Note that tcp_v{4,6}_fill_cb() do not set it. Very litle impact on overall code size/complexity: $ scripts/bloat-o-meter -t vmlinux.old vmlinux.new add/remove: 0/0 grow/shrink: 2/1 up/down: 8/-15 (-7) Function old new delta tcp_v6_rcv 3038 3042 +4 tcp_v4_rcv 3035 3039 +4 tcp_conn_request 2938 2923 -15 Total: Before=24436060, After=24436053, chg -0.00%	9.4	More Details
CVE-2026-63764	lmdeploy's OpenAI-compatible API server contains a server-side request forgery vulnerability that allows unauthenticated attackers to access internal services and cloud metadata endpoints by supplying a crafted image_url that redirects to internal targets. Attackers can send a POST request to the chat completions endpoint with an image_url pointing to an attacker-controlled server that responds with an HTTP 302 redirect to internal addresses such as loopback or instance-metadata endpoints, bypassing the initial URL safety check because redirects are followed without re-validating each hop through the safety guard.	9.3	More Details
CVE-2026-65057	Keep (commit 91c75e0) contains a server-side request forgery vulnerability that allows unauthenticated attackers to make the backend issue arbitrary HTTP requests by supplying attacker-controlled host values to the unprotected healthcheck endpoint. Attackers can send a crafted JSON payload with a malicious host parameter to cause the backend to issue outbound requests to internal services or cloud metadata endpoints, enabling theft of cloud credentials and internal network reconnaissance.	9.3	More Details
CVE-2026-65048	Ninja Forms plugin for WordPress versions 3.10.4 through 3.14.9 contains an unauthenticated stored cross-site scripting vulnerability in the Repeatable Fieldset feature where parseSubmissionIndex() accepts arbitrary strings as submission indexes without numeric validation, and admin_form_element() interpolates the index directly into HTML without escaping. An unauthenticated attacker can submit a public form with a crafted repeater child key containing malicious script payloads, which execute in an administrator's browser when viewing submissions in the WordPress admin panel, enabling session-cookie theft, creation of administrator accounts, installation of malicious plugins, and arbitrary modification of site content.	9.3	More Details
CVE-2026-64825	Home Assistant Core before 2026.6.0 contains a path traversal vulnerability that allows unauthenticated attackers to write arbitrary files to any directory on the host filesystem by uploading a crafted backup archive during the initial onboarding window. Attackers can manipulate the 'name' field inside the uploaded archive's backup.json to supply an absolute path, causing pathlib.Path.__truediv__ to discard the configured backup directory prefix and write attacker-controlled content to arbitrary locations, with full filesystem access when the process runs as root.	9.3	More Details

CVE-2026-65049	Ninja Forms plugin version 3.14.8 and prior for WordPress Multisite contains an incorrect authorization vulnerability that allows a subsite Administrator to trigger network-wide deletion of all Ninja Forms data by exploiting a site-scoped capability check combined with unsafe multisite migration defaults. Attackers can send a crafted POST request to the admin-ajax.php endpoint with the nf_delete_all_data action and a per-site nonce to invoke migration routines that unconditionally iterate all blogs via switch_to_blog(), dropping all nf3_* tables and clearing options and transients across every subsite in the network without requiring super-admin or network-admin privileges.	9.3	More Details
CVE-2026-64018	In the Linux kernel, the following vulnerability has been resolved: net: mana: validate rx_req_idx to prevent out-of-bounds array access In mana_hwc_rx_event_handler(), rx_req_idx is derived from sge->address in DMA-coherent memory. In Confidential VMs (SEV-SNP/TDX), this memory is shared unencrypted and HW can modify WQE contents at any time. No bounds check exists on rx_req_idx, which can lead to an out-of-bounds access into reqs[]. Add bounds check on rx_req_idx in mana_hwc_rx_event_handler() before using it to index the reqs[] array.	9.3	More Details
CVE-2026-15091	IBM Engineering AI Hub 1.0.0, 1.1.0, and 1.2.0 could allow a remote attacker to execute arbitrary scripts due to improper neutralization of input during web page generation.	9.3	More Details
CVE-2026-63089	WireGuard Easy through 15.3.0, fixed in commit 66b292b, contains a cryptographically weak one-time link token generation vulnerability that allows unauthenticated network attackers to recover WireGuard peer credentials by brute-forcing a keyspace of at most 1000 candidate tokens per client ID, as the token is computed using CRC32 over a random value constrained to 0-999. Attackers can enumerate candidate tokens against the unauthenticated /cnf/:oneTimeLink route, which lacks rate limiting and does not validate token expiration, to obtain a peer's PrivateKey and PresharedKey and impersonate that peer on the VPN network.	9.3	More Details
CVE-2026-61736	LightRAG provides simple and fast retrieval-augmented generation. Prior to 1.5.4, the server defaults to CORS_ORIGINS=* combined with allow_credentials=True in lightrag/api/lightrag_server.py, causing Starlette CORSMiddleware to effectively whitelist every origin for credentialed cross-origin requests. Any malicious website visited by an authenticated LightRAG user can silently make authenticated API requests, exfiltrating documents and knowledge graph data or performing destructive actions such as deleting the document store. This vulnerability is fixed in 1.5.4.	9.3	More Details
CVE-2026-52842	Lightpanda is a headless browser designed for AI and automation. Prior to 0.3.1, Lightpanda searched for @ across the entire URL string instead of only the authority component when computing a page origin, so a URL such as `http://attacker.com/@victim.com/` was fetched from attacker.com but treated as `http://victim.com`, allowing a complete Same-Origin Policy bypass. This issue is fixed in version 0.3.1.	9.3	More Details
CVE-2026-52843	Lightpanda is a headless browser designed for AI and automation. Prior to 0.2.9, Lightpanda fetch() and XMLHttpRequest unconditionally attached session cookies to every HTTP request, ignoring credentials: omit, credentials: same-origin, credentials: include, and XMLHttpRequest.withCredentials, allowing an attacker-controlled origin in a Lightpanda session to issue authenticated cross-origin requests against a victim origin. This issue is fixed in version 0.2.9.	9.3	More Details
CVE-2026-64080	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_ffa: Snapshot notifier callbacks under lock Both notification handlers currently look up a notifier callback under notify_lock, drop the lock, and then dereference the returned notifier entry. A concurrent unregister can delete and free that entry in the gap, leaving the handler to dereference stale memory. Copy the callback pointer and callback data while notify_lock is still held and invoke the callback only after the lock is dropped. This keeps the existing callback execution model while removing the use-after-free window in both the framework and non-framework notification paths.	9.3	More Details
CVE-2026-39878	Chamilo LMS versions 1.11.38 and earlier contain a stored cross-site scripting vulnerability in the user registration form that allows any unauthenticated attacker to execute arbitrary JavaScript in an administrator's browser session, leading to full platform admin account takeover. This has been patched in 1.11.40.	9.3	More Details
CVE-2026-64034	In the Linux kernel, the following vulnerability has been resolved: net: mana: Fix TOCTOU double-fetch of hwc_msg_id from DMA buffer In mana_hwc_rx_event_handler(), resp->response.hwc_msg_id is read from DMA-coherent memory and bounds-checked, then mana_hwc_handle_resp() re-reads the same field from the same DMA buffer for test_bit() and pointer arithmetic. DMA-coherent memory is mapped uncacheable on x86 and is shared, unencrypted, in Confidential VMs (SEV-SNP/TDX), so each load goes directly to host-visible memory. A H/W can modify the value between the check and the use, bypassing the bounds validation. Fix this by reading hwc_msg_id exactly once using READ_ONCE() into a stack-local variable in mana_hwc_rx_event_handler(), and passing the validated value as a parameter to mana_hwc_handle_resp().	9.3	More Details
CVE-2026-	ZEBRA is a Zcash node written entirely in Rust. Prior to zebrad 5.0.0, halo2_gadgets 0.5.0, orchard 0.14.0, zcash_primitives 0.28.0, and zcashd 6.20.0, the variable-base scalar multiplication gadget in halo2_gadgets/src/ecc/chip/mul/incomplete.rs used assign_advice() for the base point without a copy constraint tying it to the actual base, allowing a malicious prover to produce a valid proof for an Orchard	9.3	More

54496	Action with an under-constrained base point and bypass the diversified-address-integrity check that binds pk_d, g_d, ivk, the nullifier (nf), and the spend validating key (ak) to the note being spent. This issue is fixed in zebrad 5.0.0, halo2_gadgets 0.5.0, orchard 0.14.0, zcash_primitives 0.28.0, and zcashd 6.20.0.		Details
CVE-2026-63940	In the Linux kernel, the following vulnerability has been resolved: KVM: SEV: Ignore Port I/O requests of length '0' Explicitly ignore Port I/O requests of length '0' (or count '0'), so that setting up the software scratch area (and other code) doesn't have to worry about underflowing the length, and to allow for WARNing on trying to configure the scratch area with len==0.	9.3	More Details
CVE-2026-63939	In the Linux kernel, the following vulnerability has been resolved: KVM: SEV: Compute the correct max length of the in-GHCB scratch area When setting the length of the GHCB scratch area, and the area is in the GHCB shared buffer, set the effective length of the scratch area to the max possible size given the start of the guest-provided pointer, and the end of the shared buffer. The code was "fine" when first introduced, as KVM doesn't consult the length of the buffer when emulating MMIO, because the passed in @len always specifies the *max* size required. But for PSC requests, the incoming @len is just the minimum length (to process the header), and KVM needs to know the full size of the scratch area to avoid buffer overflows (spoiler alert). Opportunistically rename @len => @min_len to better reflect its role.	9.3	More Details
CVE-2026-63938	In the Linux kernel, the following vulnerability has been resolved: KVM: SEV: Check PSC request indices against the actual size of the buffer When processing Page State Change (PSC) requests, validate the PSC buffer against the effective size of the scratch area, which could be less than the maximum size if the guest provided a pointer that isn't exactly at the start of the GHCB shared buffer.	9.3	More Details
CVE-2026-49445	Cilium is a networking, observability, and security solution. Prior to 1.17.14, 1.18.8, and 1.19.2, when Cilium L7 functionality is enabled, the embedded or standalone Envoy instance creates a world-accessible admin.sock on cluster nodes, allowing a local attacker to access Envoy admin endpoints, expose TLS secrets, disrupt cluster traffic, or terminate Envoy. This issue is fixed in versions 1.17.14, 1.18.8, and 1.19.2.	9.2	More Details
CVE-2026-62241	clawvet self-hosted API server (apps/api) before 0.7.5 hard-codes a fallback JWT secret ('clawvet-dev-secret-change-me') in auth.ts and ships it as the default in .env.example. Because GET /api/v1/scans returns scan records containing userId values without authentication, a remote unauthenticated attacker can harvest a victim's userId, forge a valid HS256 cg_session cookie offline using the known secret, and call GET /api/v1/auth/me to obtain the victim's email address, subscription plan, and secret apiKey. The published clawvet npm package (CLI only) is not affected.	9.1	More Details
CVE-2026-57075	YAML::Syck versions before 1.47 for Perl allow an out-of-bounds read via a signed-char lookup-table index in syck_base64dec. The base64 decoder in the bundled libsyck indexes the 256-entry static table b64_xtable with a signed char, so any !!binary byte >= 0x80 sign-extends to a negative index and reads before the table. The decoder receives the raw bytes of any !!binary node, a standard YAML type not gated by $LoadBlessed or $LoadCode, so it is reached on the default Load path. Any caller that runs Load or LoadFile on an untrusted document containing a !!binary scalar with a high-bit byte triggers the read, and the value read can surface in the decoded result.	9.1	More Details
CVE-2026-44231	RT is an open source, enterprise-grade issue and ticket tracking system. Versions prior to 5.0.10, 6.0.0 and above, prior to 6.0.3 contain an information disclosure and privilege escalation vulnerability in the REST 2.0 API. A privileged (non-administrative) user can obtain authentication credentials belonging to other users — including users with administrative privileges — and use those credentials to read data as those users via RT's feed endpoints. The same request that exposes the credentials also rotates them, invalidating previously-distributed feed URLs across the instance. This issue has been fixed in versions 5.0.10 and 6.0.3.	9.1	More Details
CVE-2026-28321	SolarWinds Serv-U is affected by a broken access control vulnerability that could allow arbitrary file read and write, which can then be used to escalate privileges and execute code as root. A domain administrator access is required, and the impact is lower in Windows installations.	9.1	More Details
CVE-2026-28317	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to privilege escalation. This issue requires domain administrator access. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28314	SolarWinds Serv-U is affected by an insecure direct object reference vulnerability that leads to an account takeover. User authentication is required. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28316	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to privilege escalation to a system administrator with the ability to execute commands as the root user. This issue requires a domain account with administrator access. The impact is lower in Windows deployments.	9.1	More Details
	Yamcs is a mission control framework. Prior to 5.12.7, a server-side code injection vulnerability existed in the Yamcs algorithm evaluation engine org.yamcs.algorithms.JavaExprAlgorithmExecutionFactory, which		

CVE-2026-44632	dynamically compiled and evaluated user-controlled algorithm text through the Janino compiler without enforcing a secure sandbox, so an authenticated user with the ChangeMissionDatabase privilege could override an existing algorithm's text via the mission database REST API and inject Java code (for example using java.lang.Runtime) to achieve remote code execution on the underlying host operating system. This issue is fixed in versions 5.12.7 and 5.13.0, which disable algorithm editing by default.	9.1	More Details
CVE-2026-57074	XML::Bare versions through 0.53 for Perl have an unbounded character lookahead. The parserc_parse function attempts to check for multicharacter strings such as "<![CDATA" or element terminators such as ">" without checking that the offsets are within the buffer. Truncated strings such as "<a/" can trigger an out-of-bounds read.	9.1	More Details
CVE-2026-57073	HTML::Bare versions through 0.04 for Perl have an unbounded character lookahead. The parserc_parse function attempts to check for multicharacter strings such as "<![CDATA" or element terminators such as ">" without checking that the offsets are within the buffer. Truncated strings such as "<a/" can trigger an out-of-bounds read. Note that the latest version available on CPAN is version 0.02. Newer versions are available on the git repository.	9.1	More Details
CVE-2026-46621	Yamcs is a mission control framework. Prior to 5.12.7, the Yamcs script evaluation engine for Python algorithms dynamically compiled and evaluated user-controlled algorithm text using Jython through the JSR-223 ScriptEngine API without enforcing a secure sandbox, so an authenticated user with the ChangeMissionDatabase privilege could override an existing Python algorithm's logic through the mission database REST API and import and execute arbitrary Java classes such as java.lang.Runtime to achieve remote code execution on the underlying host operating system. This issue is fixed in versions 5.12.7 and 5.13.0, which disable algorithm editing by default.	9.1	More Details
CVE-2026-45568	zrok is software for sharing web services, files, and network resources. Prior to 2.0.3, zrok's Python SDK ProxyShare Flask proxy route accepts an absolute URL in the request path and passes it to urllib.parse.urljoin, allowing the requested path to replace the configured target host and causing requests.request to return a server-side response from an attacker-chosen URL. This issue is fixed in version 2.0.3.	9.1	More Details
CVE-2026-28312	SolarWinds Serv-U is affected by a privilege escalation vulnerability. This would elevate a group’s access to system administrator and allow code execution as root. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-14890	SGLang uses an expert-parallel backup subsystem that exposes a ZeroMQ PULL socket on a routable network interface that does not contain authentication or deserialization safeguards, allowing an attacker to provide a malicious pickle file that results in unauthenticated remote code execution when the feature is enabled and the service is reachable over the network.	9.1	More Details
CVE-2026-26718	A Cross-Site Request Forgery (CSRF) vulnerability exists in the xxl-job-admin web application v.3.0.0 that allows an attacker to perform unauthorized modifications to Glue IDE shell scripts. The affected endpoint lacks proper CSRF token validation and accepts arbitrary HTTP methods via a permissive request mapping	9.1	More Details
CVE-2026-53512	Better Auth is an authentication and authorization library for TypeScript. Prior to 1.6.11, the legacy oidcProvider and mcp plugins expose OAuth token endpoints whose refresh_token grant authenticates only possession of the bound refreshToken row and matching client_id, without verifying the confidential client's client_secret, allowing an attacker with a valid refresh_token to mint access tokens and rotated refresh tokens through /api/auth/oauth2/token or /api/auth/mcp/token. The @better-auth/oauth-provider package is not affected. This issue is fixed in version 1.6.11.	9.1	More Details
CVE-2026-43637	Cornac before 2.6.0 contains a path traversal (Tar Slip) vulnerability that allows attackers to write arbitrary files outside the intended cache directory by supplying a crafted TAR archive containing ../ sequences, absolute paths, or symlink/hardlink entries to the _extract_archive() function in cornac/utils/download.py. Attackers can trigger this vulnerability through the built-in dataset loaders, which automatically download and extract archives, causing archive.extractall() to write files to arbitrary locations on the filesystem accessible to the running process.	9.1	More Details
CVE-2026-28313	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to SMTP hijacking leading to arbitrary account takeover. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28307	SolarWinds Serv-U is affected by a privilege escalation vulnerability that allows a domain user group to be elevated into an administrator group. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28310	SolarWinds Serv-U is affected by a privilege escalation vulnerability that allows a domain administrator to escalate their user type to that of a system administrator. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-16406	Mitigation bypass in the Networking component. This vulnerability was fixed in Firefox 153.	9.1	More Details
CVE-2026-	The Kirki WordPress plugin before 6.0.12 does not validate a user-supplied URL before requesting it server-side, allowing unauthenticated attackers to make the site issue HTTP requests to arbitrary hosts		More

13147	(Server-Side Request Forgery).	9.1	Details
CVE-2026-62415	The Joomla extension Membership Pro prior version 4.6.2 did by default allow unauthenticated users to upload media assets.	9.1	More Details
CVE-2026-64609	Out-of-bounds read via sun.misc.Unsafe in Apache Fory. When out-of-band zero-copy deserialization is used, readAlignedVarUint() can read beyond the bounds of the underlying buffer. Out-of-band zero-copy deserialization is an opt-in feature; applications that do not use it are not affected. This issue affects Apache Fory (formerly Apache Fury): from 0.5.0 before 1.4.0. Versions before 0.11.0 were published under the Maven coordinates org.apache.fury:fury-core. Users are recommended to upgrade to version 1.4.0, which fixes the issue.	9.1	More Details
CVE-2026-63992	In the Linux kernel, the following vulnerability has been resolved: tunnels: do not assume transport header in iptunnel_pmtud_check_icmp() In some cases, iptunnel_pmtud_check_icmp() can be called while skb transport header is not set. This triggers an out-of-bound access, because (typeof(skb->transport_header))~0U is 65535. Access the icmp header based on IPv4 network header, after making sure icmp->type is present in skb linear part. Note that iptunnel_pmtud_check_icmpv6()) is fine.	9.1	More Details
CVE-2026-16359	Incorrect boundary conditions in the Audio/Video: GMP component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	9.1	More Details
CVE-2026-52199	An issue in Generic OEM UZ801_v2.1 4G LTE Router V3.4.3 allows a remote attacker to execute arbitrary code via the sbin/adbd component	9.1	More Details
CVE-2026-28309	SolarWinds Serv-U is affected by a broken access control vulnerability that allows a domain administrator to create system administrator accounts. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-42168	django-pyas2 through 1.2.3 is vulnerable to OS command injection via the cmd_receive and cmd_send fields on the Partner model. These fields are passed directly to os.system() in pyas2/utils.py without sanitization, allowing an authenticated admin user to execute arbitrary commands on the server when an AS2 message is received or sent.	9.1	More Details
CVE-2026-16394	Mitigation bypass in the DOM: Security component. This vulnerability was fixed in Firefox 153.	9.1	More Details
CVE-2025-51677	An issue was discovered in openRISC OR1200 commit 83ac6b. An output mismatch between the RTL and the netlist of the or1200 cpu output port can lead to unexpected behavior.	9.1	More Details
CVE-2026-12694	Missing Authorization vulnerability in Vimesoft Inc. Enterprise Video Platform allows Accessing Functionality Not Properly Constrained by ACLs. This issue affects Enterprise Video Platform: from 3.11.0.0 before 3.25.0.	9.1	More Details
CVE-2024-23564	HCL Aftermarket EPC is affected by Business Logic Vulnerability using which a non valid user of the application can obtain passwords from the server and redirect them to their own email address by manipulating the server's response. The application includes checks in the initial requests to verify the validity of the provided UserId, but similar validation is not applied to Email requests when sending passwords to user emails.	9.1	More Details
CVE-2026-28302	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to privilege escalation and remote code execution as root. This issue requires group administrator access. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28304	SolarWinds Serv-U is affected by a remote code execution vulnerability that, when exploited, can allow the arbitrary execution of code remotely as root. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28305	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to remote code execution as root. A domain account with admin privileges and read and write access to the home directory is required. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28306	SolarWinds Serv-U is affected by a privilege escalation vulnerability that allows a domain administrator to elevate their privileges to a system administrator. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-28308	SolarWinds Serv-U is affected by an insecure direct object reference (IDOR) vulnerability that can lead to remote code execution. Domain administrator access is required. The impact is lower in Windows deployments.	9.1	More Details
CVE-2026-	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: UI Framework). Supported versions that are affected are 13.5 and 24.1. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. While the vulnerability is in Oracle Enterprise Manager Base Platform, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Enterprise	9.1	More

46989	Manager Base Platform accessible data as well as unauthorized update, insert or delete access to some of Oracle Enterprise Manager Base Platform accessible data and unauthorized ability to cause a partial denial of service (partial DOS) of Oracle Enterprise Manager Base Platform. CVSS 3.1 Base Score 9.1 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:L/A:L).		Details
CVE-2026-11386	An input validation and injection vulnerability exists in Canonical ubuntu-pro-client (formerly ubuntu-advantage-tools). The client constructs APT source files (such as /etc/apt/sources.list.d/ubuntu-.list or their DEB822 equivalents) using data received directly from the contract server response via the directives.suites[] and directives.aptURL fields. Because the client utilizes Python's str.format() to write these files without performing escaping, validation, or newline character filtering, a malicious or tampered contract response containing embedded newline (\n) characters can successfully inject arbitrary, attacker-controlled deb configuration lines into root-owned APT sources. When combined with the unvalidated additionalPackages[] field—which is passed positionally into a root-executed apt-get install command—an attacker capable of spoofing or manipulating the contract response (e.g., via a compromised internal infrastructure, an intercepted connection utilizing a trusted CA, or local logical bugs) can force the client to fetch and install malicious packages. This ultimately leads to arbitrary code execution with root privileges on the affected system. This component is preinstalled on supported Ubuntu Server releases and auto-attaches by default on cloud provider Ubuntu Pro images.	9.0	More Details
CVE-2026-62378	RustFS Console is a web management console for the RustFS distributed file system. From 0.1.7 until 0.1.10, the RustFS Console components/object/preview-modal.tsx and components/object/pdf-viewer.tsx extension-based PDF preview path can render HTML content uploaded as .pdf, allowing stored cross-site scripting in the management console and exposure of administrator AccessKeyId, SecretAccessKey, and SessionToken values. This is caused by a regression of CVE-2026-27822. This vulnerability is fixed in 0.1.10.	9.0	More Details
CVE-2026-64106	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: vgic-its: Reject restored DTE with out-of-range num_eventid_bits Userspace can restore an ITS Device Table Entry whose Size field encodes more EventID bits than the virtual ITS supports. The live MAPD path rejects that state, but vgic_its_restore_dte() accepts it and stores the out-of-range value in dev->num_eventid_bits. Reject restored DTEs with num_eventid_bits > VITS_TYPER_IDBITS before allocating the device. This mirrors the MAPD check and prevents the restored state from reaching vgic_its_restore_itt(), where the unchecked value can be converted into an oversized scan_its_table() range.	9.0	More Details
CVE-2026-12701	A path traversal vulnerability was found in pulpcore. The relative_path_validator function only verifies that content paths do not begin with "/" but fails to block directory traversal sequences such as "../" anywhere in the path. An authenticated administrator can craft a relative_path containing embedded traversal sequences (e.g., "looking/normal/../../../../etc/shadow") that escapes the intended export directory during FilesystemExport operations. Because the file content is also user-controlled (uploaded artifact), this allows arbitrary file write to any location writable by the Pulp service user, potentially leading to service compromise or further system exploitation.	9.0	More Details
CVE-2026-35198	HeyForm is an open-source form builder. Prior to version 3.0.0-rc.7, a stored cross-site scripting (XSS) vulnerability in the form builder allows a low-privileged team member to inject malicious JavaScript that executes when a team owner views the form, leading to complete account takeover through privilege escalation. Version 3.0.0-rc.7 contains a patch for the issue.	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2026-15103	The WPFunnels – Funnel Builder for WooCommerce with Checkout & One Click Upsell plugin for WordPress is vulnerable to Privilege Escalation via arbitrary option update in all versions up to, and including, 3.12.8. This is due to the `update_settings()` REST callback failing to validate the `group_id` path parameter against an allowlist of permitted option names before passing it directly to `get_option()` and `update_option()`, allowing the built-in `wp_user_roles` option — which satisfies the route's loose `[\w-]+` regex — to be targeted. This makes it possible for authenticated attackers with the `wpf_manage_funnels` capability and above to elevate their privileges to administrator by writing a crafted role definition containing arbitrary capabilities into the `wp_user_roles` option, thereby granting any WordPress role full site administrator access. The `wpf_manage_funnels` capability is typically assigned to the Funnel Manager custom role created by the plugin, meaning this role is the minimum required to exploit the vulnerability.	8.8	More Details
CVE-2026-59851	A flaw was found in libssh. On servers with GSSAPIKeyExchange enabled, the gssapi-keyex path does not verify whether the authenticated Kerberos principal is authorized for the requested local user, allowing authenticated clients to log in as arbitrary users.	8.8	More Details
CVE-2026-	OpenClaw versions before 2026.6.5 contain an authentication bypass vulnerability that allows lower-trust callers to reach admin-scoped tools. Attackers can perform actions requiring stronger authorization by	8.8	More

62207	exploiting insufficient policy checks on configured input paths.		Details
CVE-2026-63919	In the Linux kernel, the following vulnerability has been resolved: xfrm: input: hold netns during deferred transport reinjection Transport-mode reinjection stores a struct net pointer in skb->cb and uses it later from xfrm_trans_reinject(). That pointer must stay valid until the deferred callback runs. Take a netns reference when queueing deferred reinjection work and drop it after the callback completes. Use maybe_get_net() so the queueing path does not revive a namespace that is already being torn down. This keeps the existing workqueue design and fixes the netns lifetime handling in one place for all users of xfrm_trans_queue_net().	8.8	More Details
CVE-2026-16097	A vulnerability was found in Shibby Tomato 1.28. This vulnerability affects the function sub_42537C of the component Scheduler Name Handler. The manipulation of the argument a1 results in stack-based buffer overflow. It is possible to launch the attack remotely. This project is superseded by FreshTomato.	8.8	More Details
CVE-2026-16096	A vulnerability has been found in Shibby Tomato 1.28 RT-N5x MIPSR2 Build 124. This affects the function sub_40BB50 of the file /proc/webmon_recent_domains. The manipulation leads to stack-based buffer overflow. It is possible to initiate the attack remotely. This project is superseded by FreshTomato.	8.8	More Details
CVE-2026-16095	A flaw has been found in Shibby Tomato 1.28 RT-N5x MIPSR2 Build 124. Affected by this issue is the function setup_conntrack of the file /sbin/rc. Executing a manipulation of the argument ct_tcp_timeout can lead to out-of-bounds write. The attack may be performed from remote. This project is superseded by FreshTomato.	8.8	More Details
CVE-2026-11767	The Free Builder for Elementor WordPress plugin before 1.6.7 does not sanitise submitted contact form field values before storing them and outputting them in the admin dashboard, allowing unauthenticated attackers to perform Stored Cross-Site Scripting attacks that execute when a logged-in administrator views the form submissions.	8.8	More Details
CVE-2026-47871	VMware Avi Load Balancer contains a directory traversal vulnerability. Flaws in file path validation allow malicious, authenticated network users to perform directory traversal attacks. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	8.8	More Details
CVE-2026-5674	A flaw was found in PipeWire, a multimedia server. This vulnerability allows an attacker to escape sandboxed applications, such as Flatpak, by exploiting PipeWire's PulseAudio compatibility layer. An attacker with minimal permissions within a sandboxed environment can load a malicious library, leading to arbitrary code execution outside the sandbox and potential compromise of the user's system.	8.8	More Details
CVE-2026-63921	In the Linux kernel, the following vulnerability has been resolved: ip6: vti: Use ip6_tnl.net in vti6_siocdevprivate(). After patch 1/2 in this series, vti6_update() unlinks and relinks the tunnel through t->net. vti6_siocdevprivate() still uses dev_net(dev) for the collision lookup. For a tunnel moved through IFLA_NET_NS_FD, dev_net(dev) is the new netns, not t->net. SIOCCHGTUNNEL on a migrated tunnel then runs: net = dev_net(dev) /* migrated netns */ t = vti6_locate(net, &p1, false) /* misses target in t->net */ ... t = netdev_priv(dev) vti6_update(t, &p1, false) /* mutates t->net's hash */ A caller in the migrated netns picks params that match a tunnel in the creation netns. The lookup in dev_net(dev) finds nothing. vti6_update() prepends the migrated tunnel at the head of the creation netns hash bucket for those params. Later lookups in the creation netns resolve to the migrated device. xfrm receive delivers the matched packets through a device the caller controls. Reachable from an unprivileged user namespace (unshare --user --map-root-user --net). Cross tenant scope on container hosts. Switch the SIOCCHGTUNNEL path on a non fallback device to use t->net for the lookup. The lookup now matches the netns vti6_update() operates on. Also add ns_capable(self->net->user_ns, CAP_NET_ADMIN) before the lookup. The check at the top of the case is against dev_net(dev)->user_ns, which after migration is the attacker's netns. A caller there can pick params absent from self->net, the lookup returns NULL, t becomes self, and vti6_update() inserts the device into the creation netns hash. The new check requires CAP_NET_ADMIN in the creation netns user_ns too. SIOCADDTUNNEL and SIOCCHGTUNNEL on the fallback device keep dev_net(dev), which equals init_net there.	8.8	More Details
CVE-2026-62217	OpenClaw 2026.5.14-beta.1 before 2026.5.27 contain an authorization flaw in the QQBot exec approvals feature. When the feature is enabled and reachable, a lower-trust caller or configured input path could execute or persist actions beyond the caller's intended authorization, allowing non-allowlisted senders to perform unauthorized operations.	8.8	More Details
CVE-2026-64124	In the Linux kernel, the following vulnerability has been resolved: net: devmem: reject dma-buf bind with non-page-aligned size or SG length net_devmem_bind_dmabuf() trusts dmabuf->size and sg_dma_len() to be PAGE_SIZE multiples without checking: - tx_vec is sized dmabuf->size / PAGE_SIZE, and net_devmem_get_niov_at() only bounds-checks virt_addr < dmabuf->size before indexing tx_vec[virt_addr / PAGE_SIZE]. With size = N*PAGE_SIZE + r (1 <= r < PAGE_SIZE), sendmsg() at iov_base = N*PAGE_SIZE passes the bound check and reads tx_vec[N] -- one past. - owner->area.num_niovs = len / PAGE_SIZE while gen_pool_add_owner() covers the full byte len, so a non-page-multiple non-final sg desyncs num_niovs from the gen_pool region for every later sg, on both RX and TX. dma-buf does not require page-aligned sizes, so the bind path has to enforce what its own indexing assumes. Reject both with -EINVAL. The size check is TX-only (only tx_vec is sized off dmabuf->size); the SG-length check covers both directions.	8.8	More Details

CVE-2026-62218	OpenClaw 2026.1.20 before 2026.5.27 contain an authorization bypass vulnerability in the device.pair.approve feature that allows lower-trust callers to bypass role-management checks. Attackers can perform actions requiring stronger authorization by reaching the affected feature through configured input paths.	8.8	More Details
CVE-2026-25039	Parsec is a cloud-based application for simple and cryptographically secure file sharing. The application does not sanitize the workspace name, creating a vulnerability if that workspace name is a UNC path. When creating mountpoint in the windows filesystem to mount the workspace of an organization, the application does not sanitize the workspace name. The cause issue if the workspace name evaluate to a UNC path since it's allowed for the name to containt `\` char. If the UNC path is invalid (or the targeted resource is not available) the application become unresponsive otherwise the system will interact with the mounted UNC path allowing the attacker to retrieve to [`NTLM`] hash.	8.8	More Details
CVE-2026-63923	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: validate body pcifunc in rvu_mbox_handler_rep_event_notify rvu_mbox_handler_rep_event_notify() in drivers/net/ethernet/marvell/octeontx2/af/rvu_rep.c queues a sender-controlled REP_EVENT_NOTIFY request body verbatim, and rvu_rep_up_notify() then forwards event->pcifunc (the nested body field, distinct from the AF-normalised header pcifunc) into rvu_get_pfvf(), rvu_get_pf() and the AF->PF mailbox device index without any bounds check. A VF attached to a PF that has been put into switchdev representor mode reaches this path: the VF mailbox handler otx2_pfvf_mbox_handler() forwards every message id including MBOX_MSG_REP_EVENT_NOTIFY to AF without an allowlist, and the AF dispatcher rewrites only msg->pcifunc, leaving struct rep_event::pcifunc attacker-controlled. The sibling rvu_mbox_handler_esw_cfg() refuses requests whose header pcifunc is not rvu->rep_pcifunc; this handler has no equivalent gate. An out-of-range body pcifunc selects an &rvu->pf[]/&rvu->hwvf[] element past the allocated array and, for RVU_EVENT_MAC_ADDR_CHANGE, turns into a six-byte attacker-chosen OOB ether_addr_copy() target inside the queued worker; KASAN reports a slab-out-of-bounds write in rvu_rep_wq_handler. Reject malformed requests at the handler entry by gating on is_pf_func_valid(), which is already the canonical PF/VF range check in this driver; expose it via rvu.h so callers in rvu_rep.c can use it instead of open-coding the same range arithmetic.	8.8	More Details
CVE-2026-63807	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/mmu: Ensure hugepage is in by slot before checking max mapping level When recovering hugepages in the shadow MMU, verify that the base gfn of the shadow page is actually contained within the target memslot, *before* querying the max mapping level given the shadow page's gfn. Failure to pre-check the validity of the gfn can lead to an out-of-bounds access to the slot's lpage_info (which typically manifests as a host #PF because the lpage_info is vmalloc'd) if the guest creates a hugepage mapping (in its PTEs) that extends "below" the bounds of a memslot. When faulting in memory for a guest, and the size of the guest mapping is greater than KVM's (current) max mapping, then KVM will create a "direct" shadow page (direct in that there are no gPTEs to shadow, and so the target gfn is a direct calculation given the base gfn of the shadow page). The hugepage recovery flow looks for such direct shadow pages, as forcing 4KiB mappings when dirty logging generates the guest > host mapping size case. When the 4KiB restriction is lifted, then KVM can replace the shadow page with a hugepage. But if KVM originally used a smaller mapping than the guest because the range of memory covered by the guest hugepage exceeds the bounds of a memslot, then KVM will link a direct shadow page with a gfn that is outside the bounds of the memslot being used to fault in memory. The rmap entry added for the leaf mapping is correct and within bounds, but the gfn of the leaf SPTE's parent shadow page will be out of bounds. BUG: unable to handle page fault for address: ffffc90000806ffc #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 100000067 P4D 100000067 PUD 1002a7067 PMD 10612f067 PTE 0 Oops: Oops: 0000 [#1] SMP CPU: 13 UID: 1000 PID: 757 Comm: mmu_stress_test Not tainted 7.1.0-rc1-48ce1e26eace-x86_pir_to_irr_comments-vm #341 PREEMPT Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 0.0.0 02/06/2015 RIP: 0010:kvm_mmu_max_mapping_level+0x79/0x2b0 [kvm] Call Trace: <TASK> kvm_mmu_recover_huge_pages+0x21b/0x320 [kvm] kvm_set_memslot+0x1ee/0x590 [kvm] kvm_set_memory_region.part.0+0x3a1/0x4d0 [kvm] kvm_vm_ioctl+0x9bf/0x15d0 [kvm] __x64_sys_ioctl+0x8a/0xd0 do_syscall_64+0xb7/0xbb0 entry_SYSCALL_64_after_hwframe+0x4b/0x53 RIP: 0033:0x7f21c0f1a9bf </TASK> Don't bother pre-checking the bounds of the potential hugepage, i.e. don't check that e.g. sp->gfn + KVM_PAGES_PER_HPAGE(sp->role.level + 1) is also within the memslot, as the checks performed by kvm_mmu_max_mapping_level() are a superset of the basic bounds checks. I.e. pre-checking the full range would be a dubious micro-optimization.	8.8	More Details
CVE-2026-63866	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: Clear wcid pointer in mt7996_mac_sta_deinit_link() Clear WCID pointer removing the sta link in mt7996_mac_sta_deinit_link routine.	8.8	More Details
CVE-2026-63865	In the Linux kernel, the following vulnerability has been resolved: bpf: Drop task_to_inode and inet_conn_established from lsm sleepable hooks bpf_lsm_task_to_inode() is called under rcu_read_lock() and bpf_lsm_inet_conn_established() is called from softirq context, so neither hook can be used by sleepable LSM programs.	8.8	More Details
CVE-	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a heap-based buffer overflow vulnerability within the virtual channel forwarding mechanism. When forwarding data from a remote client to the internal channel server, the xrdp process utilizes a fixed-size buffer without adequate bounds checking on		

2026-44178	the incoming payload. An authenticated remote attacker can exploit this flaw by sending a specially crafted virtual channel message that exceeds the buffer capacity, leading to heap memory corruption. This may result in a denial of service or the execution of arbitrary code with the privileges of the xrdp process. This issue has been fixed in version 0.10.6.1.	8.8	More Details
CVE-2026-47399	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Prior to version 0.1.4, the workspace-scoped REST routes contain a systemic object-level authorization flaw that allows an authenticated user from one workspace to access, modify, and delete objects belonging to another workspace by supplying the victim object's global UUID. The affected pattern appears in workspace-scoped routes such as agents, projects, issues, and comments. The route layer verifies that the caller is a member of the `workspace_id` provided in the URL, but the service layer later resolves the target object by global object ID only. It does not verify that the resolved object actually belongs to the workspace in the URL. As a result, a valid member of `workspace_attacker` can call a route under `/api/v1/workspaces/{workspace_attacker}/...` while supplying an object UUID from `workspace_victim`. The server authorizes the request based on membership in `workspace_attacker`, then fetches or mutates the victim object by global UUID. This breaks the platform's workspace isolation boundary. PraisonAI Platform version 0.1.4 patches the issue.	8.8	More Details
CVE-2026-47405	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have a broken workspace authorization check that allows any authenticated low-privilege workspace member to escalate their own role to `owner`. The issue is caused by privileged workspace-management routes using the shared dependency `require_workspace_member(...)` without requiring `admin` or `owner`. The dependency defaults to `min_role="member"`, so routes that should be administrative are accessible to ordinary workspace members. As a result, a normal workspace member can promote their own account from `member` to `owner`; add arbitrary users as `owner` or `admin`; change other members' roles; remove legitimate owners or members; take over workspace membership completely; and/or perform destructive workspace operations after escalation. This is a broken access control / vertical privilege escalation vulnerability. PraisonAI Platform version 0.1.4 patches the issue.	8.8	More Details
CVE-2026-64138	In the Linux kernel, the following vulnerability has been resolved: ksmbd: validate SID in parent security descriptor during ACL inheritance Introduce smb_validate_ntsd_sid() helper to safely validate Owner SID and Group SID inside the NT Security Descriptor (smb_ntsd) retrieved from the parent directory.	8.8	More Details
CVE-2026-62223	OpenClaw before 2026.5.18 contain an authorization bypass vulnerability in the device-pair approval feature that allows lower-trust callers to execute actions beyond their intended authorization. Attackers can exploit misconfigured input paths to execute or persist unauthorized actions when the affected feature is enabled and reachable.	8.8	More Details
CVE-2026-44880	A buffer overflow vulnerability was found in the command line interface of AOS-CX. Successful exploitation of these vulnerabilities could allow an remote low-privileged user to execute arbitrary code as a privileged user on the underlying operating system.	8.8	More Details
CVE-2026-63937	In the Linux kernel, the following vulnerability has been resolved: KVM: SEV: Use READ_ONCE() when reading entries/indices from PSC buffer Use READ_ONCE() when reading entries/indices from the guest-accessible Page State Change buffer to defend against TOCTOU bugs. Don't bother with READ_ONCE()/WRITE_ONCE() for cases where KVM is writing (and not consuming the result!), as the guest isn't supposed to touch the buffer while it's being processed. I.e. using READ_ONCE() is all about protecting against misbehaving guests.	8.8	More Details
CVE-2026-62203	OpenClaw versions before 2026.6.6 contain an environment variable filtering vulnerability in host exec that fails to properly sanitize rustup startup variables. Attackers with lower-trust caller access or configured input paths can execute or persist actions beyond their intended authorization level.	8.8	More Details
CVE-2026-63917	In the Linux kernel, the following vulnerability has been resolved: ip6: vti: Use ip6_tnl.net in vti6_changelink(). ip netns add ns1 ip netns add ns2 ip -n ns1 link add vti6_test type vti6 remote ::1 local ::2 key 7 ip -n ns1 link set vti6_test netns ns2 ip -n ns2 link set vti6_test type vti6 remote ::3 local ::4 key 9 ip netns del ns2 ip netns del ns1 [132.495484] ------------[cut here]------------ [132.497609] kernel BUG at net/core/dev.c:12376! Commit 61220ab34948 ("vti6: Enable namespace changing") dropped NETIF_F_NETNS_LOCAL from vti6 devices. A vti6 tunnel can then move through IFLA_NET_NS_FD. After the move dev_net(dev) points at the new netns while t->net stays at the creation netns. vti6_changelink() and vti6_update() still use dev_net(dev) and dev_net(t->dev). They unlink from one per netns hash and relink into another. The creation netns is left with a stale entry. cleanup_net() of that netns later walks freed memory. Reachable from an unprivileged user namespace (unshare --user --map-root-user --net). Cross tenant scope on container hosts.	8.8	More Details
CVE-2026-62202	OpenClaw versions 2026.6.1 before 2026.6.9 contain a privilege escalation vulnerability in isolated cron jobs that allows lower-trust callers to regain denied execution tools. Attackers can execute or persist actions beyond their intended authorization by leveraging misconfigured input paths in the affected cron feature.	8.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: capture fast-RX rate before mesh reuses skb->cb ieee80211_invoke_fast_rx() reads RX status through IEEE80211_SKB_RXCB(skb), which aliases the same skb->cb storage that ieee80211_rx_mesh_data() reuses as IEEE80211_TX_INFO. In the unicast forward path, mesh_data does: info = IEEE80211_SKB_CB(fwd_skb); memset(info, 0, sizeof(*info)); on the same skb the caller still names via rx->skb, then either queues the skb for TX (success) or kfree_skb()'s it	8.8	More

64117	(no-route) before returning RX_QUEUED. The caller's RX_QUEUED arm then calls sta_stats_encode_rate(status) on memory that is either zeroed (success path) or freed (no-route path). The latter is KASAN slab-use-after-free in ieee80211_prepare_and_rx_handle. Fix by encoding the rate from status before invoking ieee80211_rx_mesh_data(), so the RX_QUEUED arm consumes a value captured while status was still backed by valid memory.		Details
CVE-2026-64096	In the Linux kernel, the following vulnerability has been resolved: batman-adv: mcast: fix use-after-free in orig_node RCU release batadv_mcast_purge_orig() removes entries from RCU-protected hlists but does not wait for an RCU grace period before returning. Concurrent RCU readers may still accesses references to those entries at the point of removal. RCU-protected readers trying to operate on entries like orig->mcast_want_all_ipv6_node will then access already freed memory. Fix this by moving batadv_mcast_purge_orig() to batadv_orig_node_release(), just before the call_rcu() invocation. This ensures RCU readers that were active at purge time have drained before the orig_node memory is reclaimed.	8.8	More Details
CVE-2026-55242	ERPNext is a free and open source Enterprise Resource Planning tool. Prior to 15.111.0 and 16.22.0, an authenticated user with a standard operational role can trigger server-side template injection through a configuration field, resulting in unauthorized disclosure of data outside the user's normal permission scope. This issue is fixed in versions 15.111.0 and 16.22.0.	8.8	More Details
CVE-2026-64088	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: fix negative tt_buff_len batadv_orig_node::tt_buff_len was declared as s16, but the field is never intended to hold a negative value. When a value greater than 32767 is assigned, it wraps to a negative signed integer. In batadv_send_other_tt_response(), tt_buff_len is temporarily widened to s32. The incorrectly negative s16 value propagates into the s32, causing batadv_tt_prepare_tvlv_global_data() to allocate a full sized buffer but populates only a small portion of it with the collected changeset. All remaining bits are kept uninitialized. Using an u16 avoids this type confusion and ensures that no (negative) sign extension is performed in batadv_send_other_tt_response().	8.8	More Details
CVE-2026-63108	Roo Code through 3.54.0 contains a command injection vulnerability in the auto-approve execute feature that allows attackers to bypass allowlist/denylist enforcement by nesting command substitutions inside parameter expansion defaults. The command parser in parse-command.ts replaces parameter expansions with opaque placeholders before extracting command substitutions, causing the containsDangerousSubstitution guard to miss nested payloads, which are then auto-approved based on the outer allowlisted command prefix and executed by the shell via execa, enabling arbitrary command execution.	8.8	More Details
CVE-2026-63885	In the Linux kernel, the following vulnerability has been resolved: drm/gem: fix race between change_handle and handle_delete drm_gem_change_handle_ioctl leaves the old handle live in the IDR during the window between spin_unlock(table_lock) and the final spin_lock(table_lock). A concurrent drm_gem_handle_delete on the old handle succeeds in this window, decrements handle_count to 0, and frees the GEM object while the new handle's IDR entry still references it. NULL the old handle's IDR entry before dropping table_lock so that any concurrent GEM_CLOSE on the old handle sees NULL and returns -EINVAL. Restore the old entry on the prime-bookkeeping error path.	8.8	More Details
CVE-2026-53375	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vce: Prevent partial address patches In the case that only one of lo/hi is valid, the patching could result in a bad address written to in FW.	8.8	More Details
CVE-2026-53374	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: zero-initialize GART table on allocation GART TLB is flushed after unmapping but not after mapping. Since amdgpu_bo_create_kernel() does not zero-initialize the buffer, when a single PTE is written the TLB may speculatively load other uninitialized entries from the same cacheline. Those garbage entries can appear valid, and a subsequent write to another PTE in the same cacheline may cause the GPU to use a stale garbage PTE from the TLB. Fix this by calling memset_io() to zero-initialize the GART table with gart_pte_flags immediately after allocation. Using AMDGPU_GEM_CREATE_VRAM_CLEARED, SDMA-based clear will not work since SDMA needs GART to be initialized to work. (cherry picked from commit d9af8263b82b6eaa60c5718e0c6631c5037e4b24)	8.8	More Details
CVE-2026-45805	Penpot is an open-source design tool for design and code collaboration. Prior to 2.15.0, Penpot MCP's mcp/packages/server/src/ReplServer.ts bound the ReplServer to 0.0.0.0:4403 and exposed an unauthenticated /execute endpoint that passed the code field to PluginBridge.executePluginTask(), allowing anyone on the network to execute JavaScript on the server. This issue is fixed in version 2.15.0.	8.8	More Details
CVE-2026-53593	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.224, the denylist that neutralizes dangerous file uploads (`Helper::$restricted_extensions`) is incomplete: it does not cover the `.pht` extension. The authenticated upload endpoint `POST /uploads/upload` (`SecureController@upload`) stores files with their original extension into the web-accessible directory `storage/app/public/uploads/` (served at `/storage/uploads/`). On the standard Apache + `libapache2-mod-php` deployment, the default handler `<FilesMatch ".+\.ph(ar\|p[3457]?\|t\|tml)$">` executes `.pht`, so **any authenticated agent can upload a `.pht` web shell and run arbitrary commands as the web-server user** (`www-data`). This is a direct bypass of the fix for CVE-2025-48471, which added `phtml`/`phar` but not `pht` (nor `phtm`, `phps`). Version 1.8.224 contains an updated fix.	8.8	More Details

CVE-2026-12341	This vulnerability impacts all versions of IdentityIQ and allows an unauthenticated attacker unauthorized access to protected APIs and data due to improper validation of OAuth bearer tokens.	8.8	More Details
CVE-2026-62312	9Router is an AI router & token saver. Prior to 0.5.2, 9Router allows a remote authenticated attacker to achieve arbitrary code execution on the host operating system by combining a Host header bypass of localhost-only routes with unvalidated MCP plugin args passed to child_process.spawn(), allowing malicious custom plugins to execute commands through /api/mcp//sse. This issue is fixed in version 0.5.2.	8.8	More Details
CVE-2026-20150	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20150 are related to improper access control that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-284.	8.8	More Details
CVE-2026-64109	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix UAF read of tail->len in unix_stream_data_wait() unix_stream_data_wait() does skb_peek_tail(&sk->sk_receive_queue) without holding any lock that prevents SKBs on that queue from being dequeued and freed. This has been the case since commit 79f632c71bea ("unix/stream: fix peeking with an offset larger than data in queue"). The first consequence of this is that the pointer comparison `tail != last` can be false even if `last` semantically refers to an already-freed SKB while `tail` is a new SKB allocated at the same address; which can cause unix_stream_data_wait() to wrongly keep blocking after new data has arrived, but only in a weird scenario where a peeking recv() and a normal recv() on the same socket are racing, which is probably not a real problem. But since commit 2b514574f7e8 ("net: af_unix: implement splice for stream af_unix sockets"), `tail` is actually dereferenced, which can cause UAF in the following race scenario (where test_setup() runs single-threaded, and afterwards, test_thread1() and test_thread2() run concurrently in two threads: ``` static int socks[2]; void test_setup(void) { socketpair(AF_UNIX, SOCK_STREAM, 0, socks); send(socks[1], "A", 1, 0); int peekoff = 1; setsockopt(socks[0], SOL_SOCKET, SO_PEEK_OFF, &peekoff, sizeof(peekoff)); } void test_thread1(void) { char dummy; recv(socks[0], &dummy, 1, MSG_PEEK); } void test_thread2(void) { char dummy; recv(socks[0], &dummy, 1, 0); shutdown(socks[1], SHUT_WR); } ``` when racing like this: ``` thread1 thread2 unix_stream_read_generic mutex_lock(&u->iolock) skb_peek(&sk->sk_receive_queue) skb_peek_next(skb, &sk->sk_receive_queue) mutex_unlock(&u->iolock) unix_stream_read_generic unix_state_lock(sk) skb_peek(&sk->sk_receive_queue) unix_state_unlock(sk) unix_stream_data_wait unix_state_lock(sk) tail = skb_peek_tail(&sk->sk_receive_queue) spin_lock(&sk->sk_receive_queue.lock) __skb_unlink(skb, &sk->sk_receive_queue) spin_unlock(&sk->sk_receive_queue.lock) consume_skb(skb) [frees the SKB] `tail != last`: false `tail`: true `tail->len != last_len` ***UAF*** ``` Fix the UAF by removing the read of tail->len; checking tail->len would only make sense if SKBs in the receive queue of a UNIX socket could grow, which can no longer happen. Kuniyuki explained: > When commit 869e7c62486e ("net: af_unix: implement stream sendpage > support") added sendpage() support, data could be appended to the last > skb in the receiver's queue. > > That's why we needed to check if the length of the last skb was changed > while waiting for new data in unix_stream_data_wait(). > > However, commit a0dbf5f818f9 ("af_unix: Support MSG_SPLICE_PAGES") and > commit 57d44a354a43 ("unix: Convert unix_stream_sendpage() to use > MSG_SPLICE_PAGES") refactored sendmsg(), and now data is always added > to a new skb. That means this fix is not suitable for kernels before 6.5.	8.8	More Details
CVE-2025-45868	LogicalDOC Enterprise up to and for v9.1.1 is vulnerable to blind SQL injection in the ComparisonServlet component, allowing authenticated user to manipulate SQL queries via crafted input.	8.8	More Details
CVE-2026-63796	In the Linux kernel, the following vulnerability has been resolved: ocfs2: reject oversized group bitmap descriptors ocfs2_validate_gd_parent() only bounds bg_bits against the parent allocator's chain geometry. A malicious descriptor can still claim a bg_size/bg_bits pair that exceeds the bitmap bytes that physically fit in the group descriptor block, so later bitmap scans and bit updates can run past bg_bitmap. Add a physical-cap check based on ocfs2_group_bitmap_size() for the parent allocator type and reject descriptors whose bg_size or bg_bits exceed that capacity. Keep the existing chain geometry check so both the on-disk bitmap layout and the allocator metadata must agree before the descriptor is used. Validation reproduced this kernel report: KASAN use-after-free in _find_next_bit+0x7f/0xc0 Read of size 8 Call trace: dump_stack_lvl+0x66/0xa0 (?:?) print_report+0xd0/0x630 (?:?) _find_next_bit+0x7f/0xc0 (?:?) srso_alias_return_thunk+0x5/0xfbef5 (?:?) __virt_addr_valid+0x188/0x2f0 (?:?) kasan_report+0xe4/0x120 (?:?) ocfs2_find_max_contig_free_bits+0x35/0x70 (fs/ocfs2/suballoc.c:1375) ocfs2_block_group_set_bits+0x472/0x4b0 (fs/ocfs2/suballoc.c:1457) ocfs2_cluster_group_search+0x16b/0x440 (fs/ocfs2/suballoc.c:86) ocfs2_bg_discontig_fix_result+0x1ef/0x230 (fs/ocfs2/suballoc.c:1786) ocfs2_search_chain+0x8f8/0x10a0 (fs/ocfs2/suballoc.c:1886) get_page_from_freelist+0x70e/0x2370 (?:?) lock_release+0xc6/0x290 (?:?) do_raw_spin_unlock+0x9a/0x100 (?:?) kasan_unpoison+0x27/0x60 (?:?) __bfs+0x147/0x240 (?:?) get_page_from_freelist+0x83d/0x2370 (?:?) ocfs2_claim_suballoc_bits+0x38c/0xe70 (fs/ocfs2/suballoc.c:96) sched_domains_numa_masks_clear+0x70/0xd0 (?:?) check_irq_usage+0xe8/0xb70 (?:?) __ocfs2_claim_clusters+0x18d/0x4c0 (fs/ocfs2/suballoc.c:2497) check_path+0x24/0x50 (?:?) rcu_is_watching+0x20/0x50 (?:?) check_prev_add+0xfd/0xd00 (?:?) ocfs2_add_clusters_in_btree+0x17d/0x810 (fs/ocfs2/suballoc.c:?) __folio_batch_add_and_move+0x1f5/0x3d0 (?:?) ocfs2_add_inode_data+0xd9/0x120 (fs/ocfs2/suballoc.c:?) filemap_add_folio+0x105/0x1f0 (?:?)	8.8	More Details

	ocfs2_write_begin_nolock+0x29f7/0x2f80 (fs/ocfs2/suballoc.c:3043) ocfs2_read_inode_block+0xb5/0x110 (fs/ocfs2/suballoc.c:?) down_write+0xf5/0x180 (?:?) ocfs2_write_begin+0x180/0x240 (fs/ocfs2/suballoc.c:?) __mark_inode_dirty+0x758/0x9a0 (?:?) inode_to_bdi+0x41/0x90 (?:?) balance_dirty_pages_ratelimited_flags+0xf8/0x1d0 (?:?) generic_perform_write+0x252/0x440 (?:?) mnt_put_write_access_file+0x16/0x70 (?:?) file_update_time_flags+0xe4/0x200 (?:?) ocfs2_file_write_iter+0x80a/0x1320 (fs/ocfs2/suballoc.c:?) lock_acquire+0x184/0x2f0 (?:?) ksys_write+0xd2/0x170 (?:?) apparmor_file_permission+0xf5/0x310 (?:?) read_zero+0x8d/0x140 (?:?) lock_is_held_type+0x8f/0x100 (?:?)		
CVE-2026-63915	In the Linux kernel, the following vulnerability has been resolved: nfc: hci: fix out-of-bounds read in HCP header parsing Both nfc_hci_recv_from_llc() and nci_hci_data_received_cb() read packet->header from skb->data at function entry without first checking that the buffer holds at least one byte. A malicious NFC peer can send a 0-byte HCP frame that passes through the SHDLC layer and reaches these functions, causing an out-of-bounds heap read of packet->header. The same 0-byte frame, if queued as a non-final fragment, also causes the reassembly loop to underflow msg_len to UINT_MAX, triggering skb_over_panic() when the reassembled skb is written. Fix this by adding a pskb_may_pull() check at the entry of each function before packet->header is first accessed. The existing pskb_may_pull() checks before the reassembled hcp_skb is cast to struct hcp_packet remain in place to guard the 2-byte HCP message header.	8.8	More Details
CVE-2026-64115	In the Linux kernel, the following vulnerability has been resolved: vsock/vmci: fix UAF when peer resets connection during handshake vmci_transport_recv_connecting_server() returned err = 0 for a peer RST in its default switch arm: err = pkt->type == VMCI_TRANSPORT_PACKET_TYPE_RST ? 0 : -EINVAL; That made vmci_transport_recv_listen() skip vsock_remove_pending(), leaving the pending socket on the listener's pending_links with sk_state = TCP_CLOSE while destroy: still dropped the explicit reference taken before schedule_delayed_work(). One second later vsock_pending_work() observed is_pending=true and performed full cleanup: vsock_remove_pending() then the two trailing sock_put(sk) calls -- the first reached refcount 0 and __sk_freed the socket, and the second wrote into the freed object: BUG: KASAN: slab-use-after-free in refcount_warn_saturate Write of size 4 at addr ffff88800b1cac80 by task kworker Workqueue: events vsock_pending_work Treat peer RST like any other unexpected packet type (err = -EINVAL). All destroy: arms now return err < 0, so vmci_transport_recv_listen() removes pending from pending_links synchronously and vsock_pending_work() takes the is_pending=false / !rejected branch, dropping only its own work reference. This also closes the multi-packet race Sashiko reported on v2: pending is removed from the list before any subsequent packet can find it. The pre-existing sk_acceptq_removed() gap on the err < 0 path of vmci_transport_recv_listen() that Sashiko also noted is not introduced or changed by this patch. Tested on lts-6.12.79 with KASAN: 52/100 unpatched -> 0/100 patched.	8.8	More Details
CVE-2026-63916	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: Fix OOB write in wacom_hid_set_device_mode() wacom_hid_set_device_mode() currently assumes that the HID_DG_INPUTMODE usage is always located in the first field (field[0]) of the feature report. However, a device can specify HID_DG_INPUTMODE in a different field. If HID_DG_INPUTMODE is in a field other than the first one and the first field has a report_count smaller than the usage_index of HID_DG_INPUTMODE, this leads to an out-of-bounds write to r->field[0]->value. Fix this by storing the field index of HID_DG_INPUTMODE in 'struct hid_data' during feature mapping. In wacom_hid_set_device_mode(), use this stored field index to access the correct field and add bounds checks to ensure both the field index and the value index are within valid ranges before writing.	8.8	More Details
CVE-2026-11826	OpenPLC_v3 contains a heap-based buffer overflow in the getData() function in webserver/core/modbus_master.cpp. getData() reads characters between two delimiters into a caller-supplied buffer with no size parameter and no bounds check. In parseConfig() the function is invoked with the 100-byte heap-allocated MB_device.dev_name field. An authenticated attacker with access to the OpenPLC web interface can send a crafted HTTP POST to the /modbus endpoint with an oversized device_name value; the value is persisted to mbconfig.cfg and parsed on load, overflowing dev_name and overwriting adjacent struct fields (protocol at offset 108, dev_address at offset 109, ip_port at offset 210). A 200-byte payload writes 100 bytes past the allocation. The result is heap corruption leading to runtime crash and denial of service of the PLC process control loop, with attacker-controlled overwrite of adjacent configuration fields. The upstream repository was archived on 2026-04-04 and no fix is expected; the vendor has confirmed the issue does not affect OpenPLC Runtime v4.	8.8	More Details
CVE-2026-63757	SurrealDB versions before 3.1.0 contain a session hijacking vulnerability where the HTTP /rpc sessions method returns attached session UUIDs without authentication and accepts arbitrary session fields with no ownership verification. Unauthenticated attackers can enumerate session UUIDs and impersonate authenticated sessions to read, write, delete data and escalate privileges.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: tipc: fix slab-use-after-free Read in tipc_aead_decrypt_done tipc_aead_decrypt() goes straight from tipc_bearer_hold(b) to crypto_aead_decrypt(req) without taking a reference on the netns, unlike the encrypt path. When crypto_aead_decrypt() is offloaded asynchronously (e.g. the SIMD aead wrapper queuing to cryptd), the cryptd worker runs tipc_aead_decrypt_done() later. If the bearer's netns is torn down in the meantime, cleanup_net() -> tipc_exit_net() -> tipc_crypto_stop() frees the per-netns tipc_crypto, and the completion then reads it: tipc_aead_decrypt_done() dereferences aead->crypto->stats and aead->crypto->net, and tipc_crypto_rcv_complete() dereferences aead->crypto->aead[] and the node table -- reading freed memory.		

CVE-2026-63801	Decoded KASAN splat (v7.1-rc7, CONFIG_KASAN_INLINE + TIPC + TIPC_CRYPTO): BUG: KASAN: slab-use-after-free in tipc_aead_decrypt_done (net/tipc/crypto.c:999) Read of size 8 at addr ffff8881056258a8 by task kworker/u16:2/51 Workqueue: events_unbound Call Trace: tipc_aead_decrypt_done (net/tipc/crypto.c:999) process_one_work (kernel/workqueue.c:3314) worker_thread (kernel/workqueue.c:3397 kernel/workqueue.c:3478) kthread (kernel/kthread.c:436) ret_from_fork (arch/x86/kernel/process.c:158) ret_from_fork_asm (arch/x86/entry/entry_64.S:245) Allocated by task 169: __kasan_kmalloc (mm/kasan/common.c:398 mm/kasan/common.c:415) tipc_crypto_start (net/tipc/crypto.c:1502) tipc_init_net (net/tipc/core.c:72) ops_init (net/core/net_namespace.c:137) setup_net (net/core/net_namespace.c:446) copy_net_ns (net/core/net_namespace.c:579) create_new_namespaces (kernel/nsproxy.c:132) __x64_sys_unshare (kernel/fork.c:3316) do_syscall_64 (arch/x86/entry/syscall_64.c:63) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121) Freed by task 8: kfree (mm/slub.c:6566) tipc_exit_net (net/tipc/core.c:119) cleanup_net (net/core/net_namespace.c:704) process_one_work (kernel/workqueue.c:3314) kthread (kernel/kthread.c:436) This is the same class of bug that commit e279024617134 ("net/tipc: fix slab-use-after-free Read in tipc_aead_encrypt_done") fixed for the encrypt side. The encrypt path takes maybe_get_net(aead->crypto->net) before crypto_aead_encrypt() and drops it with put_net() on the synchronous return paths and in tipc_aead_encrypt_done(); the -EINPROGRESS/-EBUSY return keeps the reference for the async callback to release. The decrypt path was left without the equivalent guard. Mirror the encrypt-side fix on the decrypt path: take a net reference before crypto_aead_decrypt() (failing with -ENODEV and the matching bearer put if it cannot be acquired), keep it across the -EINPROGRESS/-EBUSY async return, and drop it with put_net() on the synchronous success/error return and at the end of tipc_aead_decrypt_done(). Reproduced under KASAN on v7.1-rc7: a UDP bearer with a cluster key is flooded with crafted encrypted frames from an unknown peer (driving the cluster-key decrypt path) while the bearer's netns is repeatedly torn down. The completion must run asynchronously to outlive tipc_crypto_stop(); on x86 the stock aesni gcm(aes) now decrypts synchronously, so the async path was exercised via cryptd offload. The unguarded aead->crypto dereference in tipc_aead_decrypt_done() is the unpatched upstream path; tipc_aead_decrypt() still lacks maybe_get_net(aead->crypto->net), so the completion can outlive the free on any config where crypto_aead_decrypt() goes async. Found by 0sec automated security-research tooling (https://0sec.ai).	8.8	More Details
CVE-2024-58362	SurrealDB before 1.5.5 (and 2.0.0-beta before 2.0.0-beta.3) accepts an arbitrary object in the signin and signup operations of the RPC API without recursively validating it for non-computed values. When a record access method defines a SIGNIN or SIGNUP query and the RPC API is exposed to untrusted users, an unauthenticated attacker can encode a binary object containing a subquery using the bincode serialization format and supply it in place of credentials. The subquery is then executed within the database owner's SIGNIN/SIGNUP query under a system user session with the editor role, allowing the attacker to select, create, update, and delete non-IAM resources (though not view the query results directly, and not affect IAM resources, which require the owner role).	8.8	More Details
CVE-2023-54366	SurrealDB before 1.0.1 sets default table permissions to FULL instead of NONE, allowing SELECT, CREATE, UPDATE, and DELETE operations on tables without explicit permissions. Attackers with database access or unauthenticated users on publicly exposed instances can perform unrestricted operations on unprotected tables within their authorization scope.	8.8	More Details
CVE-2026-61457	The Grav API plugin (getgrav/grav-plugin-api) before 1.0.3 contains a file upload extension bypass in the API media controller. HandlesMediaUploads::validateFileExtension() inspects only the final file extension via pathinfo($filename, PATHINFO_EXTENSION), so a user with api.media.write permission can upload a file with a double extension such as shell.php.jpg to bypass the dangerous extensions blocklist. The web server may then execute the file as PHP, resulting in remote code execution.	8.8	More Details
CVE-2026-63863	In the Linux kernel, the following vulnerability has been resolved: drm/gpusvm: Fix unbalanced unlock in drm_gpusvm_scan_mm() There is a unbalanced lock/unlock to gpusvm notifier lock: [931.045868] ==================================== [931.046509] WARNING: bad unlock balance detected! [931.047149] 6.19.0-rc6+xe-********** #9 Tainted: G U [931.048150] ------------------------------------ [931.048790] kworker/u5:0/51 is trying to release lock (&gpusvm->notifier_lock) at: [931.049801] [<ffffffffa090c0d8>] drm_gpusvm_scan_mm+0x188/0x460 [drm_gpusvm_helper] [931.050802] but there are no more locks to release! [931.051463] The drm_gpusvm_notifier_unlock() sits under err_free label and the first jump to err_free is just before calling the drm_gpusvm_notifier_lock() causing unbalanced unlock.	8.8	More Details
CVE-2026-63941	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Correctly cap ZCR_EL2 provided by a guest hypervisor ZCR_EL2 can be updated by a VHE guest hypervisor either using ZCR_EL2 (which traps) or ZCR_EL1 (which does not trap). KVM handles both in different way: - on ZCR_EL2 trap, ZCR_EL2.LEN is immediately capped at the VM's own VL limit. This has the potential to break existing SW that relies on the full LEN field to be stateful. - on ZCR_EL1 access, we do absolutely nothing. On restoring the SVE context for an L2 guest, we directly restore the guest hypervisor's view of ZCR_EL2 into the physical ZCR_EL2. If the guest's view of the register was updated using the ZCR_EL2 accessor, the value has already been sanitised (with the caveat mentioned above). But if the guest used ZCR_EL1, the raw value is written into the HW, and the L2 guest can now access VLs that it shouldn't. Fix all the above by moving the VL capping to the restore points, ensuring that: - the HW is always programmed with a capped value, irrespective of the accessor being used, - the ZCR_EL2.LEN field is always completely stateful, irrespective of the accessor being used. Additionally, move ZCR_EL2 to be a sanitised register, ensuring that only the LEN field is actually stateful. This requires some creative construction of the RES0 mask, as the sysreg generation	8.8	More Details

	script does not yet generate RAZ/WI fields. [maz: rewrote commit message, tidy up access_zcr_el2()]		
CVE-2026-10081	The Unlimited Elements For Elementor WordPress plugin before 2.0.11 does not sanitize or escape Google review content fetched from the Serp API before rendering it in the Google Reviews widget output, allowing unauthenticated attackers who submit a malicious review on the targeted business's Google listing to deliver Stored XSS to any visitor (including administrators) of any WP page displaying that Place ID's reviews.	8.8	More Details
CVE-2026-7667	IBM Langflow OSS 1.0.0 through 1.10.0 allows an authenticated attacker to create a malicious flow pointing to an attacker-controlled URL that returns a specially crafted Content-Disposition header (e.g., filename="../../../target/path"), enabling arbitrary file write operations with attacker-controlled content to any path accessible by the Langflow process.	8.8	More Details
CVE-2026-64030	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: bounds-check link_id in ieee80211_ml_epcs IEEE80211_MLE_STA_EPCS_CONTROL_LINK_ID is 0x000f, so link_id extracted from a PRIO_ACCESS ML element PER_STA_PROFILE subelement can be 0..15. sdata->link[] has IEEE80211_MLD_MAX_NUM_LINKS (15) entries (indices 0..14), making index 15 out-of-bounds. A connected WiFi 7 AP can trigger this by sending an EPCS Enable Response action frame with a PER_STA_PROFILE subelement where link_id = 15. The unsolicited-notification path (dialog_token = 0) is reachable any time EPCS is already enabled, without any prior client request. sdata->link[15] reads into the first word of sdata->activate_links_work (a wiphy_work whose embedded list_head is non-NULL after INIT_LIST_HEAD), so the NULL check on the result does not catch the invalid access. The garbage pointer is then passed to ieee80211_sta_wmm_params(), which dereferences link->sdata and crashes the kernel. The same class of bug was fixed for ieee80211_ml_reconfiguration() by commit 162d331d833d ("wifi: mac80211: bounds-check link_id in ieee80211_ml_reconfiguration").	8.8	More Details
CVE-2026-63947	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: HIDP: fix missing length checks in hidp_input_report() hidp_input_report() reads keyboard and mouse payload data from an skb without first verifying that skb->len contains enough data. hidp_recv_intr_frame() pulls the 1-byte HIDP header before dispatching to hidp_input_report(). If a paired device sends a truncated packet, the handler reads beyond the valid skb data, resulting in an out-of-bounds read of skb data. The OOB bytes may be interpreted as phantom key presses or spurious mouse movement. Replace the open-coded length tracking and pointer arithmetic with skb_pull_data() calls. skb_pull_data() returns NULL if the requested bytes are not present, eliminating the need for a manual size variable and the separate skb->len guard.	8.8	More Details
CVE-2026-14499	IBM Langflow OSS 1.0.0 through 1.10.1 Langflow could allow an authenticated user to execute arbitrary commands with elevated privileges on the system due to improper validation of user supplied input in the Python Interpreter component.	8.8	More Details
CVE-2026-63829	In the Linux kernel, the following vulnerability has been resolved: net: ip_gre: require CAP_NET_ADMIN in the device netns for changelink A tunnel changelink() operates on at most two netns, dev_net(dev) and the tunnel link netns t->net. They differ once the device is created in or moved to a netns other than the one the request runs in. The rtnl changelink path checks CAP_NET_ADMIN only against dev_net(dev), so a caller privileged there but not in t->net can rewrite a tunnel that lives in t->net. Add rtnl_dev_link_net_capable() next to rtnl_get_net_ns_capable() in net/core/rtnetlink.c. It requires CAP_NET_ADMIN in the link netns and is skipped when the link netns is dev_net(dev), where the rtnl path already checked it. The other patches in this series use the same helper. Gate ipgre_changelink() and erspan_changelink() with it, at the top of the op before any attribute is parsed, because the parsers update live tunnel fields first. ipgre_netlink_parms() sets t->collect_md before ip_tunnel_changelink() runs. Commit 8b484efd5cb4 ("ip6: vti: Use ip6_tnl.net in vti6_siocdevprivate().") added the same check on the ioctl path. This adds it on RTM_NEWLINK.	8.8	More Details
CVE-2026-63831	In the Linux kernel, the following vulnerability has been resolved: mac802154: llsec: add skb_cow_data() before in-place crypto llsec_do_encrypt_unauth(), llsec_do_encrypt_auth(), llsec_do_decrypt_unauth(), and llsec_do_decrypt_auth() all perform in-place cryptographic transformations on skb data. They build a scatterlist with sg_init_one() pointing into the skb's linear data area and then pass the same scatterlist as both src and dst to the crypto API (e.g. crypto_skcipher_encrypt/decrypt, crypto_aead_encrypt/decrypt). On the RX path, __ieee802154_rx_handle_packet() clones the received skb before handing it to each subscriber via ieee802154_subif_frame(). The cloned skb shares the same underlying data buffer via reference counting. When llsec_do_decrypt() subsequently modifies this shared buffer in place, it corrupts data that other clones -- potentially belonging to other sockets or subsystems -- still reference. On the TX path, similar data sharing can occur when an skb's head has been cloned (skb_cloned() returns true). The fix is to call skb_cow_data() before performing any in-place crypto operation. skb_cow_data() ensures that the skb's data area is not shared: if the skb head is cloned or the data spans multiple fragments, it copies the data into a private buffer that can be safely modified in place. This is the same pattern used by: - ESP (net/ipv4/esp4.c, net/ipv6/esp6.c) - MACsec (drivers/net/macsec.c) - WireGuard (drivers/net/wireguard/receive.c) - TIPC (net/tipc/crypto.c) Without this guard, in-place crypto on shared skb data leads to: - Silent data corruption of other skb clones - Use-after-free when the crypto API scatterwalk writes through a page that has already been freed by another clone's kfree_skb() - Kernel crashes under concurrent 802.15.4 traffic with security enabled (KASAN/KMSAN reports slab-use-after-free) Found by 0sec (https://0sec.ai) using automated source analysis.	8.8	More Details
CVE-	The Redux Framework WordPress plugin before 4.5.13 does not restrict which user meta keys can be written when saving custom profile fields, allowing users with at least the Subscriber role to escalate their privileges		More

2026-12525	to Administrator by submitting a crafted value while updating their own profile, on sites where the Redux Framework WordPress plugin before 4.5.13's user-profile (Users extension) feature is enabled.	8.8	Details
CVE-2026-63832	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: add wcid publish check in mt76_sta_add Since mt7925_mac_sta_add publishes wcid, add publish check in mt76_sta_add to avoid reinitializing the wcid->poll_list. Found dev->sta_poll_list corruption when using mt7925 and 7.1-rc4. According to the corruption information, prev->next was changed to itself. wlan0: disconnect from AP 90:fb:5d:94:8b:e3 for new auth to 90:fb:5d:94:8b:e2 wlan0: authenticate with 90:fb:5d:94:8b:e2 (local address=84:9e:56:9c:7e:6b) wlan0: send auth to 90:fb:5d:94:8b:e2 (try 1/3) slab kmalloc-8k start ffff8c80958a6000 pointer offset 4160 size 8192 list_add corruption. prev->next should be next (ffff8c808a7488f8), but was ffff8c80958a7040. (prev=ffff8c80958a7040). mt76_wcid_add_poll+0x95/0xd0 [mt76] mt7925_mac_add_txs.part.0+0xa5/0xe0 [mt7925_common] mt7925_rx_check+0xa7/0xc0 [mt7925_common] mt76_dma_rx_poll+0x50d/0x790 [mt76] mt792x_poll_rx+0x52/0xe0 [mt792x_lib]	8.8	More Details
CVE-2026-35152	A SQL Injection vulnerability exists in Apache Fineract's Report Execution API (runreports endpoint) in versions up to and including 1.14.0. Report parameter values are incorporated into the generated SQL query without sufficient validation, allowing an authenticated user with permission to run reports to inject arbitrary SQL via crafted parameter values. This can be leveraged to perform unauthorized access to data beyond what the report was designed to expose. Users are recommended to upgrade to a version containing the fix.	8.8	More Details
CVE-2026-64881	The audit file upload handler does not sanitize filenames, allowing shell metacharacters to flow into system command execution. This input validation failure enables command injection when chained with a related vulnerability.	8.8	More Details
CVE-2026-15804	The HCM developed by MetaGuru has a SQL Injection vulnerability. Authenticated remote attackers can inject SQL commands via specific parameters, thereby compromising the confidentiality, integrity, and availability of database data.	8.8	More Details
CVE-2026-13352	The Paid Membership Plugin, Ecommerce, User Registration Form, Login Form, User Profile & Restrict Content – ProfilePress plugin for WordPress is vulnerable to Arbitrary File Upload in all versions up to, and including, 4.16.18 via the allowed_mime_types function. This is due to the unconditional registration of an upload_mimes filter that adds executable file extensions (.exe, .apk, .msi) to the global WordPress MIME allowlist, without scoping the expansion to digital-product upload contexts. This makes it possible for authenticated attackers, with author-level access and above, to upload files that may be executable, which makes remote code execution possible. This filter is registered globally on every request regardless of whether the digital products feature is configured or in use, meaning the expanded MIME allowlist affects all WordPress upload contexts site-wide.	8.8	More Details
CVE-2026-16248	A vulnerability was found in Tenda AC10 16.03.10.09_multi_TDE01. This issue affects the function fromAdvSetLanip of the file /goform/AdvSetLanip of the component httpd/netctrl. The manipulation of the argument GetValue/SetValue results in stack-based buffer overflow. The attack may be performed from remote. The exploit has been made public and could be used.	8.8	More Details
CVE-2026-64010	In the Linux kernel, the following vulnerability has been resolved: nfc: llcp: Fix use-after-free race in nfc_llcp_recv_cc() A race condition exists in the NFC LLCP connection state machine where the connection acceptance packet (CC) can be processed concurrently with socket release. This can lead to a use-after-free of the socket object. When nfc_llcp_recv_cc() moves the socket from the connecting_sockets list to the sockets list, it does so without holding the socket lock. If llcp_sock_release() is executing concurrently, it might have already unlinked the socket and dropped its references, which can result in nfc_llcp_recv_cc() linking a freed socket into the live list. Fix this by holding lock_sock() during the state transition and list movement in nfc_llcp_recv_cc(). After acquiring the lock, check if the socket is still hashed to ensure it hasn't already been unlinked and marked for destruction by the release path. This aligns the locking pattern with recv_hdlc() and recv_disc().	8.8	More Details
CVE-2026-63974	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_sync: Set HCI_CMD_DRAIN_WORKQUEUE during device close Since hci_dev_close_sync() can now be called during the reset path, we should also set HCI_CMD_DRAIN_WORKQUEUE. This avoids queuing timeouts while the hdev workqueue is being drained.	8.8	More Details
CVE-2026-58195	Agentic-Flow is an AI agent orchestration platform. Prior to 2.0.14, agentic-flow MCP server tools in src/mcp/standalone-stdio.ts, src/mcp/fastmcp/servers/claude-flow-sdk.ts, src/mcp/fastmcp/servers/stdio-full.ts, src/mcp/fastmcp/servers/http-streaming-updated.ts, src/mcp/fastmcp/servers/http-sse.ts, src/mcp/fastmcp/servers/poc-stdio.ts, src/mcp/fastmcp/tools/agent/{execute,list,parallel}.ts, src/mcp/fastmcp/tools/swarm/orchestrate.ts, and src/mcp/fastmcp/tools/hooks/pretrain.ts interpolated attacker-influenceable tool parameters such as agent, task, name, language, and agentdb directly into shell command strings passed to execSync(), allowing arbitrary OS command execution with the privileges of the MCP server user. This issue is fixed in version 2.0.14.	8.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: Fix possible crash on l2cap_ecred_conn_rsp If dcid is received for an already-assigned destination CID the spec requires that both channels to be discarded, but calling l2cap_chan_del may invalidate the tmp cursor created by	8.8	More

63975	list_for_each_entry_safe and in fact it is the wrong procedure as the chan->dcid may be assigned previously it really needs to be disconnected. Calling l2cap_chan_clone directly may still lead to l2cap_chan_del so instead schedule l2cap_chan_timeout with delay 0 to close the channel asynchronously.		Details
CVE-2026-13741	The Digits: WordPress Mobile Number Signup and Login plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 9.1.0.5. This is due to missing authorization and role validation in the `dig_update_wpwc_custom_fields()` function. This makes it possible for authenticated attackers, with Subscriber-level access and above, to escalate their privileges to Administrator by submitting a forged `digits_reg_userrole` value during profile update, granted the site administrator has configured the built-in DIGITS User Role field.	8.8	More Details
CVE-2026-63976	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: l2cap: clear chan->ident on ECRED reconfiguration success l2cap_ecred_reconf_rsp() returns early on success without clearing chan->ident. Every other L2CAP response handler (l2cap_ecred_conn_rsp, l2cap_le_connect_rsp, l2cap_config_rsp) clears chan->ident after a successful transaction to prevent the channel from matching subsequent responses with the recycled ident value. A remote attacker that completed a reconfiguration as the peer can replay a failure response with the stale ident, causing the kernel to match and destroy the already-established channel via l2cap_chan_del(chan, ECONNRESET). Clear chan->ident for all matching channels on success, and harden the failure path by using l2cap_chan_hold_unless_zero() consistent with other L2CAP handlers (l2cap_le_command_rej, __l2cap_get_chan_by_ident).	8.8	More Details
CVE-2026-15005	The Loco Translate plugin for WordPress is vulnerable to Cross-Site Request Forgery in all versions up to, and including, 2.8.5. This is due to missing or incorrect nonce validation on the execTemplate function. This makes it possible for unauthenticated attackers to execute arbitrary PHP code on the server by supplying a php://filter stream wrapper URI as the 'template' parameter, which bypasses path validation and is passed directly to the include sink in execTemplate() via a forged request granted they can trick a site administrator into performing an action such as clicking on a link.	8.8	More Details
CVE-2026-60025	The Joomla extension Events Booking prior version 5.8.0 had an frontend file upload endpoint that lacked CSRF protection.	8.8	More Details
CVE-2026-63093	Cursor for Windows version 3.2.16 contains a binary planting vulnerability that allows remote attackers to achieve arbitrary code execution by placing a malicious git.exe file in the repository root directory. When a developer clones and opens a crafted repository, Cursor automatically resolves and executes the workspace-resident git.exe during IDE startup and on a recurring timed cadence without any user interaction, running the malicious binary under the privileges of the current user.	8.8	More Details
CVE-2026-7755	IBM Langflow OSS 1.0.0 through 1.10.0 Langflow could allow remote code execution due to incomplete validation enforcement on MCP server configuration files.	8.8	More Details
CVE-2026-63946	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: fix UAF in iso_recv_frame iso_recv_frame reads conn->sk under iso_conn_lock but releases the lock before using sk, with no reference held. A concurrent iso_sock_kill() can free sk in that window, causing use-after-free on sk->sk_state and sock_queue_rcv_skb(). Fix by replacing the bare pointer read with iso_sock_hold(conn), which calls sock_hold() while the spinlock is held, atomically elevating the refcount before the lock drops. Add a drop_put label so sock_put() is called on all exit paths where the hold succeeded.	8.8	More Details
CVE-2026-8056	IBM Langflow OSS 1.0.0 through 1.10.0 allows authenticated users to override component parameters at runtime via the API. A critical security flaw exists in the parameter filtering mechanism within the `apply_tweaks()` function.	8.8	More Details
CVE-2026-40501	Cherry Studio versions 1.2.2 through 1.9.12, fixed in commit 1518530, contain a remote code execution vulnerability in SearchService that allows remote attackers to execute arbitrary code by delivering malicious JavaScript through controlled search provider content loaded into an Electron BrowserWindow configured with nodeIntegration enabled and contextIsolation disabled. Attackers who control a search engine provider, individual search result pages, or provider settings pages can execute JavaScript with full Node.js privileges, gaining access to fs, child_process, os, and process.env under the operating-system account of the Cherry Studio process.	8.8	More Details
CVE-2026-64042	In the Linux kernel, the following vulnerability has been resolved: vfio/pci: Check BAR resources before exporting a DMABUF A DMABUF exports access to BAR resources and, although they are requested at startup time, we need to ensure they really were reserved before exporting. Otherwise, it's possible to access unreserved resources through the export. Add a check to the DMABUF-creation path.	8.8	More Details
CVE-2026-21824	HCL Commerce contains an privilege escalation vulnerability that could allow denial of service, disclosure of user personal data, and performing of unauthorized administrative operations.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_sync: fix UAF in hci_le_create_cis_sync hci_le_create_cis_sync() dereferences conn->conn_timeout after releasing both		

CVE-2026-63944	rcu_read_lock() and hci_dev_lock(hdev). The conn pointer was obtained from an RCU-protected iteration over hdev->conn_hash.list and is not valid once these locks are dropped. A concurrent disconnect can free the hci_conn between the unlock and the dereference, causing a use-after-free read. The cancellation mechanism in hci_conn_del() cannot prevent this because hci_le_create_cis_pending() queues hci_create_cis_sync with data=NULL: hci_cmd_sync_queue(hdev, hci_create_cis_sync, NULL, NULL); While hci_conn_del() dequeues with data=conn: hci_cmd_sync_dequeue(hdev, NULL, conn, NULL); Since NULL != conn, the lookup in _hci_cmd_sync_lookup_entry() never matches, and the pending work item is not cancelled. Fix this by saving conn->conn_timeout into a local variable while the locks are still held, so the stale conn pointer is never dereferenced after unlock. This is the same class of bug as the one fixed by commit 035c25007c9e ("Bluetooth: hci_sync: Fix UAF on le_read_features_complete") which addressed the identical pattern in a different function. This vulnerability was identified using 0sec.ai, an open-source automated security auditing platform (https://github.com/0sec-labs).	8.8	More Details
CVE-2026-63090	ProFTPD before 1.3.9c and 1.3.10rc3 contains a heap-based buffer overflow vulnerability in the mod_sftp module that allows authenticated low-privilege attackers to achieve arbitrary code execution by sending crafted SFTP packet fragments exceeding the 16 KB reassembly buffer in the fxp.c component. Attackers can supply oversized fragments to trigger an incorrectly conditioned reallocation, corrupt pool freelist metadata, overwrite the root_fs BSS global pointer to reference a fake filesystem struct, and redirect pr_fsio_stat() to system() via a crafted RENAME request.	8.8	More Details
CVE-2026-62228	OpenClaw before 2026.6.5 contain an authorization bypass vulnerability in node exec approvals that allows lower-trust callers to execute actions beyond their intended authorization by using different gateway and node environments. Attackers can exploit mismatched environment configurations to persist or execute actions that exceed the caller's approved permissions.	8.8	More Details
CVE-2026-64153	In the Linux kernel, the following vulnerability has been resolved: drm/msm: Fix iommu_map_sgtable() return value check and avoid WARN Commit "iommu: return full error code from iommu_map_sg[_atomic]()" changed iommu_map_sgtable() to return an ssize_t and negative values in error cases, rather than a size_t and a zero. Store the return value in the appropriate type and in case of error, return it rather than WARNing. Patchwork: https://patchwork.freedesktop.org/patch/719685/	8.8	More Details
CVE-2026-62229	OpenClaw before 2026.5.18 contain an authorization bypass vulnerability in exec allowlist glob matching that allows lower-trust callers to execute actions beyond intended authorization. Attackers can craft input paths that traverse the allowlist glob patterns to execute or persist unauthorized actions when the affected feature is enabled.	8.8	More Details
CVE-2026-64178	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: bnep: Fix UAF read of dev->name bnep_add_connection() needs to keep holding the bnep_session_sem while reading dev->name (just like bnep_get_connlist() does); otherwise the bnep_session() thread can concurrently free the net_device, which can for example be triggered by a concurrent bnep_del_connection(). (This UAF is fairly uninteresting from a security perspective; calling bnep_add_connection() requires passing a capable(CAP_NET_ADMIN) check. It also requires completely tearing down a netdev during a fairly tight race window.)	8.8	More Details
CVE-2026-63085	Axelor Open Platform versions 8.x prior to 8.2.2 contains an authorization bypass vulnerability that allows authenticated non-admin users to escalate privileges by exploiting unenforced field restrictions on nested relational save operations. Attackers can modify sensitive User record fields such as roles and group by submitting changes through a related entity's save path, bypassing the USER_RESTRICTED_FIELDS control and causing the JPA persistence layer to flush attacker-supplied admin role and group assignments on commit.	8.8	More Details
CVE-2026-55084	DHIS2 is a flexible information system for data capture, management, validation, analytics and visualization. A SQL injection vulnerability was identified in the SqlView API endpoint of the DHIS2 application in the `filter` parameter used by the `/api/sqlViews/{viewId}/data.json` endpoint. An authenticated user with access to a SqlView can inject arbitrary SQL queries inside the `filter` parameter by abusing an expression executed by PostgreSQL and its output is reflected inside the application error message. This behavior enables attackers to extract arbitrary database content using error-based SQL injection. Affected versions include: 2.37, 2.38, 2.39, 2.40.x before 2.40.11.1/2.40.12, 2.41.x before 2.41.8.2, 2.42.x before 2.42.5.1, 2.43.0 before 2.43.0.1, 2.44 development branch before PR #24162 Patched versions include: 2.37-EOS (2026-06-09), 2.38-EOS (2026-06-09), 2.39-EOS (2026-06-09), 2.40.11.1, 2.40.12, 2.41.8.2, 2.42.5.1, 2.43.0.1, 2.44 development branch after PR #24162	8.8	More Details
CVE-2026-62233	grav-plugin-api before 1.0.6 fails to validate super-admin status in createApiKey, generate2fa, and disable2fa endpoints, allowing non-super api.users.write managers to escalate to super-admin. Attackers can mint API keys bound to super-admin accounts or strip 2FA from super-admin users to achieve full instance takeover.	8.8	More Details
CVE-2026-64093	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tp_meter: directly shut down timer on cleanup batadv_tp_sender_cleanup() was calling timer_delete_sync() followed by timer_delete() to guard against the timer handler re-arming itself between the two calls. This double-deletion hack relied on the sending status being set to 0 to suppress re-arming. Replace both calls with a single timer_shutdown_sync(). This function both waits for any running timer callback to complete (like timer_delete_sync()) and permanently disarms the timer so it cannot be re-armed afterwards, making re-arming prevention unconditional and self-documenting. The re-arming property is also required because	8.8	More Details

	otherwise: 1. context 0 (batadv_tp_recv_ack()) checks in batadv_tp_reset_sender_timer() if sending is still 1 -> it is 2. context 1 changes in batadv_tp_sender_shutdown() sending to 0 and in this process forces the kthread to stop timer in batadv_tp_sender_cleanup() 3. context 0 continues in batadv_tp_reset_sender_timer() and rearms the timer -> but the reference for it is already gone		
CVE-2026-58655	The bundled Grav Flex Objects plugin (getgrav/grav-plugin-flex-objects) before 1.4.0 contains a stored server-side template injection vulnerability. When rendering dynamic collection or object titles, the plugin passes user-controlled frontmatter values (page.header.flex.collection.title or page.header.flex.object.title) to Twig's template_from_string(), causing them to be evaluated as Twig code rather than treated as text. This path bypasses Grav's Security::cleanDangerousTwig() sanitization. An attacker who can control the title frontmatter of a publicly reachable Flex Objects page can achieve arbitrary Twig execution and escalate to remote command execution via access to internal Grav services such as the scheduler.	8.8	More Details
CVE-2026-57996	phpMyFAQ before 4.1.5 contains a privilege escalation vulnerability in the user/add API endpoint that allows non-SuperAdmin administrators to create SuperAdmin accounts. A delegated administrator with USER_ADD/EDIT/DELETE permissions can call POST /admin/api/user/add with isSuperAdmin: true and attacker-chosen credentials to create a SuperAdmin account, then authenticate as that account to achieve full instance takeover.	8.8	More Details
CVE-2026-54498	view_component is a framework for building reusable, testable, and encapsulated view components in Ruby on Rails. From 4.0.0 until 4.12.0, ViewComponent::Base#around_render can return HTML-unsafe strings that bypass the escaping behavior applied to normal #call return values. This creates an XSS risk when downstream applications use around_render to wrap, replace, instrument, or conditionally return content that includes user-controlled data, and ViewComponent::Collection#render_in can amplify the issue by joining per-item results and marking the entire output html_safe, converting raw unsafe output into an ActiveSupport::SafeBuffer. This issue is fixed in version 4.12.0.	8.7	More Details
CVE-2026-47869	VMware Avi Load Balancer contains a remote code execution vulnerability. A malicious authenticated user with network access may be able to inject and execute code. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	8.7	More Details
CVE-2026-45270	CI4MS is a CodeIgniter 4-based content management system skeleton. Prior to version 0.31.9.0, the `Pages` backend module registers the `html_purify` validation rule on language-keyed page content but persists the raw, un-purified POST value into the database. The public renderer for pages (`Home::index()` → `app/Views/templates/default/pages.php`) emits `$pageInfo->content` without `esc()`, yielding stored XSS that fires for every public visitor of the affected page — including administrators. Because pages may be promoted to the site home page, the payload can be served at `/` and reach every visitor of the site. Version 0.31.9.0 patches the issue.	8.7	More Details
CVE-2026-15724	In Progress ShareFile Storage Zones Controller versions prior to 5.12.5 and 6.0.2, an authenticated administrative user can exploit a path traversal vulnerability to read arbitrary files from the server filesystem, write files to arbitrary directories, or determine whether specific files exist on the server.	8.7	More Details
CVE-2026-44739	Pimcore is an Open Source Data & Experience Management Platform. Prior to 11.5.17 (LTS) and 12.3.6, the columnConfigAction endpoint in bundles/CustomReportsBundle/src/Controller/Reports/CustomReportController.php passes malicious SQL configuration through CustomReportController:columnConfigAction, SqlAdapter::getColumns, SqlAdapter::buildQueryString, and Db::fetchAssociative(), allowing an attacker with the reports_config permission to use arbitrary SELECT queries, UNION statements, dangerous database functions, and error-based SQL injection to exfiltrate or manipulate database data. This issue is fixed in versions 11.5.17 (LTS) and 12.3.6.	8.7	More Details
CVE-2026-64104	In the Linux kernel, the following vulnerability has been resolved: virt: sev-guest: Explicitly leak pages in unknown state When set_memory_{encrypted,decrypted}() fail, the user cannot know at which point the function failed, meaning that the pages are left in an unknown state from the point of view of the caller. Since the pages may be left in an unencrypted state, they are not suitable for general use, and cannot be returned safely to the buddy allocator. Avoid the issue by never freeing the pages, and then do the proper accounting by calling snp_leak_pages().	8.7	More Details
CVE-2026-47867	VMware Avi Load Balancer contains a remote code execution vulnerability. A malicious user with network access may be able to access the Avi Control plane and execute code remotely. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	8.7	More Details
CVE-2026-15631	Impact: @fastify/http-proxy versions from 9.4.0 up to and including 11.5.0 fail to validate the resolved WebSocket destination path against the configured rewrite prefix. The WebSocket routing path in WebSocketProxy.findUpstream resolves the destination via the WHATWG URL constructor, which collapses dot segments, so a crafted upgrade request with path traversal sequences can escape the rewrite prefix and reach upstream endpoints that were not meant to be exposed by the proxy. This is a variant of CVE-2021-21322 in a code path that never went through the HTTP fix in fastify/reply-from. Exploitation requires a non-normalizing WebSocket client, since browsers and the ws package normalize the request path before	8.7	More Details

	sending, but raw HTTP clients or downstream proxies that forward the request target unchanged make the attack reachable in production topologies. Patches: upgrade to @fastify/http-proxy 11.6.0. Workarounds: none.		
CVE-2026-16158	Impact: @fastify/reply-from versions from 8.3.1 up to but not including 12.6.4 build the internal URL cache key by concatenating the destination and source path without a delimiter. Different destination and source pairs can therefore produce the same key while resolving to different upstream URLs. When getUpstream selects an upstream from request data, a URL cached for one upstream can be reused for a request intended for another upstream, causing cross-upstream data access and modification. The default configuration is affected. Setting disableCache to true prevents the behavior. Patches: upgrade to @fastify/reply-from 12.6.4. Workarounds: pass disableCache: true when registering the plugin.	8.7	More Details
CVE-2026-57206	SimpleChat is a secure AI conversation application with personal and group workspaces for document-grounded interactions. Prior to 0.241.206, several plugin validation routes in application/single_app/plugin_validation_endpoint.py, including `POST /api/admin/plugins/test-instantiation`, `GET /api/admin/plugins/health-check/<plugin_name>`, `POST /api/admin/plugins/repair/<plugin_name>`, and `POST /api/plugins/validate`, relied on @swagger_route(security=get_auth_security()) documentation without enforcing @login_required, @user_required, or @admin_required at runtime, allowing unauthenticated or unauthorized clients to invoke plugin validation, health, and repair behavior. This issue is fixed in version 0.241.206.	8.6	More Details
CVE-2026-64623	Network-AI before 5.13.4 contains an improper cryptographic signature verification vulnerability in APSAdapter where the default local verifier accepts any non-empty string as valid. Unauthenticated attackers can submit forged APS delegation payloads with arbitrary scopes to bypass signature verification and obtain signed permission-grant tokens for sensitive resources including SHELL_EXEC.	8.6	More Details
CVE-2026-48795	AdonisJS is a TypeScript-first web framework. From 10.1.3 until 10.1.5 and 11.0.3, AdonisJS @adonisjs/bodyparser incompletely fixed CVE-2026-25754 because nested multipart field payloads such as user.__proto__.polluted and constructor.prototype still caused lodash _.set() via @poppinss/utils to create plain intermediate objects and pollute Object.prototype. This issue is fixed in versions 10.1.5 and 11.0.3.	8.6	More Details
CVE-2026-11349	The Modern Event Calendar Pro WordPress plugin before 7.34.0, Modern Events Calendar Lite WordPress plugin before 7.34.0 do not sanitise and escape a request parameter before using it in a SQL statement, through an AJAX action available to unauthenticated users, leading to an unauthenticated SQL injection vulnerability that allows attackers to extract sensitive data from the database.	8.6	More Details
CVE-2026-53591	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.223, an unauthenticated attacker can inject messages into any existing support conversation by sending a single email to the helpdesk's public address with a crafted `In-Reply-To` header. No credentials, tokens, or prior access are required. The injected message is rendered in the agent UI as a legitimate customer reply, the conversation is automatically reopened, and the `last_reply_from` field is set to the attacker's identity. Version 1.8.223 contains a fix.	8.6	More Details
CVE-2026-61436	PraisonAI before 4.6.78 fails to verify Svix webhook signatures in AgentMail webhook mode, allowing unauthenticated attackers to forge message.received events. Attackers can send crafted JSON payloads to the webhook endpoint to invoke configured agents with arbitrary sender addresses and message content.	8.6	More Details
CVE-2026-44023	Docling Core defines core data types and transformations for the document processing application Docling. In versions 1.5.0 and above, prior to 2.74.1, docling-core did not sufficiently restrict remote request destinations and could resolve a server-provided Content-Disposition to a local path in an unsafe manner. In applications that accept untrusted URLs, this could allow SSRF attacks targeting local files outside the user-defined cache directory. This issue has been fixed in version 2.74.1.	8.6	More Details
CVE-2026-63429	HeyForm is an open-source form builder. Prior to version 3.0.0-rc.9, `POST /api/upload` has no authentication guard, no global guard, no form-context validation, no `openToken` requirement, and no session cookie check. Any anonymous internet user can upload files (PDF, DOC/DOCX, XLS/XLSX, CSV, TXT, MP4, images, etc., up to 10 MB) and receive a permanent public URL on the HeyForm domain. The endpoint is used by both authenticated form creators and unauthenticated form submitters; because no form-context binding exists, every request to it is anonymously accepted. Version 3.0.0-rc.9 contains a patch for the issue.	8.6	More Details
CVE-2026-61836	Directus is a real-time API and App dashboard for managing SQL database content. Prior to 12.0.0, when response caching is enabled, the cache-key derivation in api/src/utils/get-cache-key.ts includes version, path, query, and accountability.user but omits authorization context such as share, role, roles, admin, app, and policies. Directus share tokens and anonymous requests can both reduce to user null, so different shares or anonymous clients requesting the same URL and query can receive a permission-filtered cached response without permission re-evaluation. This issue is fixed in version 12.0.0.	8.6	More Details
CVE-2026-12512	The Quotes llama WordPress plugin before 3.1.6 does not properly sanitize and escape a user-supplied parameter before using it in a SQL query, allowing unauthenticated attackers to perform UNION-based SQL injection and read arbitrary data from the database, including password hashes.	8.6	More Details
CVE-	stoatchat before 0.14.0 contains a server-side request forgery (SSRF) vulnerability that allows unauthenticated network-accessible attackers to bypass the DNS-based IP blocklist by exploiting incomplete		More

2026-63088	address validation in the url_is_blacklisted function, which inspects only the first resolved address while the underlying HTTP client iterates all cached addresses.	8.6	Details
CVE-2026-63306	stoatchat before 0.13.5 contains an unauthenticated server-side request forgery vulnerability in the /proxy and /embed endpoints that accept arbitrary URLs without DNS resolution filtering or private IP range validation. Attackers can enumerate internal services, fingerprint applications, and reach instance metadata endpoints by supplying malicious URLs or leveraging redirect chains to access internal infrastructure.	8.6	More Details
CVE-2026-63086	text-generation-inference through 3.3.7 contains a server-side request forgery (SSRF) vulnerability in the OpenAI-compatible multimodal chat completions endpoint that allows unauthenticated network attackers to coerce the server into issuing arbitrary HTTP GET requests by supplying a crafted image_url value in chat message content. The fetch_image function in router/src/validation.rs performs no validation of private, loopback, link-local, or cloud metadata target addresses, and the reqwest HTTP client follows redirects by default, enabling attackers to bypass scheme checks via redirect chains to reach internal services and cloud instance-metadata endpoints for internal port scanning and credential theft.	8.6	More Details
CVE-2026-15583	A confused-deputy flaw in Grafana MCP Server allows an unauthenticated remote attacker to exfiltrate the server's environment-configured Grafana service-account token by supplying a crafted X-Grafana-URL request header. This also enables SSRF against arbitrary internal services, including cloud metadata endpoints.	8.6	More Details
CVE-2024-58366	SurrealDB before 1.1.1 contains a format string vulnerability in the rquickjs Exception::throw_type function when scripting is enabled. Attackers with scripting privileges can supply format string sequences in error inputs to read arbitrary memory or execute code with SurrealDB process privileges.	8.5	More Details
CVE-2026-61430	PraisonAI before 1.6.78 contains a server-side request forgery vulnerability in the web_crawl tool that validates hostnames at check time but re-resolves them at connection time without IP pinning. Attackers can use DNS rebinding to bypass SSRF protection and retrieve internal HTTP response bodies from private or loopback services.	8.5	More Details
CVE-2026-47198	Paymenter is a free and open-source webshop solution for management of hosting services. In versions prior to 1.5.1, the checkout component improperly filters URL-writable properties, allowing authenticated users to inject arbitrary key-value pairs into server provisioning parameters. Because bundled server extensions prioritize these user-supplied properties over administrator-defined configurations, a regular user can override hosting plans and resource limits at checkout without special privileges. The Checkout Livewire component's $checkoutConfig property exposed via URL query parameters, only validating keys explicitly defined by an extension's configuration method, allowing any undefined injected keys to bypass validation entirely. These unsanitized keys are then stored directly in the database by the cart component and later passed to server extensions during provisioning, enabling user-injected data to override intended administrator settings. Depending on the active extension, this leads to unauthorized overrides of core resource limits (such as CPU, RAM, storage, or package tiers). No administrative privileges are required to exploit this vulnerability. This issue has been fixed in version 1.5.1.	8.5	More Details
CVE-2026-62226	OpenClaw 2026.3.28 before 2026.5.19 contain an authorization bypass vulnerability in the browser act route that fails to properly validate current-tab URL checks. Attackers with lower-trust access or configured input paths can perform actions requiring stronger authorization or policy checks.	8.5	More Details
CVE-2026-55234	Wekan is open source kanban built with Meteor. Prior to 9.37, Wekan DDP update allow rules in server/permissions/cards.js, server/permissions/lists.js, and server/permissions/swimlanes.js authorize against the stored source boardId and do not validate a new boardId in the update modifier. Any authenticated user with write access to their own board can call /cards/update, /lists/update, or /swimlanes/update to move cards, lists, or swimlanes into a private board they are not a member of. This issue is fixed in version 9.37.	8.5	More Details
CVE-2026-64151	In the Linux kernel, the following vulnerability has been resolved: iommupt: Check for missing PAGE_SIZE in the pgsize_bitmap Sashiko pointed out that the driver could drop PAGE_SIZE from the pgsize_bitmap. That is technically allowed but nothing does it, and such an iommu_domain would not be used with the DMA API today. Still, it is against the design and it is trivial to fix up. Lift the PT_WARN_ON to the if branch and just skip the fast path.	8.4	More Details
CVE-2026-63818	In the Linux kernel, the following vulnerability has been resolved: f2fs: validate orphan inode entry count f2fs_recover_orphan_inodes() trusts the orphan block entry_count when replaying orphan inodes from the checkpoint pack. A corrupted entry_count larger than F2FS_ORPHANS_PER_BLOCK makes the recovery loop read past the ino[] array and interpret footer or following data as inode numbers. On a crafted image, mounting an unpatched kernel can drive orphan recovery into f2fs_bug_on() and panic the kernel. Validate entry_count before consuming entries so corrupted checkpoint data fails the mount with -EFSCORRUPTED and requests fsck instead. Set ERROR_INCONSISTENT_ORPHAN as well, so the corruption reason can be recorded in the superblock s_errors[] field. This gives fsck a persistent hint even though mount-time orphan recovery failure may leave no chance to persist SBI_NEED_FSCK through a checkpoint.	8.4	More Details
	In the Linux kernel, the following vulnerability has been resolved: bpf: sockmap: fix tail fragment offset in bpf_msg_push_data When bpf_msg_push_data() inserts data in the middle of a scatterlist entry, it splits the		

CVE-2026-63926	original entry into a left fragment and a right fragment. The right fragment offset is page-local, but the code advances it with `start`, which is the message-global insertion point. For inserts into a non-first SG entry, this over-advances the offset and leaves the split layout inconsistent. Advance the right fragment offset by the fragment-local delta, `start - offset`, which matches the length removed from the front of the original entry.	8.4	More Details
CVE-2026-64045	In the Linux kernel, the following vulnerability has been resolved: ovpn: tcp - use cached peer pointer in ovpn_tcp_close() ovpn_tcp_close() loads the ovpn_socket via rcu_dereference_sk_user_data() under rcu_read_lock(), takes a reference on sock->peer, caches the peer pointer in a local, and drops the read lock. It then passes sock->peer (rather than the cached local) to ovpn_peer_del(), re-dereferencing the ovpn_socket after the RCU read section has ended. Unlike ovpn_tcp_sendmsg(), which uses the same "load under RCU, use after unlock" pattern but is protected by lock_sock() held across the function, ovpn_tcp_close() runs without the socket lock: inet_release() invokes sk_prot->close() without taking lock_sock first. ovpn_socket_release() can therefore complete its kref_put -> detach -> synchronize_rcu -> kfree(sock) sequence concurrently, in the window after ovpn_tcp_close() drops rcu_read_lock() but before it dereferences sock->peer. The synchronize_rcu() in ovpn_socket_release() protects readers that use the dereferenced pointer inside the RCU read section, not those that escape the pointer to a local and use it afterwards. A reproducer follows the pattern of commit 94560267d6c4 ("ovpn: tcp - don't deref NULL sk_socket member after tcp_close()"): trigger a peer removal (keepalive expiration or netlink OVPN_CMD_DEL_PEER) at the same moment userspace closes the TCP fd. That commit fixed the detach-side of the same race window; this one fixes the close-side at a different victim. Tighten the entry block to read sock->peer exactly once into the cached peer local, and route all subsequent uses (the hold check, the ovpn_peer_del() call, and the prot->close() invocation) through that local. sock->peer is only ever written once in ovpn_socket_new() under lock_sock(), before rcu_assign_sk_user_data() publishes the ovpn_socket, and is never reassigned afterwards - but the previous multi-read pattern made that invariant implicit rather than explicit. The same multi-read shape exists in ovpn_tcp_recvmsg(), ovpn_tcp_sendmsg(), ovpn_tcp_data_ready() and ovpn_tcp_write_space(); those will be cleaned up via a dedicated helper in a follow-up net-next series.	8.4	More Details
CVE-2026-64824	Home Assistant Core before 2026.7.0 contains a path traversal vulnerability in the backup-restore function that allows attackers to write files to arbitrary absolute filesystem paths by supplying a crafted tar archive with a SYMTYPE entry containing a benign member name paired with an absolute linkname pointing outside the extraction directory. Because the official Docker image runs the Home Assistant process as root and the subsequent regular-file entry is written through the unvalidated symlink, attackers can achieve remote code execution by overwriting auto-imported Python paths such as site-packages/sitecustomize.py or custom component directories.	8.4	More Details
CVE-2026-63860	In the Linux kernel, the following vulnerability has been resolved: RDMA/core: Prefer NLA_NUL_STRING These attributes are evaluated as c-string (passed to strcmp), but NLA_STRING doesn't check for the presence of a \0 terminator. Either this needs to switch to nla_strcmp() and needs to adjust printf fmt specifier to not use plain %s, or this needs to use NLA_NUL_STRING. As the code has been this way for long time, it seems to me that userspace does include the terminating nul, even tough its not enforced so far, and thus NLA_NUL_STRING use is the simpler solution.	8.4	More Details
CVE-2026-63952	In the Linux kernel, the following vulnerability has been resolved: memfd: deny writeable mappings when implying SEAL_WRITE When SEAL_EXEC is added, SEAL_WRITE is implied to make W^X. But the implied seal is set after the check that makes sure the memfd can not have any writable mappings. This means one can use SEAL_EXEC to apply SEAL_WRITE while having writeable mappings. This breaks the contract that SEAL_WRITE provides and can be used by an attacker to pass a memfd that appears to be write sealed but can still be modified arbitrarily. Fix this by adding the implied seals before the call for mapping_deny_writable() is done.	8.4	More Details
CVE-2026-63828	In the Linux kernel, the following vulnerability has been resolved: apparmor: mediate the implicit connect of TCP fast open sendmsg sendmsg()/sendto() with MSG_FASTOPEN is a combination of connect(2) and write(2): it opens the connection in the SYN. apparmor_socket_sendmsg() only checks AA_MAY_SEND, so a profile that grants send but denies connect lets a confined task open an outbound TCP/MPTCP connection that connect(2) would have refused, bypassing connect mediation. Mediate the implicit connect when MSG_FASTOPEN is set and a destination is supplied. Add it to apparmor_socket_sendmsg() (not the shared aa_sock_msg_perm() helper, which recvmsg also uses) and call aa_sk_perm() directly, mirroring the selinux and tomoyo fixes. sk_is_tcp() does not cover MPTCP fast open, so the SOCK_STREAM/IPPROTO_MPTCP arm is explicit.	8.4	More Details
CVE-2026-15226	A sandbox confinement bypass vulnerability exists in Canonical snapd within its internal execution environment compiler (snap-confine). The default seccomp security templates generated by the engine to restrict system calls do not filter or reject process operations capable of creating or manipulating file execution flags with set-user-ID attributes. Consequently, an application running within a strictly confined snap environment can successfully compile or drop binaries and apply setuid properties to them. If a compromised or malicious process inside the snap sandbox executes these generated setuid binaries, it can potentially circumvent architectural sandboxing assumptions, drop intended restriction policies, or execute privileged actions inside the container namespace that should otherwise be strictly blocked. The vulnerability has been resolved by hardening the seccomp template engine to block the execution and creation of setuid executables by sandboxed snap processes.	8.4	More Details

CVE-2026-64118	In the Linux kernel, the following vulnerability has been resolved: qed: fix double free in qed_cxt_tables_alloc() If one of the later PF or VF CID bitmap allocations fails, qed_cid_map_alloc() jumps to cid_map_fail and frees the previously allocated CID bitmaps before returning an error. qed_cxt_tables_alloc() then calls qed_cxt_mngr_free(), which invokes qed_cid_map_free() again. Fix this by setting each CID bitmap pointer to NULL after bitmap_free() to avoid double free. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1-rc3. Runtime reproduction was not attempted because exercising the failing allocation path requires device-specific setup.	8.4	More Details
CVE-2026-64877	An authenticated non-admin user can exploit a SQL injection flaw in the ticketing REST API to access sensitive data stored in the appliance database.	8.4	More Details
CVE-2026-63864	In the Linux kernel, the following vulnerability has been resolved: bpf: Propagate error from visit_tailcall_insn Commit e40f5a6bf88a ("bpf: correct stack liveness for tail calls") added visit_tailcall_insn() but did not check its return value.	8.4	More Details
CVE-2026-63797	In the Linux kernel, the following vulnerability has been resolved: rpmsg: char: Fix use-after-free on probe error path rpmsg_chrdev_probe() stores the newly allocated eptdev in the default endpoint's priv pointer before calling rpmsg_chrdev_eptdev_add(). If rpmsg_chrdev_eptdev_add() then fails, its error path frees eptdev while the default endpoint may still dispatch callbacks with the stale priv pointer. Avoid publishing eptdev through the default endpoint until rpmsg_chrdev_eptdev_add() succeeds. Messages received before the priv pointer is published should be ignored by rpmsg_ept_cb(). Flow-control updates can hit rpmsg_ept_flow_cb() in the same window, so make both callbacks return success when priv is NULL.	8.4	More Details
CVE-2026-28220	Wazuh is a free and open source platform used for threat prevention, detection, and response. Prior to version 4.14.5, issues in the Cluster Distributed API (DAPI) handling allow a cluster peer, or any actor able to authenticate to the cluster channel using the shared cluster key, to make the master node deserialize an attacker-controlled callable and execute it under an attacker-controlled RBAC context. The cluster code in `framework/wazuh/core/cluster/common.py` deserializes JSON with `as_wazuh_object()`, which resolves any callable whose top-level package is wazuh or api (an overly broad allowlist controlled only by `ALLOWED_CALLABLES_PACKAGES`), and DAPI requests handled in framework/wazuh/core/cluster/dapi/dapi.py accept a client-supplied rbac_permissions value that `run_local()` applies as the global RBAC context, so supplying an rbac_mode of black causes authorization checks for expose_resources-protected functions to pass without any legitimate permission assignment. Combined, these allow privileged administrative actions on the master node such as arbitrary file writes under WAZUH_PATH, creation of new API users, and tampering with security.yaml, and can be chained into full manager compromise. This issue has been fixed in version 4.14.5.	8.4	More Details
CVE-2026-61446	PraisonAI (praisonaiagents) before 1.6.78 contains a remote code execution vulnerability in the plugin manager, which loads and executes arbitrary Python (.py) files from project-level and user-home .praisonai/plugins/ directories using importlib spec_from_file_location() and exec_module() without code signing, integrity verification, or sandboxing. An attacker who can write a malicious .py file to a plugin directory (for example via path traversal, a supply chain attack, or a compromised dependency) achieves arbitrary code execution when the plugin system initializes.	8.4	More Details
CVE-2026-53369	In the Linux kernel, the following vulnerability has been resolved: udf: reject descriptors with oversized CRC length udf_read_tagged() skips CRC verification when descCRCLength + sizeof(struct tag) exceeds the block size. A crafted UDF image can set descCRCLength to an oversized value to bypass CRC validation entirely; the descriptor is then accepted based solely on the 8-bit tag checksum, which is trivially recomputable. Reject such descriptors instead of silently accepting them. A legitimate single-block descriptor should never have a CRC length that exceeds the block.	8.4	More Details
CVE-2026-64081	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_ffa: Validate framework notification message layout Framework notifications carry an indirect message in the shared RX buffer. Validate the reported offset and size before using them, reject zero-length payloads, and ensure that any non-header payload starts at the UUID field rather than in the middle of the message header. Use the validated offset and size values for both kmemdup() and the UUID parsing path so malformed firmware data cannot drive an out-of-bounds read or an oversized allocation.	8.4	More Details
CVE-2026-63906	In the Linux kernel, the following vulnerability has been resolved: usb: musb: omap2430: Fix use-after-free in omap2430_probe() In omap2430_probe(), of_node_put(np) is called prematurely before the last access to np, leading to a use-after-free if the node's reference count drops to zero. Move the of_node_put() calls after the last use of np in both the success and error paths.	8.4	More Details
	In the Linux kernel, the following vulnerability has been resolved: f2fs: bound i_inline_xattr_size for non-inline-xattr inodes When the flexible_inline_xattr feature is enabled, do_read_inode() loads the on-disk i_inline_xattr_size unconditionally: if (f2fs_sb_has_flexible_inline_xattr(sbi)) fi->i_inline_xattr_size = le16_to_cpu(ri->i_inline_xattr_size); but sanity_check_inode() only range-checks it when the inode also has the FI_INLINE_XATTR flag set. An inode that carries an inline dentry or inline data but not FI_INLINE_XATTR -- the normal layout for an inline directory -- therefore keeps a fully attacker-controlled i_inline_xattr_size from a		

CVE-2026-63815	crafted image. get_inline_xattr_addrs() returns that value with no flag gating, so it feeds the inode geometry: MAX_INLINE_DATA() = 4 * (CUR_ADDRS_PER_INODE - i_inline_xattr_size - 1) NR_INLINE_DENTRY() = MAX_INLINE_DATA() * BITS_PER_BYTE / (...) addrs_per_page() = CUR_ADDRS_PER_INODE - i_inline_xattr_size A large i_inline_xattr_size drives MAX_INLINE_DATA() and NR_INLINE_DENTRY() negative, so make_dentry_ptr_inline() sets d->max (int) to a negative value. The inline directory walk then compares an unsigned long bit_pos against that negative d->max, which is promoted to a huge unsigned bound, and reads far past the inline area: while (bit_pos < d->max) /* fs/f2fs/dir.c */ ... test_bit_le(bit_pos, d->bitmap) / d->dentry[bit_pos] ... Mounting a crafted image and reading such a directory triggers an out-of-bounds read in f2fs_fill_dentries(); the same underflow also corrupts ADDRS_PER_INODE for regular files. Validate i_inline_xattr_size against MAX_INLINE_XATTR_SIZE whenever the flexible_inline_xattr feature is enabled -- i.e. whenever the value is loaded from disk and consumed -- and keep the lower MIN_INLINE_XATTR_SIZE bound gated on inodes that actually carry an inline xattr, so legitimate inodes with i_inline_xattr_size == 0 are still accepted.	8.4	More Details
CVE-2026-47415	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. The issue CRUD endpoints (`GET / PATCH / DELETE /workspaces/{workspace_id}/issues/{issue_id}`) gate access on `require_workspace_member(workspace_id)` only, then resolve `issue_id` through `IssueService.get(issue_id)` which is a primary-key lookup with no workspace constraint. A user who is a member of any workspace `W1` can read, modify, or delete issues that belong to a different workspace `W2`. PraisonAI Platform version 0.1.4 patches the issue.	8.3	More Details
CVE-2026-10673	The Zephyr ADIN2111/ADIN1110 10BASE-T1S/T1L Ethernet driver (drivers/ethernet/eth_adin2111.c) reassembles received Ethernet frames in OPEN Alliance (OA) SPI mode by copying device-supplied 64-byte data chunks into a fixed static buffer ctx->buf of size CONFIG_ETH_ADIN2111_BUFFER_SIZE (default 1524 bytes). In eth_adin2111_oa_data_read(), each valid chunk was memcpy'd into ctx->buf[ctx->scur] and the write cursor scur advanced, with no check that scur + len stayed within the buffer. The number of chunks (up to 255, from the BUFSTS RCA field) and the per-chunk length are taken entirely from the frame data received off the wire; the cursor is only reset on a start-of-frame chunk. An attacker on the single-pair Ethernet segment can therefore send a frame whose reassembled size exceeds the configured buffer, causing the driver's RX offload thread to write attacker-controlled frame bytes past the end of the static buffer into adjacent driver/kernel memory (up to roughly 14.8 KB in the worst case). This is a remotely/adjacently reachable out-of-bounds write (CWE-787) that can corrupt memory and cause denial of service or potentially code execution. The defect was introduced when OA SPI support was added (commit 0ca8b0756b1) and shipped in releases v3.7.0 through v4.4.0. The fix adds a bounds check that drops the oversized frame and resets the cursor before the copy.	8.3	More Details
CVE-2026-47419	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an* Insecure Direct Object Reference. The agent CRUD endpoints (`GET / PATCH / DELETE /workspaces/{workspace_id}/agents/{agent_id}`) gate access on `require_workspace_member(workspace_id)` only, then resolve `agent_id` through `AgentService.get(agent_id)` which is a primary-key lookup with no workspace constraint. A user who is a member of any workspace `W1` can read, modify, or delete agents that belong to a different workspace `W2` by guessing or harvesting an agent UUID and calling `…/workspaces/W1/agents/<W2-agent-id>`. PraisonAI Platform version 0.1.4 patches the issue.	8.3	More Details
CVE-2026-20296	In Splunk Enterprise versions below 10.4.1, 10.2.5, 10.0.8, and 9.4.13, and Splunk Cloud Platform versions below 10.5.2605.0, 10.4.2604.7, 10.3.2512.16, 10.2.2510.18, and 10.1.2507.24, an attacker could trick a user that holds a role with the `list_deployment_server` capability into running arbitrary Search Processing Language (SPL) searches on their behalf as `splunk-system-user`, allowing for access to stored credentials and indexed data. The vulnerability is possible because Deployment Server endpoints in Splunk Web do not validate Cross-Site Request Forgery (CSRF) tokens on GET requests, and caller-supplied input is not correctly neutralized before it is placed into an SPL search.	8.3	More Details
CVE-2026-53516	Better Auth is an authentication and authorization library for TypeScript. Prior to 1.6.11, Better Auth's OAuth callback auto-link gate in handleOAuthUserInfo accepts implicit account linking when the OAuth provider asserts email_verified: true without requiring the local user row's emailVerified field to also be true, allowing an attacker who pre-registers a victim email through /sign-up/email to bind the victim's OAuth identity to the attacker's account. The same primitive affects one-tap, and emailAndPassword.requireEmailVerification: true does not mitigate the link-time verification change. This issue is fixed in version 1.6.11.	8.3	More Details
CVE-2026-56400	open-webui before 0.3.14 contains a cross-origin resource sharing misconfiguration allowing arbitrary origins with allow_origins=* and authenticated requests to the /api/v1/functions endpoint. Attackers can execute arbitrary code on the openwebui instance by crafting malicious cross-site requests from attacker-controlled websites when an admin user visits them.	8.3	More Details
CVE-2026-47866	VMware Avi Load Balancer contains an authorization bypass vulnerability. A malicious actor on the network can access a limited subset of the Avi Control Plane without proper authorization. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	8.3	More Details

CVE-2026-55723	When NGINX Ingress Controller is configured with Custom Resource Definitions (CRDs) or Ingress annotations, an injection vulnerability exists in the configuration generator of NGINX Ingress Controller. Multiple user-controllable fields are written into the generated NGINX configuration without sanitization. An authenticated attacker with permission to create or modify these CRDs or annotations may craft values that inject arbitrary NGINX configuration directives. Impact: An authenticated attacker granted write access to NGINX Ingress Controller CRDs or Ingress annotations through the Kubernetes API may be able to inject arbitrary NGINX configuration directives, create or delete files, or disable services. There is no data plane exposure; this is a control plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.3	More Details
CVE-2026-59258	immich before 3.0.3 contains a broken access control vulnerability in the PUT /albums/:id/user/:userId endpoint that allows shared album editors to modify member roles without owner-only restrictions. Attackers with editor access can demote the album owner to editor and promote themselves to owner in sequential requests, gaining full control including deletion and eviction capabilities.	8.3	More Details
CVE-2026-47158	Vaultwarden is a Bitwarden-compatible server written in Rust. Prior to 1.36.0, Vaultwarden's SSO authorization flow did not bind the OAuth state parameter accepted by /connect/authorize to the initiating browser session, allowed attacker-controlled PKCE parameters, and left SsoAuth records intact after failed token exchange, allowing an unauthenticated attacker to induce IdP authentication and redeem tokens for a fully authenticated session. This issue is fixed in version 1.36.0.	8.3	More Details
CVE-2026-62349	TDengine is an open source, time-series database optimized for Internet of Things devices. In 3.4.1.6 and earlier, source/libs/parser/src/parUtil.c trimString() checks space for only one byte before processing SQL string escape sequences \%, _, or \x, allowing a one-byte out-of-bounds write to the stack buffer tmpTokenBuf that can cause denial of service and potentially remote code execution. This issue is fixed in version 3.4.1.14.	8.3	More Details
CVE-2026-47688	FOG is a free open-source cloning/imaging/rescue suite/inventory management system. Prior to versions 1.5.10.1832 and 1.6.0-beta.2313, the `clearAES` and `clearPMTasks` methods in `FOGPage` can be invoked by an unauthenticated attacker via a single HTTP GET request through the public `client` node endpoint. This allows remote wiping of host AES encryption credentials and deletion of all power management scheduled tasks, with no login, session, or CSRF token required. Versions 1.5.10.1832 and 1.6.0-beta.2313 fix the issue.	8.2	More Details
CVE-2026-61435	PraisonAI before 4.6.78 contains an authentication bypass in the Call API agent invocation endpoints (src/praisonai/praisonai/api/agent_invoke.py) when PRAISONAI_CALL_AUTH=disabled is configured. The safeguard intended to restrict the disabled-auth opt-out to localhost binding derives the bind host from request.url.hostname, which is taken from the client-controlled HTTP Host header. A remote, unauthenticated attacker who can reach the service over the network can send a spoofed 'Host: 127.0.0.1' header to bypass the localhost-only restriction and list (GET /api/v1/agents) and invoke (POST /api/v1/agents/{agent_id}/invoke) registered agents without authentication.	8.2	More Details
CVE-2026-58658	GPUStack through 2.2.1, fixed in commit 4e20551, contains an unauthenticated information disclosure vulnerability that allows unauthenticated attackers to access sensitive inference logs and modify worker configuration by exploiting unprotected /serveLogs and /debug endpoints on the worker port. Attackers can enumerate model instance IDs to stream serving logs containing prompts and completions, change log levels, and read memory profiling data without any authentication.	8.2	More Details
CVE-2026-65056	mcp-webresearch 0.1.7 contains a server-side request forgery vulnerability that allows attackers to access internal network services by supplying loopback, link-local, or cloud metadata addresses to the visit_page tool, which only validates the URL protocol without filtering private or reserved IP ranges. Attackers can steer the LLM-controlled URL argument through prompt injection to navigate the server's Playwright browser to internal endpoints such as cloud instance metadata services, causing the server to return sensitive internal page content including credentials into the model context.	8.2	More Details
CVE-2026-35147	HCL DFXServer is affected by a Broken Authentication vulnerability via direct API access. The application fails to verify the user's authentication status when accessing specific API endpoints, allowing an unauthenticated attacker to interact with the APIs and perform unauthorized actions without valid credentials.	8.2	More Details
CVE-2026-13410	Dancer::Plugin::Auth::Google versions through 0.07 for Perl have TLS verification disabled. The default user agent is initialised with SSL_verify_mode explicitly disabled. An attacker with network man-in-the-middle (MITM) capability between the Dancer application and googleapis.com can intercept the OAuth2 token exchange and userinfo fetch, return a forged access_token and user profile, and be logged in to the Dancer application as any Google user.	8.2	More Details
CVE-2026-13577	Dancer2 versions through 2.1.0 for Perl generate insecure session ids when CSPRNG modules are unavailable. Dancer2::Core::Role::SessionFactory::generate_id silently falls back to a built-in rand-derived session id when both Math::Random::ISAAC::XS and Crypt::URandom are unavailable. The fallback session id is generated from a SHA-1 hash of a call to the built-in rand function, the absolute path of the Dancer2::Core::Role::SessionFactory module, an internal counter, the process id, the module instance memory address, and a shuffled string of characters (using the List::Util::shuffle function, which also uses the built-in rand function). These are all low-entropy and easily guessed sources. The built-in rand() function is seeded with 32-bits and considered unsuitable for security applications. Predictable session ids could allow an	8.2	More Details

	attacker to gain access to systems.		
CVE-2026-35149	HCL DFXServer is affected by an Authentication Bypass vulnerability via server response manipulation. An unauthorized user without valid credentials can exploit this flaw by intercepting and altering the server's authentication responses, allowing them to gain unauthorized access to the application without verification.	8.2	More Details
CVE-2026-60005	NGINX Plus and NGINX Open Source have a vulnerability in the ngx_http_slice_module module. When the slice directive and unnamed regex captures are configured or when a background cache update happens, unauthenticated attackers can send requests that may cause uninitialized memory access in the NGINX worker process, leading to limited disclosure of memory or a restart. Impact: This vulnerability may allow remote, unauthenticated attackers to have limited control to disclose memory contents or restart the NGINX worker process. There is no control plane exposure; this is a data plane issue only. Note: The ngx_http_slice_module module is not enabled by default; it's enabled with the --with-http_slice_module configuration parameter. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.2	More Details
CVE-2026-12382	A flaw was found in the AAP Gateway Envoy proxy configuration. The non-mTLS route to EDA event streams does not remove the Subject HTTP header from client requests, despite the source code defining requestHeadersToRemove for this header. An unauthenticated remote attacker can inject a spoofed Subject header matching a legitimate client certificate DN to bypass mTLS authentication and inject arbitrary events into protected EDA event streams.	8.2	More Details
CVE-2026-45325	Gestor de Oferta is a web application for managing mobility service offerings. Prior to 20260509.0340.15, @tmlmobilidade/utils has a prototype pollution vulnerability in setValueAtPath() in packages/utils/src/generic/value-at-path.ts because unsafe path segments are not blocked. This issue is fixed in version 20260509.0340.15.	8.2	More Details
CVE-2026-63913	In the Linux kernel, the following vulnerability has been resolved: netfilter: conntrack: tcp: do not force CLOSE on invalid-seq RST without direction check An unintended behavior in the TCP conntrack state machine allows a connection to be forced into the CLOSE state using an RST packet with an invalid sequence number. Specifically, after a SYN packet is observed, an RST with an invalid SEQ can transition the conntrack entry to TCP_CONNTRACK_CLOSE, regardless of whether the RST corresponds to the expected reply direction. The relevant code path assumes the RST is a response to an outgoing SYN, but does not validate packet direction or ensure that a matching SYN was actually sent in the opposite direction. As a result, a crafted packet sequence consisting of a SYN followed by an invalid-sequence RST can prematurely terminate an active NAT entry. This makes connection teardown easier than intended. So, tighten the state transition logic to ensure that RST-triggered CLOSE transitions only occur when the RST is a valid response to a previously observed SYN in the correct direction.	8.2	More Details
CVE-2026-10130	QueryWeaver contains an authentication bypass vulnerability that allows unauthenticated attackers to obtain valid session tokens for existing accounts by submitting a signup request with a known victim email address. The signup route unconditionally creates and links a new token to the matching Identity via a Cypher MERGE operation before checking whether the email belongs to an existing account, causing the server to return a valid authenticated session token for the victim's identity without requiring any prior credentials or user interaction.	8.2	More Details
CVE-2026-47255	AgenticMail gives AI agents real email addresses and phone numbers. @agenticmail/api prior to version 0.9.32 and @agenticmail/core prior to version 0.9.10 had weakness related to validation and and binding of inactive-agent hour filtering; storage SQL identifier validation; metadata-backed ownership checks for raw storage SQL; blocking direct storage metadata access through raw SQL; fail-closed outbound worker secret handling; SMTP envelope/header control-character validation before command construction; and TLS certificate verification as the default for MailSender with an explicit opt-out for local development. @agenticmail/api prior to version 0.9.32 and @agenticmail/core prior to version 0.9.10 are patched.	8.2	More Details
CVE-2026-63867	In the Linux kernel, the following vulnerability has been resolved: mptcp: close TOCTOU race while computing rcv_wnd The MPTCP output path access locklessly the MPTCP-level ack_seq in multiple times, using possibly different values for the data_ack in the DSS option and to compute the announced rcv wnd for the same packet. Refactor the cote to avoid inconsistencies which may confuse the peer. Also ensure that the MPTCP level rcv wnd is updated only when the egress packet actually contains a DSS ack.	8.2	More Details
CVE-2026-49998	Centrifugo is an open-source scalable real-time messaging server. Prior to 6.8.1, Centrifugo dynamic JWKS endpoint verification could reuse a key for one allowed issuer to verify a JWT for another allowed issuer because the JWKS cache and singleflight lookup were keyed only by JWT header kid, not by the resolved JWKS endpoint, issuer, audience, or trust-domain namespace, affecting client.token.jwks_public_endpoint, client.subscription_token.jwks_public_endpoint, internal/jwks/cache.go, and internal/jwks/manager.go. This issue is fixed in version 6.8.1.	8.2	More Details
CVE-2026-46485	Dashy is a self-hostable personal dashboard. Prior to 4.0.8, Dashy deployments using OIDC can allow unauthenticated users or non-admin authenticated users to write changes to the main config.yaml through the config-saving functionality despite configured permissions, allowing unauthorized modification of dashboard configuration and potential service disruption. This issue is fixed in version 4.0.8.	8.2	More Details

CVE-2026-41521	xrdp is an open source RDP server. Versions 0.10.6 and prior contain an integer overflow vulnerability when processing screen update messages within the vnc-any connection mode. A malicious remote VNC server can send crafted image dimensions that cause an integer overflow during memory buffer size calculation, resulting in an undersized allocation. Subsequent processing of the incoming image data using the original oversized parameters leads to an out-of-bounds read. An unauthenticated remote attacker could exploit this flaw to disclose sensitive information from the heap memory or cause a denial of service (DoS) via a process crash. This issue has been fixed in version 0.10.6.1.	8.2	More Details
CVE-2026-46415	The Caddy Defender plugin is a middleware for Caddy that allows users to block or manipulate requests based on the client's IP address. Prior to version 0.10.1, Caddy Defender used `r.RemoteAddr` when evaluating whether a request should be blocked. `RemoteAddr` is the address of the immediate peer connected to Caddy. In deployments where Caddy is behind a trusted proxy, CDN, or load balancer, the immediate peer is usually the proxy, not the original client. Caddy resolves the original client address into its `client_ip` request variable after applying the configured `trusted_proxies` policy, but Defender did not use that value. As a result, clients from blocked IP ranges could bypass Defender when accessing Caddy through a trusted proxy whose own IP address was not blocked. This affects deployments that use Defender behind trusted proxies and expect it to enforce blocking based on the real client IP. The issue is fixed in version 0.10.1 by making Defender prefer Caddys resolved `client_ip` request variable when it is available. Defender falls back to `RemoteAddr` only when Caddy has not provided a resolved client IP. There is no complete workaround in affected Defender versions for deployments that rely on Caddy's trusted proxy client IP resolution. Until upgrading, affected users should enforce equivalent IP blocking at the trusted proxy, CDN, load balancer, firewall, or other edge layer before traffic reaches Caddy. Deployments where Caddy receives traffic directly from clients, without an intermediate trusted proxy, are not affected by this bypass.	8.2	More Details
CVE-2026-1609	A flaw was found in Keycloak. When the JSON Web Token (JWT) authorization grant preview feature is enabled and a user account is disabled, Keycloak fails to validate the user’s disabled status during JWT authorization grant processing. A remote attacker with low privileges can exploit this improper access control vulnerability by presenting a valid assertion token from an external identity provider to obtain a JWT for a disabled user. This allows unauthorized access to sensitive resources.	8.1	More Details
CVE-2026-59950	The MCP Python SDK, called mcp on PyPI, is a Python implementation of the Model Context Protocol (MCP). Prior to 1.28.1, the deprecated mcp.server.websocket.websocket_server transport accepted WebSocket handshakes without applying Host or Origin header validation, leaving no SDK-level way to restrict which origins could connect to applications that exposed that transport. This issue is fixed in version 1.28.1.	8.1	More Details
CVE-2026-63756	SurrealDB versions before 3.1.0 contain a time-of-check/time-of-use race condition in the HTTP /rpc endpoint that allows unauthenticated requests to inherit authenticated session state. Unauthenticated attackers can send concurrent requests to the /rpc endpoint while legitimate authenticated traffic is active to execute operations with hijacked user privileges.	8.1	More Details
CVE-2026-63889	In the Linux kernel, the following vulnerability has been resolved: scsi: scsi_transport_fc: Widen FPIN pname walker counter to u32 An adjacent Fibre Channel fabric actor that can deliver an FPIN ELS frame to an lpfc or qla2xxx Linux initiator can trigger a non-return in the generic FC transport. This is not a local userspace or IP network path; the attacker must be able to inject fabric traffic, for example as a compromised switch or fabric controller, or as a same-zone N_Port on a fabric that permits source spoofing. The Link-Integrity and Peer-Congestion FPIN walkers used a u8 loop counter against the 32-bit on-wire pname_count field, and did not bound pname_count by the descriptor body already validated by the TLV walker. A pname_count of 256 therefore wraps the counter and keeps the loop condition true indefinitely. Factor the shared pname_list[] walk into one helper, widen the counter to u32, and clamp pname_count against the entries that fit in the descriptor body before iterating.	8.1	More Details
CVE-2026-63735	SurrealDB versions before 3.2.0 fail to validate namespace and database scope in custom API routes, allowing authenticated users to invoke endpoints in different namespaces/databases. Attackers with valid credentials for any namespace/database can access custom API endpoints in other tenants by specifying the target scope in the URL path, reading sensitive data or triggering unintended operations.	8.1	More Details
CVE-2026-63909	In the Linux kernel, the following vulnerability has been resolved: ksmbd: OOB read regression in smb_check_perm_dacl() ACE-walk loops Commit d07b26f39246 ("ksmbd: require minimum ACE size in smb_check_perm_dacl()") introduced a transposed bounds check: if (offsetof(struct smb_ace, sid) + aces_size < CIFS_SID_BASE_SIZE) Since offsetof(..sid) is 8 and CIFS_SID_BASE_SIZE is 8, this evaluates to `aces_size < 0`. Because `aces_size` is always non-negative, this check becomes dead code and never breaks the loop. Worse, that commit removed the old 4-byte guard, meaning the loop now reads `ace->size` (offset 2) even when `aces_size` is 0-3 bytes. This re-opens a 2-byte heap out-of-bounds (OOB) read past the pntsd allocation during subsequent SMB2_CREATE operations. Fix this by properly transposing the comparison to require at least 16 bytes (8-byte offset + 8-byte SID base), matching the correct form used in smb_inherit_dacl().	8.1	More Details
CVE-2026-13142	The Social Login, Passkeys, Magic Link & Email OTP WordPress plugin before 1.4.1 does not enforce rate limiting or a working attempt lockout on its passwordless email one-time-password verification, and stores the short numeric codes in plaintext, allowing an unauthenticated attacker who knows a registered email address to brute-force the code and log in as that user, including an administrator, leading to full site takeover.	8.1	More Details

CVE-2026-15008	The Uncanny Automator – Easy Automation, Integration, Webhooks & Workflow Builder Plugin plugin for WordPress is vulnerable to arbitrary file deletion due to insufficient file path validation in the fr_token function in all versions up to, and including, 7.3.1.4. This makes it possible for unauthenticated attackers to delete arbitrary files on the server, which can easily lead to remote code execution when the right file is deleted (such as wp-config.php). Exploitation requires a Forminator form connected to an Uncanny Automator recipe configured for 'Everyone', allowing unauthenticated form submissions to supply the malicious serialized payload; a gadget chain is present within the plugin via the Action_Helpers_Email __destruct() method, meaning no external gadget library is required.	8.1	More Details
CVE-2021-27137	An issue was discovered in router/upnp/src/ssdp.c in DD-WRT before 45724. An unsafe strcpy in the UPnP handling functionality allows an unauthenticated remote attacker to send a request that would overflow an internal fixed buffer. Exploitation requires the DD-WRT user to enable UPnP (which is off by default, and only listens on internal interfaces by default). This occurs in ssdp_msearch (reachable by an M-SEARCH request).	8.1	More Details
CVE-2026-63925	In the Linux kernel, the following vulnerability has been resolved: macsec: fix replay protection at XPN lower-PN wrap In macsec_post_decrypt(), when pn is U32_MAX, pn + 1 overflows u32 to 0 and the first branch never fires. If next_pn_halves.lower is also in the upper half, pn_same_half(pn, lower) is true and the XPN else-if does not fire either, leaving next_pn_halves unchanged. An attacker that captures the legitimate frame carrying pn == 0xFFFFFFFF on an XPN association can then replay it indefinitely, since lowest_pn never rises above the captured pn and macsec_decrypt() reconstructs the same IV. Extend the XPN else-if to also fire when pn + 1 wraps to 0, so receipt of pn == U32_MAX advances next_pn_halves to (upper + 1, 0).	8.1	More Details
CVE-2026-63893	In the Linux kernel, the following vulnerability has been resolved: thunderbolt: property: Reject u32 wrap in tb_property_entry_valid() entry->value is u32 and entry->length is u16; the sum is performed in u32 and wraps. A malicious XDomain peer can pick value = 0xffffff00, length = 0x100 so the sum 0x100000000 wraps to 0 and passes the > block_len check. tb_property_parse() then passes entry->value to parse_dwdata() as a dword offset into the property block, reading attacker-directed memory far past the allocation. For TEXT-typed entries with the "deviceid" or "vendorid" keys this lands in xd->device_name / xd->vendor_name and is readable back via the per-XDomain device_name / vendor_name sysfs attributes; the leak is NUL-bounded (kstrdup() stops at the first zero byte) and untargeted (the attacker picks a delta, not an absolute address). DATA-typed entries are parsed into property->value.data but not generically surfaced to userspace. Use check_add_overflow() so a wrapped sum is rejected.	8.1	More Details
CVE-2026-64176	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mvm: fix driver-set TX rates on old devices On old devices such as 7265D, rates are still encoded in version 1 format, which doesn't use the CCK/OFDM rate index (0-3/0-7) but rather their PLCP value (e.g. 10 for 1 Mbps CCK rate.) While introducing v3 rates, I changed the driver from internally handling v1 rates and converting to v2, to internally handling v3 and converting to v1 or v2 according to the firmware. I accordingly changed the code in iwl_mvm_mac80211_idx_to_hwrate() to no longer have different values for different APIs. This was correct. However, I later reverted this part of the change, because it was reported that I had broken beacon rates, causing a FW assert/crash. This caused TX_CMD rates to be set incorrectly, potentially causing a warning when reported back from the device as having been used. Fix this (hopefully correctly now) by handling beacon rates in the TX_CMD that's embedded in the beacon template command separately. Restore iwl_mvm_mac80211_idx_to_hwrate() to return only the rate index, not PLCP value, fixing the real TX_CMD.	8.1	More Details
CVE-2026-63305	AVideo through 29.0 contains an OS command injection vulnerability in the ffmpeg.json.php endpoint where notifyCode and callback parameters are concatenated into a shell command without escaping. Attackers who can craft a valid encrypted payload can inject arbitrary shell metacharacters into these fields to execute OS commands as the web-server user.	8.1	More Details
CVE-2026-63304	AVideo through 29.0 contains an OS command injection vulnerability in plugin/API/standAlone/functions.php where the listFFmpegProcesses() function interpolates unsanitized keyword parameters inside single quotes without escaping. Attackers who can craft a valid encrypted codeToExec payload can break out of the single-quoted grep context and execute arbitrary OS commands as the web-server user.	8.1	More Details
CVE-2025-60357	AhnLab EPP Management v1.0.14.32-6249 was discovered to contain a NoSQL injection vulnerability via the eventlog/agentEvent/list endpoint.	8.1	More Details
CVE-2026-12585	The Abandoned Cart Lite for WooCommerce WordPress plugin before 6.8.2 does not protect the integrity of its cart-recovery tokens or bind them to the requesting account, allowing unauthenticated attackers to forge a recovery link that logs them in as another user when the automatic-login option is enabled.	8.1	More Details
CVE-2026-62418	Low-privileged authenticated Server-Side Request Forgery (SSRF) vulnerability in Apache Syncope via Connectors and Resources check. This issue affects Apache Syncope: from 3.0.0-M0 through 3.0.16, from 4.0.0-M0 Through 4.0.6, from 4.1.0-M0 through 4.1.1. Users are recommended to upgrade to version 4.0.7 / 4.1.2, which fix this issue.	8.1	More Details
CVE-	The Grav API plugin (getgrav/grav-plugin-api) before 1.0.6 contains an authorization bypass: API keys can be created with a restricted scopes array, but the ApiKeyAuthenticator class never reads or enforces these		

2026-62231	scopes. It loads and returns the owning user's full account object, so a key created with limited scopes (e.g. read-only) can perform any write, delete, or administrative operation the owning user is authorized for. Fixed in 1.0.6.	8.1	More Details
CVE-2026-13473	IBM Storage Protect Client 8.1.0.0 through 8.1.27.0, 8.1.27.1, and 8.2.0.0 through 8.2.1.0 IBM Storage Protect is vulnerable to a heap-based buffer overflow, caused by improper bounds checking. A remote attacker could overflow a buffer and execute arbitrary code on the system or cause the server to crash.	8.1	More Details
CVE-2026-44019	Docling Core defines core data types and transformations for the document processing application Docling. In versions 2.5.0 and above, prior to 2.74.1, docling-core could allow local file:// image references and accepted inline data: content without a decoded-size limit. In applications that accept untrusted image references, this may allow access to local files readable by the process or excessive memory use from large inline payloads. This issue has been fixed in version 2.74.1.	8.1	More Details
CVE-2026-62209	OpenClaw versions 2026.5.10-beta.1 before 2026.6.5 contain an authorization bypass in the ClickClack agent-mode dispatch feature, which could ignore the toolsAllow policy check. When the affected feature is enabled and reachable, a lower-trust caller or configured input path could perform actions that should have required a stronger authorization or policy check.	8.1	More Details
CVE-2026-45260	Pimcore is an Open Source Data & Experience Management Platform. Prior to 11.5.17 (LTS) and 12.3.7, Pimcore's WebDAV asset endpoint exposes a MOVE operation through /asset/webdav{path} without an authentication plugin in bundles/CoreBundle/src/Controller/WebDavController.php, and models/Asset/WebDAV/Tree.php performs asset mutation and deletion through models/Asset.php before checking a current Pimcore user or the rename, delete, create, or publish permissions, allowing unauthorized asset deletion, moves, or overwrites. This issue is fixed in versions 11.5.17 (LTS) and 12.3.7.	8.1	More Details
CVE-2026-62685	File Browser is a file managing interface for uploading, deleting, previewing, renaming, and editing files within a specified directory. Prior to 2.63.17, File Browser builds new user scopes from usernames passed through cleanUsername() when Signup=true and CreateUserDir=true, but the many-to-one normalization can collapse usernames such as team/one, team one, and team-one to the same home directory without checking whether the resulting scope is already taken, allowing a second registrant to gain full read and write access to another user's files. This issue is fixed in version 2.63.17.	8.1	More Details
CVE-2026-32821	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, any authenticated API user who has their own access token can ask the collection API to evaluate permissions as a different user by supplying `user_email`. If the target user has collections, this can expose those collections through the API. In V4, once a collection id is known, the same controller also offers `add_item` and `remove_item` routes without any object-level `authorize!` checks, creating a likely cross-user modification path. This is patched in version 26.06.08.	8.1	More Details
CVE-2026-13445	IBM Langflow OSS 1.0.0 through 1.10.1 can allow an authenticated attacker to exploit the SaveToFile component to read and modify another user's uploaded files by specifying absolute paths pointing to victim storage locations. In append mode, the attacker's workflow reads victim file contents, appends attacker-controlled data, and uploads a copy containing victim data to the attacker's namespace (confidentiality breach). In overwrite mode, the attacker can replace victim file contents with arbitrary data (integrity breach). This breaks the storage ownership boundary between users.	8.1	More Details
CVE-2026-62234	Grav before 2.0.4 fails to restrict cURL protocols in webhook dispatch, allowing authenticated users with api.webhooks.write permission to create webhooks with file://, dict://, or gopher:// URLs. Attackers can trigger webhook events to read local files, access process information, or pivot to internal services via unrestricted protocol handlers.	8.1	More Details
CVE-2026-61443	PraisonAI before 1.6.78 contains a remote code execution vulnerability in SkillTools.run_skill_script() that executes scripts without path containment validation. Attackers can supply absolute file paths to execute arbitrary scripts from any filesystem location, including those outside the intended working directory.	8.1	More Details
CVE-2026-56287	A boolean-based SQL Injection vulnerability exists in Apache Fineract's Client Search API (GET /api/v1/clients) in versions up to and including 1.14.0. The orderBy and sortOrder request parameters are concatenated into a SQL query without sufficient validation, allowing an authenticated user with permission to view clients to inject arbitrary SQL via a crafted orderBy value. This can be leveraged to perform blind boolean-based data extraction and, on MySQL/MariaDB, to disclose arbitrary files readable by the database process via the LOAD_FILE() function. Users are recommended to upgrade to a version containing the fix	8.1	More Details
CVE-2026-47418	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. The project CRUD endpoints (`GET / PATCH / DELETE /workspaces/{workspace_id}/projects/{project_id}` and `GET .../{project_id}/stats`) gate access on `require_workspace_member(workspace_id)` only, then resolve `project_id` through `ProjectService.get(project_id)` / `update(project_id, ...)` / `delete(project_id)` / `get_stats(project_id)`. None of these calls thread `workspace_id` through to constrain the lookup. A user who is a member of any workspace `W1` can read, modify, delete, or read stats for projects that belong to a different workspace `W2`. PraisonAI Platform version 0.1.4 patches the issue.	8.1	More Details

CVE-2026-47417	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. The comment endpoints (`POST /workspaces/{workspace_id}/issues/{issue_id}/comments` and `GET .../comments`) gate access on `require_workspace_member(workspace_id)` only, then call `CommentService.create(issue_id=issue_id, ...)` and `CommentService.list_for_issue(issue_id)` without verifying that `issue_id` belongs to `workspace_id`. A user who is a member of any workspace `W1` can read every comment on, and post new comments to, any issue in any other workspace `W2`. PraisonAI Platform version 0.1.4 patches the issue.	8.1	More Details
CVE-2026-47412	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an authorization bypass enabling destructive action. The `DELETE /workspaces/{workspace_id}` endpoint is gated only by `require_workspace_member(workspace_id)` (default `min_role="member"`). Any member of the workspace can issue a single DELETE to wipe the entire workspace, including every project, issue, comment, agent, label, and member record (cascading via the foreign-key relationships). There is no owner-role gate, no confirmation token, no soft-delete window, no recovery path. PraisonAI Platform version 0.1.4 patches the issue.	8.1	More Details
CVE-2026-47409	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an authorization bypass enabling owner lockout. The `DELETE /workspaces/{workspace_id}/members/{user_id}` endpoint is gated only by `require_workspace_member(workspace_id)` (default `min_role="member"`). Any member can remove any other member, including the workspace owner, using a single DELETE. There is no caller-role check, no target-role check, no "cannot remove last owner" guard. PraisonAI Platform version 0.1.4 patches the issue.	8.1	More Details
CVE-2026-53390	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix out-of-bounds read in smb_check_perm_dacl() The permission-check ACE walk in smb_check_perm_dacl() validates the ACE header size and caps sid.num_subauth at SID_MAX_SUB_AUTHORITIES, but it never checks that ace->size is actually large enough to contain num_subauth sub-authorities before compare_sids() dereferences them. CIFS_SID_BASE_SIZE covers the SID header up to but excluding the sub_auth[] array, and offsetof(struct smb_ace, sid) is the ACE header, so the existing guards only guarantee the 8-byte SID base, i.e. zero sub-authorities. compare_sids() then reads ace->sid.sub_auth[i] for i < min(local_sid->num_subauth, ace->sid.num_subauth). The local comparison SIDs (sid_everyone, sid_unix_NFS_mode, and the id_to_sid() result) always have at least one sub-authority, and an attacker controls the ACE revision and authority bytes (which lie within the in-bounds SID base), so they can match one of those SIDs and force the sub_auth read. A crafted ACE with size == 16 and num_subauth >= 1 placed at the tail of the security descriptor therefore causes a heap out-of-bounds read of up to SID_MAX_SUB_AUTHORITIES * sizeof(__le32) bytes past the pntsd allocation. The security descriptor is loaded by ksmbd_vfs_get_sd_xattr() into a buffer sized exactly to the on-disk data (kzalloc(sd_size) in ndr_decode_v4_ntacl()), so the read lands past the allocation. The malformed descriptor can be stored verbatim via SMB2_SET_INFO (the DACL is not normalised before being written to the security.NTACL xattr) and the read fires on a subsequent SMB2_CREATE access check, making this reachable by an authenticated client on a share that uses ACL xattrs. Add the missing num_subauth-versus-ace_size check, mirroring the identical guards already present in the sibling parsers parse_dacl() and smb_inherit_dacl().	8.1	More Details
CVE-2026-61873	Grav before 9.1.8 contains an arbitrary file write vulnerability in the Form plugin's process.save.filename parameter, which is validated against path traversal before Twig processing but never re-validated after rendering. Attackers can submit form data containing path traversal sequences that are processed through Twig templates, allowing them to write arbitrary files including PHP webshells to the web root or other sensitive directories.	8.1	More Details
CVE-2026-47406	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. The dependency endpoints (`POST/GET /workspaces/{workspace_id}/issues/{issue_id}/dependencies` and `DELETE .../dependencies/{dep_id}`) gate access on `require_workspace_member(workspace_id)` only, then dispatch to `DependencyService` calls that take URL/body-supplied issue and dependency IDs without verifying any of them belong to the membership-checked workspace. Most damaging: `create_dependency` accepts `body.depends_on_issue_id` from the request body — that ID is checked against nothing — letting an attacker create a "blocks" or "related" link between any two issues anywhere in the database. PraisonAI Platform version 0.1.4 patches the issue.	8.1	More Details
CVE-2026-20156	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20156 are related to improper restriction of operations within the bounds of a memory buffer that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-119.	8.1	More Details
CVE-2026-57821	A SQL Injection vulnerability exists in Apache Fineract's Office Search API (GET /api/v1/offices) in versions up to and including 1.14.0. The orderBy request parameter is concatenated into a SQL query without sufficient validation, allowing an authenticated user with permission to view offices to inject arbitrary SQL via a crafted orderBy value. This is a bypass of the ColumnValidator fix introduced for CVE-2024-32838, which does not detect bare subqueries in the ORDER BY position. This can be leveraged to perform time-based blind SQL injection for data exfiltration. Because the injected query blocks the database connection for its full duration,	8.1	More Details

	concurrent exploitation can exhaust the application's database connection pool, resulting in denial of service for other users. Users are recommended to upgrade to a version containing the fix.		
CVE-2026-47398	PraisonAI is a multi-agent teams system. The v4.6.32 chokepoint refactor (which patched CVE-2026-44334 / GHSA-xcmw-grxf-wjhj) added the PRAISONAI_ALLOW_LOCAL_TOOLS env-var gate to the tool_override.py sinks. However, two additional spec.loader.exec_module call sites in praisonai/agents_generator.py were missed and remain completely unguarded in versions prior to 4.6.40. Both functions accept a module_path parameter sourced from YAML configuration and execute it without validation, signature checking, or the env-var gate. Version 4.6.40 fixes the issue.	8.1	More Details
CVE-2026-11961	The User Registration & Membership WordPress plugin before 5.2.3 does not validate that the membership tier submitted during public registration is one of the tiers allowed by the registration form before assigning that tier's associated user role, allowing unauthenticated users to register into an arbitrary published membership tier and obtain its role — up to administrator when such a tier exists.	8.1	More Details
CVE-2026-10678	The MCTP-over-I2C+GPIO target binding in Zephyr (subsys/pmci/mctp/mctp_i2c_gpio_target.c) processes pseudo-register writes from an I2C bus master byte-by-byte in mctp_i2c_gpio_target_write_received() without validating the order or the receive buffer. In the affected versions the MCTP_I2C_GPIO_RX_MSG_ADDR (data) handler dereferences and writes through b->rx_pkt without checking that the receive buffer was allocated: a controller that selects the data register and writes a byte without first sending the length register (which is what allocates the buffer) causes a write of an attacker-chosen byte through a NULL/unallocated mctp_pktbuf pointer (i.e. into a small attacker-advanceable offset above address 0), producing memory corruption or a hard fault. The same handler also performs a write-then-check bounds test, allowing a one-byte heap overflow at data[255] when more than 255 data bytes are sent. Because the I2C target callback is invoked with raw bytes supplied by whatever device is the bus master and the binding performs no authentication, a malicious or malfunctioning controller on the bus can trigger these without any prior protocol state, leading to memory corruption and/or denial of service on the target device. The vulnerable code was introduced when the I2C+GPIO target binding was added and shipped in Zephyr v4.3.0 and v4.4.0. The fix defers allocation to the first data byte with a NULL check, treats a missing length as a zero-sized packet rejected by libmctp, and moves the bounds check before the store.	8.1	More Details
CVE-2026-63094	SigNoz through 0.133.0 contains an open redirect vulnerability in the SSO authentication flow that allows unauthenticated attackers to steal session tokens from any user on instances configured with Google OAuth, SAML, or OIDC. Attackers can call the unauthenticated sessions context endpoint with a ref parameter pointing to an attacker-controlled host, deliver the resulting crafted login URL to a victim, and receive the victim's access and refresh tokens when they complete SSO authentication.	8.1	More Details
CVE-2026-43978	wger is a free, open-source workout and fitness manager. In versions prior to 2.6, a gym trainer can escalate their session to any higher-privileged account (gym manager, general manager) by chaining two calls to the trainer-login endpoint. Once a trainer performs a legitimate switch into a low-privileged user, the session flag trainer.identity is set and this flag alone bypasses the permission check on all subsequent trainer-login calls. This grants full gym administration capabilities including viewing all member data, modifying contracts, managing gym configuration, and accessing other trainers' and managers' personal information. This issue has been fixed in version 2.6.	8.1	More Details
CVE-2026-46351	BigBlueButton is an open-source virtual classroom. Prior to 3.0.21, bbb-web generated conference sessionToken values with insufficiently secure randomness in bbb-common-web/src/main/java/org/bigbluebutton/api/Util.java and bigbluebutton-web/grails-app/controllers/org/bigbluebutton/web/controllers/ApiController.groovy, allowing a session user to predict other users' conference session tokens and impersonate them. This issue is fixed in version 3.0.21.	8.1	More Details
CVE-2026-47129	NextCRM is open-source customer relationship management (CRM) software. Versions prior to 0.12.0 have a Broken Access Control (BAC) vulnerability in the `activateUser` and `deactivateUser` Next.js Server Actions of NextCRM. The application fails to verify if the requesting user holds the `admin` role. Consequently, any authenticated user (even those with the lowest `member` or `viewer` roles) can arbitrarily activate or deactivate any user account in the system, including the main administrator. Version 0.12.0 fixes the issue.	8.1	More Details
CVE-2026-9323	The urwid web display backend (urwid/display/web.py) generates web session identifiers (urwid_id) in Screen.start() by concatenating two random.randrange(10**9) calls that use Python's Mersenne Twister PRNG, which is not cryptographically secure. Each call consumes approximately 30 bits of PRNG state, and the Mersenne Twister internal state is approximately 19,937 bits, so an attacker who observes approximately 334 session IDs (for example via the X-Urwid-ID HTTP response header) can fully reconstruct the internal state and predict all past and future session IDs (Path B). The same identifier is also used as the filename of a FIFO created in the world-listable /tmp directory (for example /tmp/urwid375487765176907690.in), so any local user on the host can list /tmp to enumerate active session tokens directly (Path A). With a valid session ID, an attacker can read the victim's terminal screen via the polling endpoint, inject keystrokes into the victim's session (yielding OS-level code execution with the session owner's privileges if the session runs a shell), and inject exit sequences or flood the FIFO to terminate or crash the session. A prior Bandit S311 warning on this usage was suppressed with # noqa: S311 rather than fixed	8.1	More Details
CVE-	BigBlueButton is an open-source virtual classroom. Prior to 3.0.21, bbb-web checksum validation could be		

2026-46353	bypassed when a presentationUploadExternalUrl parameter was supplied to API request handling in CreateMeeting.java and ValidationService.java, allowing a user to send valid requests to some endpoints without a checksum. This issue is fixed in version 3.0.21.	8.1	More Details
CVE-2026-55173	WWBN AVideo is an open source video platform. Versions 29.0 and below remain vulnerable to OS command injection because the fix for CVE-2026-33482 was incomplete and still does not neutralize a single & (the shell background operator). CVE-2026-33482 reported that sanitizeFFmpegCommand() (plugin/API/standAlone/functions.php) failed to strip $(...) command substitution, allowing OS command injection at the execAsync() sh -c sink. The fix (commit 25c8ab90) added $, (,), {, }, \n, \r to the denylist character class and a str_replace('&&', '', ...), but did not account for the single &. ffmpeg.json.php builds the command from _decryptString(getInput('codeToExecEncrypted')). This is the same threat model the original advisory accepted (“an attacker who can craft a valid encrypted payload can achieve arbitrary command execution on the standalone encoder server”) and the same CVSS basis (AV:N/AC:H/PR:N). Multiple &-separated commands can be chained (e.g. download + execute). Redirect-based payloads are blocked by the > strip, but command execution (e.g. & curl http://attacker/..., & nc ..., dropping/running a file) is not. This issue has been patched by this commit: https://github.com/WWBN/AVideo/commit/c1cfa2bea8a351a1d07f5758f82887403e3abf1f.	8.1	More Details
CVE-2026-12281	The Shibboleth WordPress plugin before 2.5.4 does not fail closed when its HTTP header identity mode is enabled without an anti-spoofing key, treating any request that carries identity headers as an authenticated session without verifying them. On a deployment where untrusted client headers reach the application, an unauthenticated attacker can log in with forged identity headers and, when automatic account creation and the default administrator role mapping are enabled, create and sign in as a new administrator. Exploitation requires the non-default HTTP header attribute mode, an empty or absent spoof key, automatic account creation enabled, and a deployment that does not strip untrusted client headers before they reach the application.	8.1	More Details
CVE-2026-42533	A vulnerability exists in NGINX Plus and NGINX Open Source when a map directive uses regex matching and a string expression references the map's regex capture variables before referencing the map output variable. Alternatively, the same result could be achieved by using a non-cacheable variable in a string expression under certain conditions. An unauthenticated attacker along with conditions beyond their control can exploit this vulnerability by sending crafted HTTP requests. This may cause a heap buffer overflow in the NGINX worker process leading to a restart. Additionally, attackers can execute code on systems with Address Space Layout Randomization (ASLR) disabled or when the attacker can bypass ASLR. Impact: This vulnerability may allow remote attackers to cause a denial-of-service (DoS) on the NGINX system or to possibly trigger a code execution. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details
CVE-2026-58660	Kanboard through 1.2.52, fixed in commit 564cc30, BoardAjaxController save() method (used by the kanban board drag-and-drop endpoint) validates the caller's role on the attacker-supplied project_id but never verifies that the supplied task_id actually belongs to that project. Because task identifiers are sequential integers shared across the entire instance, any authenticated user who is a member of at least one project can enumerate and move (corrupt/hide) tasks belonging to any other project on the same instance, including private projects they have no membership or role on.	8.1	More Details
CVE-2026-13448	IBM Langflow OSS 1.0.0 through 1.10.1 Lanflow OSS contains an unauthenticated remote code execution vulnerability in the public flow build endpoint (/api/v1/build_public_tmp/{flow_id}/flow). The vulnerability stems from an incomplete denylist in the validate_public_flow_no_code_execution() function that fails to block several code-execution agent components including OpenDsStarAgent, CodeActAgentSmolagents, and CSVAgent.	8.1	More Details
CVE-2026-53518	Better Auth is an authentication and authorization library for TypeScript. From 1.6.0 until 1.6.11, the @better-auth/oauth-provider POST /oauth2/token endpoint for the authorization_code grant redeems a single-use authorization code through a non-atomic find-then-delete sequence, allowing two concurrent requests to pass the read step and mint independent access tokens, refresh tokens, and ID tokens; legacy /oauth2/token and /mcp/token paths in oidc-provider and mcp plugins share the same primitive. This issue is fixed in version 1.6.11.	8.1	More Details
CVE-2026-53517	Better Auth is an authentication and authorization library for TypeScript. From 1.4.8-beta.7 until 1.6.11, the @better-auth/oauth-provider POST /oauth2/token endpoint on the refresh_token grant performs a non-atomic read, validate, revoke, and mint sequence on the oauthRefreshToken row, allowing concurrent requests with the same parent refresh token to pass the revoked check and create forked refresh-token families; the vulnerable range also includes embedded better-auth plugin versions before 1.6.0. This issue is fixed in version 1.6.11.	8.1	More Details
CVE-2026-62215	OpenClaw versions before 2026.6.5 contain an authentication bypass vulnerability in HTTP Canvas responses that allows lower-trust callers to forge trusted A2UI actions. Attackers can perform actions requiring stronger authorization by submitting crafted requests through configured input paths, bypassing intended policy checks.	8.0	More Details
	Pimcore is an Open Source Data & Experience Management Platform. Prior to 11.5.17 (LTS) and 12.3.7,		

CVE-2026-45162	multiple Pimcore locations call PHP's unserialize() on data from database columns and filesystem files without the allowed_classes restriction, including lib/Tool/Authentication.php, models/Site/Dao.php, models/DataObject/ClassDefinition/CustomLayout/Dao.php, models/Tool/TmpStore/Dao.php, models/Asset/WebDAV/Service.php, and admin-ui-classic-bundle/src/Helper/Dashboard.php, enabling object injection and remote code execution if an attacker can control the serialized data source. This issue is fixed in versions 11.5.17 (LTS) and 12.3.7.	8.0	More Details
CVE-2026-47237	Kubeflow Community Distribution helps users to install Kubeflow Platform in popular Kubernetes clusters. Prior to version 26.03-rc.1, a Kubeflow setup based on the official manifests or most other packaged Kubeflow distributions is vulnerable to authorization token stealing from any user of the Kubeflow UI or APIs, such as the Dashboard, Pipelines API, or Notebooks. With this token, the attacker can take over the user's account and the data that is processed by that user. The attacker needs a valid user with the ``kubeflow-edit`` role / Contributor role in a random Kubeflow namespace to perform this attack. This is given if _Automatic Profile Creation_ is enabled. Version 26.03-rc.1 fixes the issue.	8.0	More Details
CVE-2026-55626	xrdp is an open source RDP server. In versions 0.10.6 and prior, when an authenticated user session is initialized using the Xvnc backend over UNIX domain sockets, the Xvnc process is launched with insufficient authentication mechanisms. A local authenticated attacker could exploit this vulnerability to bypass intended session isolation, allowing them to unauthorizedly view or control the active desktop sessions of other users on the same system. Users using other backends, such as xorgxrdp or Xvnc over TCP sockets, are not affected. This issue has been fixed in version 0.10.6.1.	8.0	More Details
CVE-2026-46923	Vulnerability in the Oracle Public Sector Financials (International) product of Oracle E-Business Suite (component: Authorization). Supported versions that are affected are 12.2.3-12.2.15. Difficult to exploit vulnerability allows high privileged attacker with network access via HTTP to compromise Oracle Public Sector Financials (International). While the vulnerability is in Oracle Public Sector Financials (International), attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle Public Sector Financials (International). CVSS 3.1 Base Score 8.0 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:H/UI:N/S:C/C:H/I:H/A:H).	8.0	More Details
CVE-2026-62222	OpenClaw before 2026.5.22 contain a vulnerability in setup-mode discovery that allows loading of untrusted workspace plugins. Attackers with lower-trust caller access or control over configured input paths can execute or persist actions beyond their intended authorization level.	7.8	More Details
CVE-2026-47868	VMware Avi Load Balancer contains a local privilege escalation vulnerability. A malicious user with local access may be able to escalate their privileges to run code as root. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	7.8	More Details
CVE-2026-9147	uproot dynamically generates Python class source code from ROOT TStreamerInfo records in a file and compiles it at runtime. Some file-controlled streamer metadata fields (for example, streamer element names) are interpolated into the generated Python source without safe quoting via repr() or the !r format specifier. An attacker who can supply a crafted ROOT file can place Python expression-breaking content into a streamer metadata field. When uproot generates and invokes the corresponding reader method, the injected Python expression is evaluated in the context of the process opening the file, resulting in arbitrary Python code execution in applications that open or process attacker-controlled ROOT files with affected uproot code paths.	7.8	More Details
CVE-2026-63918	In the Linux kernel, the following vulnerability has been resolved: l2tp: use refcount_inc_not_zero in l2tp_session_get_by_ifname A reader in l2tp_session_get_by_ifname() can return a pointer to a session whose refcount has reached zero. The getter takes its reference with plain refcount_inc(), but every other session getter in the same file (l2tp_v2_session_get, l2tp_v3_session_get, and the corresponding _get_next variants) uses refcount_inc_not_zero() because the IDR/RCU lookup can race with refcount_dec_and_test() -> l2tp_session_free() -> kfree_rcu(). The ifname getter is the only outlier; the inconsistency was raised on-list after 979c017803c4 ("l2tp: use list_del_rcu in l2tp_session_unhash"). A reader inside rcu_read_lock_bh() that matches session->ifname can be preempted between the strcmp() and the refcount_inc(). If the last reference drops on another CPU in that window, the reader's refcount_inc() runs on a counter that has reached zero. refcount_t catches the addition-on-zero, prints "refcount_t: addition on 0; use-after-free", saturates the counter, and returns the saturated pointer to the caller. Session memory is held live by the in-flight RCU read section, but the kfree_rcu() callback queued from l2tp_session_free() will free it once the grace period closes; a caller that dereferences the returned session past that point hits a slab-use-after-free. On PREEMPT_RT local_bh_disable() is a per-CPU sleeping lock and the preemption window is real; on stock PREEMPT kernels local_bh_disable() is a preempt_count increment that closes the cross-CPU race in practice (see below). Use refcount_inc_not_zero() and continue the list walk on failure, matching the other session getters in the file. The ifname getter is the only session getter in net/l2tp/ that still uses the bare refcount_inc() pattern; this change restores file-internal consistency. The success path is unchanged.	7.8	More Details
CVE-2026-63844	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v4.0.3 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 2f6afc97d259d530f4f86c7743efbc573a8da927)	7.8	More Details

CVE-2026-63945	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: serialize iso_sock_clear_timer with socket lock iso_sock_close() calls iso_sock_clear_timer() before acquiring lock_sock(sk). iso_sock_clear_timer() reads iso_pi(sk)->conn twice without the socket lock held: if (!iso_pi(sk)->conn) return; cancel_delayed_work(&iso_pi(sk)->conn->timeout_work); Concurrently, iso_conn_del() executes under lock_sock(sk) and calls iso_chan_del(), which sets iso_pi(sk)->conn to NULL and may result in the final reference to the connection being dropped: CPU0 CPU1 ---- ---- iso_sock_clear_timer() if (conn != NULL) ... lock_sock(sk) iso_chan_del() iso_pi(sk)->conn = NULL cancel_delayed_work(conn) /* NULL deref or UAF */ iso_pi(sk)->conn is not stable across the unlock window, causing a NULL pointer dereference or use-after-free. Serialize iso_sock_clear_timer() with the socket lock by moving it inside lock_sock()/release_sock(), matching the pattern used in iso_conn_del() and all other call sites.	7.8	More Details
CVE-2026-63927	In the Linux kernel, the following vulnerability has been resolved: usb: dwc2: Fix use after free in debug code We're not allowed to dereference "urb" after calling usb_hcd_giveback_urb() so save the urb->status ahead of time.	7.8	More Details
CVE-2026-63930	In the Linux kernel, the following vulnerability has been resolved: iio: buffer: hw-consumer: fix use-after-free in error path In the err_put_buffers cleanup path of iio_hw_consumer_alloc(), the code was using list_for_each_entry() to iterate through buffers while calling iio_buffer_put() which can free the current buffer if refcount drops to 0. The list_for_each_entry() loop macro then evaluates buf->head.next to continue iteration, accessing the freed buffer. Fix this by using list_for_each_entry_safe().	7.8	More Details
CVE-2026-63942	In the Linux kernel, the following vulnerability has been resolved: parport: Fix race between port and client registration The parport subsystem registers port devices before they are fully initialised, resulting in a race condition where client drivers such as lp can attach to ports that are not completely initialised or even being torn down. When the port and client drivers are built as modules and loaded around the same time during boot, this occasionally results in a crash. I was able to make this happen reliably in a VM with a PC-style parallel port by patching parport_pc to fail probing: > --- a/drivers/parport/parport_pc.c > +++ b/drivers/parport/parport_pc.c > @@ -2069,7 +2069,7 @@ static struct parport *__parport_pc_probe_port(unsigned long int base, > if (!p) > goto out3; > > - base_res = request_region(base, 3, p->name); > + base_res = NULL; > if (!base_res) > goto out4; > and then running: while true; do modprobe lp & modprobe parport_pc wait rmmod lp parport_pc done for a few seconds. In the long term I think port registration should be changed to put the call to device_add() inside parport_announce_port(), but since the latter currently cannot fail this will require changing all port drivers. For now, add a flag to indicate whether a port has been "announced" and only try to attach client drivers to ports when the flag is set.	7.8	More Details
CVE-2026-53411	A time-of-check to time-of-use (TOCTOU) race condition in the installation and uninstallation process of certain Zoom Clients for Windows could allow an authenticated local user to escalate privileges.	7.8	More Details
CVE-2026-48373	Acrobat Reader is affected by a Heap-based Buffer Overflow vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-63950	In the Linux kernel, the following vulnerability has been resolved: mm/rmap: initialize nr_pages to 1 at loop start in try_to_unmap_one Initialize nr_pages to 1 at the start of each loop iteration, like folio_referenced_one() does. Without this, nr_pages computed by a previous folio_unmap_pte_batch() call can be reused on a later iteration that does not run folio_unmap_pte_batch() again. mmap a 64K large folio with MAP_ANONYMOUS | MAP_DROPPABLE, then call madvise(MADV_FREE), then make the last page device-exclusive via HMM_DMIRROR_EXCLUSIVE. Trigger node reclaim through sysfs. Now, in try_to_unmap_one(), we will first clear the first 15 out of 16 entries mapping the lazyfree folio. This will set nr_pages to 15. In the next pvmw walk, this nr_pages gets reused on a device-exclusive pte, thus potentially corrupting folio refcount/mapcount. At the moment, I have a userspace program which can make the kernel spit out a trace, but the blow up is in folio_referenced_one(), because there are existing bugs in the interaction between device-private and rmap (which too I am investigating). I did a one liner kernel change to avoid going into folio_referenced_one(), and the kernel blows up at folio_remove_rmap_ptes in try_to_unmap_one which is what I wanted. Note that the bug is there not since file folio batching but lazyfree folio batching, since device-exclusive only works for anonymous folios. Userspace visible effect is simply kernel crashing somewhere due to refcount/mapcount corruption.	7.8	More Details
CVE-2026-63951	In the Linux kernel, the following vulnerability has been resolved: zram: fix use-after-free in zram_writeback_endio A crash was observed in zram_writeback_endio due to a NULL pointer dereference in wake_up. The root cause is a race condition between the bio completion handler (zram_writeback_endio) and the writeback task. In zram_writeback_endio, wake_up() is called on &wb_ctl->done_wait after releasing wb_ctl->done_lock. This creates a race window where the writeback task can see num_inflight become 0, return, and free wb_ctl before zram_writeback_endio calls wake_up(). CPU 0 (zram_writeback_endio) CPU 1 (writeback_store) ========== ========== zram_writeback_slots zram_submit_wb_request zram_submit_wb_request wait_event(wb_ctl->done_wait) spin_lock(&wb_ctl->done_lock); list_add(&req->entry, &wb_ctl->done_reqs); spin_unlock(&wb_ctl->done_lock); wake_up(&wb_ctl->done_wait); zram_complete_done_reqs spin_lock(&wb_ctl->done_lock); list_add(&req->entry, &wb_ctl->done_reqs);	7.8	More Details

	spin_unlock(&wb_ctl->done_lock); while (num_inflight) > 0) spin_lock(&wb_ctl->done_lock); list_del(&req->entry); spin_unlock(&wb_ctl->done_lock); // num_inflight becomes 0 atomic_dec(num_inflight); // Leave zram_writeback_slots // Free wb_ctl release_wb_ctl(wb_ctl); // UAF crash! wake_up(&wb_ctl->done_wait); This patch fixes this race by using RCU. By protecting wb_ctl with rcu_read_lock() in zram_writeback_endio and using kfree_rcu() to free it, we ensure that wb_ctl remains valid during the execution of zram_writeback_endio.		
CVE-2026-63954	In the Linux kernel, the following vulnerability has been resolved: hpfs: fix a crash if hpfs_map_dnode_bitmap fails If hpfs_map_dnode_bitmap fails, the code would call hpfs_brelse4 on uninitialized quad buffer head, causing a crash.	7.8	More Details
CVE-2026-63970	In the Linux kernel, the following vulnerability has been resolved: vsock/virtio: bind uarg before filling zerocopy skb virtio_transport_send_pkt_info() allocates or reuses the zerocopy uarg before entering the send loop, but virtio_transport_alloc_skb() still fills the skb before it inherits that uarg. When fixed-buffer vectored zerocopy hits MAX_SKB_FRAGS, io_sg_from_iter() may partially attach managed frags and return -EMSGSIZE. The rollback path call kfree_skb() to free an skb that carries SKBFL_MANAGED_FRAG_REFS but no uarg, so skb_release_data() falls through to ordinary frag unref. Pass the uarg into virtio_transport_alloc_skb() and bind it immediately before virtio_transport_fill_skb(). This keeps control or no-payload skbs untouched while ensuring success and rollback share one lifetime rule.	7.8	More Details
CVE-2026-63971	In the Linux kernel, the following vulnerability has been resolved: sctp: fix race between sctp_wait_for_connect and peeloff sctp_wait_for_connect() drops and re-acquires the socket lock while waiting for the association to reach ESTABLISHED state. During this window, another thread can peeloff the association to a new socket via getsockopt(SCTP_SOCKOPT_PEELOFF), changing asoc->base.sk. After re-acquiring the old socket lock, sctp_wait_for_connect() returns success without noticing the migration — the caller then accesses the association under the wrong lock in sctp_datamsg_from_user(). Add the same sk != asoc->base.sk check that sctp_wait_for_sndbuf() already has, returning an error if the association was migrated while we slept.	7.8	More Details
CVE-2026-9762	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to remote code execution when jdbc url is under user control.	7.8	More Details
CVE-2026-63977	In the Linux kernel, the following vulnerability has been resolved: dpll: zl3073x: use __dpll_device_change_ntf() and remove change_work The change_work was introduced to send device change notifications from DPLL device callbacks without deadlocking on dpll_lock, since the callbacks are already invoked under that lock. Now that __dpll_device_change_ntf() is exported for callers that already hold dpll_lock, use it directly and remove the change_work infrastructure entirely. This eliminates a race condition where change_work could be re-scheduled after cancel_work_sync() during device teardown, potentially causing the handler to dereference a freed or NULL dpll_dev pointer.	7.8	More Details
CVE-2026-63985	In the Linux kernel, the following vulnerability has been resolved: ethtool: eeprom: add more safeties to EEPROM Netlink fallback The Netlink fallback path for reading module EEPROM (fallback_set_params()) validates that offset < eeprom_len, but does not check that offset + length stays within eeprom_len. The ioctl equivalent (ethtool_get_any_eeprom() in ioctl.c) has always enforced both bounds: if (eeprom.offset + eeprom.len > total_len) return -EINVAL; This could lead to surprises in both drivers and device FW. Add the missing offset + length validation to fallback_set_params(), mirroring the ioctl. Similarly - ethtool core in general, and ethtool_get_any_eeprom() in particular tries to zero-init all buffers passed to the drivers to avoid any extra work of zeroing things out. eeprom_fallback() uses a plain kmalloc(), change it to zalloc.	7.8	More Details
CVE-2026-57860	ForgeCode (tailcallhq/forgecode), an AI pair-programming CLI, automatically loads and executes the MCP servers defined in a repository's .mcp.json file on startup without user confirmation. A malicious repository can supply a crafted .mcp.json whose mcpServers entries specify arbitrary command and args values (for example, command: bash with args: ['-c', 'touch /tmp/pwned']). When a user runs the forge CLI inside a cloned untrusted repository, the specified commands are spawned with the invoking user's privileges, resulting in arbitrary code execution. This provides a reliable initial-access and persistence primitive against developers who evaluate untrusted repositories with ForgeCode.	7.8	More Details
CVE-2026-57076	YAML::Syck versions before 1.47 for Perl allow a heap use-after-free via an anchor name reused as an anchors-table key in syck_hdlr_add_anchor. In the bundled libsyck an anchor name allocated by syck_strndup is stored both as node->anchor, freed when the node is freed, and as the key in the parser's anchors table. Freeing the node frees the shared key, and a later anchor redefinition makes st_delete compare against the freed key, so st_strcmp reads freed heap memory. Anchors are a standard YAML feature and need no special flags, so this is reached on the default Load path. Any caller that runs Load or LoadFile on an untrusted document that redefines an anchor reaches the read of freed memory.	7.8	More Details
CVE-2026-53386	In the Linux kernel, the following vulnerability has been resolved: iio: adc: ti-ads1298: add bounds check to pga_settings index ads1298_pga_settings has 7 elements but ADS1298_MASK_CH_PGA can yield values 0-7. If it yields a value >= 7, this causes an out-of-bounds array access. Add a bounds check and return -EINVAL if the index is out of range. Note that the remaining value b111 is reserved so should not be seen in a correctly functioning system.	7.8	More Details

CVE-2026-63911	In the Linux kernel, the following vulnerability has been resolved: xfrm: iptfs: reset runtime state when cloning SAs iptfs_clone_state() clones the IPTFS mode data with kmemdup(). This copies runtime objects which must not be shared with the original SA, including the embedded sk_buff_head, hrtimers, spinlock, and in-flight reassembly/reorder state. If xfrm_state_migrate() fails after clone_state() but before the later init_state() call has reinitialized those fields, the cloned state can be destroyed by xfrm_state_gc_task() with list and timer state copied from the original SA. With queued packets this lets the clone splice and free skbs owned by the original IPTFS queue, leading to use-after-free and double-free reports in iptfs_destroy_state() and skb release paths. Reinitialize the clone's runtime state before publishing it through x->mode_data. Because clone_state() now publishes a destroyable mode_data object before init_state(), take the mode callback module reference there. Avoid taking it again from __iptfs_init_state() for the same object.	7.8	More Details
CVE-2026-63814	In the Linux kernel, the following vulnerability has been resolved: f2fs: validate ACL entry sizes in f2fs_acl_from_disk() f2fs_acl_count() only validates the aggregate ACL xattr length. A malformed ACL can still place ACL_USER or ACL_GROUP in a slot that only contains struct f2fs_acl_entry_short bytes, and f2fs_acl_from_disk() then reads entry->e_id before verifying that a full entry fits. Require a short entry before reading e_tag and e_perm, and require a full entry before reading e_id for ACL_USER and ACL_GROUP. Return -EFSCORRUPTED from these new truncated-entry checks, while keeping the pre-existing -EINVAL paths unchanged. Validation reproduced this kernel report: KASAN slab-out-of-bounds in __f2fs_get_acl+0x6fb/0x7e0 RIP: 0033:0x7f4b835ea7aa The buggy address belongs to the object at ffff888114589960 which belongs to the cache kmalloc-8 of size 8 The buggy address is located 0 bytes to the right of allocated 8-byte region [ffff888114589960, ffff888114589968) Read of size 4 Call trace: dump_stack_lvl+0x66/0xa0 (?:?) print_report+0xce/0x630 (?:?) __f2fs_get_acl+0x6fb/0x7e0 (fs/f2fs/acl.c:169) srso_alias_return_thunk+0x5/0xfbef5 (?:?) __virt_addr_valid+0x224/0x430 (?:?) kasan_report+0xe0/0x110 (?:?) __f2fs_get_acl+0x5/0x7e0 (fs/f2fs/acl.c:169) __get_acl+0x281/0x380 (?:?) vfs_get_acl+0x10b/0x190 (?:?) do_get_acl+0x2a/0x410 (?:?) do_get_acl+0x9/0x410 (?:?) do_getxattr+0xe8/0x260 (?:?) filename_getxattr+0xd1/0x140 (?:?) do_getname+0x2d/0x2d0 (?:?) path_getxattrat+0x16c/0x200 (?:?) lock_release+0xc8/0x290 (?:?) cgroup_update_frozen+0x9d/0x320 (?:?) lockdep_hardirqs_on_prepare+0xea/0x1a0 (?:?) trace_hardirqs_on+0x1a/0x170 (?:?) _raw_spin_unlock_irq+0x28/0x50 (?:?) do_syscall_64+0x115/0x6a0 (arch/x86/entry/syscall_64.c:87) entry_SYSCALL_64_after_hwframe+0x77/0x7f (?:?)	7.8	More Details
CVE-2026-63987	In the Linux kernel, the following vulnerability has been resolved: ethtool: coalesce: cap profile updates at NET_DIM_PARAMS_NUM_PROFILES ethnl_update_profile() walks the ETHTOOL_A_PROFILE_IRQ_MODERATION nest list with an index 'i' and writes new_profile[i++] without bounding i. The destination is kmemdup()'d at NET_DIM_PARAMS_NUM_PROFILES entries (5), but the Netlink nest count is entirely user-controlled. Netlink policies do not have support for constraining the number of nested entries (or number of multi-attr entries).	7.8	More Details
CVE-2026-63858	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: add hook transactions for device deletions Restore the flag that indicates that the hook is going away, ie. NFT_HOOK_REMOVE, but add a new transaction object to track deletion of hooks without altering the basechain/flowtable hook_list during the preparation phase. The existing approach that moves the hook from the basechain/flowtable hook_list to transaction hook_list breaks netlink dump path readers of this RCU-protected list. It should be possible use an array for nft_trans_hook to store the deleted hooks to compact the representation but I am not expecting many hook object, specially now that wildcard support for devices is in place. Note that the nft_trans_chain_hooks() list contains a list of struct nft_trans_hook objects for DELCHAIN and DELFLOWTABLE commands, while this list stores struct nft_hook objects for NEWCHAIN and NEWFLOWTABLE. Note that new commands can be updated to use nft_trans_hook for consistency. This patch also adapts the event notification path to deal with the list of hook transactions.	7.8	More Details
CVE-2026-63819	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to do sanity check on f2fs_get_node_folio_ra() kernel BUG at fs/f2fs/file.c:845! Oops: invalid opcode: 0000 [#1] SMP KASAN NOPTI CPU: 0 UID: 0 PID: 5336 Comm: syz.0.0 Not tainted syzkaller #0 PREEMPT(full) Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 RIP: 0010:f2fs_do_truncate_blocks+0x1115/0x1140 fs/f2fs/file.c:845 Code: fc fc 90 0f 0b e8 8b 9d 9a fd 90 0f 0b e8 83 9d 9a fd 48 89 df 48 c7 c6 60 d1 1a 8c e8 54 f1 fc fc 90 0f 0b e8 6c 9d 9a fd 90 <0f> 0b e8 64 9d 9a fd 90 0f 0b 90 e9 93 fd ff ff e8 56 9d 9a fd 90 RSP: 0018:ffffc9000e4474c0 EFLAGS: 00010283 RAX: ffffffff842b1d34 RBX: 0000000000000003 RCX: 0000000000100000 RDX: ffffc9000f03a000 RSI: 0000000000035503 RDI: 0000000000035504 RBP: ffffc9000e447608 R08: ffff8880123b0000 R09: 0000000000000002 R10: 00000000fffffffe R11: 0000000000000002 R12: 0000000000000001 R13: 0000000000000000 R14: 1ffff92001c88ea0 R15: 00000000ffff039c FS: 00007f7e02ee36c0(0000) GS:ffff88808c887000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007ff0305c4000 CR3: 0000000012d4c000 CR4: 0000000000352ef0 Call Trace: <TASK> f2fs_truncate_blocks+0x10a/0x300 fs/f2fs/file.c:882 f2fs_truncate+0x471/0x7c0 fs/f2fs/file.c:940 f2fs_evict_inode+0xa3f/0x1ac0 fs/f2fs/inode.c:907 evict+0x61e/0xb10 fs/inode.c:841 f2fs_fill_super+0x5f43/0x78f0 fs/f2fs/super.c:5224 get_tree_bdev_flags+0x431/0x4f0 fs/super.c:1694 vfs_get_tree+0x92/0x2a0 fs/super.c:1754 fc_mount fs/namespace.c:1193 [inline] do_new_mount_fc fs/namespace.c:3758 [inline] do_new_mount+0x341/0xd30 fs/namespace.c:3834 do_mount fs/namespace.c:4167 [inline] __do_sys_mount fs/namespace.c:4383 [inline] __se_sys_mount+0x31d/0x420 fs/namespace.c:4360 do_syscall_x64 arch/x86/entry/syscall_64.c:63 [inline] do_syscall_64+0x15f/0xf80 arch/x86/entry/syscall_64.c:94 entry_SYSCALL_64_after_hwframe+0x77/0x7f count = ADDRS_PER_PAGE(dn.node_folio, inode); count -= dn.ofs_in_node; f2fs_bug_on(sbi, count < 0); The fuzz test	7.8	More Details

	will trigger above bug_on in f2fs. The root cause should be: in the corrupted inode, there is a direct node which has the same ino and nid in its footer, so in f2fs_do_truncate_blocks(), after f2fs_get_dnode_of_data() finds such dnode: 1) ADDRS_PER_PAGE(dn.node_folio, inode) will return 923 2) once dn.ofs_in_node points to addr[923, 1017] Then it will trigger the system panic. Let's introduce NODE_TYPE_NON_IXNODE to indicate current node should not be an inode or xattr node, and then use it in below path to detect inconsistent node chain in inode mapping table: - f2fs_do_truncate_blocks - f2fs_get_dnode_of_data - f2fs_get_node_folio_ra - __get_node_folio - f2fs_sanity_check_node_footer - case NODE_TYPE_NON_IXNODE -> check whether it is inode|xnode		
CVE-2026-63856	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v2.0 enc/dec rings VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit e2b5499fca55f1a32960a311bbb62e35891eaf73)	7.8	More Details
CVE-2026-63855	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v2.5 enc/dec rings VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit efc9dd5590894109bce9a0bfe1fa5592dd6b20b1)	7.8	More Details
CVE-2026-63823	In the Linux kernel, the following vulnerability has been resolved: keys: Pin request_key_auth payload in instantiate paths A: request_key() B: KEYCTL_INSTANTIATE_IOV ========== ========== create auth key store rka in auth key wait for helper get auth key load rka from auth key copy user payload sleep on #PF helper completed detach and free rka destroy auth key wake up use rka->target_key **USE-AFTER-FREE** Give request_key_auth payloads a refcount. Take a payload reference while authkey->sem stabilizes the payload and revocation state. Hold that reference across the instantiate and reject paths. Drop the auth key owning reference from revoke and destroy. [jarkko: Replaced the first two paragraphs of text with an actual concurrency scenario.]	7.8	More Details
CVE-2026-63824	In the Linux kernel, the following vulnerability has been resolved: KEYS: fix overflow in keyctl_pkey_params_get_2() The length for the internal output buffer is calculated incorrectly, which can result overflow when a too small buffer is provided. Fix the bug by allocating internal output with the size of the maximum length of the cryptographic primitive instead of caller provided size.	7.8	More Details
CVE-2026-63827	In the Linux kernel, the following vulnerability has been resolved: apparmor: fix use-after-free in rawdata dedup loop aa_replace_profiles() walks ns->rawdata_list to dedup the incoming policy blob against entries already attached to existing profiles. Per the kernel-doc on struct aa_loaddata, list membership does not hold a reference: profiles hold pcount, and when the last pcount drops, do_ploaddata_rmfs() is queued on a workqueue that takes ns->lock and removes the entry. Between dropping the last pcount and the workqueue running, an entry remains on the list with pcount == 0. aa_get_profile_loaddata() is an unconditional kref_get() on pcount, so when the dedup loop hits such an entry, refcount hardening reports refcount_t: addition on 0; use-after-free. inside aa_replace_profiles(), and the poisoned counter then trips "saturated" and "underflow" warnings on the subsequent uses of the same loaddata. Before commit a0b7091c4de4 ("apparmor: fix race on rawdata dereference") the dedup path used a get_unless_zero-style helper on a single counter, so the existing "if (tmp)" guard was meaningful. The split-refcount refactor introduced aa_get_profile_loaddata(), which has plain kref_get() semantics, and the guard quietly became a no-op. Introduce aa_get_profile_loaddata_not0(), matching the existing _not0 convention used by aa_get_profile_not0(), and use it for the rawdata_list dedup lookup so dying entries are skipped. Reproduced on x86_64 with v7.1-rc5 in QEMU+KVM running Ubuntu 24.04 + stress-ng 0.17.06: stress-ng --apparmor 1 --klog-check --timeout 60s Without this patch the three refcount_t warnings fire within a few seconds. With it the same 60 s run is clean. Coverage is a smoke-test only; a longer soak with CONFIG_KASAN, CONFIG_KCSAN and CONFIG_PROVE_LOCKING would be welcome from anyone with the cycles.	7.8	More Details
CVE-2026-63854	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v3.0 enc/dec rings VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 663bed3c7b8b9a7624b0d95d300ddae034ad0614)	7.8	More Details
CVE-2026-63853	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v4.0 enc ring VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit fd852c048b46f9825e904a4f3f4538fe9d8827d9)	7.8	More Details
CVE-2026-63840	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v5.3.0 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 86ac011ae234c03fb872f4945913391ea1d8862e)	7.8	More Details
CVE-2026-63841	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v5.0.1 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 742a98e2e81702df8fe1b1eccee5223220a03dc2)	7.8	More Details
CVE-2026-63842	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v5.0.0 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 0f43893d3cd478fa57836697525b338817c9c23d)	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v4.0.3 enc ring VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS	7.8	More

63852	submissions with user fences. (cherry picked from commit ff1a5a125c5a70c328806b9bc01d7d942cf3f9aa)		Details
CVE-2026-63851	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v4.0.5 enc ring VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 084d94ac93707bdda07efb5cee786f632de4219b)	7.8	More Details
CVE-2026-63850	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v5.0.0 enc ring VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 49b1fbbb5a071197ee71e2d70959b1cb29bdc317)	7.8	More Details
CVE-2026-63849	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn: set no_user_fence for VCN v5.0.1 enc ring VCN encoder and decoder rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit e16be95a2c3ee712b142cb27d2dca0b461181359)	7.8	More Details
CVE-2026-63848	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v2.0 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 96179da0c6b059eb31706a0abe8dd6381c533143)	7.8	More Details
CVE-2026-63847	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v2.5 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 3216a7f4e2642bda5fd14f57586e835ae9202587)	7.8	More Details
CVE-2026-63843	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v4.0.5 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit f05d0a4f21fc720116d6e238f23308b199891058)	7.8	More Details
CVE-2026-63846	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v3.0 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 4d7d774f100efb5089c86a1fb8c5bf47c63fc9ef)	7.8	More Details
CVE-2026-63816	In the Linux kernel, the following vulnerability has been resolved: f2fs: atomic: fix UAF issue on f2fs_inode_info.atomic_inode - ioctl(F2FS_IOC_GARBAGE_COLLECT_RANGE) - shrink - f2fs_gc - gc_data_segment - ra_data_block(cow_inode) - mapping = F2FS_I(inode)->atomic_inode->i_mapping : f2fs_is_cow_file(cow_inode) is true - f2fs_evict_inode(atomic_inode) - clear_inode_flag(fi->cow_inode, FI_COW_FILE) - F2FS_I(fi->cow_inode)->atomic_inode = NULL ... - truncate_inode_pages_final(atomic_inode) - f2fs_grab_cache_folio(mapping) : create folio in atomic_inode->mapping - clear_inode(atomic_inode) - BUG_ON(atomic_inode->i_data.nrpages) We need to add a reference on fi->atomic_inode before using its mapping field during garbage collection, otherwise, it will cause UAF issue.	7.8	More Details
CVE-2026-63813	In the Linux kernel, the following vulnerability has been resolved: Revert "f2fs: remove non-uptodate folio from the page cache in move_data_block" This reverts commit 9609dd704725a40cd63d915f2ab6c44248a44598. The kernel panics are keeping to be reported especially when the f2fs partition get almost full. By investigation, we find that the reason is one f2fs page got freed to buddy without being deleted from LRU and the root cause is the race happened in [2] which is enrolled by this commit. There are 3 race processes in this scenario, please find below for their main activities. The changed code in move_data_block() lets the GC path evict the tail-end folio from the page cache through folio_end_dropbehind(). Once folio_unmap_invalidate() removes the folio from mapping->i_pages, the page-cache references for all pages in the folio are dropped. The folio is then kept alive only by temporary external references, which allows a later split to operate on a folio whose subpages are no longer protected by page-cache references. After the page-cache references are gone, split_folio_to_order() can split the big folio into individual pages and put the resulting subpages back on the LRU. For tail pages beyond EOF, split removes them from the page cache and drops their page-cache references. A tail page can then remain on the LRU with PG_lru set while holding only the split caller's temporary reference. When free_folio_and_swap_cache() drops that final reference, the page enters the final folio_put() release path. In parallel, folio_isolate_lru() can observe the same tail page with a non-zero refcount and PG_lru set. It clears PG_lru before taking its own reference. If this races with the final folio_put() from the split path, __folio_put() sees PG_lru already cleared and skips lruvec_del_folio(). The page is then freed back to the allocator while its lru links are still present in the LRU list. A later LRU operation on a neighboring page detects the stale link and reports list corruption. [1] [22.486082] list_del corruption. next->prev should be ffffffffec10e0ac8, but was dead000000000122. (next=ffffffffec10e0a88) [22.486130] ------------[cut here]------------ [22.486134] kernel BUG at lib/list_debug.c:67! [22.486141] Internal error: Oops - BUG: 00000000f2000800 [#1] SMP [22.488502] Tainted: [W]=WARN, [O]=OOT_MODULE [22.488506] Hardware name: Spreadtrum UMS9230 1H10 SoC (DT) [22.488511] pstate: 604000c5 (nZCv daIF +PAN -UAO -TCO -DIT -SSBS BTYPE=--) [22.488517] pc : __list_del_entry_valid_or_report+0x14c/0x154 [22.488531] lr : __list_del_entry_valid_or_report+0x14c/0x154 [22.488539] sp : ffffffc08006b830 [22.488542] x29: ffffffc08006b868 x28: 0000000000003020 x27: 0000000000000000 [22.488553] x26: 0000000000000000 x25: 0000000000000004 x24: ffffffffec10e0ac0 [22.488564] x23: 00000000000000e8 x22: 0000000000000024 x21: dead000000000122 [22.488574] x20: ffffffffec10e0a88 x19: ffffffffec10e0ac8 x18: ffffffc080061060 [22.488585] x17: 20747562202c3863 x16: 6130653031636566 x15: 0000000000000058 [22.488595] x14: 0000000000000004 x13: ffffff80f91e0000 x12: 0000000000000003 [22.488605] x11: 0000000000000003 x10: 0000000000000001 x9 : ffe85721f0e25f00 [22.488615] x8 : ffe85721f0e25f00 x7 : 0000000000000000 x6 : 6c65645f7473696c [	7.8	More Details

	22.488625] x5 : ffffffed39b23026 x4 : 0000000000000000 x3 : 0000000000000010 [22.488636] x2 : 0000000000000000 x1 : 0000000000000000 x0 : 000000000000006d [22.488647] Call trace: [22.488651] __list_del_entry_valid_or_report+0x14c/0x154 (P) [22.488661] __folio_put+0x2bc/0x434 [22.488670] folio_put+0x28/0x58 [22.488678] do_garbage_collect+0x1a34/0x2584 [22.488689] f2fs_gc+0x230/0x9b4 [22.488697] f2fs_fallocate+0xb90/0xdf4 [22.488706] vfs_fallocate+0x1b4/0x2bc [22.488716] __arm64_sys_fallocate+0x44/0x78 [22.488725] invoke_syscall+0x58/0xe4 [22.488732] do_el0_svc+0x48/0xdc [22.488739] el0 ---truncated---		
CVE-2026-53373	In the Linux kernel, the following vulnerability has been resolved: mm/vma: do not try to unmap a VMA if mmap_prepare() invoked from mmap() The mmap_prepare hook functionality includes the ability to invoke mmap_prepare() from the mmap() hook of existing 'stacked' drivers, that is ones which are capable of calling the mmap hooks of other drivers/file systems (e.g. overlayfs, shm). As part of the mmap_prepare action functionality, we deal with errors by unmapping the VMA should one arise. This works in the usual mmap_prepare case, as we invoke this action at the last moment, when the VMA is established in the maple tree. However, the mmap() hook passes a not-fully-established VMA pointer to the caller (which is the motivation behind the mmap_prepare() work), which is detached. So attempting to unmap a VMA in this state will be problematic, with the most obvious symptom being a warning in vma_mark_detached(), because the VMA is already detached. It's also unncessary - the mmap() handler will clean up the VMA on error. So to fix this issue, this patch propagates whether or not an mmap action is being completed via the compatibility layer or directly. If the former, then we do not attempt VMA cleanup, if the latter, then we do. This patch also updates the userland VMA tests to reflect the change.	7.8	More Details
CVE-2026-63812	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix incorrect FI_NO_EXTENT handling in __destroy_extent_node() When __destroy_extent_node() sets the inode flag FI_NO_EXTENT, it does not reset the length of the largest extent to 0 and update the inode folio. Since modifications to the extent tree are disallowed afterward, the cached largest extent may become stale. This can trigger the following error in xfstests generic/388: F2FS-fs (dm-0): sanity_check_extent_cache: inode (ino=1761) extent info [220057, 57, 6] is incorrect, run fsck to fix In the f2fs_drop_inode path, __destroy_extent_node() does not need to guarantee that et->node_cnt is 0, because concurrency with writeback is expected in this path, and writeback may update the extent cache. This patch reverts commit ed78aeebef05 ("f2fs: fix node_cnt race between extent node destroy and writeback"), and remove the unnecessary zero check of et->node_cnt.	7.8	More Details
CVE-2026-63910	In the Linux kernel, the following vulnerability has been resolved: dma-buf: fix UAF in dma_buf_fd() tracepoint Once FD_ADD() returns, the fd is live in the file descriptor table and a thread sharing that table can close() it before DMA_BUF_TRACE() runs. The close drops the last reference, __fput() frees the dma_buf, and the tracepoint then dereferences dmabuf to take dmabuf->name_lock -- slab-use-after-free. Split FD_ADD() back into get_unused_fd_flags() + fd_install() and emit the tracepoint between them. While the fdtable slot is reserved with a NULL file pointer, a racing close() returns -EBADF without entering __fput(), so the dma_buf stays alive across the trace. Same approach as commit 2d76319c4cbb ("dma-buf: fix UAF in dma_buf_put() tracepoint"). This undoes the FD_ADD() conversion done in commit 34dfce523c90 ("dma: convert dma_buf_fd() to FD_ADD()"); FD_ADD() has no place to hook the tracepoint safely.	7.8	More Details
CVE-2026-53380	In the Linux kernel, the following vulnerability has been resolved: media: rzv2h-ivc: Fix concurrent buffer list access The list of buffers (`rzv2h_ivc::buffers.queue`) is protected by a spinlock (`rzv2h_ivc::buffers.lock`). However, in `rzv2h_ivc_transfer_buffer()`, which runs in a separate workqueue, the `list_del()` call is executed without holding the spinlock, which makes it possible for the list to be concurrently modified Fix that by removing a buffer from the list in the lock protected section. [assign ivc->buffers.curr in critical section as reported by Barnabas]	7.8	More Details
CVE-2026-53381	In the Linux kernel, the following vulnerability has been resolved: virtiofs: fix UAF on submount umount iput() called from fuse_release_end() can Oops if the super block has already been destroyed. Normally this is prevented by waiting for num_waiting to go down to zero before commencing with super block shutdown. This only works, however, for the last submount instance, as the wait counter is per connection, not per superblock. Revert to using synchronous release requests for the auto_submounts case, which is virtiofs only at this time.	7.8	More Details
CVE-2026-63845	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/jpeg: set no_user_fence for JPEG v4.0 ring JPEG rings do not support 64-bit user fence writes, reject CS submissions with user fences. (cherry picked from commit 8d0cac9478a3f046279c657d6a2545de49ae675a)	7.8	More Details
CVE-2026-53409	Improper Privilege Management in Zoom Rooms for Windows before version 7.1.0 may allow an authenticated user to conduct an escalation of privilege via local access.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: serialize DMABUF cancel against request completion ffs_epfile_dmabuf_io_complete() calls usb_ep_free_request() on the completed request but leaves priv->req, the back-pointer that ffs_dmabuf_transfer() set on submission, pointing at the freed memory. A later FUNCTIONFS_DMABUF_DETACH ioctl or ffs_epfile_release() on the close path still sees priv->req non-NULL under ffs->eps_lock: if (priv->ep && priv->req) usb_ep_dequeue(priv->ep, priv->req); so usb_ep_dequeue() is called on a freed usb_request. On dummy_hcd the dequeue path only walks a live queue and pointer-compares, so the freed pointer reads without faulting and KASAN requires an explicit check at		

CVE-2026-63894	the FunctionFS call site to surface the use-after-free. On SG-capable in-tree UDCs the dequeue path dereferences the supplied request immediately: * chipidea's ep_dequeue() does container_of(req, struct ci_hw_req, req) and reads hwreq->req.status before acquiring its own lock. * cdnsp's cdnsp_gadget_ep_dequeue() reads request->status first. The narrower option of clearing priv->req via cmpxchg() in the completion does not close the race: the completion runs without eps_lock, so a cancel path holding eps_lock can still observe priv->req non-NULL, race a concurrent completion that clears and frees, and pass the freed pointer to usb_ep_dequeue(). A slightly longer fix that moves the free into the cleanup work is needed. Same class of lifetime race as the recent usbip-vudc timer fix [1]. Take eps_lock in the sole place that mutates priv->req from the callback direction by moving usb_ep_free_request() out of the completion into ffs_dmabuf_cleanup(), the existing work handler scheduled by ffs_dmabuf_signal_done() on ffs->io_completion_wq. Clear priv->req there under eps_lock before freeing, and only clear if priv->req still names our request (a subsequent ffs_dmabuf_transfer() on the same attachment may have queued a new one). This keeps the existing dummy_hcd sync-dequeue invariant: the completion callback is still invoked by the UDC without eps_lock held (dummy_hcd drops its own lock before calling the callback), and the callback now takes no f_fs lock at all. Serialization against the cancel path happens in cleanup, which runs from the workqueue with no f_fs lock held on entry. The priv ref count protects the containing ffs_dmabuf_priv: ffs_dmabuf_transfer() takes a ref via ffs_dmabuf_get(), cleanup drops it via ffs_dmabuf_put(), so priv stays live for the cleanup even after the cancel path's list_del + ffs_dmabuf_put. The ffs_dmabuf_transfer() error path no longer frees usb_req inline: fence->req and fence->ep are set before usb_ep_queue(), so ffs_dmabuf_cleanup() (scheduled by the error-path ffs_dmabuf_signal_done()) owns the free regardless of whether the queue succeeded. Reproduced under KASAN on both detach and close paths against dummy_hcd with an observability hook (kasan_check_byte(priv->req) immediately before usb_ep_dequeue) at the two FunctionFS cancel sites to surface the stale-pointer access; the hook is not part of this patch. The KASAN allocator / free stacks in the captured splats identify the same request: alloc in dummy_alloc_request, free in dummy_timer, fault reached from ffs_epfile_release (close) and from the FUNCTIONFS_DMABUF_DETACH ioctl (detach). With the patch applied, both paths are silent under the same hook. The bug is reached from the FunctionFS device node, which in real deployments is owned by the privileged gadget daemon (adbd, UMS, composite gadget services, etc.); it is not reachable from unprivileged userspace or from a USB host on the cable. FunctionFS mounts default to GLOBAL_ROOT_UID, but the filesystem supports uid=, gid=, and fmode= delegation to a non-root gadget daemon, so on real deployments the attacker may be a less-privileged service rather than root.	7.8	More Details
CVE-2026-53388	In the Linux kernel, the following vulnerability has been resolved: fuse: re-lock request before replacing page cache folio fuse_try_move_folio() unlocks the request on entry but does not re-lock it on the success path. This means fuse_chan_abort() can end the request and free the fuse_io_args (eg fuse_readpages_end()) while the subsequent copy chain logic after fuse_try_move_folio() accesses the fuse_io_args, leading to use-after-free issues. Fix this by calling lock_request() before replace_page_cache_folio(). This ensures the request is locked on the success path which will prevent the fuse_io_args from being freed while the later copying logic runs, and also ensures that the ap->folios[i]->mapping is never null since ap->folios[i] will always point to the newfolio after replace_page_cache_folio().	7.8	More Details
CVE-2026-53389	In the Linux kernel, the following vulnerability has been resolved: net/tcp-ao: fix use-after-free of key in del_async path In tcp_ao_delete_key(), the del_async path skips the current_key and rnext_key validity checks present in the synchronous path, assuming these pointers are always NULL on LISTEN sockets. However, if a key was added with set_current=1/set_rnext=1 while the socket was in CLOSE state, current_key and rnext_key will be non-NULL after listen() transitions the socket to LISTEN. When such a key is deleted with del_async=1, hlist_del_rcu() and call_rcu() free the key without clearing the dangling pointers. After the RCU grace period, getsockopt(TCP_AO_INFO) dereferences current_key->sndid and rnext_key->rcvid from freed slab memory. Clear current_key and rnext_key in the del_async path when they reference the key being deleted.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: drm/i915: Fix potential UAF in TTM object purge TLDR: The bo->ttm object might be changed by calling ttm_bo_validate(), move casting it to an i915_tt object later to actually get the right pointer. A user reported hitting the following bug under heavy use on DG2: [26620.095550] Oops: general protection fault, probably for non-canonical address 0xa56b6b6b6b6b6b8b: 0000 1 SMP NOPTI [26620.095556] CPU: 2 UID: 0 PID: 631 Comm: Xorg Not tainted 6.18.8 #1 PREEMPT(lazy) [26620.095558] Hardware name: ASRock B850M Steel Legend WiFi/B850M Steel Legend WiFi, BIOS 3.50 09/18/2025 [26620.095559] RIP: 0010:i915_ttm_purge+0x84/0x100 [i915] [26620.095604] Code: 00 00 00 48 8d 54 24 10 48 89 e6 48 89 fb e8 83 aa ae ff 85 c0 75 6f 48 83 bb a8 01 00 00 00 74 2c 48 8b 45 78 48 85 c0 74 23 <48> 8b 78 20 48 c7 c2 ff ff ff ff 31 f6 e8 7a 73 e3 e0 48 8b 7d 78 [26620.095605] RSP: 0018:ffffc90005fd7430 EFLAGS: 00010282 [26620.095607] RAX: a56b6b6b6b6b6b6b RBX: ffff8881f46c3dc0 RCX: 0000000000000000 [26620.095608] RDX: 0000000000000000 RSI: 0000000000000246 RDI: 00000000ffffffff [26620.095609] RBP: ffff888289610f00 R08: 0000000000000001 R09: ffff88823b022000 [26620.095609] R10: ffff888103029b28 R11: ffff8881fc7f3800 R12: ffff88810b6150d0 [26620.095609] R13: ffff888289610f00 R14: 0000000000000000 R15: ffff8881f46c3dc0 [26620.095610] FS: 00007f1004d86900(0000) GS:ffff88901c858000(0000) knlGS:0000000000000000 [26620.095611] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [26620.095611] CR2: 00007f0fdf489000 CR3: 000000035b0c1000 CR4: 0000000000750ef0 [26620.095612] PKRU: 55555554 [26620.095612] Call Trace: [26620.095615] <TASK> [26620.095615] i915_ttm_move+0x2b9/0x420 [i915] [26620.095642] ? ttm_tt_init+0x65/0x80 [ttm] [26620.095644] ? i915_ttm_tt_create+0xc6/0x150 [i915] [26620.095667] ttm_bo_handle_move_mem+0xb6/0x160 [ttm]		

CVE-2026-63884	[26620.095669] ttm_bo_evict+0x100/0x150 [ttm] [26620.095671] ? preempt_count_add+0x64/0xa0 [26620.095673] ? _raw_spin_lock+0xe/0x30 [26620.095675] ? _raw_spin_unlock+0xd/0x30 [26620.095675] ? i915_gem_object_evictable+0xb7/0xd0 [i915] [26620.095704] ttm_bo_evict_cb+0x6e/0xd0 [ttm] [26620.095705] ttm_lru_walk_for_evict+0xa6/0x200 [ttm] [26620.095708] ttm_bo_alloc_resource+0x185/0x4f0 [ttm] [26620.095709] ? init_object+0x62/0xd0 [26620.095712] ttm_bo_validate+0x7a/0x180 [ttm] [26620.095713] ? _raw_spin_unlock_irqrestore+0x16/0x30 [26620.095714] __i915_ttm_get_pages+0xb0/0x170 [i915] [26620.095737] i915_ttm_get_pages+0x9f/0x150 [i915] [26620.095759] ? i915_gem_do_execbuffer+0xedc/0x2b40 [i915] [26620.095786] ? alloc_debug_processing+0xd0/0x100 [26620.095787] ? _raw_spin_unlock_irqrestore+0x16/0x30 [26620.095788] ? i915_vma_instance+0xa0/0x4e0 [i915] [26620.095822] __i915_gem_object_get_pages+0x2f/0x40 [i915] [26620.095848] i915_vma_pin_ww+0x706/0x980 [i915] [26620.095875] ? i915_gem_do_execbuffer+0xedc/0x2b40 [i915] [26620.095904] eb_validate_vmas+0x170/0xa00 [i915] [26620.095930] i915_gem_do_execbuffer+0x1201/0x2b40 [i915] [26620.095953] ? alloc_debug_processing+0xd0/0x100 [26620.095954] ? _raw_spin_unlock_irqrestore+0x16/0x30 [26620.095955] ? i915_gem_execbuffer2_ioctl+0xc9/0x240 [i915] [26620.095977] ? __wake_up_sync_key+0x32/0x50 [26620.095979] ? i915_gem_execbuffer2_ioctl+0xc9/0x240 [i915] [26620.096001] ? __slab_alloc.isra.0+0x67/0xc0 [26620.096003] i915_gem_execbuffer2_ioctl+0x11a/0x240 [i915] Results from decode_stacktrace.sh pointed to dereference of a file pointer field of a i915 TTM page vector container associated with an object being purged on eviction. That path is taken when the object is marked as no longer needed. Code analysis revealed a possibility of the i915 TTM page vector container being replaced with a new instance inside a function that purges content of the object, should it be still busy. That function is called, indirectly via a more general function that changes the object's placement and caching policy, ---truncated---	7.8	More Details
CVE-2026-53400	In the Linux kernel, the following vulnerability has been resolved: i2c: core: fix adapter registration race Adapters can be looked up based on their id using i2c_get_adapter() which takes a reference to the embedded struct device. Make sure that the adapter (including its struct device) has been initialised before adding it to the IDR to avoid accessing uninitialised data which could, for example, lead to NULL-pointer dereferences or use-after-free. Note that the i2c-dev chardev, which is registered from a bus notifier, currently uses i2c_get_adapter() so the adapter needs to be added to the IDR before registration.	7.8	More Details
CVE-2026-63793	In the Linux kernel, the following vulnerability has been resolved: ntfs: serialize volume label accesses Protect vol->volume_label with a mutex and snaphost the label before copy_to_user. This prevent a use-after-free when FS_IOC_SETFSLABEL replaces the vol->volume_label and FS_IOC_GETTSLABEL reads it concurrently.	7.8	More Details
CVE-2026-63881	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: fix a vulnerability of integer overflow in kfd debugger get_queue_ids() computes array_size = num_queues * sizeof(uint32_t), which could overflow on 32-bit size_t build. using array_size() instead, it saturates to SIZE_MAX on overflow. (cherry picked from commit 2d57a0475f085c08b49312dfd8edcb461845f285)	7.8	More Details
CVE-2026-63879	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix amdgpu_hmm_range_get_pages The notifier sequence must only be read once or otherwise we could work with invalid pages. While at it also fix the coding style, e.g. drop the pre-initialized return value and use the common define for 2G range. (cherry picked from commit c08972f555945cda57b0adb72272a37910153390)	7.8	More Details
CVE-2026-63799	In the Linux kernel, the following vulnerability has been resolved: sched/mmcid: Fix OOB clear_bit when CID is MM_CID_UNSET in fixup path In mm_cid_fixup_cpus_to_tasks(), when rq->curr has the target mm and mm_cid.active is set, the CID is checked with cid_in_transit() before setting the transition bit. In per-CPU mode a newly forked or exec'd task can be running with mm_cid.cid == MM_CID_UNSET because CIDs are assigned lazily on schedule-in. With cid_in_transit() the guard passes for MM_CID_UNSET (no transit bit), converts it to MM_CID_UNSET | MM_CID_TRANSIT and stores it back; later mm_cid_schedout() feeds this to clear_bit() with MM_CID_UNSET as the bit number, triggering an out-of-bounds write. Symptoms: this is genuine memory corruption, but a bounded out-of-bounds write, not an arbitrary one. MM_CID_UNSET is the fixed sentinel BIT(31), so once the bad value reaches mm_cid_schedout() the cid_from_transit_cid() strip leaves MM_CID_UNSET, which fails the "cid < max_cids" convergence test and falls into mm_drop_cid() -> clear_bit(MM_CID_UNSET, mm_cidmask(mm)). The cid bitmap is embedded in the mm_struct slab object (after cpu_bitmap and mm_cpus_allowed) and is only num_possible_cpus() bits wide, so clearing bit 31 is a deterministic OOB bit-clear at a fixed offset of 2^31 / 8 == 256 MiB past the bitmap base. The address is not attacker-influenced (fixed sentinel -> fixed offset) and the op only clears a single bit; what sits 256 MiB further along the direct map is whatever kernel object happens to live there, so this corrupts one bit of unpredictable kernel memory -- it is not an arbitrary-address or arbitrary-value write. It triggers only in per-CPU CID mode, when a CPU is running an active task of the target mm whose cid is still MM_CID_UNSET -- the fork()/execve() window before that task's next schedule-in assigns it a real CID -- and a per-CPU -> per-task fixup walks over it (the mode fallback driven by a thread exit, sched_mm_cid_exit(), or by the deferred max_cids recompute in mm_cid_work_fn()). In practice syzkaller surfaced it as a KASAN use-after-free reported in __schedule -> mm_cid_switch_to, where the offending clear_bit() is inlined via mm_cid_schedout() -> mm_drop_cid(). Guard the transition-bit assignment against MM_CID_UNSET, in addition to the existing cid_in_transit() check, so the bit is only set on a genuine task-owned CID. A CPU-owned (MM_CID_ONCPU) CID of a running active task is handled by the cid_on_cpu(pcp->cid) branch above and never reaches this path, so excluding MM_CID_UNSET (and the already-transitioning case) is sufficient.	7.8	More Details

CVE-2026-63875	In the Linux kernel, the following vulnerability has been resolved: arm64: tlb: Flush walk cache when unsharing PMD tables When huge_pmd_unshare() is called to unshare a PMD table, the tlb_unshare_pmd_ptdesc() function sets tlb->unshared_tables=true but the aarch64 tlb_flush() only checked tlb->freed_tables to determine whether to use TLBF_NONE (vae1is, invalidates walk cache) or TLBF_NOWALKCACHE (vale1is, leaf-only). This caused the stale PMD page table entry to remain in the walk cache after unshare, potentially leading to incorrect page table walks. Fix by including unshared_tables in the check, so that when unsharing tables, TLBF_NONE is used and the walk cache is properly invalidated. Here is the detailed distinction between vae1is and vale1is: | Instruction Combination | Actual Invalidation Scope | | ---------- | ----------| | `VAE1IS` + TTL=`0` | All entries at all levels (full invalidation) | | `VAE1IS` + TTL=`2` (L2) | Non-leaf at Level 0/1 + leaf at Level 2 | | `VALE1IS` + TTL=`0` | Leaf entries at all levels (non-leaf not cleared) | | `VALE1IS` + TTL=`2` (L2) | Leaf entry at Level 2 only |	7.8	More Details
CVE-2026-63802	In the Linux kernel, the following vulnerability has been resolved: blk-cgroup: fix UAF in __blkcg_rstat_flush() When multiple blkgs in the same blkcg are released concurrently, a use-after-free can occur. The race happens when one blkg's __blkcg_rstat_flush() removes another blkg's iostat entries via llist_del_all(). The second blkg sees an empty list and proceeds to free itself while the first is still iterating over its entries. Move the flush from __blkg_release() (RCU callback) to blkg_release() (before call_rcu). This ensures the RCU grace period waits for any concurrent flush's rcu_read_lock() section to complete before freeing.	7.8	More Details
CVE-2026-63874	In the Linux kernel, the following vulnerability has been resolved: net: mctp: usb: fix race between urb completion and rx_retry cancellation It's possible that sequencing between setting ->stopped and cancelling the rx_retry work (in ndo_stop) could leave us with an urb queued: T1: ndo_stop T2: rx_retry_work ---------- ---------- LD: ->stopped => false ST: ->stopped <= true usb_kill_urb() mctp_usb_rx_queue() usb_submit_urb() cancel_delayed_work_sync() That urb completion can then re-schedule rx_retry_work. Strenghen the sequencing between the stop (preventing another requeue) and the cancel by updating both atomically under a new rx lock. After setting ->rx_stopped, and cancelling pending work, we know that the requeue cannot occur, so all that's left is killing any pending urb.	7.8	More Details
CVE-2026-63870	In the Linux kernel, the following vulnerability has been resolved: ieee802154: 6lowpan: only accept IPv6 packets in lowpan_xmit() The aoe driver (or similar) generates a non-IPv6 packet (e.g., ETH_P_AOE) and queues it for transmission via dev_queue_xmit() on a 6LoWPAN interface (configured by the user or test case). Since the packet is not IPv6, the 6LoWPAN header_ops->create function (lowpan_header_create or header_create) returns early without initializing the lowpan_addr_info structure in the skb headroom. In the transmit function (lowpan_xmit), the driver calls lowpan_header (or setup_header) which unconditionally copies and uses the lowpan_addr_info from the headroom, which contains uninitialized data. Fix this by dropping non IPv6 packets. A similar fix is needed in net/bluetooth/6lowpan.c bt_xmit().	7.8	More Details
CVE-2026-63803	In the Linux kernel, the following vulnerability has been resolved: hdlc_ppp: sync per-proto timers before freeing hdlc state Each PPP control protocol (LCP/IPCP/IPV6CP) embedded in struct ppp registers a timer via timer_setup(). That struct ppp is the hdlc->state allocation, which detach_hdlc_protocol() frees with kfree() in both teardown paths: unregister_hdlc_device() and the re-attach inside attach_hdlc_protocol(). The ppp proto never registered a .detach callback, so detach_hdlc_protocol() performs no timer synchronization before the kfree(). The only cancel, timer_delete(&proto->timer) in ppp_cp_event(), is partial (it does not wait for a running callback) and only runs on the ->CLOSED transition; ppp_stop()/ppp_close() do not sync either. A ppp_timer callback already executing (blocked on ppp->lock) survives the kfree and then dereferences proto->state / ppp->lock in freed memory, leading to a use-after-free. Fix this by adding a .detach helper that calls timer_shutdown_sync() on every per-proto timer. detach_hdlc_protocol() invokes proto->detach(dev) before kfree(hdlc->state), so timer_shutdown_sync() now runs on both free paths. timer_shutdown_sync() is used instead of timer_delete_sync() because the keepalive path re-arms the timer through add_timer()/mod_timer() and shutdown blocks any re-activation during teardown. Initialize the per-protocol timers in ppp_ioctl() when the protocol is attached, and remove the now-redundant timer_setup() from ppp_start(), so that the timers are initialized exactly once at attach time and ppp_timer_release() never operates on uninitialized timer_list structures. attach_hdlc_protocol() uses kmalloc() (not kzalloc), so struct ppp's protos[i].timer is uninitialized garbage until the first timer_setup(); without this init-at-attach, attaching the PPP protocol without ever bringing the device up would leave timer_shutdown_sync() operating on uninitialized memory in .detach. Moving the init out of ppp_start() (which only runs on NETDEV_UP) into the attach path makes the initialization unconditional and avoids initializing the same timer_list twice. This bug was found by static analysis.	7.8	More Details
CVE-2026-63805	In the Linux kernel, the following vulnerability has been resolved: crypto: nx - fix nx_crypto_ctx_exit argument nx_crypto_ctx_shash_exit calls nx_crypto_ctx_exit with crypto_shash_ctx(...) but crypto_shash_ctx gives a nx_crypto_ctx *, not a crypto_tfm *. Fix the type in nx_crypto_ctx_exit and drop the bogus crypto_tfm_ctx call. This fixes the following oops: BUG: Unable to handle kernel data access at 0xc0403effffffffc8 Faulting instruction address: 0xc000000000396cb4 Oops: Kernel access of bad area, sig: 11 [#15] Call Trace: nx_crypto_ctx_shash_exit+0x24/0x60 crypto_shash_exit_tfm+0x28/0x40 crypto_destroy_tfm+0x98/0x140 crypto_exit_ahash_using_shash+0x20/0x40 crypto_destroy_tfm+0x98/0x140 hash_release+0x1c/0x30 alg_sock_destruct+0x38/0x60 __sk_destruct+0x48/0x2b0 af_alg_release+0x58/0xb0 __sock_release+0x68/0x150 sock_close+0x20/0x40 __fput+0x110/0x3a0 sys_close+0x48/0xa0 system_call_exception+0x140/0x2d0 system_call_common+0xf4/0x258 .. which came from hardlink(1) opportunistically using AF_ALG. The same problem exists with nx_crypto_ctx_skcipher_exit getting a context	7.8	More Details

	it wasn't expecting, but apparently nobody hit that for years.		
CVE-2026-63809	In the Linux kernel, the following vulnerability has been resolved: bpf: use kvfree() for replaced sysctl write buffer proc_sys_call_handler() allocates its temporary sysctl buffer with kvzalloc() and passes it to __cgroup_bpf_run_filter_sysctl(). Since kvzalloc() may fall back to vmalloc() for large allocations, freeing that buffer with kfree() is wrong and can corrupt memory. Use kvfree() to safely handle both kmalloc and kvzalloc()/vmalloc allocations. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1-rc5. Reproduced the bug based on v7.1-rc4 in a QEMU x86_64 guest booted with KASAN and CONFIG_FAILSLAB enabled. To exercise the replacement path, the test tree also included the accompanying fix for the stale ret == 1 check in __cgroup_bpf_run_filter_sysctl(). The reproducer confines failslab injections to the proc_sys_call_handler() range, uses stacktrace-depth=32, and injects fail-nth=1 while writing 8191 bytes to /proc/sys/kernel/domainname from a task in the target cgroup. Under that setup, fail-nth=1 triggered the fault: BUG: unable to handle page fault for address: ffffeb0200024d48 #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: Oops: 0000 SMP KASAN NOPTI CPU: 2 UID: 0 PID: 209 Comm: repro_proc_sys_ Not tainted 7.1.0-rc4-00686-g97625979a5d4 PREEMPT(lazy) Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.15.0-1 04/01/2014 RIP: 0010:kfree+0x6e/0x510 ... Call Trace: <TASK> ? __cgroup_bpf_run_filter_sysctl+0x626/0xc30 __cgroup_bpf_run_filter_sysctl+0x74d/0xc30 ? __pfx___cgroup_bpf_run_filter_sysctl+0x10/0x10 ? srso_return_thunk+0x5/0x5f ? __kvmalloc_node_noprof+0x345/0x870 ? proc_sys_call_handler+0x250/0x480 ? srso_return_thunk+0x5/0x5f proc_sys_call_handler+0x3a2/0x480 ? __pfx_proc_sys_call_handler+0x10/0x10 ? srso_return_thunk+0x5/0x5f ? selinux_file_permission+0x39f/0x500 ? srso_return_thunk+0x5/0x5f ? lock_is_held_type+0x9e/0x120 vfs_write+0x98e/0x1000 ... </TASK> With this fix applied on top of the same test setup, rerunning the reproducer with fail-nth=1 yields no corresponding Oops reports.	7.8	More Details
CVE-2026-63817	In the Linux kernel, the following vulnerability has been resolved: f2fs: validate compress cache inode only when enabled F2FS_COMPRESS_INO() uses NM_I(sbi)->max_nid as the synthetic inode number for the compressed page cache inode. That inode only exists when the compress_cache mount option is enabled. When compress_cache is disabled, max_nid is outside the valid inode range. A corrupted directory entry that points to ino == max_nid should therefore be rejected by f2fs_check_nid_range(). However, is_meta_ino() currently treats F2FS_COMPRESS_INO() as a meta inode unconditionally, so f2fs_iget() bypasses do_read_inode() and its nid range check, and instantiates a fake internal inode instead. Gate the compressed cache inode case on COMPRESS_CACHE, matching f2fs_init_compress_inode(). With compress_cache disabled, ino == max_nid now follows the normal inode path and is rejected as an out-of-range nid.	7.8	More Details
CVE-2026-53401	In the Linux kernel, the following vulnerability has been resolved: fbdev: omap2: fix use-after-free in omapfb_mmap omapfb_mmap() has a race condition with OMAPFB_SETUP_PLANE ioctl that can lead to use-after-free: The fb_mmap() entry point holds mm_lock but not lock (fb_info->lock), while ioctl handlers like OMAPFB_SETUP_PLANE hold lock but not mm_lock. This allows concurrent execution. In omapfb_mmap(): 1. rg = omapfb_get_mem_region(ofbi->region); // Get old region ref 2. start = omapfb_get_region_paddr(ofbi); // Read from NEW region 3. len = fix->smem_len; // Read from NEW region 4. vm_iomap_memory(vma, start, len); // Map NEW region memory 5. atomic_inc(&rg->map_count); // Increment OLD region! Concurrently, OMAPFB_SETUP_PLANE can: - Reassign ofbi->region = new_rg - Update fix->smem_len - OMAPFB_SETUP_MEM then checks NEW region's map_count (0!) and frees it This leaves userspace with a mapping to freed physical memory. The fix is to read all required values (start, len) from the same region reference (rg) that will have its map_count incremented, preventing the region from being freed while still mapped.	7.8	More Details
CVE-2026-63995	In the Linux kernel, the following vulnerability has been resolved: ethtool: cmis: validate start_cmd_payload_size from module The CMIS firmware update code reads start_cmd_payload_size from the module's FW Management Features CDB reply and uses it directly as the byte count for memcpy. The destination buffer is 112 bytes (ETHTOOL_CMIS_CDB_LPL_MAX_PL_LENGTH - 8). So a malicious module (or corrupted response) can cause a OOB write later on in cmis_fw_update_start_download(). Let's error out. If modules that expect longer LPL writes actually exist we should revisit. struct cmis_cdb_start_fw_download_pl's definition has to move, no change there.	7.8	More Details
CVE-2026-64057	In the Linux kernel, the following vulnerability has been resolved: afs: Fix the locking used by afs_get_link() The afs filesystem in the kernel doesn't do locking correctly for symbolic links. There are a number of problems: (1) It doesn't do any locking around afs_read_single() to prevent races between multiple ->get_link() calls, thereby allowing the possibility of leaks. (2) It doesn't use RCU barriering when accessing the buffer pointers during RCU pathwalk. (3) It can race with another thread updating the contents of the symlink if a third party updated it on the server. Fix this by the following means: (0) Move symlink handling into its own file as this makes it more complicated. (1) Take the validate_lock around afs_read_single() to prevent races between multiple ->get_link() calls. (2) Keep a separate copy of the symlink contents with an rcu_head. This is always going to be a lot smaller than a page, so it can be kmalloc'd and save quite a bit of memory. It also needs a refcount for non-RCU pathwalk. (3) Split the symlink read and write-to-cache routines in afs from those for directories. (4) Discard the I/O buffer as soon as the write-to-cache completes as this is a full page (plus a folio_queue). (5) If there's no cache, discard the I/O buffer immediately after reading and copying if there is no cache.	7.8	More Details

CVE-2026-64112	In the Linux kernel, the following vulnerability has been resolved: rbd: eliminate a race in lock_dwork draining on unmap Given how rbd_lock_add_request() and rbd_img_exclusive_lock() are written, lock_dwork may be (re)queued more than it's actually needed: for example in case a new I/O request comes in while we are in the middle of rbd_acquire_lock() on behalf of another I/O request. This is expected and with rbd_release_lock() preemptively canceling lock_dwork is benign under normal operation. A more problematic example is maybe_kick_acquire(): if (have_requests || delayed_work_pending(&rbd_dev->lock_dwork)) { dout("%s rbd_dev %p kicking lock_dwork\n", __func__, rbd_dev); mod_delayed_work(rbd_dev->task_wq, &rbd_dev->lock_dwork, 0); } It's not unrealistic for lock_dwork to get canceled right after delayed_work_pending() returns true and for mod_delayed_work() to requeue it right there anyway. This is a classic TOCTOU race. When it comes to unmapping the image, there is an implicit assumption of no self-initiated exclusive lock activity past the point of return from rbd_dev_image_unlock() which unlocks the lock if it happens to be held. This unlock is assumed to be final and lock_dwork (as well as all other exclusive lock tasks, really) isn't expected to get queued again. However, lock_dwork is canceled only in cancel_tasks_sync() (i.e. later in the unmap sequence) and on top of that the cancellation can get in effect nullified by maybe_kick_acquire(). This may result in rbd_acquire_lock() executing after rbd_dev_device_release() and rbd_dev_image_release() run and free and/or reset a bunch of things. One of the possible failure modes then is a violated rbd_assert(rbd_image_format_valid(rbd_dev->image_format)); in rbd_dev_header_info() which is called via rbd_dev_refresh() from rbd_post_acquire_action(). Redo exclusive lock task draining to provide saner semantics and try to meet the assumptions around rbd_dev_image_unlock().	7.8	More Details
CVE-2026-64108	In the Linux kernel, the following vulnerability has been resolved: cifs: Fix busy dentry used after unmounting Since commit 340cea84f691c ("cifs: open files should not hold ref on superblock"), cifs file only holds the dentry ref_cnt, the cifs file close work(cfile->deferred) could be executed after unmounting, which will trigger a warning in generic_shutdown_super: BUG: Dentry 00000000a14a6845{i=c,n=file} still in use (1) [unmount of cifs cifs] The detailed processs is: process A process B kworker fd = open(PATH) vfs_open file->__f_path = *path // dentry->d_lockref.count = 1 cifs_open cifs_new_fileinfo cfile->dentry = dget(dentry) // dentry->d_lockref.count = 2 close(fd) __fput cifs_close queue_delayed_work(deferredclose_wq, cfile->deferred) dput(dentry) // dentry->d_lockref.count = 1 smb2_deferred_work_close _cifsFileInfo_put list_del(&cifs_file->flist) umount cleanup_mnt deactivate_super cifs_kill_sb cifs_close_all_deferred_files_sb cifs_close_all_deferred_files // cannot find cfile, skip _cifsFileInfo_put kill_anon_super generic_shutdown_super shrink_dcache_for_umount umount_check WARN ! // dentry->d_lockref.count = 1 cifsFileInfo_put_final dput(cifs_file->dentry) // dentry->d_lockref.count = 0 Fix it by flushing 'deferredclose_wq' before calling kill_anon_super. Fetch a reproducer in https://bugzilla.kernel.org/show_bug.cgi?id=221548.	7.8	More Details
CVE-2026-64099	In the Linux kernel, the following vulnerability has been resolved: drm/v3d: Fix use-after-free of CPU job query arrays on error path The CPU job ioctl's fail label calls kvfree() on cpu_job's timestamp and performance query arrays after v3d_job_cleanup(), which drops the job's last reference and frees cpu_job. Reading cpu_job at that point is a use-after-free. Also, on the early v3d_job_init() failure path, it is a NULL dereference, since v3d_job_deallocate() zeroes the local pointer. In the success path, the arrays are released from the scheduler's .free_job callback, but on the error path, they are freed manually, as the job was never pushed to the scheduler. While the success path deals with this correctly, the fail path doesn't. On top of that, the manual kvfree() calls only free the array storage; they don't drm_syncobj_put() the per-query syncobjs that v3d_timestamp_query_info_free() and v3d_performance_query_info_free() release on the success path. So the same fail path that triggers the use-after-free also leaks one syncobj reference per query. Unify the CPU job teardown into the CPU job's kref destructor, mirroring v3d_render_job_free(). The scheduler's .free_job slot reverts to the generic v3d_sched_job_free() and the fail label drops the manual kvfree() calls, leaving a single teardown path that is reached from both the scheduler and the ioctl error path. That removes the use-after-free, the NULL dereference, and the syncobj leak by construction.	7.8	More Details
CVE-2026-64098	In the Linux kernel, the following vulnerability has been resolved: drm/virtio: use uninterruptible resv lock for plane updates virtio_gpu_cursor_plane_update() and virtio_gpu_resource_flush() lock the framebuffer BO's dma_resv via virtio_gpu_array_lock_resv() and ignore its return value. The function can fail with -EINTR from dma_resv_lock_interruptible() (signal during lock wait) or with -ENOMEM from dma_resv_reserve_fences() (fence slot allocation), leaving the resv lock not held. The queue path then walks the object array and calls dma_resv_add_fence(), which requires the lock held; with lockdep enabled this trips dma_resv_assert_held(): WARNING: drivers/dma-buf/dma-resv.c:296 at dma_resv_add_fence+0x71e/0x840 Call Trace: virtio_gpu_array_add_fence virtio_gpu_queue_ctrl_sgs virtio_gpu_queue_fenced_ctrl_buffer virtio_gpu_cursor_plane_update drm_atomic_helper_commit_planes drm_atomic_helper_commit_tail commit_tail drm_atomic_helper_commit drm_atomic_commit drm_atomic_helper_update_plane __setplane_atomic drm_mode_cursor_universal drm_mode_cursor_common drm_mode_cursor_ioctl drm_ioctl __x64_sys_ioctl Beyond the WARN, mutating the dma_resv fence list without the lock races with concurrent readers/writers and can corrupt the list. Both call sites run inside the .atomic_update plane callback, which DRM atomic helpers do not allow to fail (by the time it runs, the commit has been signed off to userspace and there is no clean rollback path). Moving the lock acquisition to .prepare_fb was rejected because the broader lock scope deadlocks against other BO locking paths in the same atomic commit. Introduce virtio_gpu_lock_one_resv_uninterruptible() that uses dma_resv_lock() instead of dma_resv_lock_interruptible(). This eliminates the -EINTR failure mode -- the realistic syzbot trigger -- without extending the lock hold across the commit. The helper locks a single BO and rejects nents > 1 with -EINVAL; both fix sites lock exactly one BO. Use it from virtio_gpu_cursor_plane_update() and	7.8	More Details

	virtio_gpu_resource_flush(); check the return value to handle the remaining -ENOMEM case from dma_resv_reserve_fences() by freeing the objs and skipping the plane update for that frame. The framebuffer BOs touched here are not shared with other contexts and lock contention is expected to be brief, so the loss of signal-interruptibility is acceptable. Other callers of virtio_gpu_array_lock_resv() (the ioctl paths) continue to use the interruptible variant. The bug was reported by syzbot, triggered via fault injection (fail_nth) on the DRM_IOCTL_MODE_CURSOR path, which forces the -ENOMEM branch in dma_resv_reserve_fences().		
CVE-2026-64097	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Validate GPIO pin LUT table size before iterating [Why&How] The GPIO pin table parsers in get_gpio_i2c_info() and bios_parser_get_gpio_pin_info() derive an element count from the VBIOS table_header.structuresize field, then iterate over gpio_pin[] entries. However, GET_IMAGE() only validates that the table header itself fits within the BIOS image. If the VBIOS reports a structuresize larger than the actual mapped data, the loop reads past the end of the BIOS image, causing an out-of-bounds read. Fix this by calling bios_get_image() to validate that the full claimed structuresize is accessible within the BIOS image before entering the loop in both functions. (cherry picked from commit ba5e95b43b773ae1bf1f66ee6b31eb774e65afe3)	7.8	More Details
CVE-2026-64086	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) include PEC byte in pmbus_block_xfer read buffer adm1266_pmbus_block_xfer() sets up the read transaction with .buf = data->read_buf, .len = ADM1266_PMBUS_BLOCK_MAX + 2, but read_buf in struct adm1266_data is declared as u8 read_buf[ADM1266_PMBUS_BLOCK_MAX + 1]; For a max-length block response (length byte = 255 + up to 1 PEC byte), the i2c controller is told to write 257 bytes into a 256-byte buffer, putting one byte past the end of read_buf. The same response also makes the subsequent PEC compare if (crc != msgs[1].buf[msgs[1].buf[0] + 1]) read a byte beyond the array. Bump the read_buf declaration to ADM1266_PMBUS_BLOCK_MAX + 2 so the buffer can hold the length byte, up to 255 payload bytes, and the PEC byte the i2c_msg length already accounts for.	7.8	More Details
CVE-2026-64084	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) cap PDIO scan in get_multiple at ADM1266_PDIO_NR adm1266_gpio_get_multiple() iterates the PDIO portion of the caller-supplied mask using for_each_set_bit_from(gpio_nr, mask, ADM1266_GPIO_NR + ADM1266_PDIO_STATUS) { ... } where ADM1266_PDIO_STATUS is the PMBus command code (0xE9, i.e. 233), not the number of PDIO pins. The intended upper bound is ADM1266_GPIO_NR + ADM1266_PDIO_NR = 25. gpiolib hands in a mask sized for gc.ngpio (= 25 bits on this chip), so the iteration walks find_next_bit() up to 242, reading up to 217 extra bits (a handful of unsigned-long words: four on 64-bit, seven on 32-bit) of whatever lives past the end of the mask in the caller's stack. Any incidental set bit in that range then drives a set_bit(gpio_nr, bits) call that writes past the end of the caller-supplied bits array too -- both out-of-bounds. Substitute ADM1266_PDIO_NR for the constant so the scan stops at the last real PDIO bit.	7.8	More Details
CVE-2026-64082	In the Linux kernel, the following vulnerability has been resolved: riscv: Fix register corruption from uninitialized cregs on error compat_riscv_gpr_set() calls cregs_to_regs() unconditionally, even when user_regset_copyin() fails. Since cregs is an uninitialized stack variable, a copyin failure causes uninitialized stack data to be written into the target task's pt_regs, corrupting its register state and potentially leaking kernel stack contents. compat_restore_sigcontext() has the same issue: it calls cregs_to_regs() even when __copy_from_user() fails, leading to the same corruption of the signal-returning task's register state on error. Only call cregs_to_regs() when the user copy succeeds.	7.8	More Details
CVE-2026-64078	In the Linux kernel, the following vulnerability has been resolved: netfilter: x_tables: add and use xtables_unregister_table_exit Previous change added xtables_unregister_table_pre_exit to detach the table from the packetpath and to unlink it from the active table list. In case of rmmod, userspace that is doing set/getsockopt for this table will not be able to re-instantiate the table: 1. The larval table has been removed already 2. existing instantiated table is no longer on the xt pernet table list. This adds the second stage helper: unlink the table from the dying list, free the hook ops (if any) and do the audit notification. It replaces xt_unregister_table().	7.8	More Details
CVE-2026-15809	A flaw was found in CRI-O. The fix for a previous vulnerability (CVE-2022-4318) was incorrect, allowing it to be bypassed. An attacker capable of setting environment variables on a container can inject a newline character into the HOME environment variable. This issue allows the addition of arbitrary lines into /etc/passwd by use of a specially crafted environment variable.	7.8	More Details
CVE-2026-64077	In the Linux kernel, the following vulnerability has been resolved: netfilter: ebtables: move to two-stage removal scheme Like previous patches for x_tables, follow same pattern in ebtables. We can't reuse xt helpers: ebt_table struct layout is incompatible. table->ops assignment is now done while still holding the ebt mutex to make sure we never expose partially-filled table struct.	7.8	More Details
CVE-2026-64076	In the Linux kernel, the following vulnerability has been resolved: netfilter: bridge: eb_tables: close module init race sashiko reports for unrelated patch: Does the core ebtables initialization in ebtables.c suffer from a similar race? Once nf_register_sockopt() completes, the sockopts are exposed globally. sockopt has to be registered last, just like in ip/ip6/arptables.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: fs/statmount: fix slab out-of-bounds write in statmount_mnt_idmap statmount_mnt_idmap() writes one mapping with seq_printf() and then manually advances seq->count to include the NUL separator. If seq_printf() overflows, seq_set_overflow() sets seq-		More

2026-64074	>count to seq->size. The manual seq->count++ changes this to seq->size + 1. seq_has_overflowed() then no longer detects the overflow. The corrupted count returns to statmount_string(), which later executes: seq->buf[seq->count++] = '\0'; This causes a 1-byte NULL out-of-bounds write on the dynamically allocated seq buffer. Fix this by checking for overflow immediately after seq_printf().	7.8	Details
CVE-2026-8933	A local privilege escalation vulnerability exists in snap-confine, a set-capabilities core component used internally by Canonical snapd to construct the secure execution environment for snap applications. This vulnerability uniquely affects versions of snap-confine configured with set-capabilities (rather than standard set-uid-root installations). Due to a flaw in how privilege boundaries or security sandboxes are initialized when the binary runs under limited ambient capabilities, a local, unprivileged attacker can exploit this behavior to bypass intended restrictions and execute arbitrary code. Successful exploitation allows the local user to elevate their privileges to full root authority.	7.8	More Details
CVE-2026-64073	In the Linux kernel, the following vulnerability has been resolved: irq_work: Fix use-after-free in irq_work_single() on PREEMPT_RT On PREEMPT_RT, non-HARD irq_work runs in per-CPU kthreads via run_irq_workd(), so irq_work_sync() uses rcuwait() to wait for BUSY==0. After irq_work_single() clears BUSY via atomic_cmpxchg(), it still dereferences @work for irq_work_is_hard() and rcuwait_wake_up(). An irq_work_sync() caller on another CPU that enters after BUSY is cleared can observe BUSY==0 immediately, return, and free the work before those accesses complete — causing a use-after-free. Fix this by wrapping run_irq_workd() in guard(rcu)() so that the entire irq_work_single() execution is within an RCU read-side critical section. Then add synchronize_rcu() in irq_work_sync() after rcuwait_wait_event() to ensure the caller waits for the RCU grace period before returning, preventing premature frees.	7.8	More Details
CVE-2026-64114	In the Linux kernel, the following vulnerability has been resolved: ipv4: raw: reject IP_HDRINCL packets with ihl < 5 raw_send_hdrinc() validates that the caller-supplied IPv4 header fits within the message length: iphlen = iph->ihl * 4; err = -EINVAL; if (iphlen > length) goto error_free; if (iphlen >= sizeof(*iph)) { /* fix up saddr, tot_len, id, csum, transport_header */ } It does not, however, reject ihl < 5. For such a packet the "if (iphlen >= sizeof(*iph))" branch is skipped, leaving the crafted iphdr untouched, but the packet is still handed to __ip_local_out() and onward. Downstream consumers that read iph->ihl assume a sane value: net/ipv4/ah4.c:ah_output() in particular subtracts sizeof(struct iphdr) from top_iph->ihl * 4 and passes the (signed-int-negative, then cast to size_t) result to memcpy(), producing an OOB access of length close to SIZE_MAX and a host kernel panic. An IPv4 header with ihl < 5 is malformed by definition (RFC 791: "Internet Header Length is the length of the internet header in 32 bit words ... Note that the minimum value for a correct header is 5."). The kernel should not be willing to inject such a packet into its own output path. Reject "iphlen < sizeof(*iph)" alongside the existing "iphlen > length" check. This matches the principle that locally constructed packets that re-enter the IP stack must pass the same basic sanity tests that a foreign packet would be subjected to. Once this lands, the "if (iphlen >= sizeof(*iph))" wrapper around the fixup branch becomes redundant; left in place to keep the patch minimal and backport-friendly. A follow-up can unwrap it. Note that commit 86f4c90a1c5c ("ipv4, ipv6: ensure raw socket message is big enough to hold an IP header") ensures the message buffer is large enough to hold an iphdr, but does not constrain the self-reported iph->ihl. Reachability: the malformed packet source is any caller with CAP_NET_RAW, including an unprivileged process in a user+net namespace on a kernel with CONFIG_USER_NS=y. The reproduced AH crash also requires a matching xfrm AH policy on the outgoing route; a container granted CAP_NET_ADMIN can install that state and policy in its netns. Loopback bypasses xfrm_output, so the trigger uses a real netdev. Reproduced on UML + KASAN: kernel-mode fault at addr 0x0 with memcpy_orig at the crash site. Same shape reproduces inside a rootless Docker container with --cap-add NET_ADMIN on a stock distro kernel.	7.8	More Details
CVE-2026-58558	Permission control vulnerability in the file system. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	7.8	More Details
CVE-2026-64123	In the Linux kernel, the following vulnerability has been resolved: net: hsr: defer node table free until after RCU readers HSR node-list and node-status generic-netlink operations run under rcu_read_lock(). They walk hsr->node_db through hsr_get_next_node() and hsr_get_node_data(), but RTM_DELLINK teardown removes the same node table with plain list_del() and frees each node immediately. That lets a generic-netlink reader hold a struct hsr_node pointer across hsr_dellink(). In a KASAN build, widening the reader window after hsr_get_next_node() obtains the node reproduces a slab-use-after-free when the reader copies node->macaddress_A; the freeing stack is hsr_del_nodes() from hsr_dellink(). Use list_del_rcu() and defer the free through the existing hsr_free_node_rcu() callback. This matches the lifetime rule used by the HSR prune paths, which already delete nodes with list_del_rcu() and call_rcu().	7.8	More Details
CVE-2026-15895	OS command injection in the npm package loading component in AWS jsii-diff before 1.131.0 might allow context-dependent attackers to execute arbitrary commands via crafted package specifiers passed to the npm: source argument. To mitigate this issue, users should upgrade to jsii-diff v1.131.0 or later.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: ethtool: cmis: require exact CDB reply length Malicious SFP module could respond with rpl_len longer than what cmis_cdb_process_reply() expected, leading to OOB writes. Malicious HW is a bit theoretical but some modules may just be buggy and/or the reads may occasionally get corrupted, so let's protect the kernel. The existing check protects from short replies. We need to protect from long ones, too. All callers that pass a non-zero rpl_exp_len cast the reply payload to a fixed-layout struct and read fields at fixed offsets, with no version negotiation or short-reply		

2026-63996	handling: - cmis_cdb_validate_password() - cmis_cdb_module_features_get() - cmis_fw_update_fw_mng_features_get() so let's assume that responses longer than expected do not have to be handled gracefully here. Add a warning message to make the debug easier in case my understanding is wrong... Note that page_data->length (argument of kmalloc) comes from last arg to ethtool_cmis_page_init() which is rpl_exp_len. Note2 that AIs also like to point out overflows in args->req.payload itself (which is a fixed-size 120 B buffer, on the stack), but callers should be reading structs defined by the standard, so protecting from requests for more data than max seem like defensive programming.	7.8	More Details
CVE-2026-46709	Tabby (formerly Terminus) is a highly configurable terminal emulator. Prior to 1.0.234, Tabby inserts dropped file paths from tabby-electron/src/pathDrop.ts into the active shell without neutralizing command substitution metacharacters such as $(...) and `...`, so the incomplete CVE-2026-45038 fix for control characters still allows code execution when the victim presses Enter. This issue is fixed in version 1.0.234.	7.8	More Details
CVE-2026-52349	Directory Traversal vulnerability in Menyoo 2.0 Versions before commit 729aa48: fixed in commit 729aa48 allows a local attacker to execute arbitrary code via the Spooner file management, VehicleSpawner save/folder/rename functionality, WeaponOptions save/folder/rename functionality, PedComponentChanger create folder/createfile/rename functionality.	7.8	More Details
CVE-2026-56687	Dell ThinOS 10, versions prior to 2605_10.2100, contain an Obsolete Feature in UI vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Unauthorized access.	7.8	More Details
CVE-2026-58659	PyTorch Lightning through 2.6.5, fixed in commit d710d68, contains a remote code execution vulnerability in the _load_state function that imports and executes attacker-controlled module names from checkpoint _instantiator hyperparameters. Attackers can craft malicious checkpoint files that bypass weights_only=True protections to execute arbitrary code when LightningModule.load_from_checkpoint is called.	7.8	More Details
CVE-2026-40633	Dell PowerScale OneFS versions 9.5.0.0 through 9.10.1.7, versions 9.11.0.0 through 9.13.0.2 contains an Insertion of Sensitive Information into Log File vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Information disclosure.	7.8	More Details
CVE-2026-40952	CVE-2026-40952 is a privilege misconfiguration in the Secure Access installer for the Windows client and server prior to version 14.55. Attackers with local access to the client or server can use it to elevate privileges to Administrator when Secure Access is installed in a non-default location.	7.8	More Details
CVE-2026-48389	DNG SDK versions 1.7.1 2536 and earlier are affected by a Stack-based Buffer Overflow vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-64624	FreeRDP before 3.28.0 treats lines beginning with forward slash in RDP files as raw command-line options, exposing the entire CLI parser surface to untrusted files. Attackers can craft malicious RDP files with /rdp2tcp, /cert:ignore, or /drive options to execute arbitrary commands, bypass certificate validation, or expose local filesystems without user interaction.	7.8	More Details
CVE-2026-64181	In the Linux kernel, the following vulnerability has been resolved: mm: fix __vm_normal_page() to handle missing support for pmd_special()/pud_special() On x86 32-bit with THP enabled, zap_huge_pmd() is seen to generate a "WARNING: mm/memory.c:735 at __vm_normal_page+0x6a/0x7d", from the VM_WARN_ON_ONCE(is_zero_pfn(pfn) || is_huge_zero_pfn(pfn)); followed by "BUG: Bad rss-counter state"s, then later "BUG: Bad page state"s when reclaim gets to call shrink_huge_zero_folio_scan(). It's as if the _PAGE_SPECIAL bit never got set in the huge_zero pmd: and indeed, whereas pte_special() and pte_mkspecial() are subject to a dedicated CONFIG_ARCH_HAS_PTE_SPECIAL, pmd_special() and pmd_mkspecial() are subject to CONFIG_ARCH_SUPPORTS_PMD_PFNMAP, which is never enabled on any 32-bit architecture. While the problem was exposed through commit d80a9cb1a64a ("mm/huge_memory: add and use normal_or_softleaf_folio_pmd()"), it was an oversight in commit af38538801c6 ("mm/memory: factor out common code from vm_normal_page_*()") and would result in other problems: * huge zero folio accounted in smaps, pagemap (PAGE_IS_FILE) and numamaps as file-backed THP * folio_walk_start() returning the folio even without FW_ZEROPAGE set. Callers seem to tolerate that, though. ... and triggering the VM_WARN_ON_ONE(), although never reported so far. To fix it, teach vm_normal_page_pmd()/vm_normal_page_pud() to consider whether pmd_special/pud_special is actually implemented.	7.8	More Details
CVE-2026-15905	Use after free in Aura in Google Chrome prior to 150.0.7871.128 allowed a local attacker to potentially exploit heap corruption via a malicious file. (Chromium security severity: High)	7.8	More Details
CVE-2026-64152	In the Linux kernel, the following vulnerability has been resolved: iommu: Handle unmap error when iommu_debug is enabled Sashiko noticed a latent bug where the map error flow called iommu_unmap() which calls iommu_debug_unmap_begin()/iommu_debug_unmap_end() however since this is an error path the map flow never actually established the original iommu_debug_map() it will malfunction. Lift the unmap error handling into iommu_map_nosync() and reorder it so the trace_map()/iommu_debug_map() records the partial mapping and then immediately unmaps it. This avoid creating the unbalanced tracking and provides saner tracing instead of a unmap unmatched to any map.	7.8	More Details

CVE-2026-64137	In the Linux kernel, the following vulnerability has been resolved: smb: client: require net admin for CIFS SWN netlink CIFS_GENL_CMD_SWN_NOTIFY is the userspace witness-notify command. The intended sender is the cifs.witness helper, but the generic-netlink operation currently has no capability flag, so any local process can send RESOURCE_CHANGE or CLIENT_MOVE notifications to the in-kernel witness handler. The same family exposes CIFS_GENL_MCGRP_SWN without multicast-group capability flags. Register messages sent to that group include the witness registration id and, for NTLM-authenticated mounts, the username, domain, and password attributes copied from the CIFS session. An unprivileged local process should not be able to join that group and receive those messages. Require CAP_NET_ADMIN for incoming SWN_NOTIFY commands with GENL_ADMIN_PERM, and require CAP_NET_ADMIN over the network namespace for joining the SWN multicast group with GENL_MCAST_CAP_NET_ADMIN. The cifs.witness service runs with the privileges needed for both operations.	7.8	More Details
CVE-2026-64134	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcm: Don't setup bogus iov_iter for silencing At transition to the iov_iter for PCM data transfer, we blindly applied the iov_iter setup also for silencing (i.e. data = NULL), and it leads to a calculation of bogus iov_iter. Fortunately this didn't cause troubles on most of architectures but it goes wrong on RISC-V now, causing a NULL dereference. Handle the NULL data case to treat the silencing in interleaved_copy() for addressing the bug above. noninterleaved_copy() has already the NULL data handling, so it doesn't need changes.	7.8	More Details
CVE-2026-64058	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix netfs_read_folio() to wait on writeback Fix netfs_read_folio() to wait for an ongoing writeback to complete so that it can trust the dirty flag and whatever is attached to folio->private (folio->private may get cleaned up by the collector before it clears the writeback flag).	7.8	More Details
CVE-2026-64133	In the Linux kernel, the following vulnerability has been resolved: ALSA: asihpi: Fix potential OOB array access at reading cache find_control() to retrieve a cached info accesses the array with the given index blindly, which may lead to an OOB array access. Add a sanity check for avoiding it.	7.8	More Details
CVE-2026-64053	In the Linux kernel, the following vulnerability has been resolved: block: don't overwrite bip_vcnt in bio_integrity_copy_user() bio_integrity_add_page() already sets bip_vcnt to 1 for the bounce segment. Overwriting it with nr_vecs breaks bip_vcnt <= bip_max_vcnt on WRITE (bip_max_vcnt is 1), so the gap-merge checks in block/blk.h read past the bip_vec[] flex array. On READ the read is in bounds but lands on a saved user bvec instead of the bounce. The line was added for split propagation, but bio_integrity_clone() doesn't copy bip_vcnt and BIP_CLONE_FLAGS excludes BIP_COPY_USER.	7.8	More Details
CVE-2026-64015	In the Linux kernel, the following vulnerability has been resolved: security/keys: fix missed RCU read section on lookup Nicholas Carlini reports that the keyring code calls assoc_array_find() in find_key_to_update() without holding the RCU read lock, while the assoc_array_gc() code really is designed around removing the node from the tree and then freeing it after an RCU grace-period. The regular key handling doesn't see this because holding the keyring semaphore hides any lifetime issues, but the persistent key handling uses a different model. Instead of extending the keyring locking, just do the simple RCU locking that the assoc_array was designed for.	7.8	More Details
CVE-2026-64029	In the Linux kernel, the following vulnerability has been resolved: ALSA: seq: Serialize UMP output teardown with event_input seq_ump_process_event() borrows client->out_rfile.output without synchronizing with the first-open and last-close transition in seq_ump_client_open() and seq_ump_client_close(). The last output unuse can therefore drop opened[STR_OUT] to zero and release the rawmidi file while an in-flight event_input callback is still inside snd_rawmidi_kernel_write(). That leaves the rawmidi substream runtime exposed to teardown before the write path has taken its own buffer reference. Add a per-client rwlock for the event_input-visible output file. Publish a newly opened output file under the write side, and hold the read side from the output lookup through snd_rawmidi_kernel_write(). The last output close copies and clears the visible output file under the write side, then drops the lock and releases the saved rawmidi file. Use IRQ-safe rwlock guards because event_input can also be reached from atomic sequencer delivery. The buggy scenario involves two paths, with each column showing the order within that path: path A label: event_input path path B label: last unuse path 1. seq_ump_process_event() reads 1. seq_ump_client_close() client->out_rfile.output. drops opened[STR_OUT] to zero. 2. snd_rawmidi_kernel_write1() 2. snd_rawmidi_kernel_release() has not yet pinned runtime. closes the output file. 3. The writer continues using 3. close_substream() frees the borrowed substream. substream->runtime. This keeps the output substream and runtime alive for the full event_input write while keeping rawmidi release outside the rwlock. KASAN reproduced this as a slab-use-after-free in snd_rawmidi_kernel_write1(), with allocation through seq_ump_use()/snd_seq_port_connect() and free through seq_ump_unuse()/snd_seq_port_disconnect(). Validation reproduced this kernel report: KASAN slab-use-after-free in snd_rawmidi_kernel_write1+0x9d/0x400 RIP: 0033:0x7f5528af837f Read of size 8 Call trace: dump_stack_lvl+0x73/0xb0 (?:?) print_report+0xd1/0x650 (?:?) srso_alias_return_thunk+0x5/0xfbef5 (?:?) __virt_addr_valid+0x1a7/0x340 (?:?) kasan_complete_mode_report_info+0x64/0x200 (?:?) kasan_report+0xf7/0x130 (?:?) snd_rawmidi_kernel_write1+0x9d/0x400 (?:?) __asan_load8+0x82/0xb0 (?:?) update_stack_state+0x1ef/0x2d0 (?:?) snd_rawmidi_kernel_write+0x1a/0x20 (?:?) seq_ump_process_event+0xd4/0x120 (sound/core/seq/seq_ump_client.c:82) __snd_seq_deliver_single_event+0x8a/0xe0 (?:?) snd_seq_deliver_from_ump+0x2b2/0xd60 (?:?) lock_acquire+0x14e/0x2e0 (?:?) find_held_lock+0x31/0x90 (?:?) snd_seq_port_use_ptr+0xa6/0xe0 (?:?) __kasan_check_write+0x18/0x20 (?:?) do_raw_read_unlock+0x32/0xa0 (?:?) _raw_read_unlock+0x26/0x50 (?:?) snd_seq_deliver_single_event+0x45c/0x4b0 (?:?) snd_seq_deliver_event+0x10d/0x1b0 (?:?)	7.8	More Details

	snd_seq_client_enqueue_event+0x192/0x240 (?:?) snd_seq_write+0x2cd/0x450 (?:?) apparmor_file_permission+0x20/0x30 (?:?) security_file_permission+0x51/0x60 (?:?) vfs_write+0x1ce/0x850 (?:?) __fget_files+0x12b/0x220 (?:?) lock_release+0xc8/0x2a0 (?:?) __rcu_read_unlock+0x74/0x2d0 (?:?) __fget_files+0x135/0x220 (?:?) ksys_write+0x15a/0x180 (?:?) rcu_is_watching+0x24/0x60 (?:?) __x64_sys_write+0x46/0x60 (?:?) x64_sys_call+0x7d/0x20d0 (?:?) do_syscall_64+0xc1/0x360 (arch/x86/entry/syscall_64.c:87) entry_SYSCALL_64_after_hwframe+0x77/0x7f (?:?)		
CVE-2026-53366	In the Linux kernel, the following vulnerability has been resolved: ipv4: account for fraggap on the paged allocation path In __ip_append_data(), when the paged-allocation branch is taken, alloclen and pagedlen are computed as alloclen = fragheaderlen + transhdrlen; pagedlen = datalen - transhdrlen; datalen already includes fraggap, but the fraggap bytes carried over from the previous skb are copied into the new skb's linear area at offset transhdrlen by the subsequent skb_copy_and_csum_bits(). The linear area is therefore undersized by fraggap bytes while pagedlen is overstated by the same amount. The non-paged branch sets alloclen to fraglen, which already accounts for fraggap because datalen does. Bring the paged branch in line by adding fraggap to alloclen and subtracting it from pagedlen. After this adjustment, copy no longer collapses to -fraggap on the paged path, so remove the stale comment describing that old arithmetic.	7.8	More Details
CVE-2026-64027	In the Linux kernel, the following vulnerability has been resolved: net: shaper: rework the VALID marking (again) Recent commit changed the semantics from NOT_VALID to VALID. I didn't realize that the flags are not stored atomically with the entry in XArray. There's still a race of reader observing a VALID mark for a slot, getting interrupted, writer replacing the entry with a different one, reader continuing, fetching the entry which is now a different pointer than the pointer for which VALID was meant. The biggest consequence of this is that we may see a UAF since net_shaper_rollback() assumed that entries without VALID can be freed without observing RCU. Looks like the XArray marks are buying us nothing at this point. Let's convert the code to an explicit valid field. The smp_load_acquire() / smp_store_release() barriers are marginally cleaner.	7.8	More Details
CVE-2026-64026	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix DATA decrypt vs splice() by copying data to buffer in recvmsg This improves the fix for CVE-2026-43500. Fix the pagecache corruption from in-place decryption of a DATA packet transmitted locally by splice() by getting rid of the packet sharing in the I/O thread and unconditionally extracting the packet content into a bounce buffer in which the buffer is decrypted. recvmsg() (or the kernel equivalent) then copies the data from the bounce buffer to the destination buffer. The sk_buff then remains unmodified. This has an additional advantage in that the packet is then arranged in the buffer with the correct alignment required for the crypto algorithms to process directly. The performance of the crypto does seem to be a little faster and, surprisingly, the unencrypted performance doesn't seem to change much - possibly due to removing complexity from the I/O thread. Yet another advantage is that the I/O thread doesn't have to copy packets which would slow down packet distribution, ACK generation, etc.. The buffer belongs to the call and is allocated initially at 2K, sufficiently large to hold a whole jumbo subpacket, but the buffer will be increased in size if needed. However, to take this work, MSG_PEEK may cause a later packet to be decrypted into the buffer, in which case the earlier one will need re-decrypting for a subsequent recvmsg(). Note that rx_pkt_offset may legitimately see 0 as a valid offset now, so switch to using USHRT_MAX to indicate an invalid offset. Note also that I would generally prefer to replace the buffers of the current sk_buff with a new kmalloc'd buffer of the right size, ditching the old data and frags as this makes the handling of MSG_PEEK easier and removes the re-decryption issue, but this looks like quite a complicated thing to achieve. skb_morph() looks half way to what I want, but I don't want to have to allocate a new sk_buff.	7.8	More Details
CVE-2026-64023	In the Linux kernel, the following vulnerability has been resolved: gpio: aggregator: fix a potential use-after-free On error we free aggr->lookups->dev_id before removing the entry from the lookup table. If a concurrent thread calls gpiod_find() before we remove the entry, it could iterate over the list and call gpiod_match_lookup_table() which unconditionally dereferences dev_id when calling strcmp(). Reverse the order of cleanup.	7.8	More Details
CVE-2026-64017	In the Linux kernel, the following vulnerability has been resolved: blk-mq: pop cached request if it is usable When submitting a bio to blk-mq, if the task should sleep after peeking a cached request, but before it pops it, the plug flushes and calls blk_mq_free_plug_rqs, freeing the cached_rqs. This creates a use-after-free bug. Fix this by popping the cached request before any possible blocking calls if it is suitable for use. Popping this request first holds a queue reference, so avoid any serialization races with queue freezes and can safely proceed with dispatching that request to the driver. This potentially increases a timing window from when a driver wants to freeze its queue to when requests stop being dispatched. That scenario is off the fast path though, and drivers need to appropriately handle requests during a freeze request anyway. The downside is the popped element needs to be individually freed when we performed a bio plug merge. The cached request would have had to be freed later anyway, but this patch does it inline with building the plug list instead of after flushing it.	7.8	More Details
CVE-2026-61433	PraisonAI before 4.6.78 fails to safely encode deployment configuration values when generating Python source code for API servers. Attackers can inject arbitrary Python expressions through the deploy.api.host and agents_file configuration parameters that execute when the generated server starts or handles requests.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: bridge: mcast: Fix a possible use-after-free when removing a bridge port When per-VLAN multicast snooping is enabled, the bridge iterates over all the		

CVE-2026-64032	bridge ports, disables the per-port multicast context on each port and enables the per-{port, VLAN} multicast contexts instead. The reverse happens when per-VLAN multicast snooping is disabled. When global multicast snooping is enabled, the bridge iterates over all the bridge ports and enables the per-port multicast context on each port. The reverse happens when multicast snooping is disabled. The above scheme can result in a situation where both types of contexts (per-port and per-{port, VLAN}) are enabled on a single bridge port: # ip link add name br1 up type bridge mcast_snooping 1 mcast_querier 1 vlan_filtering 1 # ip link add name dummy1 up master br1 type dummy # ip link set dev br1 type bridge mcast_vlan_snooping 1 # ip link set dev br1 type bridge mcast_snooping 0 # ip link set dev br1 type bridge mcast_snooping 1 This is not intended and it is a problem since the commit cited below. Prior to this commit, when removing a bridge port, br_multicast_disable_port() would disable the per-port multicast context and the per-{port, VLAN} multicast contexts would get disabled when flushing VLANs. After this commit, br_multicast_disable_port() only disables the per-port multicast context if per-VLAN multicast snooping is disabled. If both types of contexts were enabled on the port when it was removed, the per-port multicast context would remain enabled when freeing the bridge port, leading to a use-after-free [1]. Fix by preventing the bridge from enabling / disabling the per-port multicast contexts when toggling global multicast snooping if per-VLAN multicast snooping is enabled. [1] ODEBUG: free active (active state 0) object: ffff88810f8bda78 object type: timer_list hint: br_ip6_multicast_port_query_expired (net/bridge/br_multicast.c:1927) WARNING: lib/debugobjects.c:629 at debug_print_object+0x1b1/0x3e0, CPU#5: swapper/5/0 [...] Call Trace: <IRQ> __debug_check_no_obj_freed (lib/debugobjects.c:1116) kfree (mm/slub.c:2620 mm/slub.c:6250 mm/slub.c:6565) kobject_cleanup (lib/kobject.c:689) rcu_do_batch (kernel/rcu/tree.c:2617) rcu_core (kernel/rcu/tree.c:2869) handle_softirqs (kernel/softirq.c:622) __irq_exit_rcu (kernel/softirq.c:656 kernel/softirq.c:496 kernel/softirq.c:735) irq_exit_rcu (kernel/softirq.c:752) sysvec_apic_timer_interrupt (arch/x86/kernel/apic/apic.c:1061 (discriminator 47) arch/x86/kernel/apic/apic.c:1061 (discriminator 47)) </IRQ>	7.8	More Details
CVE-2026-64011	In the Linux kernel, the following vulnerability has been resolved: nfc: llcp: Fix use-after-free in llcp_sock_release() llcp_sock_release() unconditionally unlinks the socket from the local sockets list. However, if the socket is still in connecting state, it is on the connecting list. Fix this by checking the socket state and unlinking from the correct list.	7.8	More Details
CVE-2026-64051	In the Linux kernel, the following vulnerability has been resolved: accel/qaic: Add overflow check to remap_pfn_range during mmap The call to remap_pfn_range in qaic_gem_object_mmap is susceptible to (re)mapping beyond the VMA if the BO is too large. This can cause use after free issues when munmap() unmaps only the VMA region and not the additional mappings. To prevent this, check the remaining size of the VMA before remapping and truncate the remapped length if sg->length is too large. [jhugo: fix braces from checkpatch --strict]	7.8	More Details
CVE-2026-64008	In the Linux kernel, the following vulnerability has been resolved: accel/rocket: fix UAF via dangling GEM handle in create_bo rocket_ioctl_create_bo() inserts a GEM handle into the file's IDR via drm_gem_handle_create() early on, then performs several operations that can fail (sgt allocation, drm_mm insert, iommu_map). If any fail after the handle is live, the error path calls drm_gem_shmem_object_free() which kfree's the object without removing the handle from the IDR. This leaves a dangling handle pointing to freed slab memory. Any subsequent ioctl using that handle (PREP_BO, FINI_BO, SUBMIT) calls drm_gem_object_lookup() and dereferences freed memory (UAF). Fix by moving drm_gem_handle_create() to after all fallible operations succeed, matching the pattern used by panfrost, lima, and etnaviv. Also fix drm_mm_insert_node_generic() whose return value was silently overwritten by iommu_map_sgtable() on the next line. Add the missing error check. [tomeu: Move handle creation to the very end]	7.8	More Details
CVE-2026-64005	In the Linux kernel, the following vulnerability has been resolved: net/smc: Do not re-initialize smc hashtables INIT_HLIST_HEAD(&smc_v*_hashinfo.ht) are called after smc_nl_init(), proto_register() and sock_register(). This can lead to smc_v*_hashinfo.ht being reset even though hash entries already exist and are being used, possibly resulting in a corrupted list. Remove unnecessary and dangerous re-initialisation of smc_v*_hashinfo.ht in smc_init(); it is implicitly initialised to zero anyhow. Add HLIST_HEAD_INIT to the definitions for clarity.	7.8	More Details
CVE-2026-64004	In the Linux kernel, the following vulnerability has been resolved: net/iucv: fix locking in .getsockopt Mirror iucv_sock_setsockopt() and wrap the whole switch in lock_sock()/release_sock(). The pre-existing SO_MSGLIMIT-only lock becomes redundant and is removed. Any AF_IUCV HIPER user can potentially crash the kernel by racing recvmsg() with getsockopt(SO_MSGSIZE): the SO_MSGSIZE arm dereferences iucv->hs_dev->mtu after iucv_sock_close() (called from the racing recvmsg()) has set hs_dev to NULL, producing a NULL pointer dereference oops.	7.8	More Details
CVE-2026-64002	In the Linux kernel, the following vulnerability has been resolved: ipv4: free net->ipv4.sysctl_local_reserved_ports after unregister_net_sysctl_table() ipv4_sysctl_exit_net() is currently freeing net->ipv4.sysctl_local_reserved_ports too soon. Only after unregister_net_sysctl_table() we can be sure no threads can possibly use the sysctls, including /proc/sys/net/ipv4/ip_local_reserved_ports.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: erofs: fix managed cache race for unaligned extents After unaligned compressed extents were introduced, the following race could occur: [Thread 1] [Thread 2] (z_erofs_fill_bio_vec) <handle a Z_EROFS_PREALLOCATED_FOLIO folio> ... filemap_add_folio (1) (z_erofs_bind_cache) <the same folio is found..> .. .. folio_attach_private (2) filemap_add_folio (3) again Since (1) is executed but (2) hasn't been executed yet, it's possible that another thread finds the same managed		More

2026-64031	folio in z_erofs_bind_cache() for a different pcluster and calls filemap_add_folio() again since folio->private is still Z_EROFS_PREALLOCATED_FOLIO. Fix this by explicitly clearing folio->private before making the folio visible in the managed cache so that another pcluster can simply wait on the locked managed folio as what we did for other shared cases [1]. This only impacts unaligned data compression (`-E48bit` with zstd, for example). [1] Commit 9e2f9d34dd12 ("erofs: handle overlapped pclusters out of crafted images properly") was originally introduced to handle crafted overlapped extents, but it addresses unaligned extents as well.	7.8	Details
CVE-2026-64009	In the Linux kernel, the following vulnerability has been resolved: xfrm: Check for underflow in xfrm_state_mtu Leo Lin reported OOB write issue in esp component: xfrm_state_mtu() returns u32 but performs its arithmetic in unsigned modulo-2^32 space using an attacker-influenced "header_len + authsize + net_adj" subtracted from a small "mtu" argument. A nobody user can install an IPv4 ESP tunnel SA with a large authentication key (XFRMA_ALG_AUTH_TRUNC, e.g. hmac(sha512), 64-byte key, 64-byte trunc), configure a small interface MTU (68 bytes), and set XFRMA_TFCPAD to a large value. When a single UDP datagram is then sent through the tunnel, xfrm_state_mtu() underflows to a near-2^32 value, and esp_output() consumes it as a signed int via: padto = min(x->tfcpad, xfrm_state_mtu(x, mtu_cached)) esp.tfclen = padto - skb->len (assigned to int) esp.tfclen ends up negative (e.g. -207). It is sign-extended to size_t when passed to memset() inside esp_output_fill_trailer(), producing a ~16 EB write of zeroes at skb_tail_pointer(skb). KASAN logs it as "Write of size 18446744073709551537 at addr ffff888...". Check for underflow and return 1. This causes the sendmsg attempt to fail with ENETUNREACH.	7.8	More Details
CVE-2026-16493	A flaw was found in ansible-core. The _extract_collection_from_git() function in ansible-core's concrete_artifact_manager.py constructs git clone commands without a '--' (end-of-options) separator before user-supplied URLs when installing collections from git sources. An attacker who provides a crafted collection source URI containing git argument injection payloads can achieve arbitrary command execution when a user runs 'ansible-galaxy collection install' with the malicious source. This is an incomplete fix for CVE-2026-11332, which hardened the role install path but missed the equivalent collection install code path.	7.8	More Details
CVE-2026-64044	In the Linux kernel, the following vulnerability has been resolved: ovpn: respect peer refcount in CMD_NEW_PEER error path ovpn_nl_peer_new_doit()'s error path calls ovpn_peer_release() directly rather than ovpn_peer_put(), bypassing the kref. The accompanying comment ("peer was not yet hashed, thus it is not used in any context") holds for UDP but not for TCP. For UDP, the ovpn_socket union uses the .ovpn arm and never points back at a peer; UDP encap_recv looks up peers via the not-yet-populated hashtables, so the new peer is unreachable until ovpn_peer_add() publishes it. For TCP, ovpn_socket_new() sets ovpn_sock->peer and ovpn_tcp_socket_attach() publishes ovpn_sock via rcu_assign_sk_user_data(). From that moment until ovpn_socket_release() detaches in the error path, the TCP fd is fully wired: userspace recvmsg / sendmsg / close / poll on the fd, as well as the strparser-driven ovpn_tcp_rcv() path, can reach the peer through sk_user_data -> ovpn_sock->peer and bump its refcount via ovpn_peer_hold(). ovpn_tcp_socket_wait_finish() (called inside ovpn_socket_release()) drains strparser and the tx work, but does not synchronize with userspace syscall callers that already hold a peer reference. If ovpn_nl_peer_modify() or ovpn_peer_add() returns an error while such a caller is in flight - notably an ovpn_tcp_recvmsg() blocked in __skb_recv_datagram() on peer->tcp.user_queue - the direct ovpn_peer_release() destroys the peer while the caller still holds the reference, and the eventual ovpn_peer_put() from that caller operates on freed memory. Replace the direct destructor call with ovpn_peer_put() so the kref correctly defers destruction until the last reference is dropped. In the common case where no concurrent user is present, behaviour is unchanged: the kref hits zero immediately and ovpn_peer_release_kref() runs the same destructor. With this conversion ovpn_peer_release() has no callers outside peer.c - ovpn_peer_release_kref() in the same translation unit is the only remaining user - so make it static and drop its declaration from peer.h.	7.8	More Details
CVE-2026-64036	In the Linux kernel, the following vulnerability has been resolved: cgroup/rstat: validate cpu before css_rstat_cpu() access css_rstat_updated() is exposed as a BPF kfunc and accepts a caller-provided cpu argument. The function uses cpu for per-cpu rstat lookups without checking whether it refers to a valid possible CPU. A BPF iter/cgroup program with CAP_BPF and CAP_PERFMON can pass an invalid cpu value. On an unfixed UBSCAN_BOUNDS test kernel, cpu == 0x7fffffff triggers: UBSAN: array-index-out-of-bounds in kernel/cgroup/rstat.c:31:9 index 2147483647 is out of range for type 'long unsigned int [64]' Call Trace: css_rstat_updated bpf_iter_run_prog cgroup_iter_seq_show bpf_seq_read Add cpu validation to the BPF-facing css_rstat_updated() kfunc and move the common implementation to __css_rstat_updated() for in-kernel callers.	7.8	More Details
CVE-2026-3842	A flaw was found in QEMU. This vulnerability allows a local attacker within a guest virtual machine to write data beyond its allocated memory. This occurs when cpu_physical_memory_map() returns a shorter length than expected, leading to an out-of-bounds write. Successful exploitation could result in unauthorized access to guest memory or corruption of heap-allocated objects, potentially causing information disclosure, data integrity issues, or a denial of service.	7.8	More Details
CVE-2026-64050	In the Linux kernel, the following vulnerability has been resolved: drm/msm/dpu: don't mix devm and drmm functions Mixing devm and drmm functions will result in a use-after-free on msm driver teardown if userspace keeps a reference on the drm device: The WB connector data will be destroyed because of the use of devm_kzalloc()), while the usersoace still can try interacting with the WB connector (which uses drmm_ functions). Change dpu_writeback_init() to use drmm_. Patchwork: https://patchwork.freedesktop.org/patch/722656/	7.8	More Details

CVE-2026-64041	In the Linux kernel, the following vulnerability has been resolved: ASoC: codecs: fs210x: fix possible buffer overflow In fs210x_effect_scene_info(), a string was copied like this: strscpy(DST, SRC, strlen(SRC) + 1); A buffer overflow would happen if strlen(SRC) >= sizeof(DST). Actually, strscpy() must be used this way: strscpy(DST, SRC, sizeof(DST)); strscpy(DST, SRC); // defaults to sizeof(DST)	7.8	More Details
CVE-2026-62201	OpenClaw versions before 2026.6.6 contain a network policy bypass vulnerability in the sandbox exec-server that allows lower-trust callers to reach internal network destinations blocked by OpenClaw policy. Attackers can send HTTP requests through the exec-server to access network resources that should have been restricted by configured policies.	7.7	More Details
CVE-2026-53514	Better Auth is an authentication and authorization library for TypeScript. Prior to 1.6.11, and in 1.6.14 and later when invitation IDs can be obtained outside the invited mailbox and requireEmailVerificationOnInvitation: true is not enabled, the organization plugin's acceptInvitation, rejectInvitation, getInvitation, and listUserInvitations recipient endpoints use session.user.email and an invitation ID without sufficient verified-email ownership proof, allowing a user with an unverified session for the invited email address to accept an organization invitation after obtaining the invitation ID. This issue is fixed for the original default behavior in version 1.6.11, while 1.6.14 restored compatibility for built-in opaque invitation IDs and leaves affected configurations requiring secure options.	7.7	More Details
CVE-2026-45806	Penpot is an open-source design tool for design and code collaboration. Prior to 2.15.0, Penpot's remote image import passed the user-controlled url from frontend/src/app/main/data/workspace/media.cljs into the backend RPC method :create-file-media-object-from-url in backend/src/app/rpc/commands/media.clj, where media/download-image in backend/src/app/media.clj used the shared HTTP client without destination filtering, allowing an authenticated file editor to reach internal-only endpoints. This issue is fixed in version 2.15.0.	7.7	More Details
CVE-2026-48799	Postiz is an AI social media scheduling tool. Prior to 2.21.8, Postiz fails to verify Nowpayments IPN callback authenticity against the payment provider shared secret and reads the target subscription identifier from the untrusted request body, allowing a low-privileged account to grant arbitrary organizations lifetime PRO subscriptions without payment. This issue is fixed in version 2.21.8.	7.7	More Details
CVE-2026-62227	OpenClaw 2026.4.14 before 2026.5.26 contain a server-side request forgery vulnerability in browser snapshot routes that fail to validate post-navigation destinations. Attackers with lower-trust access can bypass OpenClaw policy checks to reach network destinations that should have been blocked.	7.7	More Details
CVE-2026-61835	Directus is a real-time API and App dashboard for managing SQL database content. Prior to 12.0.0, the SSRF protection on Directus's file-import-from-URL feature can be bypassed using the address 0.0.0.0 because api/src/request/is-denied-ip.ts treats 0.0.0.0 as a keyword for local interfaces but never blocks the literal address itself. On Linux and macOS, connecting to 0.0.0.0 reaches localhost, so an authenticated user with file-upload rights can make the server fetch internal services through the /files/import endpoint and retrieve the response as a downloadable file. This issue is fixed in version 12.0.0.	7.7	More Details
CVE-2026-47164	Vaultwarden is a Bitwarden-compatible server written in Rust. Prior to 1.36.0, Vaultwarden's SSO login flow checked the IdP email_verified claim only for new-user creation and not when SSO_SIGNUPS_MATCH_EMAIL=true linked an IdP identity to an existing local account, allowing an attacker-controlled IdP identity asserting a victim email address to bind to and authenticate as that account. This issue is fixed in version 1.36.0.	7.7	More Details
CVE-2026-57077	YAML::Syck versions before 1.47 for Perl allow an out-of-bounds read via an unbounded newline scan in newline_len. In the bundled libsyck newline_len and is_newline dereference the scan pointer, and the following byte for a "\r\n" pair, with no NUL-terminator or bounds check. During block-scalar lexing at a document boundary the scan runs one byte past the heap lexer buffer. This is an incomplete fix of CVE-2025-11683, on a lexer path the earlier fix did not cover. Any caller that runs Load or LoadFile on an untrusted document with a block scalar at a document boundary reaches the over-read.	7.7	More Details
CVE-2026-46555	WhatsApp MCP Server is a Model Context Protocol (MCP) server for WhatsApp, enabling Claude to read and send WhatsApp messages. Prior to version 0.2.1, the `whatsapp-bridge` HTTP API listens on `127.0.0.1:8080` without authentication and without Host header validation, and the `/api/send` endpoint accepts an absolute `media_path` parameter without confining it to a safe directory. Combined, these issues allow any local process running as the same user as the bridge to send WhatsApp messages from the paired account without authorization; the same caller to read arbitrary files readable by the user (e.g. SSH private keys, browser session data, source code, dotfiles) and exfiltrate them as WhatsApp document attachments; and/or a remote attacker to trigger the same operations via DNS rebinding from a webpage the user visits, since no Host header validation is performed. In MCP environments, "local caller" extends beyond processes the user explicitly launched — sibling MCP servers, IDE extensions, and tool-triggered flows running in the user's session can act as the effective caller. This issue is fixed in whatsapp-mcp v0.2.1 and corresponding Docker images / release artifacts. Users should upgrade immediately. The fix introduces bearer token authentication on the bridge HTTP API (configured via environment variable, required on all requests, validated with constant-time comparison); host header allow-list validation to prevent DNS rebinding; and confinement of `media_path` to a configured directory, with rejection of absolute paths outside the root and path traversal sequences. This is a breaking change for clients of the bridge API. For users who cannot immediately upgrade: Stop the bridge, or block loopback access to port 8080, when the bridge is not actively in use; avoid	7.7	More Details

	running the bridge alongside untrusted MCP servers, browser extensions, or other untrusted local processes; avoid browsing untrusted sites while the bridge is running (DNS rebinding mitigation); and/or run the bridge under a dedicated user account or in a sandbox/container with no access to sensitive files.		
CVE-2026-14251	A flaw was found in the OpenShift GitOps operator. The ClusterRole reconciler does not validate resource ownership when reconciling ClusterRole objects. A namespace-scoped Argo CD instance can trigger deletion of a ClusterRole owned by a cluster-scoped Argo CD instance by crafting a name collision, resulting in a denial of service.	7.7	More Details
CVE-2026-61644	FastGPT is a knowledge-based AI application platform. From 4.14.17 until 4.15.0-beta5, the POST /api/core/chat/record/getCollectionQuote endpoint authenticates the caller's chat and collection context, but the initialId center-node lookup is not bound to that authorized context. A low-privileged tenant user can call the endpoint with valid attacker-owned appId, chatId, chatItemDataId, and collectionId values while supplying another tenant's dataset data id as initialId, causing the response to include foreign dataset quote or full-text content. This issue is fixed in version 4.15.0-beta5.	7.7	More Details
CVE-2026-63739	SurrealDB before 3.1.5 contains an arbitrary file read vulnerability in the DEFINE ANALYZER mapper filter that allows database users with EDITOR or OWNER roles to read files accessible to the SurrealDB process. Attackers can specify arbitrary file paths in the mapper filter and retrieve file contents through query error messages when the SURREAL_FILE_ALLOWLIST is empty or not configured.	7.7	More Details
CVE-2026-64039	In the Linux kernel, the following vulnerability has been resolved: drm/msm/snapshot: fix dumping of the unaligned regions The snapshotting code internally aligns data segment to 16 bytes. This works fine for DPU code (where most of the regions are aligned), but fails for snapshotting of the DSI data (because DSI data region is shifted by 4 bytes). Fix the code by removing length alignment and by accurately printing last registers in the region. While reworking the code also fix the 16x memory overallocation in msm_disp_state_dump_regs(). Patchwork: https://patchwork.freedesktop.org/patch/725449/	7.7	More Details
CVE-2026-46987	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: Application Service Level Mgmt). Supported versions that are affected are 13.5 and 24.1. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. While the vulnerability is in Oracle Enterprise Manager Base Platform, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Enterprise Manager Base Platform accessible data. CVSS 3.1 Base Score 7.7 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:N/A:N).	7.7	More Details
CVE-2026-63107	LimeSurvey through 6.17.10 and 7.0.4 contains a server-side request forgery vulnerability in the REST API survey template endpoint that allows authenticated users to cause the server to issue arbitrary HTTP requests by supplying a manipulated Host header. Attackers can exploit the unsanitized use of the HTTP Host header in the getTemplateData() function to reach internal network services, cloud metadata endpoints, and extract sensitive credentials such as IAM tokens from instance metadata services.	7.7	More Details
CVE-2026-45313	Sandboxie-Plus is an open source sandbox-based isolation software for Windows. Prior to 1.17.6, GuiServer::WndHookRegisterSlave in Sandboxie/core/svc/GuiServer.cpp stores attacker-supplied hthread and hproc fields from a GUI_WND_HOOK_REGISTER request without validating that the thread belongs to the sandboxed process or that the function pointer is in the caller address space, and GuiServer::WndHookNotifySlave then calls OpenThread(THREAD_SET_CONTEXT, FALSE, whk->hthread) and QueueUserAPC((PAPCFUNC)whk->hproc, hThread, (ULONG_PTR)req->threadid) as SYSTEM, allowing a sandboxed process to execute arbitrary code in an unsandboxed host process. This issue is fixed in version 1.17.6.	7.7	More Details
CVE-2026-7754	IBM Langflow OSS 1.0.0 through 1.10.0 Langflow 1.9.0 could allow server-side request forgery (SSRF) due to insecure default configuration and incomplete enforcement of the SSRF protection mechanism.	7.7	More Details
CVE-2026-54910	FileBrowser Quantum is a free, self-hosted, web-based file manager. Prior to version 1.4.3-beta, the `subtitlesHandler` endpoint (`GET /api/media/subtitles`) accepts two user-controlled query parameters: `path` and `name`, both of which are used in filesystem operations without sanitization, creating two independent path traversal vectors. The primary vector is the `path` parameter: it is passed directly to `idx.GetRealPath()` without calling `SanitizeUserPath()`, allowing an attacker to escape the storage root and set `parentDir` to any directory on the host. No existing anchor file is required. The secondary vector is the `name` parameter: it is joined with `parentDir` via `filepath.Join(parentDir, name)` without stripping directory components, allowing traversal relative to any resolved `parentDir`. Any authenticated user (regardless of role or permissions) can exploit either vector to read any text file readable by the server process, including `/etc/passwd`, SSH keys, database credentials, and JWT signing keys. Version 1.4.3-beta patches the issue.	7.7	More Details
CVE-2026-	HyperDX before 2.31.0 contains a server-side request forgery vulnerability that allows authenticated team members to direct the server to arbitrary internal destinations by supplying a caller-controlled host parameter to the ClickHouse proxy test endpoint with no URL validation or allowlist enforcement. Attackers	7.7	More Details

63731	can exploit the reflected error responses from the endpoint to disclose internal service response bodies, enabling access to internal APIs, container services, and cloud provider metadata endpoints.		
CVE-2026-63769	Huginn through 2022.08.18 contains a server-side request forgery vulnerability in the fetch_url method of ScenarioImport that allows authenticated users to make arbitrary HTTP requests by submitting crafted URLs. Attackers can probe internal network services, enumerate ports via error signatures, and access cloud metadata endpoints to retrieve sensitive credentials.	7.7	More Details
CVE-2026-46701	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.4.5, the MCP SSE server defaults to an empty secret (`process.env['NETWORK_AI_MCP_SECRET'] ?? ''` at `bin/mcp-server.ts:89`), which causes `_isAuthorized` (`lib/mcp-transport-sse.ts:254`) to return `true` unconditionally for every request — no `Authorization` header is required. Simultaneously, `_handleRequest` sets `Access-Control-Allow-Origin: *` (`lib/mcp-transport-sse.ts:272`) on every response, so a cross-origin browser fetch can read the result without restriction. An unauthenticated attacker who can lure a user to a malicious web page can invoke all 22 exposed MCP tools — including `config_set`, `agent_spawn`, and `blackboard_write` — against a default-configured localhost server. Version 5.4.5 patches the issue.	7.6	More Details
CVE-2026-55544	NextCRM is open-source customer relationship management (CRM) software. In version 0.12.1, the MCP campaign tools expose campaign read and write operations over the network using user-generated Bearer API tokens (`nxtc__...`). The application has an authorization model that restricts normal users to campaigns they created, but multiple MCP campaign handlers ignore the authenticated user ID and query or mutate campaigns only by object ID. As a result, a low-privileged authenticated user with a valid MCP API token can enumerate all campaigns, read campaign details, update or delete campaigns owned by other users, modify campaign templates and steps, and potentially trigger or pause campaign delivery. Version 0.12.2 fixes the issue.	7.6	More Details
CVE-2026-45337	Better Auth is an authentication and authorization library for TypeScript. From 1.6.0 until 1.6.11, the deviceAuthorization plugin treats any authenticated session as the owner of any pending device code because GET /device does not claim the row and POST /device/approve and POST /device/deny short-circuit when userId is unset, allowing an authenticated attacker who learns a valid user_code to bind the polling device to the attacker's account or deny the legitimate flow. This issue is fixed in version 1.6.11.	7.6	More Details
CVE-2026-47414	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. Five label endpoints — `PATCH /workspaces/{workspace_id}/labels/{label_id}`, `DELETE .../labels/{label_id}`, `POST .../issues/{issue_id}/labels/{label_id}`, `DELETE .../issues/{issue_id}/labels/{label_id}`, `GET .../issues/{issue_id}/labels` — gate access on `require_workspace_member(workspace_id)` only and pass URL-supplied `label_id` and `issue_id` straight through to `LabelService` without verifying either belongs to the workspace. PraisonAI Platform version 0.1.4 patches the issue.	7.6	More Details
CVE-2026-63869	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: limit injected antenna index in ieee80211_parse_tx_radiotap When parsing the radiotap header of an injected frame, ieee80211_parse_tx_radiotap() uses the IEEE80211_RADIOTAP_ANTENNA value directly as a shift count: info->control.antennas |= BIT(*iterator.this_arg); *iterator.this_arg is an 8-bit value taken straight from the frame supplied by userspace, so BIT() can be asked to shift by up to 255. That is undefined behaviour on the unsigned long and is reported by UBSAN: UBSAN: shift-out-of-bounds in net/mac80211/tx.c:2174:30 shift exponent 235 is too large for 64-bit type 'unsigned long' Call Trace: ieee80211_parse_tx_radiotap+0xadb/0x1950 net/mac80211/tx.c:2174 ieee80211_monitor_start_xmit+0xb1f/0x1250 net/mac80211/tx.c:2451 ... packet_sendmsg+0x3eb6/0x50f0 net/packet/af_packet.c:3109 info->control.antennas is a 2-bit bitmap (u8 antennas:2), so only antenna indices 0 and 1 can ever be represented. Ignore any larger value instead of shifting out of bounds.	7.6	More Details
CVE-2026-54560	Cloudreve is a self-hosted file management and sharing system. From 4.12.0 until 4.16.1, Cloudreve's OAuth access tokens are issued without the OAuth client_id claim, so the JWT verifier does not load token scopes into request context and RequiredScopes treats the request like non-scoped session authentication, allowing a low-scope OAuth access token to call APIs requiring higher scopes such as file, share, workflow, user setting, WebDAV account, and potentially admin scopes. This issue is fixed in version 4.16.1.	7.6	More Details
CVE-2026-50147	Metabase is an open-source business intelligence and embedded analytics tool. From 1.57.0 until 1.57.19.1, 1.58.14.1, 1.59.10, and 1.60.4, an attacker who can configure a Metabase database connection can read arbitrary files from the Metabase server's filesystem by adding unsafe JDBC parameters to a MySQL or MariaDB connection, causing the driver to read files from the Metabase host and expose the contents through queries against the connected database or through validation error messages. This issue is fixed in versions 1.57.19.1, 1.58.14.1, 1.59.10, and 1.60.4.	7.6	More Details
	The Classic (BR/EDR) L2CAP signaling handlers l2cap_br_conf_req() and l2cap_br_conf_rsp() in subsys/bluetooth/host/classic/l2cap_br.c validated the minimum command size against buf->len (the bytes remaining in the whole received PDU) instead of len (the per-command data length from the L2CAP signaling header). Because multiple signaling commands can be packed into one PDU, buf->len may exceed a command's len. An attacker can send a CONF_REQ command with a header length smaller than the configuration-request structure (e.g. 0), followed by another command so that buf->len still satisfies the		

CVE-2026-10680	check. The check then passes incorrectly and opt_len = len - sizeof(*req) underflows the uint16_t to a near-0xFFFF value. The configuration-option loop, which lacks an opt_len-versus-buf->len guard, then walks far past the end of the pooled ACL receive buffer using net_buf pull primitives that perform no runtime bounds check, producing an out-of-bounds read of host memory and, when the out-of-bounds option bytes encode an MTU or flush-timeout option, an out-of-bounds write. The BR/EDR signaling channel is processed before pairing/encryption and an L2CAP channel to an L0 service such as SDP can be opened without pairing, so an unauthenticated peer within radio range that can establish an ACL connection can trigger the flaw, leading to memory corruption and denial of service (host/device crash). The defect is present in released versions including v4.4.0. The fix validates against len instead of buf->len in both handlers.	7.6	More Details
CVE-2026-52870	The MCP Python SDK, called mcp on PyPI, is a Python implementation of the Model Context Protocol (MCP). From 1.23.0 until 1.27.2, default handlers installed by server.experimental.enable_tasks() for tasks/list, tasks/get, tasks/result, and tasks/cancel operate only on task identifiers without recording the session that created each task, allowing any connected client to enumerate, read results from, consume messages for, or cancel other clients' tasks. This issue is fixed in version 1.27.2.	7.6	More Details
CVE-2026-11575	The PhonePe Payment Solutions WordPress plugin before 3.1.0 does not properly verify the authenticity of incoming payment callbacks: the secret used to validate the callback signature is empty on sites configured through the current setup flow, so the expected signature reduces to an unkeyed hash of the request body that anyone can compute. This allows unauthenticated attackers to forge a payment-success notification and mark unpaid WooCommerce orders as paid without any payment being made.	7.5	More Details
CVE-2024-51316	The Tenda TX9 V22.03.02.20 firmware has a denial of service vulnerability in the update_dev_name function of the file /goform/SetOnlineDevName	7.5	More Details
CVE-2026-47690	MeltanoHub is the source code for hub.meltano.com, the central place for Meltano plugins. Versions of the repo prior to commit 923820de8f64d753951fbbd54f7282a3d5f75173 were vulnerable to exfiltration of `GITHUB_TOKEN` with write permissions to the repository. The vulnerable workflow used pull_request_target, which runs in the context of the base repository with access to secrets. Commit 923820de8f64d753951fbbd54f7282a3d5f75173 fixes the issue. No known workarounds are available.	7.5	More Details
CVE-2026-59117	Integer overflow or wraparound in Windows Terminal allows an unauthorized attacker to execute code over a network.	7.5	More Details
CVE-2026-44891	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, io.netty.handler.codec.stomp.StompSubframeDecoder fails to limit the total number of headers or their cumulative size per frame, and the maxLineLength parameter only restricts individual header lines. An attacker can send a large number of short headers that are accumulated in memory inside DefaultStompHeadersSubframe until the JVM throws an OutOfMemoryError, causing denial of service for servers exposing a STOMP endpoint based on StompSubframeDecoder. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.	7.5	More Details
CVE-2026-56452	Path traversal in the sshd-scp component of Apache MINA SSHD. Apache MINA SSHD is a Java library for client-side and server-side SSH. The implementation of receiving files or directories via SCP did not validate filenames in SCP "C" or "D" commands. A malicious sender could send filenames containing paths, resulting in files to be written in attacker-controlled places. The issue affects only * applications that use no longer supported Apache MINA SSHD versions < 2.0.0 and use the SCP functions to receive files, * or applications using sshd-scp in Apache MINA SSHD >= 2.0.0 to receive files. Applications using Apache MINA SSHD >= 2.0.0 not using sshd-scp are not affected. The issue is fixed in Apache MINA 2.19.0 and 3.0.0-M5. Affected applications are advised to upgrade to these versions.	7.5	More Details
CVE-2025-51678	An issue was discovered in RISC-V PicoRV32 commit 87c89a. A mismatch in the PCPI INSN and memory address can lead to unexpected behavior.	7.5	More Details
CVE-2026-9171	IBM PowerVM Novalink are vulnerable to a denial of service, caused by sending a specially-crafted request. A remote attacker could exploit this vulnerability to cause the server to consume memory resources.	7.5	More Details
CVE-2026-50274	Datadog dd-trace-go is a Go client library for Datadog application performance monitoring, profiling, and security monitoring. Prior to 2.8.1, Datadog tracing libraries that implement W3C baggage propagation parse incoming baggage HTTP headers without enforcing DD_TRACE_BAGGAGE_MAX_ITEMS or DD_TRACE_BAGGAGE_MAX_BYTES limits on the extract path. A remote, unauthenticated attacker can send a request whose baggage header contains an arbitrarily large number of comma-separated key-value pairs or a single very large value, causing unbounded CPU and memory consumption and enabling a remote denial of service against HTTP services with baggage propagation enabled. This issue is fixed in version 2.8.1.	7.5	More Details
CVE-2026-	FileCodeBox before 2.4 contains a rate-limit bypass vulnerability in the IPRateLimit class that allows unauthenticated attackers to circumvent request throttling by supplying attacker-controlled X-Real-IP and X-Forwarded-For headers without verification of trusted reverse proxy origin. Attackers can supply unique	7.5	More Details

64619	spoofed IP values on each request to enumerate all possible share codes and retrieve other users' files without authentication.		
CVE-2026-53383	In the Linux kernel, the following vulnerability has been resolved: ksmbd: reject non-VALID session in compound request branch smb2_check_user_session() takes a shortcut for any operation that is not the first in a COMPOUND request: it reuses work->sess (the session bound by the first operation) and validates only the SessionId, then returns "valid". It never re-checks work->sess->state == SMB2_SESSION_VALID, and a SessionId of 0xFFFFFFFFFFFFFFFF (ULLONG_MAX, the MS-SMB2 related-operation value) skips even the id comparison. The standalone path (ksmbd_session_lookup_all() plus the SESSION_SETUP state machine) does enforce the VALID state; the compound branch bypasses all of it. A SESSION_SETUP carrying only an NTLM Type-1 (NtLmNegotiate) blob publishes a fresh SMB2_SESSION_IN_PROGRESS session whose sess->user is still NULL (->user is assigned later, by ntlm_authenticate()). Used as operation 1 of a COMPOUND with operation 2 = TREE_CONNECT (related, SessionId=ULLONG_MAX, \\host\IPC$), the tree-connect then runs on that IN_PROGRESS session and reaches ksmbd_ipc_tree_connect_request(), which dereferences user_name(sess->user) with sess->user == NULL (transport_ipc.c:687/701/704) -> remote NULL-pointer dereference and a kernel Oops that wedges the ksmbd worker for all clients. Reject any non-first compound operation that lands on a session which is not SMB2_SESSION_VALID, mirroring the validity the standalone lookup path enforces. SESSION_SETUP itself legitimately runs on an IN_PROGRESS session, but it is never carried as a non-first compound operation, so multi-leg authentication is unaffected by this check.	7.5	More Details
CVE-2026-61371	Microsoft AVML before 0.17.0 could follow a symlink when opening a destination output path on Unix, allowing truncation/overwrite of the symlink target. The destructive effect is performed at open-time via O_TRUNC, and can happen before full input validation completes (“truncation-before-validation”).	7.5	More Details
CVE-2026-64194	Net::DNS versions through 1.55 for Perl allow Denial of Service via deep DNS compression pointer chains. Net::DNS::DomainName::decode follows RFC 1035 compression pointers by recursing into itself with no depth limit. It is possible to construct a name which saturates the call stack (at least with larger TCP responses), leading to a potential Denial of Service. The guard `$link < $offset` prevents forward and circular chains, but still allows arbitrarily long backward chains. The per-offset cache (`$cache`) is populated at the start of each call and short-circuits only re-traverses of the same offset - the initial descent through a fresh chain still recurses at full depth. A crafted packet can chain two-byte compression pointers so that each one points two bytes earlier than the previous, producing a chain length of `offset / 2`. For the 14-bit pointer field (max offset 16383) this gives up to ~8191 recursive frames. For a TCP DNS message the limit is the 16-bit length field (~32767 frames). Perl's default C stack handles only a few thousand frames; beyond that the process receives SIGSEGV or similar, which is a denial-of-service for any application parsing untrusted DNS data. The vulnerability is triggered by `Net::DNS::Packet->new(\$wire)` i.e. any point where the library decodes a DNS message from the network.	7.5	More Details
CVE-2026-63770	Glance through 0.8.5 contains an IP address spoofing vulnerability in the authentication handler that allows unauthenticated attackers to bypass brute-force lockout protections by supplying arbitrary values in the X-Forwarded-For request header when the server proxied option is enabled. Attackers can manipulate the leftmost value of the X-Forwarded-For header to make each login attempt appear to originate from a distinct IP address, preventing the per-IP failed-login counter from reaching the lockout threshold and enabling unlimited credential guessing against the authentication endpoint.	7.5	More Details
CVE-2026-44907	A denial of service vulnerability could be triggered by sending specially crafted HTTP requests to server function endpoints, this could lead to excessive CPU usage; affecting the following packages: react-server-dom-webpack, react-server-dom-parcel, react-server-dom-turbopack (versions 19.0.0 through 19.0.7, 19.1.0 through 19.1.8, and 19.2.0 through 19.2.7).	7.5	More Details
CVE-2026-60085	PraisonAI before 4.6.78 contains an unenforced security policy vulnerability in the default Subprocess Sandbox backend where blocked_commands, blocked_paths, blocked_imports, allow_subprocess, and allow_file_write restrictions are completely ignored. Attackers can execute arbitrary subprocess commands, read sensitive files, and perform destructive operations despite explicit security policy configuration.	7.5	More Details
CVE-2026-62386	The Grav API plugin (getgrav/grav-plugin-api) before 1.0.0-rc.16 accepts JWT access tokens through the ?token= URL query parameter on every API route (JwtAuthenticator::extractBearerToken fallback). Because tokens are embedded in URLs, they are logged verbatim in web server access logs, leaked via the Referer header, stored in browser history, and captured by upstream proxy and CDN logs, exposing valid admin access tokens. A leaked token grants unauthorized API access, including reading configuration and user data, creating admin accounts, modifying system settings, and deleting pages.	7.5	More Details
CVE-2026-15322	IBM Engineering AI Hub 1.0.0, 1.1.0, and 1.2.0 could allow a remote attacker to obtain sensitive information due to the exposure of session tokens in URLs.	7.5	More Details
CVE-2026-62230	Grav before 2.0.4 ships a default .htaccess (and reference webserver-configs/htaccess.txt) whose rules blocking access to sensitive file types (.yaml, .php, .json, etc.) lack the [NC] flag, making extension matching case-sensitive. On case-insensitive filesystems (Windows/NTFS, macOS/HFS+, or Docker volume mounts), an unauthenticated attacker can request these files with uppercase or mixed-case extensions (e.g., .YAML, .PHP) to bypass the restrictions and read sensitive configuration files that may contain API keys and credentials.	7.5	More Details

CVE-2026-16445	A flaw was found in dracut. A remote attacker on the adjacent network can exploit this vulnerability by providing specially crafted DHCP options, such as a malicious root-path, next-server, or bootfile name, to a system using dracut's NetworkManager-based initrd network module. These options are improperly handled and written into a temporary shell script without proper escaping, leading to command injection. This allows the attacker to achieve root code execution within the initramfs during system boot.	7.5	More Details
CVE-2026-16405	Information disclosure in the Networking: WebSockets component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	7.5	More Details
CVE-2026-15957	Smithy-RS is a Rust code generation and runtime framework that generates HTTP clients and servers from Smithy interface definitions, powering the AWS SDK for Rust and custom service implementations. Uncontrolled recursion in the JSON, CBOR, and XML deserializer functions emitted by Amazon smithy-rs code generation could allow remote attackers to cause a denial of service (process abort via stack exhaustion) via a small request containing deeply nested data for a recursive model shape to a generated SDK or server. To mitigate this issue, users should upgrade to aws-sdk-rust release-2026-06-02 or later. Users building custom servers with smithy-rs codegen should regenerate from smithy-rs release-2026-06-01 or later.	7.5	More Details
CVE-2026-8082	The bpost-shipping-platform WordPress plugin before 3.2.3 does not properly sanitize a parameter before using it in a SQL query during WooCommerce order submission, allowing unauthenticated attackers to perform time-based blind SQL injection on stores running this bpost-shipping-platform WordPress plugin before 3.2.3.	7.5	More Details
CVE-2026-7872	IBM Langflow OSS 1.0.0 through 1.10.0 allows an authenticated attacker to read arbitrary files including the JWT signing key and forge authentication tokens for any user.	7.5	More Details
CVE-2026-47667	CImg Library is a C++ library for image processing. Prior to version 4.0.0 in `_load_analyze()`, the header_size field is read as an `unsigned int` from the first 4 bytes of an Analyze/NIfTI file and passed directly to `new unsigned char[header_size]` without being bounded against the actual file size. A value up to ~4 GB is accepted. If the subsequent `fread` returns `short` as it will for any malformed file), the function throws a `CImgIOException` and the allocated buffer is never freed. A 6-byte crafted file is sufficient to trigger an allocation of ~1.3 GB per call, with the full allocation leaked on every error path. The issue is reachable via `load_analyze()` and the generic `load()` when the file extension is .hdr, .img, or .nii. Version 4.0.0 fixes the issue.	7.5	More Details
CVE-2026-13765	The LearnPress – WordPress LMS Plugin for Create and Sell Online Courses plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 4.4.1 via the check_answer. This makes it possible for unauthenticated attackers to extract the correct-answer markers, full option lists, explanations, and question content for any quiz question on the site — including questions belonging to paid courses the attacker is not enrolled in.	7.5	More Details
CVE-2026-55833	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, Netty SPDY header decoding continues inflating zlib-compressed header blocks after the raw header parser has exceeded `maxHeaderSize` and marked the frame truncated in `SpdyFrameCodec`, allowing a remote peer to send a small compressed `HEADERS` block that expands into much larger raw header data and causes compression-amplified CPU and allocation churn. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.	7.5	More Details
CVE-2026-50272	dd-trace is the Datadog APM client for Node.js. Prior to 5.100.0, W3C baggage propagation in packages/dd-trace/src/baggage.js and packages/dd-trace/src/opentracing/propagation/text_map.js parsed incoming baggage HTTP headers without enforcing DD_TRACE_BAGGAGE_MAX_ITEMS or DD_TRACE_BAGGAGE_MAX_BYTES on extraction. A remote, unauthenticated attacker can send a request whose baggage header contains an arbitrarily large number of comma-separated key-value pairs, or a single very large value, causing unbounded CPU and memory consumption and enabling a remote denial of service against any HTTP service with baggage propagation enabled. This issue is fixed in version 5.100.0.	7.5	More Details
CVE-2026-59173	Uncontrolled Resource Consumption vulnerability in Apache Traffic Server. This issue affects Apache Traffic Server: from 9.0.0 through 9.1.13, from 10.0.0 through 10.1.2. Users are recommended to upgrade to version 9.1.14 or 10.1.3, which fixes the issue.	7.5	More Details
CVE-2026-55831	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, Netty's SPDY SETTINGS decoder accepts a peer-declared SETTINGS entry count up to the 24-bit frame-length limit and materializes every unique setting ID in `DefaultSpdySettingsFrame`, allowing a remote SPDY/3.1 peer to send a syntactically valid roughly 2 MiB SETTINGS frame that creates 262144 map entries and amplifies network input into heap growth and ordered-map insertion work. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.	7.5	More Details
CVE-2026-51833	Xenforo 2.3.8 is vulnerable to SSRF. Attackers that have administrator privileges or are able to add/save RSS feeds can enumerate internal services (ports) or expose the original IP address of the server.	7.5	More Details

CVE-2026-56740	JLine is a Java library for handling console input. Prior to 3.30.14, 4.0.16, and 4.2.1, the JLine3 Telnet server remote-telnet module does not limit the number of environment variables a client may inject via the Telnet NEW-ENVIRON option, and TelnetIO.readNEVariables() in TelnetIO.java:1127-1180 stores each variable pair in a HashMap held by ConnectionData, allowing an unauthenticated attacker to flood unique variable pairs before the terminating IAC SE byte and exhaust JVM heap memory with an OutOfMemoryError. This issue is fixed in versions 3.30.14, 4.0.16, and 4.2.1.	7.5	More Details
CVE-2026-56741	JLine is a Java library for handling console input. Prior to 3.30.14, 4.0.16, and 4.2.1, the JLine3 Telnet server remote-telnet module does not apply an upper bound to terminal dimensions received via the Telnet NAWS option, and TelnetIO.handleNAWS() in TelnetIO.java:856-879 reads client-supplied width and height as 16-bit unsigned integers and passes values such as 65535x65535 to setTerminalGeometry(), allowing an unauthenticated remote attacker to repeatedly alternate values and trigger continuous expensive rendering work that causes CPU exhaustion and denial of service. This issue is fixed in versions 3.30.14, 4.0.16, and 4.2.1.	7.5	More Details
CVE-2026-65052	Ninja Forms WordPress plugin version 3.14.8 and prior contains an improper input validation vulnerability that allows unauthenticated attackers to inject arbitrary numeric values into form calculations and payment totals by submitting values that do not match any configured option in ListSelect or ListRadio fields. Attackers can tamper with form submission payloads to the ajax submit endpoint, causing the get_calc_value() method to fail open and return attacker-controlled values, enabling manipulation of payment amounts to zero or arbitrary figures and bypassing admin-configured pricing logic.	7.5	More Details
CVE-2026-50271	Datadog dd-trace-py is the Datadog Python APM client. Prior to 4.8.2, Datadog tracing libraries that implement W3C baggage propagation parse incoming baggage HTTP headers without enforcing DD_TRACE_BAGGAGE_MAX_ITEMS or DD_TRACE_BAGGAGE_MAX_BYTES limits on the extract path. A remote, unauthenticated attacker can send a request whose baggage header contains an arbitrarily large number of comma-separated key-value pairs or a single very large value, causing unbounded CPU and memory consumption and enabling a remote denial of service against HTTP services with baggage propagation enabled. This issue is fixed in version 4.8.2.	7.5	More Details
CVE-2026-54340	h2o is an HTTP server with support for HTTP/1.x, HTTP/2 and HTTP/3. Prior to commit 9265bdd, there is an HTTP/2 state amplification issue that combines HPACK decompression amplification with Slowloris-style stream stalling. Amplified decoded header state can be retained by stalled HTTP/2 streams, and depending on the configuration, additional limits are needed to bound decoded header state and prevent attack. This issue has been fixed by commit 9265bdd.	7.5	More Details
CVE-2026-39359	Wazuh is a free and open source platform used for threat prevention, detection, and response. In versions 4.0.0 through 4.10.3 and 4.11.0 through 4.14.4, a logic flaw affects the Wazuh Manager's enrollment daemon (authd) and synchronization daemon (remoted). The authd process allows agents to select a group during enrollment but does not filter path traversal sequences such as "..." While the manager checks for the group directory using wopendir(), the ".." sequence references the parent directory (/var/ossec/etc), allowing it to pass validation. After the malicious group is accepted and stored in the manager's global database, the remoted process uses this unchecked value to build paths for agent configuration synchronization. As a result, sensitive files from /var/ossec/etc, such as client.keys, ossec.conf, and internal certificates, are included in the agent's shared configuration stream and exposed to the attacker. This issue has been fixed in versions 4.10.4 and 4.14.5.	7.5	More Details
CVE-2026-34150	Wazuh is a free and open source platform used for threat prevention, detection, and response. In versions 1.0.0 and above, prior to 4.14.5, a heap buffer overflow in wazuh-analysisd allows an unauthenticated remote attacker to crash the Wazuh manager's analysis engine, causing complete loss of SIEM alert processing. The attack exploits the default configuration shipped in the official wazuh/wazuh-docker deployment with default configuration. An attacker can enroll with authd without a password to obtain a valid agent ID and encryption key, connect to remoted over the Wazuh agent protocol, and inject rootcheck events containing {key: value} patterns longer than 30 bytes that trigger a sprintf overflow of a 30-byte buffer in W_JSON_ParseRootcheck, corrupting the heap and crashing wazuh-analysisd so that all alert processing silently stops while the dashboard and API keep showing stale data.	7.5	More Details
CVE-2024-58368	SurrealDB versions before 1.1.0 fail to properly parse the ID, DB, and NS headers in HTTP REST API requests containing special characters. Unauthenticated attackers can send crafted HTTP requests with malformed header values to trigger an uncaught exception that crashes the server.	7.5	More Details
CVE-2026-44453	h2o is an HTTP server with support for HTTP/1.x, HTTP/2 and HTTP/3. Prior to commit 6b5370d, h2o is vulnerable to a Denial of Service attack when calling alloca under certain conditions. When serving static files, h2o builds the file path on stack, by calling alloca. The maximum size of the memory allocated using alloca can be as huge as ~600KB, which exceeds the default pthread stack size used by musl libc (128KB). If the amount of memory allocated by alloca exceeds the stack size, the h2o server crashes with a segmentation fault, while it tries to touch the guard page. This issue has been fixed by commit 6b5370d.	7.5	More Details
CVE-2026-	When an HTTP/2 profile is configured on a virtual server, undisclosed requests can cause an increase in memory resource utilization. Impact: System performance can degrade until the TMM process is either forced to restart or is manually restarted. This vulnerability allows a remote, unauthenticated attacker to cause a degradation of service that can lead to a denial-of-service (DoS) on the BIG-IP system. There is no	7.5	More Details

59762	control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.		
CVE-2026-44436	Quicly is an IETF QUIC protocol implementation intended primarily for use within the H2O HTTP server. Prior to commit 8b178e6, Quicly is vulnerable to a Denial of Service attack through connection state corruption. In QUIC Invariants, the maximum length of a Connection ID is 255 bytes, while QUIC version 1 further restricts the maximum to 20 bytes. Quicly implements QUIC version 1 and therefore its CID buffers are limited to 20 bytes. However, to be able to respond to unknown versions of QUIC, its packet decoder accepts Connection IDs of up to 255 bytes. As its CID buffers are merely 20 bytes long, Quicly must reject QUIC version 1 packets with Connection IDs longer than that. The command line tool bundled with Quicly has had that check, however the library itself lacked such enforcement. As a consequence, when used by applications that lack their own enforcement, the connection state becoming inconsistent to buffer overrun. Fortunately, the overflow stops within the allocated chunk of memory, but nevertheless, the bug leads to assertion failures. This issue has been fixed by commit 8b178e6.	7.5	More Details
CVE-2026-44435	Quicly is an IETF QUIC protocol implementation intended primarily for use within the H2O HTTP server. Prior to commit 937d0e9, an assertion failure is raised when the total number of valid handshake messages received over a CRYPTO stream of a single packet number space exceeds 32KB, causing a Denial of Service. This issue has been fixed by commit 937d0e9.	7.5	More Details
CVE-2026-51031	FlareSolverr before version 3.4.7 contains a server-side request forgery (SSRF) vulnerability in the /v1 API endpoint. This allows a remote attacker to obtain sensitive information	7.5	More Details
CVE-2026-50151	oras-go is a Go library for managing OCI artifacts. Prior to 2.6.1, registry/remote/repository.go in blobStore.completePushAfterInitialPost follows a registry-controlled Location header during monolithic blob upload and reuses the Authorization header from the initial POST request for the subsequent PUT request, allowing a malicious registry to return a cross-host Location and receive the caller's credentials at an attacker-controlled endpoint. This issue is fixed in version 2.6.1.	7.5	More Details
CVE-2026-56339	Capgo (Cap-go/capgo) before 12.128.2 contains an information disclosure vulnerability in the Supabase PostgREST SECURITY DEFINER RPC function public.rescind_invitation that allows unauthenticated attackers to enumerate organization existence. The function returns distinct error messages (NO_ORG vs NO_RIGHTS) when called with only a publishable API key, enabling attackers to discover valid organization IDs and increase the attack surface for targeted phishing or social engineering campaigns.	7.5	More Details
CVE-2026-45799	Wire provides gRPC and protocol buffers for Android, Kotlin, Swift, and Java. Prior to 6.3.0 and 7.0.0-alpha03, ByteArrayProtoReader32.skipGroup() and ProtoReader.skipGroup() in wire-runtime do not validate that a LENGTH_DELIMITED field length is non-negative before skip(), allowing a crafted protobuf varint encoding -128 as a signed Int to make skip(-128) move the internal position negative and make the next readByte() throw ArrayIndexOutOfBoundsException instead of the documented IOException or ProtocolException, which can crash services using ProtoAdapter.decode(byte[]) on untrusted payloads. This issue is fixed in versions 6.3.0 and 7.0.0-alpha03.	7.5	More Details
CVE-2026-43977	wger is a free, open-source workout and fitness manager. In versions prior to 2.6, any authenticated user can read another user's private workout session notes, exercise history, and training statistics by calling the /logs/ and /stats/ actions on a routine they do not own. The vulnerability exists in RoutineViewSet (wger/manager/api/views.py). The view defines two custom actions /logs/ and /stats/ that are intended to return data for the requesting user's own training history within a routine. However, the underlying permission check (RoutinePermission.has_object_permission) grants read access to any authenticated user when the routine has is_template=True, regardless of ownership. When the /logs/ or /stats/ actions are invoked against a routine the attacker does not own, they return the owner's private workout history, not the attacker's. This issue has been fixed in version 2.6.	7.5	More Details
CVE-2026-53994	ProFTPD mod_sftp contains a heap-based buffer overflow reachable by an authenticated SFTP user. The fxp_packet_read() function accepts the attacker-supplied 32-bit big-endian SFTP packet length without a minimum sanity check. A value of 0 causes an unsigned subtraction elsewhere in the read path to underflow to approximately 4 GB. That oversized request reaches the core memory allocator, where the rounded size is computed in size_t but passed to new_block() as a 32-bit int; the low 32 bits of 0x100000000 are 0, so new_block() returns a small (~512-byte) block while the caller is told it received ~4 GB. The subsequent fill loop then streams attacker-controlled bytes past the end of the 544-byte allocation, producing an attacker-controlled heap buffer overflow. An authenticated user can crash the per-connection ProFTPD session child on demand with a single malformed SFTP packet (packet_len=0 followed by a body greater than approximately 544 bytes), producing reliable authenticated remote denial of service. Depending on heap layout and adjacent allocations, heap metadata corruption and further consequences beyond denial of service may be possible, though only denial of service is demonstrated by the supplied proof of concept.	7.5	More Details
CVE-2026-	WWBN AVideo is an open source video platform. Versions prior to 29.0 expose .env files to unauthenticated users through the official Docker compose configuration. The official docker-compose.yml mounts the entire project root directory as the Apache document root, causing the .env file — which contains database credentials, admin passwords, and infrastructure configuration — to be served as a static file at /.env. No	7.5	More Details

33692	.htaccess rule or Apache configuration blocks access to dotfiles. Exploitation enables direct database access, admin panel takeover, and further lateral movement within the Docker network. This issue has been resolved in version 29.0.		
CVE-2026-32807	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, anyone with a DataLink UUID can fetch the attached text file directly, even if the link is expired, the caller is unauthenticated, or the normal show flow would have denied access. Because the route is public and the mailer embeds the direct file URL, any leaked, forwarded, logged, or stale email link can continue to expose the attachment.	7.5	More Details
CVE-2026-53391	In the Linux kernel, the following vulnerability has been resolved: NFSv4/pNFS: reject zero-length r_addr in nfs4_decode_mp_ds_addr nfs4_decode_mp_ds_addr() decodes the r_netid and r_addr opaques of a netaddr4 from a GETDEVICEINFO multipath-DS body, then immediately calls strrchr(buf, '.') to locate the port separator. Both decodes use xdr_stream_decode_string_dup(), and the current code checks only "nlen < 0" / "rlen < 0" before dereferencing the returned string. When the on-wire opaque has length zero, xdr_stream_decode_opaque_inline() returns 0 and xdr_stream_decode_string_dup() falls through to its "*str = NULL; return ret" tail, leaving buf NULL with a return value of 0. The "< 0" check does not catch this, and the next line is strrchr(NULL, '.'), a kernel NULL pointer dereference reachable from any pNFS-flexfile client mounted against a malicious or compromised metadata server. Reject the zero-length cases explicitly so the decoder fails with -EBADMSG (treated as a malformed GETDEVICEINFO body) instead of panicking the client.	7.5	More Details
CVE-2026-64140	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix null pointer dereference in proc_show_files() When a SMB2 client opens a file with a durable v2 handle and then issues SMB2 SESSION_LOGOFF, session_fd_check() clears fp->tcon = NULL on the reconnectable file pointer but leaves the fp registered in global_ft.idr until the durable scavenger fires (up to fp->durable_timeout seconds later). During that window any read of /proc/fs/ksmbd/files (mode 0400) panics the kernel because proc_show_files() walks global_ft.idr and unconditionally dereferences fp->tcon->id with no NULL guard. Reproducer requires only a successful SMB2 SESSION_SETUP and a share configured with 'durable handles = yes'. KASAN report on mainline 70390501d194: general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] RIP: 0010:proc_show_files+0x118/0x740 Call Trace: proc_show_files+0x118/0x740 seq_read_iter+0x4ef/0xe10 proc_reg_read_iter+0x1b7/0x280 ... Guard the dereference. A durable-disconnected fp legitimately has no tcon; report its tree id as 0 rather than oopsing.	7.5	More Details
CVE-2026-49353	9Router is an AI router & token saver. In 0.4.45 and earlier, 9Router's src/dashboardGuard.js local-only access gate used Host and Origin headers in isLocalRequest() to protect /api/mcp/*, /api/tunnel/*, and /api/cli-tools/*, allowing header spoofing in reverse proxy or tunnel deployments to reach MCP child process stdin paths.	7.5	More Details
CVE-2026-63872	In the Linux kernel, the following vulnerability has been resolved: esp: fix page frag reference leak on skb_to_sgvec failure In esp_output_tail(), when esp->inplace is false, the old skb page frags are replaced with a new page from the xfrm page_frag cache. The source scatterlist (sg) is built from the old frags before the replacement, and esp_ssg_unref() is responsible for releasing the old page references after the crypto operation completes. However, if the second skb_to_sgvec() call (which builds the destination scatterlist from the new page) fails, the code jumps to error_free which only calls kfree(tmp). The old page frag references captured in the source scatterlist are never released: 1. sg[] is built from old frags via skb_to_sgvec() (no extra get_page) 2. nr_frags is set to 1 and frag[0] is replaced with the new page 3. Second skb_to_sgvec() fails -> goto error_free 4. kfree(tmp) frees the sg[] memory but old frags are not unref'd 5. kfree_skb() only releases frag[0] (the new page), not the old ones Fix this by adding a bool parameter to esp_ssg_unref() that, when true, unconditionally unrefs the source scatterlist frags without checking req->src and req->dst, since those fields are not yet initialized by aead_request_set_crypt() at the point of the error. Existing callers pass false to preserve the original behavior. The same issue exists in both esp4 and esp6 as the code is identical.	7.5	More Details
CVE-2026-13397	HTML::Bare versions through 0.04 for Perl will hang in an infinite loop when parsing malformed attributes. The parserc_parse function never advances the attribute-parse state cursor on certain malformed attribute forms, looping forever. Nameless attributes such as "<a ='c'>" or unbalanced quotes "<a b='''''c'>" can trigger this condition. Note that the latest version available on CPAN is version 0.02. Newer versions are available on the git repository.	7.5	More Details
CVE-2026-13401	XML::Bare versions through 0.53 for Perl will hang in an infinite loop when parsing malformed attributes. The parserc_parse function never advances the attribute-parse state cursor on certain malformed attribute forms, looping forever. Nameless attributes such as "<a ='c'>" or unbalanced quotes "<a b='''''c'>" can trigger this condition.	7.5	More Details
CVE-2026-14741	HTTP::Date versions before 6.08 for Perl allow CPU exhaustion via polynomial regex backtracking in parse_date. parse_date() matches the date string against a chain of alternative regexes, and str2time() delegates to it. Several of these patterns place unbounded quantifiers next to each other before a trailing `\s*$` anchor. A valid date prefix followed by a long interior run of digits, letters, or whitespace and a single trailing byte that defeats the final match forces the engine to repartition the run, giving polynomial (about quadratic) backtracking. A header value of a few tens of kilobytes runs for tens of seconds of CPU. HTTP::Date parses timestamps such as HTTP `Date`, `Expires`, and `Last-Modified` headers, which	7.5	More Details

	commonly originate from untrusted sources. Any caller that passes an untrusted date header to str2time() or parse_date() can be driven to consume unbounded CPU, a denial of service.		
CVE-2026-45367	HAPI FHIR is a complete implementation of the HL7 FHIR standard for healthcare interoperability in Java. Prior to 6.9.7, the FHIRPathEngine implementation passes user-controlled regular expressions from matches(), matchesFull(), and replaceMatches() to Java regex operations without effective timeouts, allowing catastrophic backtracking and denial of service. This issue is fixed in version 6.9.7.	7.5	More Details
CVE-2026-45576	zrok is software for sharing web services, files, and network resources. From 0.4.23 until 2.0.3, `zrok2 copy` stores attacker-controlled WebDAV or zrok drive paths such as /../outside.txt in the source inventory and passes them to FilesystemTarget.WriteStream, allowing the sync pipeline to write files outside the selected local filesystem destination root. This issue is fixed in version 2.0.3.	7.5	More Details
CVE-2026-64141	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix null pointer dereference in compare_guid_key() session_fd_check() walks the per-inode m_op_list during durable-handle session teardown and sets op->conn = NULL for every opinfo whose conn matched the closing session's connection. The matching opinfo, however, stays linked in its per-ClientGuid lease_table_list entry's lb->lease_list because destroy_lease_table() only runs on full TCP-connection teardown, not on SESSION_LOGOFF. If the same TCP connection then negotiates a fresh session with the same ClientGuid (ClientGuid is bound to NEGOTIATE, not the session, and is unchanged across LOGOFF + SETUP) and issues a SMB2 CREATE with a lease context on a different inode, find_same_lease_key() walks lb->lease_list, reaches the stale opinfo, and calls compare_guid_key(), which unconditionally dereferences opinfo->conn->ClientGUID. The conn pointer is NULL and the kernel panics. Reproducer requires only a successful SMB2 SESSION_SETUP and a share configured with 'durable handles = yes'. KASAN report on mainline 70390501d194: general protection fault, probably for non-canonical address 0xdffffc0000000069: 0000 [#1] SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000348-0x000000000000034f] Workqueue: ksmbd-io handle_ksmbd_work RIP: 0010:bcmp+0x5b/0x230 Call Trace: compare_guid_key+0x4b/0xd0 find_same_lease_key+0x324/0x690 smb2_open+0x6aea/0x8e60 handle_ksmbd_work+0x796/0xee0 ... Faulting address 0x348 is the offset of ClientGUID within struct ksmbd_conn, confirming opinfo->conn was NULL. Read opinfo->conn once and bail out if it has been cleared by a concurrent session_fd_check(). A half-detached opinfo cannot be the owner of an active lease, so returning 0 is the correct match result.	7.5	More Details
CVE-2026-49212	Symfony UX is a JavaScript ecosystem for Symfony. From 2.8.0 until 2.36.0 and 3.1.0, the HMAC computed by Symfony\UX\LiveComponent\LiveComponentHydrator covered only sorted prop key/value pairs and did not include the component name, the slot identifier (props vs propsFromParent), or request context, allowing a signed blob minted for one component or slot to be replayed in another and set a read-only prop on a target component. This issue is fixed in versions 2.36.0 and 3.1.0.	7.5	More Details
CVE-2026-64148	In the Linux kernel, the following vulnerability has been resolved: pds_core: fix error handling in pdsc_devcmd_wait Fix two cases where pdsc_devcmd_wait() returns stale success from the completion register instead of an error: 1. FW crash: If firmware stops running, the wait loop breaks early with running=false. The condition "if ((!done || timeout) && running)" is false, so error handling is bypassed and stale status is returned. Check !running first and return -ENXIO. 2. Timeout: If a command times out, err is set to -ETIMEDOUT but then overwritten by pdsc_err_to_errno(status) which reads stale status. Return -ETIMEDOUT immediately after cleaning up. Both errors now propagate to pdsc_devcmd_locked() which queues health_work for recovery.	7.5	More Details
CVE-2026-12691	Missing authentication for critical function vulnerability in Vimesoft Inc. Enterprise Video Platform allows Authentication Bypass. This issue affects Enterprise Video Platform: from 3.11.0.0 before 3.25.0.	7.5	More Details
CVE-2026-64175	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mld: stop TX during firmware restart When iwlwifi firmware crashes (e.g., NMI_INTERRUPT_UNKNOWN on Intel BE201/Wi-Fi 7), iwl_mld_nic_error() sets mld->fw_status.in_hw_restart to true. However, iwl_mld_tx_from_txq() does not check this flag before dequeuing frames from mac80211 and pushing them to the transport layer. Since the firmware is dead, iwl_trans_tx() returns -EIO for each frame, which then gets freed immediately. Under high-throughput conditions (e.g., Tailscale UDP traffic or active SSH sessions), this creates a tight dequeue-send-fail-free loop that wastes CPU cycles and generates rapid skb allocation churn, leading to memory pressure from slab fragmentation. The RX path already has this guard (iwl_mld_rx_mpdu checks in_hw_restart at rx.c:1906), and so does the TXQ allocation worker (iwl_mld_add_txqs_wk at tx.c:156). Add the same guard to iwl_mld_tx_from_txq() to stop all TX during firmware restart. Frames left in mac80211's TXQs are naturally drained after restart completes, when queue reallocation triggers iwl_mld_tx_from_txq() via iwl_mld_add_txq_list(), or when new upper-layer traffic invokes wake_tx_queue. Tested on ASUS Zenbook 14 UX3405CA with Intel BE201 (Wi-Fi 7) on kernel 6.19.5 where the firmware crashes approximately every 10-15 minutes under Tailscale traffic.	7.5	More Details
	Meshtastic is an open source mesh networking solution. Prior to version 2.7.23.b246bcd, a single node advertising a User.long_name that contains a malformed character encoding can render other radios unusable over BLE when managed through the iOS app. The malformed name does not need to be maliciously crafted — it can arise from ordinary buffer truncation and has been observed occurring naturally in the wild. At least one code path could place a null terminator in the middle of a multibyte sequence,		

CVE-2026-42566	leaving a malformed User.long_name in the node database. The problem surfaced downstream: the iOS app enforced encoding validation and therefore cannot parse a node database once it contains a poisoned entry. This caused BLE sync to enter a fail/retry loop, resulting in loss of control over the affected device. For a typical user managing their radio with the iOS app, the device becomes effectively unusable until the poisoned node ages out of the on-device database, or unless they have an alternate management path (e.g., the Python CLI, which can be used to identify and remove the offending entries manually). Because the malformed name propagates through the mesh, the temporary presence of a single affected node can degrade BLE management for iOS users across a wide geographical area for an extended period. Less technical users have no straightforward recovery path. Starting in version 2.7.23.b246bcd, the firmware has added input sanitization and regression tests demonstrating recovery for already-poisoned devices. The apps have also taken steps to ensure more graceful handling of malformed encoding sequences as well.	7.5	More Details
CVE-2026-40957	o CVE-2026-40957 is a frameable content vulnerability in the Secure Access server login page prior to 14.55. Attackers with control of a malicious web site could use it to potentially steal credentials from an unwary administrator.	7.5	More Details
CVE-2026-12592	The SlimStat Analytics WordPress plugin before 5.5.0 does not escape a visitor-controlled geolocation value before outputting it in its admin analytics reports, allowing unauthenticated visitors to store a cross-site scripting payload that executes in the browser of an administrator who views the reports. Exploitation requires the SlimStat Analytics WordPress plugin before 5.5.0 to be configured to use the Cloudflare geolocation provider.	7.5	More Details
CVE-2026-6656	Crypt::Password versions through 0.28 for Perl are susceptible to timing attacks. The check_password method uses the built-in eq operator. This allows discrepancies in timing to be used to guess the underlying hash.	7.5	More Details
CVE-2026-64116	In the Linux kernel, the following vulnerability has been resolved: ipv6: ioam: add NULL check for idev in ipv6_hop_ioam() Reported by Sashiko: The function ipv6_hop_ioam() accesses __in6_dev_get(skb->dev)->cnf.ioam6_enabled without validating the returned idev pointer. Because addrconf_ifdown() can concurrently clear dev->ip6_ptr via RCU, __in6_dev_get() can return NULL during interface teardown, which could cause a NULL pointer dereference when processing an IOAM Hop-by-Hop option. Let's add a check and use SKB_DROP_REASON_IPV6DISABLED accordingly.	7.5	More Details
CVE-2026-7189	Insertion of sensitive information into sent data vulnerability in Proliz Software Ltd. Co. Proliz's OBS allows Accessing Functionality Not Properly Constrained by ACLs. This issue affects Proliz's OBS: before v3.6.0.	7.5	More Details
CVE-2026-36590	An issue in EMQ NanoMQ v.0.24.9 allows a remote attacker to cause a denial of service via the nni_qos_db_set function in broker_tcp.c component	7.5	More Details
CVE-2026-53598	Prompty is a markdown file format (.prompty) for LLM prompts. Prior to 2.0.0-beta.2, Prompty loaders expanded ${file:...} references in .prompty frontmatter without enforcing that resolved paths stayed within the prompt directory or allowed roots, allowing an attacker-controlled prompt file to read local files through absolute paths, .. traversal, or symlink escapes. This issue is fixed in versions 2.0.0-beta.2.	7.5	More Details
CVE-2026-63980	In the Linux kernel, the following vulnerability has been resolved: net/handshake: Use spin_lock_bh for hn_lock nvmet_tcp_state_change(), a socket callback that runs in BH context, can reach handshake_req_cancel() via nvmet_tcp_schedule_release_queue() and tls_handshake_cancel(). handshake_req_cancel() acquires hn->hn_lock with plain spin_lock(). If a process-context thread on the same CPU holds hn->hn_lock when a softirq invokes the cancel path, the lock attempt deadlocks. This is the only caller that invokes tls_handshake_cancel() from BH context; every other consumer calls it from process context. Deferring the cancel to process context in the NVMe target is not straightforward: nvmet_tcp_schedule_release_queue() must call tls_handshake_cancel() atomically with its state transition to DISCONNECTING. If the cancel were deferred, the handshake completion callback could fire in the window before the cancel runs, observe the unexpected state, and return without dropping its kref on the queue. Reworking that interlock is considerably more invasive than hardening the handshake lock. Convert all hn->hn_lock acquisitions from spin_lock/spin_unlock to spin_lock_bh/spin_unlock_bh so the lock is never taken with softirqs enabled.	7.5	More Details
CVE-2026-64003	In the Linux kernel, the following vulnerability has been resolved: scsi: core: Run queues for all non-SDEV_DEL devices from scsi_run_host_queues While a SCSI host is in a recovery state, scsi_mq_requeue_cmd() will not set the requeue list for a requeued command to be kicked in the future. The expectation is a call to scsi_run_host_queues() will kick all SCSI devices once the recovery state is cleared. However, scsi_run_host_queues() uses shost_for_each_device() which uses scsi_device_get() and so will ignore devices in a partially removed state like SDEV_CANCEL. But these devices may also have requeued requests, leaving their requests stuck from not being kicked and causing the removal process of the device to hang. scsi_run_host_queues() needs to run against more devices than the macro shost_for_each_device() allows. Instead of using the too limiting scsi_device_get() state checks, only ignore devices in SDEV_DEL state or when unable to acquire a reference. Attempt to run the queues for all other devices when scsi_run_host_queues() is called.	7.5	More Details

CVE-2026-63972	In the Linux kernel, the following vulnerability has been resolved: net: mana: Skip redundant detach on already-detached port When mana_per_port_queue_reset_work_handler() runs after a previous detach succeeded but attach failed, the port is left in a detached state with apc->tx_qp and apc->rxqs already freed. Calling mana_detach() again unconditionally leads to NULL pointer dereferences during queue teardown. Add an early exit in mana_detach() when the port is already in detached state (!netif_device_present) for non-close callers, making it safe to call idempotently. This allows the queue reset handler and other recovery paths to simply retry mana_attach() without redundant teardown.	7.5	More Details
CVE-2026-63968	In the Linux kernel, the following vulnerability has been resolved: ipv6: fix possible infinite loop in fib6_select_path() Found while auditing the same pattern Sashiko reported in rt6_fill_node() [1]. Apply the same fix as commit f8d8ce1b515a ("ipv6: fix possible infinite loop in fib6_info_uses_dev()"). Writers holding tb6_lock can list_del_rcu(&first->fib6_siblings) without waiting for RCU readers; first->fib6_siblings.next then still points into the old ring and this softirq-side walker never reaches &first->fib6_siblings as its terminator. fib6_purge_rt() always WRITE_ONCE()s first->fib6_nsiblings to 0 before list_del_rcu(), so an inside-loop check is a reliable detach signal. [1] https://sashiko.dev/#/patchset/20260526020227.4857-1-jiayuan.chen%40linux.dev	7.5	More Details
CVE-2026-63955	In the Linux kernel, the following vulnerability has been resolved: mm/vmalloc: do not trigger BUG() on BH disabled context __get_vm_area_node() currently triggers a BUG() if in_interrupt() returns true. However, in_interrupt() also reports true when BH are disabled. The bridge code can call rhashtable_lookup_insert_fast() with bottom halves disabled: __vlan_add() -> br_fdb_add_local() spin_lock_bh(&br->hash_lock); <-- Disable BH -> fdb_add_local() -> fdb_create() -> rhashtable_lookup_insert_fast() -> kvmalloc() -> vmalloc() -> __get_vm_area_node() -> BUG_ON(in_interrupt()) spin_unlock_bh(&br->hash_lock) this triggers the BUG() despite the caller not being in NMI or hard IRQ context. Replace the in_interrupt() check with in_nmi() || in_hardirq().	7.5	More Details
CVE-2026-64020	In the Linux kernel, the following vulnerability has been resolved: nvme-pci: fix dma_vecs leak on p2p memory We don't unmap P2P memory, so we don't need to track it. The dma_vec allocation was getting leaked on the completion.	7.5	More Details
CVE-2026-21729	Loki queries with large limits can cause large memory allocations which can impact the availability of the service, depending on its deployment strategy.	7.5	More Details
CVE-2026-12753	The Advance Product Search- Voice & Ajax Search for WooCommerce plugin for WordPress is vulnerable to generic SQL Injection via the 's' and 'match' parameter in all versions up to, and including, 1.4.4 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-64048	In the Linux kernel, the following vulnerability has been resolved: net/smc: reject CHID-0 ACCEPT that matches an empty ism_dev slot On the SMC-D client, slot 0 of ini->ism_dev[]/ini->ism_chid[] is reserved for an SMC-Dv1 device. smc_find_ism_v2_device_clnt() populates V2 entries starting at index 1, so when no V1 device is selected slot 0 is left in its kzalloc()'ed state with ism_dev[0] == NULL and ism_chid[0] == 0. smc_v2_determine_accepted_chid() then matches the peer's CHID against the array starting from index 0 using the CHID alone. A malicious peer replying to a SMC-Dv2-only proposal with d1.chid == 0 matches the empty slot, ini->ism_selected becomes 0, and the subsequent ism_dev[0]->lgr_lock dereference in smc_conn_create() faults at offsetof(struct smcd_dev, lgr_lock) == 0x68: BUG: KASAN: null-ptr-deref in _raw_spin_lock_bh+0x79/0xe0 Write of size 4 at addr 0000000000000068 by task exploit/144 Call Trace: _raw_spin_lock_bh smc_conn_create (net/smc/smc_core.c:1997) __smc_connect (net/smc/af_smc.c:1447) smc_connect (net/smc/af_smc.c:1720) __sys_connect __x64_sys_connect do_syscall_64 Require ism_dev[i] to be non-NULL before accepting a CHID match.	7.5	More Details
CVE-2026-48863	A flaw was found in libsolv. A stack-based buffer overflow vulnerability exists in the PGP verification component due to incorrect length handling when copying EdDSA 's' MPI into a stack buffer. A remote attacker could craft a malicious Ed25519 PGP signature with mismatched MPI lengths. Processing this crafted signature could lead to a denial of service in automated package or repository processing workflows.	7.5	More Details
CVE-2026-23538	A vulnerability was identified in the Feast Feature Server's `/ws/chat` endpoint that allows remote attackers to establish persistent WebSocket connections without any authentication. By opening a large number of simultaneous connections, an attacker can exhaust server resources—such as memory, CPU, and file descriptors—leading to a complete denial of service for legitimate users.	7.5	More Details
CVE-2026-7488	Insertion of sensitive information into sent data vulnerability in IKAS Technology Inc. E-Commerce allows Retrieve Embedded Sensitive Data. This issue affects E-Commerce: through 03062026.	7.5	More Details
CVE-2026-59861	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.0, Kiota's Ruby generator embedded OpenAPI default fields, property names, and other schema-derived strings through CodeMethodWriter.cs and SanitizeForQuotedLiteral() in Writers/StringExtensions.cs into Ruby double-quoted literals without escaping #, allowing attacker-controlled #{expr}, #$var, or #@var interpolation markers to inject arbitrary Ruby code	7.5	More Details

	into generated model classes. This issue is fixed in version 1.32.0.		
CVE-2026-59862	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.0, Kiota's Python generator let attacker-controlled enum value descriptions from x-ms-enum.values[].description flow through KiotaBuilder.SetEnumOptions into Documentation.DescriptionTemplate and PythonConventionService.RemoveInvalidDescriptionCharacters without newline sanitization, allowing generated inline comments to split and execute attacker-controlled Python code at module scope when generated modules were imported. This issue is fixed in version 1.32.0.	7.5	More Details
CVE-2026-8396	Improper restriction of XML external entity reference vulnerability in Netcad Software Inc. NetGIS allows Serialized Data External Linking. This issue affects NetGIS: from 5.0.66 before 7.2.2.	7.5	More Details
CVE-2026-53392	In the Linux kernel, the following vulnerability has been resolved: NFSv4/flexfiles: reject zero filehandle version count ff_layout_alloc_lseg() decodes the filehandle-version array count from the flexfiles layout body. The value is used as the count for kzalloc_objs(), and the current code only rejects NULL. A zero count yields ZERO_SIZE_PTR, which can be stored in dss_info->fh_versions even though later flexfiles paths assume that at least one filehandle version exists. Reject fh_count == 0 before the allocation, matching the existing zero version_count validation in the flexfiles GETDEVICEINFO parser. A QEMU/KASAN run with a malformed flexfiles layout hit: KASAN: null-ptr-deref in range [0x0000000000000010-0x0000000000000017] RIP: 0010:ff_layout_encode_ff_layoutupdate.isra.0+0x15f/0x750 ff_layout_encode_layoutreturn+0x683/0x970 nfs4_xdr_enc_layoutreturn+0x278/0x3a0 Kernel panic - not syncing: Fatal exception The patched kernel rejects the malformed layout without KASAN/oops/panic, and a valid fh_count=1 regression still opens, reads, and unmounts cleanly.	7.5	More Details
CVE-2026-49211	Symfony UX is a JavaScript ecosystem for Symfony. From 2.2.0 until 2.36.0 and 3.1.0, Symfony\UX\Autocomplete\Doctrine\EntitySearchUtil::addSearchClause() builds the LIKE expression used by the autocomplete endpoint by wrapping the client-supplied query in %...% without escaping SQL LIKE wildcards (%, _, \), allowing unauthenticated users to turn the public BaseEntityAutocompleteType endpoint into a broad matcher or blind boolean oracle against every column in default searchable_fields. This issue is fixed in versions 2.36.0 and 3.1.0.	7.5	More Details
CVE-2026-46941	Vulnerability in the Oracle Cost Management product of Oracle E-Business Suite (component: Cost Maintenance). Supported versions that are affected are 12.2.3-12.2.15. Difficult to exploit vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Cost Management. Successful attacks of this vulnerability can result in takeover of Oracle Cost Management. CVSS 3.1 Base Score 7.5 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:L/UI:N/S:U/C:H/I:H/A:H).	7.5	More Details
CVE-2026-32820	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, the documentation and static markdown renderer accepts attacker-controlled path segments and only runs them through the Rails HTML sanitizer, which does not remove directory traversal sequences. An unauthenticated attacker can traverse out of the intended `docs` or `static` directories and render arbitrary `.md` files from the application root or engine root. This is patched in version 26.06.08.	7.5	More Details
CVE-2026-52746	JSONata is a JSON query and transformation language. Prior to 2.2.0, malicious non-matching inputs to the $toMillis function can cause superlinear backtracking in the ISO-8601 validation regex, leading to denial of service in applications that evaluate user-provided JSONata expressions. This issue is fixed in version 2.2.0.	7.5	More Details
CVE-2026-15352	A vulnerability exists in the Health & Safety (HS) application of NASA's Core Flight System (cFS). The flaw allows the application to crash via segmentation fault when processing a routine Housekeeping Telemetry request, leading to denial of service.	7.5	More Details
CVE-2026-54538	xrdp is an open source RDP server. In versions 0.10.6 and prior, a n issue was discovered where the software fails to properly validate the totalLength field within the RDP protocol control header during packet reception. An unauthenticated remote attacker can exploit this vulnerability by sending a specially crafted packet that forces the xrdp process or thread into an infinite, CPU-bound loop. Because the internal pointer fails to advance and the deadlock prevention mechanism is bypassed for specific protocol data unit types, the process consumes excessive CPU resources indefinitely. This can render the xrdp service unavailable and potentially lead to system-wide resource exhaustion if multiple malicious connections are established. This issue has been fixed in version 0.10.6.1.	7.5	More Details
CVE-2026-59955	Apollo is a reliable configuration management system suitable for microservice configuration management scenarios. Prior to 2.5.2, Apollo ConfigService may allow unauthorized access to raw configuration data when AccessKey or management key authentication is enabled because requests under /configfiles/raw/{appId}/{clusterName}/{namespace} are parsed for authentication as appId raw instead of the actual path appId, causing ConfigService to look up AccessKey secrets for raw before verifying the request signature and potentially continue without signature verification for the target appId. This issue is fixed in version 2.5.2.	7.5	More Details
	Apollo is a reliable configuration management system suitable for microservice configuration management		

CVE-2026-59954	scenarios. Prior to 2.5.2, Apollo ConfigService may allow unauthorized access to configuration data when AccessKey or management key authentication is enabled because ConfigService can accept a non-canonical appId variant during authentication while downstream request handling resolves it to the protected app, including accent variants under accent-insensitive collations or trailing-space variants under PAD SPACE collations on /configs and /configfiles endpoints. This issue is fixed in version 2.5.2.	7.5	More Details
CVE-2026-48812	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.221, FreeScout's attachment download route skips token authentication for any attachment whose `token_type` is set to `1` (`TOKEN_TYPE_LEGACY`). Because this route is unauthenticated and the file path is deterministic, an unauthenticated remote attacker can download any attachment that was created by an older version of FreeScout without possessing a valid token or session. Version 1.8.221 contains a fix.	7.5	More Details
CVE-2026-45804	Diffusers is the a library for pretrained diffusion models. Prior to 0.38.0, Diffusers' DiffusionPipeline.from_pretrained flow can bypass the trust_remote_code guard because download() validates model_index.json and custom pipeline code before later loading from a cached folder that can change, allowing a Hub repository with custom .py pipeline code to execute through the custom pipeline flow without passing custom_pipeline or trust_remote_code=True. This issue is fixed in version 0.38.0.	7.5	More Details
CVE-2026-45793	Composer is a dependency Manager for the PHP language. Prior to 1.10.28, 2.2.28, and 2.9.8, Composer\IO\BaseIO::loadConfiguration() validates GitHub OAuth tokens with the regex ^[.A-Za-z0-9_]+$ and interpolates rejected tokens into an UnexpectedValueException; GitHub Actions GITHUB_TOKEN values using the ghs_<id>_<base64url-JWT> format can contain -, fail validation, and be disclosed to stderr or CI logs. This issue is fixed in versions 1.10.28, 2.2.28, and 2.9.8.	7.5	More Details
CVE-2026-20187	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20187 are related to improper handling of exceptional conditions that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-703.	7.5	More Details
CVE-2026-20158	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20158 are related to improper control of a resource through its lifetime that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-664.	7.5	More Details
CVE-2026-20157	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20157 are related to missing encryption that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-311.	7.5	More Details
CVE-2026-62309	CoreDNS is a DNS server written in Go. Prior to 1.14.4, a single 28-byte UDP datagram can crash the CoreDNS process when the proxyproto plugin is enabled because plugin/pkg/proxyproto/proxyproto.go PacketConn.ReadFrom handles a PROXY v2 header with non-UDP transport such as family byte 0x11, reassigns addr from a nil readFrom result after parseProxyProtocol errors, and calls addr.String() in the warning log before ServeDNS recovery applies. This issue is fixed in version 1.14.4.	7.5	More Details
CVE-2026-20153	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco RoomOS engineering team has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20153 are related to improper input validation that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-20.	7.5	More Details
CVE-2026-64612	A flaw was found in libcupsfilters and cups-filters. The PNG image reading function creates a libpng reader without installing an error recovery handler, causing the CUPS image filter process to abort when processing a malformed PNG file. An unauthenticated attacker could exploit this by submitting a specially crafted PNG print job, leading to denial of service of the in-flight print job.	7.5	More Details
CVE-2026-53397	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix posix_acl leak on SETACL decode failure nfsaclsvc_decode_setaclargs() and nfs3svc_decode_setaclargs() each call nfs_stream_decode_acl() twice, first for NFS_ACL and then for NFS_DFACL. Each successful call transfers ownership of a freshly allocated posix_acl into argp->acl_access or argp->acl_default. If the first call succeeds but the second fails, the decoder returns false and argp->acl_access is left dangling. ACLPROC2_SETACL.pc_release was wired to nfssvc_release_attrstat and ACLPROC3_SETACL.pc_release was wired to nfs3svc_release_fhandle. Both only call fh_put() and have no knowledge of the ACL fields on argp. The posix_acl_release() pairs sat at the out: labels inside nfsacld_proc_setacl() and nfsd3_proc_setacl(), but svc_process() skips pc_func when pc_decode returns false, so that cleanup is unreachable on decode failure: svc_process_common() pc_decode() /* decode_setaclargs: false */ /* pc_func skipped */ pc_release() /* fh_put only -- ACLs leaked */ The orphaned posix_acl is leaked for the lifetime of the server. Fix by adding nfsaclsvc_release_setacl() and nfs3svc_release_setacl(), which release both argp->acl_access and argp->acl_default in addition to fh_put(), and wiring them as pc_release for their respective SETACL procedures. pc_release runs on every path	7.5	More Details

	svc_process() takes after decode, including decode failure, so the posix_acl_release() pairs are removed from the proc functions' out: labels to keep ownership in one place. This matches the existing release_getacl() pattern used by the sibling GETACL procedures.		
CVE-2026-53395	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix dead ACL conflict guard in nfsd4_create nfsd4_create() steals create->cr_dpacl/cr_pacl into the local nfsd_attrs via the designated initializer, then immediately sets the source pointers to NULL. The subsequent conflict guard tests the already-nilled source fields, making it permanently dead code: if (create->cr_acl) { if (create->cr_dpacl || create->cr_pacl) /* always false */ When a client encodes both FATTR4_WORD0_ACL and FATTR4_WORD2_POSIX_{DEFAULT,ACCESS}_ACL in the same CREATE fattr bitmap, nfsd4_acl_to_attr() overwrites attrs.na_pacl/na_dpacl without releasing the originals, leaking two posix_acl slab objects per request. Repeated requests cause unbounded slab exhaustion. Fix by checking attrs.na_dpacl/na_pacl (the stolen values) instead of the nilled create->cr_dpacl/cr_pacl, matching the correct pattern already used in nfsd4_setattr().	7.5	More Details
CVE-2026-53394	In the Linux kernel, the following vulnerability has been resolved: nfsd: avoid leaking pre-allocated openowner on unconfirmed retry race When find_or_alloc_open_stateowner() encounters an unconfirmed owner, it calls release_openowner() and sets oo = NULL. Control then falls through past the `if (oo)` guard -- which would have freed any pre-allocated `new` -- and unconditionally executes `new = alloc_stateowner(...)`. If `new` was already allocated on a prior iteration, the pointer is silently overwritten and the previous allocation (slab object + owner name buffer) is leaked. This requires a race: two NFSv4.0 OPEN threads with the same owner string, where a concurrent thread inserts a new unconfirmed owner into the hash between retry iterations. The window is narrow but repeatable under adversarial conditions. Fix by adding `goto retry` after `oo = NULL` so the already-allocated `new` is reused on the next iteration rather than overwritten.	7.5	More Details
CVE-2026-62351	TDengine is a time-series database optimized for Internet of Things devices. Prior to 3.4.1.15, source/libs/transport/src/transComm.c transDecompressMsg() read STransCompMsg.contLen when pHead->comp == 1 without first validating that the RPC packet contained the 8-byte STransCompMsg structure, causing an unauthenticated out-of-bounds read, uncontrolled allocation, integer underflow, and server crash. This issue is fixed in version 3.4.1.15.	7.5	More Details
CVE-2026-49485	HAPI FHIR is a complete implementation of the HL7 FHIR standard for healthcare interoperability in Java. Prior to 6.9.9 and 6.9.4.2, all implementations of FHIRPathEngine accept arbitrary FHIRPath expressions and evaluate them without input validation, and the FHIRPath functions matches(), matchesFull(), and replaceMatches() pass user-controlled regular expressions to Java's Pattern.compile() and String.replaceAll() through an incomplete timeout utility. An attacker can send a resource containing an evil regex pattern that causes catastrophic backtracking, exhausting CPU resources and causing denial of service in the FHIR Validator HTTP endpoint and affected org.hl7.fhir.* modules. This issue is fixed in versions 6.9.9 and 6.9.4.2.	7.5	More Details
CVE-2026-32806	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, any authenticated user can request arbitrary partials or helper-backed render functions through /remote_render. The endpoint does not restrict which partial can be rendered and does not apply controller-specific authorization before rendering the selected view. This enables a low-privileged user to retrieve server-side rendered admin content that is otherwise hidden by navigation and route checks. On the test instance, a Standard user was able to retrieve the PostgreSQL admin dashboard stats even though /admin itself redirected away. This is patched in 26.06.08.	7.5	More Details
CVE-2026-16221	Impact: fast-uri versions from 2.3.1 through 4.1.0 (including the 3.x line up to 3.1.3 and the 2.x line up to 2.4.2) do not treat a literal backslash character (U+005C) as an authority delimiter. Node's native WHATWG URL parser, used by fetch, undici, and Node's http and https clients, normalizes the backslash to a forward slash for special schemes such as http, https, ws, wss, ftp, and file. As a result, the two parsers extract different hosts from the same input string. Applications that use fast-uri to enforce host-based policy such as allowlists, denylists, loopback or SSRF filtering, redirect validation, or outbound proxy routing before passing the same URL into Node's URL or fetch consumers can be steered to an unintended destination, including cloud metadata endpoints, loopback, or internal hosts. Patches: upgrade to fast-uri 4.1.1, 3.1.4, or 2.4.3. Workarounds: none.	7.5	More Details
CVE-2026-63747	SurrealDB versions before 3.1.0 contain a denial of service vulnerability in the RPC use handler that panics when db is set without a namespace. Unauthenticated attackers can send a malformed WebSocket message to the /rpc endpoint to crash the server process.	7.5	More Details
CVE-2026-63760	SurrealDB before 3.1.0 fails to enforce the configured recursion depth limit in the value and JSON parser when processing nested braces, brackets, or parentheses. Unauthenticated attackers can send deeply nested JSON payloads to the WebSocket /rpc endpoint to exhaust server memory and crash the process.	7.5	More Details
CVE-2026-63101	Open Event Server through 1.19.1 contains a missing authentication vulnerability that allows unauthenticated attackers to export the complete member roster of any group, including email addresses, names, join dates, and roles, by submitting requests to the group followers CSV export endpoint which lacks any authentication decorator. Attackers can enumerate sequential group IDs via brute-force, trigger an export via the unauthenticated POST endpoint, then poll the unauthenticated task status endpoint until	7.5	More Details

	completion to retrieve a download URL containing the full member CSV.		
CVE-2026-64622	Network-AI (npm: network-ai) versions 5.12.2 through 5.13.3 fail to apply the configured authorization check (checkAuth/secret) to the ApprovalInbox GET read routes, so even when an operator configures a secret, unauthenticated actors can access sensitive approval request details. The GET /approvals/?status=all, GET /approvals/:id, GET /approvals/stats, and GET /approvals/sse routes disclose full ApprovalEntry content including action/target shell-command strings, file paths, justifications, and risk levels. All responses also carry a hardcoded Access-Control-Allow-Origin: * header, enabling cross-origin disclosure from any website the operator visits. This is an incomplete fix for GHSA-mxjx-28vx-xjjj.	7.5	More Details
CVE-2026-12997	The Gravity Forms plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 2.10.4 via the 'gform_uploaded_files' parameter parameter. This makes it possible for unauthenticated attackers to read the contents of arbitrary files on the server, which can contain sensitive information. Exploitation requires the targeted form to not enforce login (so publicly accessible), which allows the unauthenticated attacker to reach the process_send_resume_link endpoint and supply an arbitrary recipient email address to receive the traversal-retrieved file as a notification attachment.	7.5	More Details
CVE-2026-62389	ws before 8.21.1 contains a memory exhaustion vulnerability in lib/receiver.js where the fragment guard only triggers when fragment count reaches maxFragments, allowing attackers to exhaust memory by sending incomplete fragmented WebSocket messages. Attackers can send a text frame with FIN=0 followed by continuation frames without completing the sequence, causing each fragment to be stored as a separate Buffer object with significant overhead, enabling denial of service through heap exhaustion.	7.5	More Details
CVE-2026-50273	Datadog .NET Tracer is a client library for Datadog APM for .NET applications. Prior to 3.43.0, Datadog tracing libraries that implement W3C baggage propagation parse incoming baggage HTTP headers without enforcing DD_TRACE_BAGGAGE_MAX_ITEMS or DD_TRACE_BAGGAGE_MAX_BYTES on extraction, allowing a remote unauthenticated attacker to send a baggage header with many comma-separated key-value pairs or one very large value and cause unbounded CPU and memory consumption in services with baggage propagation enabled. This issue is fixed in version 3.43.0.	7.5	More Details
CVE-2026-35287	Vulnerability in Oracle Application Testing Suite. The supported version that is affected is 13.3.0.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via TCP to compromise Oracle Application Testing Suite. Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Application Testing Suite accessible data. CVSS 3.1 Base Score 7.5 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N).	7.5	More Details
CVE-2026-52203	An issue in MCMS v.6.1.1 allows a remote attacker to obtain sensitive information via the source parameter.	7.5	More Details
CVE-2026-45713	Mailpit is an email testing tool and API for developers. Prior to version 1.30.0, the Mailpit SMTP server has a Server.MaxSize int field that controls the maximum allowed DATA payload size, but the field is never assigned anywhere outside test code, leaving it at Go's zero value (0 ⇒ "no limit"). The same applies to the HTTP /api/v1/send endpoint, whose request body is decoded with json.NewDecoder(r.Body) and no http.MaxBytesReader. Because Mailpit's default listeners bind [::]:1025 (SMTP) and [::]:8025 (HTTP), with no authentication required on either, a single network-reachable attacker can push an arbitrarily large message into Mailpit and watch RAM consumption spike with a ~7-10× amplification factor (raw frame → enmime envelope tree → search-text index → zstd-encoded write to SQLite). Repeating the attack — or running it concurrently from multiple connections — drives the process to OOM-kill. Version 1.30.0 contains a patch.	7.5	More Details
CVE-2026-46513	Frogman provides headless PBX control through MCP and HTTP API. Prior to 1.6.2, Frogman stored API tokens generated by Tools/CreateApiToken.php:33-36 as raw bin2hex(random_bytes(32)) strings in oc_api_tokens, and Frogman.class.php:78 authenticated the X-Frogman-Token header by comparing it with the stored raw value, allowing database read access to recover reusable active tokens at their assigned permission level, including admin. This issue is fixed in version 1.6.2.	7.4	More Details
CVE-2026-46943	Vulnerability in the Oracle Retail EFTLink product of Oracle Retail Applications (component: Core/Plugin). Supported versions that are affected are 21.0.0-25.0.0. Difficult to exploit vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle Retail EFTLink. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Retail EFTLink accessible data as well as unauthorized access to critical data or complete access to all Oracle Retail EFTLink accessible data. CVSS 3.1 Base Score 7.4 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:N).	7.4	More Details
CVE-2026-62232	Grav before 2.0.4 contains a two-factor authentication bypass vulnerability in the login plugin where the regenerate2FASecret task checks only user existence, not authorization, during the pending TOTP challenge window. Attackers who know the victim's password can call this task without a CSRF nonce to overwrite the 2FA secret with an attacker-chosen value, compute a valid TOTP code, and complete authentication while reducing 2FA to password-only protection.	7.4	More Details
	Open WebUI before 0.9.5 contains a stored cross-site scripting vulnerability in the OAuth authentication flow where the picture claim URL MIME type is inferred from file extension rather than Content-Type header,		

CVE-2026-56398	allowing SVG files to bypass the profile image validator and be stored as data URIs. Authenticated users who visit the profile image endpoint receive attacker-controlled SVG content with inline disposition and no default security headers, enabling script execution in the same origin to steal authentication tokens and achieve account takeover.	7.3	More Details
CVE-2026-47685	FOG is a free open-source cloning/imaging/rescue suite/inventory management system. Prior to versions 1.5.10.1832 and 1.6.0-beta.2313, the unauthenticated inventory service endpoint (`/service/inventory.php`) persists client-supplied values without sanitization, and the Host Management Inventory page renders all static inventory fields into HTML without output encoding, allowing stored cross-site scripting that executes in any administrator's browser. Versions 1.5.10.1832 and 1.6.0-beta.2313 fix the issue.	7.3	More Details
CVE-2026-47687	FOG is a free open-source cloning/imaging/rescue suite/inventory management system. Prior to versions 1.5.10.1832 and 1.6.0-beta.2313, the `selectForm()` helper in `fogpage.class.php` renders `<option>` labels using raw, unescaped user input. An unauthenticated attacker who knows any registered host's MAC address can POST a malicious `sysproduct` value to `/service/inventory.php`, which is stored in the database. When an administrator opens Reports > Inventory, the payload breaks out of the `<option>` element and executes arbitrary JavaScript in the admin's browser. Versions 1.5.10.1832 and 1.6.0-beta.2313 fix the issue.	7.3	More Details
CVE-2026-16014	A vulnerability was found in code-projects Hospital Bed Management System 1.0. This affects an unknown part of the component Login Form. Performing a manipulation of the argument Username results in sql injection. Remote exploitation of the attack is possible. The exploit has been made public and could be used.	7.3	More Details
CVE-2026-63358	FileGator accepts arbitrary Unix permission values via the '/chmoditems' API endpoint and passes the value directly to PHP's native 'chmod()' function through 'octdec()' conversion, with no validation. This allows an authenticated user with 'chmod' permission to upgrade their privileges to root.	7.3	More Details
CVE-2026-16484	A flaw has been found in SourceCodester Class and Exam Timetabling System 1.0. Affected by this vulnerability is an unknown functionality of the file /edit_subjecta.php. This manipulation of the argument ID causes sql injection. The attack can be initiated remotely. The exploit has been published and may be used.	7.3	More Details
CVE-2024-32386	Directory traversal vulnerability in Kerlink Kerlink Wirnet iStation 868 KerOS v.4.3.3_20200803132042 allows a remote attacker to obtain sensitive information via the SNMP update mechanism.	7.3	More Details
CVE-2026-16209	A vulnerability has been found in Gerapy up to 0.9.13. The impacted element is an unknown function of the file gerapy/server/core/views.py of the component Project Upload Endpoint. Such manipulation leads to missing authentication. The attack may be launched remotely. The exploit has been disclosed to the public and may be used. The name of the patch is bd4891c60315f17611a3b7a651ffe0fba7cfe71e. Applying a patch is advised to resolve this issue.	7.3	More Details
CVE-2026-16200	A vulnerability has been found in zevorn rt-claw up to 0.2.0. This impacts the function claw_tool_invoke of the file claw/services/swarm/swarm.c of the component RPC Handler. The manipulation leads to incorrect authorization. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16210	A vulnerability was found in newpanjing simpleui 2026.01.13. This affects the function self.get_action of the file simpleui/admin.py of the component AjaxAdmin AJAX Endpoint. Performing a manipulation results in missing authentication. Remote exploitation of the attack is possible. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16227	A security vulnerability has been detected in SourceCodester Class and Exam Timetabling System 1.0. This impacts an unknown function of the file /edit_subject.php. Such manipulation of the argument ID leads to sql injection. The attack may be performed from remote. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2026-16228	A vulnerability was detected in SourceCodester Class and Exam Timetabling System 1.0. Affected is an unknown function of the file /edit_schoolyr.php. Performing a manipulation of the argument ID results in sql injection. It is possible to initiate the attack remotely. The exploit is now public and may be used.	7.3	More Details
CVE-2026-62290	cert-manager adds certificates and certificate issuers as resource types in Kubernetes clusters, and simplifies the process of obtaining, renewing and using those certificates. From 1.18.0 until 1.19.6 and 1.20.3, Challenge resources under acme.cert-manager.io can be created directly by namespace users without admission validation tying the Challenge to an Order, owner reference, or Issuer-selected solver, allowing attacker-controlled Challenge.spec.solver values referencing a ClusterIssuer to bypass DNS01 solver selectors such as dnsZones, dnsNames, and matchLabels and cause cert-manager to use ClusterIssuer DNS credentials for attacker-selected provider settings and DNS names, including disclosure of X-Api-User and X-Api-Key headers for acme-dns. This issue is fixed in versions 1.19.6 and 1.20.3.	7.3	More Details
	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including		

CVE-2026-32825	version 25.07.3, the application accepts unlimited password guesses against both the browser login flow and the JSON login endpoint. The source code enables Devise's `:lockable` module on the user model but explicitly disables both lock and unlock strategies, and no request throttling or rate-limiting layer was identified in the Rails code. This creates a direct online password-guessing risk: - valid user accounts can be attacked continuously without temporary lockout - the same weakness is reachable through both `/users/sign_in` and `/api/v4/auth/login` - successful guessing yields a normal session cookie in the HTML flow or a fresh JWT in the API flow - the API endpoint is especially attractive for automation because it requires no CSRF token This has been patched in version 26.06.08.	7.3	More Details
CVE-2026-32824	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, a low-privileged authenticated API user can supply `forwardToUrl` and `redirectUrl` values when triggering password reset or confirmation flows. Those values are then embedded into the outgoing email workflow without host allowlisting. This creates two related abuse paths: - password reset or confirmation links can be sent to a victim with the token already attached to an attacker-controlled `forwardToUrl` - after a legitimate password reset completes, the browser is redirected to attacker-controlled `redirectUrl` In practice, this can be used for phishing, token capture, confirmation hijacking, or steering a victim from a trusted email into an attacker domain. This is patched in version 26.06.08.	7.3	More Details
CVE-2026-16252	A security flaw has been discovered in Beijing Shenzhou Shihan Technology Multimedia Integrated Business Display System 8.2.2. Impacted is an unknown function of the file /admin/system/structure/updateStructure/deflate/Insecure/Staffshinel Ds.jsp?Shine ID=aaa. The manipulation of the argument Structure_ID results in sql injection. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks.	7.3	More Details
CVE-2026-12080	A flaw was found in the QEMU Guest Agent (qga). A local unprivileged user can exploit a vulnerability in the guest-ssh-add-authorized-keys command handler by manipulating symbolic links. This can occur either through a deterministic directory-symlink bypass or a Time-of-Check to Time-of-Use (TOCTOU) file-symlink race. Successful exploitation allows the attacker to gain ownership of arbitrary root-owned files or directories, leading to root access. This vulnerability requires an external management layer (e.g., libvirt) to trigger the affected code path.	7.3	More Details
CVE-2026-64621	FreeRDP before 3.28.0 (affected 3.x through 3.27.1) contains a double-free vulnerability in freerdp_client_rdp_file_apply_to_settings() (client/common/file.c) when parsing the selectedmonitors field of a .rdp connection file. The MonitorIds array is allocated through the settings object, and a raw non-owning pointer to it is freed on the strtoul error path without clearing settings->MonitorIds, leaving it dangling; at teardown freerdp_settings_free() frees the same buffer again. An attacker who convinces a victim to open a crafted .rdp file with oversized monitor tokens can trigger a size-controlled double-free in any FreeRDP CLI client (xfreerdp/sdl-freerdp/wlfreerdp) in the default configuration.	7.3	More Details
CVE-2026-16247	In _connect.BRAIN versions prior to 5.06, the application LogPathConfig.exe is executed during setup. During this process, existing permissions on %ProgramData% are deleted and replaced, granting the Windows group Everyone full control instead of restricting access to %ProgramData%\Bizerba_connect.BRAIN or %ProgramData%\Bizerba\BCT. Starting with _connect.BRAIN 5.06, the setup no longer executes this tool.	7.3	More Details
CVE-2026-16246	In BRAIN2 versions prior to 3.09, the application LogPathConfig.exe is executed during setup. As a result, the Windows group Everyone is granted full control over %ProgramData% instead of being restricted to %ProgramData%\Bizerba\BRAIN2\. Starting with BRAIN2 3.09, the setup no longer executes this tool. However, the optional component Bizerba ScriptService still executes it. Bizerba ScriptService is being deprecated and will no longer be included starting with BRAIN2 version 3.11.	7.3	More Details
CVE-2026-64158	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix write streaming disablement if fd open O_RDWR In netfs_perform_write(), "write streaming" (the caching of dirty data in dirty but !uptodate folios) is performed to avoid the need to read data that is just going to get immediately overwritten. However, this is/will be disabled in three circumstances: if the fd is open O_RDWR, if fscache is in use (as we need to round out the blocks for DIO) or if content encryption is enabled (again for rounding out purposes). The idea behind disabling it if the fd is open O_RDWR is that we'd need to flush the write-streaming page before we could read the data, particularly through mmap. But netfs now fills in the gaps if ->read_folio() is called on the page, so that is unnecessary. Further, this doesn't actually work if a separate fd is open for reading. Fix this by removing the check for O_RDWR, thereby allowing streaming writes even when we might read. This caused a number of problems with the generic/522 xfstest, but those are now fixed.	7.3	More Details
CVE-2026-45738	Argo CD is a declarative, GitOps continuous delivery tool for Kubernetes. Prior to 3.2.12, 3.3.10, and 3.4.2, Argo CD users with application write access can set link.argocd.argoproj.io/* annotations whose pipe-separated values are rendered by ui/src/app/applications/components/application-summary/application-summary.tsx in the Summary tab URLs section as anchor href values without URL validation, allowing javascript: execution in a higher-privileged user's authenticated Argo CD origin session. This issue is fixed in versions 3.2.12, 3.3.10, and 3.4.2.	7.3	More Details
	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: MGMT: validate Add Extended Advertising Data length MGMT_OP_ADD_EXT_ADV_DATA is registered as a variable-length command, with		

CVE-2026-64126	MGMT_ADD_EXT_ADV_DATA_SIZE as the fixed header size. The handler then uses cp->adv_data_len and cp->scan_rsp_len to validate and copy cp->data, but it never checks that those bytes are part of the mgmt command payload. A short command can therefore make add_ext_adv_data() pass an out-of-bounds pointer into tlv_data_is_valid(). If the bytes beyond the command buffer are addressable, they can also be copied into the advertising instance as scan response data, where the caller can read them back via MGMT_OP_GET_ADV_INSTANCE. The trigger requires CAP_NET_ADMIN in the initial user namespace; KASAN reports an 8-byte slab-out-of-bounds read. Reject commands whose length does not match the fixed header plus both advertising data lengths before parsing cp->data.	7.3	More Details
CVE-2026-13104	A potential vulnerability was reported in Lenovo App Store, distributed exclusively in the Chinese market, that could allow a local authenticated user to execute arbitrary code with elevated privileges.	7.3	More Details
CVE-2026-13103	A potential path traversal vulnerability was reported in Lenovo App Store, distributed exclusively in the Chinese market, that could allow a local authenticated user to execute arbitrary code.	7.3	More Details
CVE-2026-63883	In the Linux kernel, the following vulnerability has been resolved: serial: qcom_geni: fix kfifo underflow when flush precedes DMA completion IRQ When uart_flush_buffer() runs before the DMA completion IRQ is delivered, the following race can occur (all steps serialized by uart_port_lock): 1. DMA starts: tx_remaining = N, kfifo contains N bytes 2. DMA completes in hardware; IRQ is pending but not yet delivered 3. uart_flush_buffer() acquires the port lock and calls kfifo_reset(), making kfifo_len() = 0 while tx_remaining remains N 4. uart_flush_buffer() releases the port lock 5. DMA IRQ fires; handle_tx_dma() acquires the port lock and calls uart_xmit_advance(uport, tx_remaining) on an empty kfifo uart_xmit_advance() increments kfifo->out by tx_remaining. Since kfifo_reset() already set both in and out to 0, out wraps past in, causing kfifo_len() to return UART_XMIT_SIZE - tx_remaining. The next start_tx_dma() call then submits a DMA transfer of stale buffer data. Fix this by snapshotting kfifo_len() at the start of handle_tx_dma() and skipping uart_xmit_advance() when fifo_len < tx_remaining, which indicates the kfifo was reset by a preceding flush.	7.3	More Details
CVE-2026-63914	In the Linux kernel, the following vulnerability has been resolved: xfrm: route MIGRATE notifications to caller's netns xfrm_send_migrate() in net/xfrm/xfrm_user.c and pfkey_send_migrate() in net/key/af_key.c both hardcode &init_net for the multicast that announces a successful XFRM_MSG_MIGRATE / SADB_X_MIGRATE. XFRM_MSG_MIGRATE arrives on a per-netns NETLINK_XFRM socket, and the rest of the xfrm/af_key netlink path was made netns-aware in 2008. The other 14 multicast paths in xfrm_user.c route their event using xs_net(x), xp_net(xp) or sock_net(skb->sk); only the migrate path was missed. Two consequences of the init_net hardcoding: 1. The notification (selector, old/new endpoint addresses, and the km_address) is delivered to listeners on init_net's XFRMNLGRP_MIGRATE / pfkey BROADCAST_ALL groups rather than on the issuing netns. An IKE daemon running in init_net therefore receives migration notifications originating from any other netns on the host. 2. An IKE daemon running inside a non-init netns and subscribed to its own XFRMNLGRP_MIGRATE / pfkey groups never receives the notification of its own migration. IKEv2 MOBIKE / address-update handling inside a netns is silently broken. Thread struct net through km_migrate() and the xfrm_mgr.migrate function pointer, drop the &init_net override in xfrm_send_migrate() and pfkey_send_migrate(), and pass the caller's net (already in scope in xfrm_migrate() via sock_net(skb->sk)) all the way down. struct xfrm_mgr is in-tree only and not exported as a stable API, so the function-pointer signature change is internal. pfkey_broadcast() is already netns-aware via net_generic(net, pfkey_net_id) since the pernet conversion. The five other pfkey_broadcast() callers in af_key.c already pass xs_net(x), sock_net(sk) or a per-netns net, so this only removes the &init_net outlier.	7.3	More Details
CVE-2026-15907	A flaw has been found in H3C SecPath F1000-C8300 up to 20260522. This impacts an unknown function of the file /webui/?g=log_fw_nbc_mail_jsondata. Executing a manipulation of the argument subject can lead to sql injection. The attack can be executed remotely. The exploit has been published and may be used. The vendor was contacted early about this disclosure and confirmed the existence of the vulnerability. A technical fix is planned to be released.	7.3	More Details
CVE-2026-56624	Improper certificate validation in Apache MINA SSHD (server-side). Apache MINA SSHD is a Java library for client-side and server-side SSH. Server-side OpenSSH user certificate validation during user authentication in an Apache MINA SSHD server did not check for the unsupported force-command or verify-required options that could be embedded in the certificate, nor did it validate these options. As a result it was possible that a user could authenticate with such a certificate that included a force-command option but still was able to execute other commands. What other command exactly would be available to the user depends on the implementation of the server. This issue is fixed in Apache MINA SSHD 2.19.0 and 3.0.0-M5. Applications are advised to upgrade to these versions. The fix rejects OpenSSH user certificates that include these options, since Apache MINA SSHD implements neither force-command nor sk-*-cert-v01@openssh.com user certificates (which are the only ones for which verify-required would make sense).	7.3	More Details
CVE-2026-16016	A vulnerability was identified in poco-ai poco-claw up to 0.5.4. This issue affects the function run_task of the file executor/app/api/v1/task.py. The manipulation of the argument callback_url leads to server-side request forgery. The attack is possible to be carried out remotely. The exploit is publicly available and might be used. The reported GitHub issue was closed automatically due to inactivity.	7.3	More Details
CVE-	A vulnerability was determined in SourceCodester Class and Exam Timetabling System 1.0/1.php. Affected by this vulnerability is an unknown functionality of the file /edit_room1.php. Executing a manipulation of the		More

2026-16154	argument ID can lead to sql injection. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized.	7.3	Details
CVE-2026-16329	A vulnerability was identified in D-Link DNS-320 1.0.2. Impacted is an unknown function of the file /photo_center/php/uploadify.php. The manipulation of the argument Malicious Handler leads to unrestricted upload. The attack may be initiated remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2026-16447	A vulnerability has been found in D-Link DNS-320 1.0.2. Impacted is an unknown function of the file /web/jquery/uploader/multi_uploadify.php. The manipulation of the argument Filedata[] leads to unrestricted upload. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used.	7.3	More Details
CVE-2026-16332	A vulnerability was detected in D-Link DNS-320 1.0.2. This impacts an unknown function of the file /mydlink/multi_uploadify.php. Performing a manipulation of the argument Filedata[] results in unrestricted upload. The attack is possible to be carried out remotely. The exploit is now public and may be used.	7.3	More Details
CVE-2026-16331	A security vulnerability has been detected in D-Link DNS-320 1.0.2. This affects an unknown function of the file /web/function/save_ajax.php. Such manipulation of the argument Malicious Handler leads to unrestricted upload. The attack can be executed remotely. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2026-16084	A weakness has been identified in Sipeed PicoClaw up to 0.2.9. This impacts the function web_fetch of the file pkg/tools/integration/web.go. This manipulation causes server-side request forgery. Remote exploitation of the attack is possible. The exploit has been made available to the public and could be used for attacks. Patch name: c15aac21fe05ee103a470e1104bc891754e83392. To fix this issue, it is recommended to deploy a patch.	7.3	More Details
CVE-2026-16330	A weakness has been identified in D-Link DNS-320 1.0.2. The impacted element is an unknown function of the file /web/jquery/uploader/uploadify.php. This manipulation of the argument https:/ucn9h68n9289.feishu.cn/wiki/JJcTwHz7aiKeq6kSItMcoeSUnMc?from=from_copylink causes unrestricted upload. Remote exploitation of the attack is possible. The exploit has been made available to the public and could be used for attacks.	7.3	More Details
CVE-2026-61438	PraisonAI before 4.6.78 contains a remote code execution vulnerability in JobWorkflowExecutor._exec_inline_python() due to insufficient AST validation of workflow script steps. Attackers can create malicious YAML workflow files with import os statements followed by os.system() calls that bypass sandbox checks and execute arbitrary OS commands with process privileges.	7.3	More Details
CVE-2026-61427	PraisonAI before 4.6.78 exposes the MCP HTTP-stream transport without authentication by default: the CLI --api-key option defaults to None, and the server only enforces Authorization/Bearer checks when an API key is configured. When an operator runs 'praisonai mcp serve --transport http-stream' without an API key, an unauthenticated client (no Authorization header, and no Origin header, which is also permitted) can initialize a session, enumerate the available tools (tools/list), and invoke tools (tools/call). Additionally, the dispatcher forwards tool-call arguments to handlers without validating them against the advertised inputSchema. The server binds to 127.0.0.1 by default, so remote exploitation requires the operator to bind to a network-accessible address (e.g., --host 0.0.0.0).	7.3	More Details
CVE-2026-16327	A vulnerability was determined in D-Link DNS-320 1.0.2. This issue affects some unknown processing of the file /web/web_file/upload.php. Executing a manipulation of the argument File can lead to unrestricted upload. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2026-16125	A vulnerability was found in zevorn rt-claw up to 0.2.0. The affected element is the function claw_net_get/claw_net_post of the file claw/services/tools/net.c of the component http_request. The manipulation of the argument url results in server-side request forgery. The attack can be executed remotely. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16126	A vulnerability was determined in zevorn rt-claw up to 0.2.0. The impacted element is the function handle_rpc_request of the file claw/services/swarm/swarm.c of the component Swarm RPC Receiver. This manipulation causes incorrect authorization. The attack is possible to be carried out remotely. The exploit has been publicly disclosed and may be utilized. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16127	A vulnerability was identified in zevorn rt-claw up to 0.2.0. This affects the function claw_net_get/claw_net_post of the file claw/tools/tool_net.c of the component http_request. Such manipulation of the argument url leads to server-side request forgery. The attack may be performed from remote. The exploit is publicly available and might be used. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16324	A vulnerability was identified in Metasoft 美特软件 MetaCRM up to 6.4.0 Beta06. The impacted element is an unknown function of the file /business/qnaire/upload.jsp. Such manipulation of the argument File leads to unrestricted upload. The attack may be launched remotely. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details

CVE-2026-16128	A security flaw has been discovered in zevorn rt-claw up to 0.2.0. This impacts the function receiver_thread of the file claw/services/swarm/swarm.c of the component http_request. Performing a manipulation results in server-side request forgery. It is possible to initiate the attack remotely. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-16152	A vulnerability was found in SourceCodester Class and Exam Timetabling System 1.0. Affected is an unknown function of the file /edit_rooma.php. Performing a manipulation of the argument ID results in sql injection. The attack is possible to be carried out remotely. The exploit has been made public and could be used.	7.3	More Details
CVE-2026-51082	A race condition between the vncproxy and vncwebsocket API calls in Proxmox Virtual Environment (PVE) 9.x pve-manager before 9.1.9 and 8.x before 8.4.19; qemu-server 9.x before 9.1.7 and 8.x before 8.4.7; and pve-container before 6.1.3 (PVE 9.x) and before 5.3.4 (PVE 8.x) allows an attacker with privileges to call "vncproxy" to hijack a VNC session that is established in parallel by a different user for a different VM.	7.2	More Details
CVE-2026-1771	The MapSVG plugin for WordPress is vulnerable to arbitrary file uploads due to missing file type validation in the SVGFile constructor in all versions up to, and including, 8.14.0 This is due to an incorrect conditional check that prevents file validation from taking place. This makes it possible for authenticated attackers, with Administrator-level access and above, to upload arbitrary files on the affected site's server which may make remote code execution possible.	7.2	More Details
CVE-2026-14448	An high privileged remote attacker can exploit an authenticated OS command injection vulnerability in the system_certificates view due to improper neutralization of special elements in an OS command. This can result in a total loss of confidentiality, availability and integrity.	7.2	More Details
CVE-2026-46981	Vulnerability in the Oracle Utilities Network Management System product of Oracle Utilities Applications (component: Mobile). Supported versions that are affected are 2.5.0.1.0-2.5.0.1.17, 2.5.0.2.0-2.5.0.2.11, 2.6.0.1.0-2.6.0.1.12, 2.6.0.2.0-2.6.0.2.8 and 25.12.0.0.0-25.12.0.0.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Utilities Network Management System. While the vulnerability is in Oracle Utilities Network Management System, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle Utilities Network Management System accessible data as well as unauthorized read access to a subset of Oracle Utilities Network Management System accessible data. CVSS 3.1 Base Score 7.2 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:L/I:L/A:N).	7.2	More Details
CVE-2026-46954	Vulnerability in the Oracle Human Resources product of Oracle E-Business Suite (component: Data Removal Tool). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows high privileged attacker with network access via HTTP to compromise Oracle Human Resources. Successful attacks of this vulnerability can result in takeover of Oracle Human Resources. CVSS 3.1 Base Score 7.2 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:H/I:H/A:H).	7.2	More Details
CVE-2026-6952	A post-authentication command injection vulnerability in the "LogServer" field of the syslog component in Zyxel AX7501-B1 firmware versions through 5.17(ABPC.7.2)C0 could allow an authenticated attacker with administrator privileges to execute OS commands on an affected device.	7.2	More Details
CVE-2026-63453	Buffer overflow vulnerabilities exist in the command line interface of AOS-CX. Successful exploitation of these vulnerabilities could allow a remote high-privileged user to execute arbitrary code as a privileged user on the underlying operating system.	7.2	More Details
CVE-2026-44878	A vulnerability in the web-based management interface of an ECOS device could allow a highly privileged, authenticated remote attacker to access the device's filesystem. Successful exploitation of this vulnerability could allow an attacker to access sensitive files and tamper with or delete system data.	7.2	More Details
CVE-2026-44879	A vulnerability in the command line interface of ECOS devices could allow a highly privileged, authenticated remote attacker to perform command injection on certain CLI commands. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	7.2	More Details
CVE-2026-63454	An authenticated path traversal vulnerability exists in AOS-CX. Successful exploitation of this vulnerability allows an attacker to copy arbitrary files to a user readable location from the command line interface of the underlying operating system, which could lead to remote code execution.	7.2	More Details
CVE-2026-46988	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: Connector Framework). Supported versions that are affected are 13.5 and 24.1. Easily exploitable vulnerability allows high privileged attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. Successful attacks of this vulnerability can result in takeover of Oracle Enterprise Manager Base Platform. CVSS 3.1 Base Score 7.2 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:H/I:H/A:H).	7.2	More Details
CVE-2026-	TDengine is an open source, time-series database optimized for Internet of Things devices. Prior to 3.4.1.15, a user with create udf privilege could upload a crafted shared library and install it as a user-defined function,	7.2	More

62350	such as eval, then execute arbitrary C code on the TDengine server side through database queries. This issue is fixed in version 3.4.1.15.		Details
CVE-2026-20297	In Splunk Enterprise versions below 10.4.1, 10.2.5, 10.0.8, 9.4.13, and 9.3.14, and Splunk Cloud Platform versions below 10.5.2605.0, 10.4.2604.6, 10.2.2510.18, and 10.1.2507.24, a user who holds a role that contains the `edit_local_apps` and `install_apps` capabilities could cause a legitimate app installation to write files outside the intended app directory, into `$SPLUNK_HOME/etc/` and its subdirectories. The vulnerability is caused by a path traversal in the app installation workflow, which does not restrict the installation path to the intended app directory.	7.2	More Details
CVE-2026-44982	CrowdSec offers crowdsourced protection against malicious IPs. From 1.5.0 until 1.7.8, pkg/appsec/request.go NewParsedRequestFromRequest allocated a request body buffer from max(r.ContentLength, 0), so HTTP/1.1 requests using Transfer-Encoding: chunked and HTTP/2 requests without a content-length header produced an empty body and caused WAF rules targeting REQUEST_BODY, BODY_ARGS, ARGS_POST, JSON, or XML to be skipped. This issue is fixed in version 1.7.8.	7.2	More Details
CVE-2026-13042	The RPB Chessboard plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Comment Content in all versions up to, and including, 8.1.2 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. WordPress's save-time kses sanitization does not mitigate this issue because the crafted payload uses only kses-allowed tags and attributes (such as an <a> element with title and href), and the dangerous attribute-breaking HTML is synthesized entirely at render time by the plugin's own comment_text filter.	7.2	More Details
CVE-2026-15395	The Kali Forms — Contact Form & Drag-and-Drop Builder plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'digitalSignature' Field Value in all versions up to, and including, 2.4.18 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The required form-submission nonce is publicly available on any page containing the form shortcode, making this exploitable by fully unauthenticated attackers without any precondition beyond the form being published.	7.2	More Details
CVE-2026-7543	The Breakdance plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'fields' parameter in versions up to, and including, 2.7.1 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	7.2	More Details
CVE-2026-64880	Unsanitized user-supplied input in report filtering parameters is concatenated directly into SQL queries without proper escaping or parameterized queries, enabling blind SQL injection and unauthorized database read access.	7.1	More Details
CVE-2026-3183	Zohocorp ManageEngine ADSelfService Plus versions before 6524 are vulnerable to Multi Factor Authentication Bypass.	7.1	More Details
CVE-2026-56623	Path traversal on Windows in Apache MINA SSHD component sshd-git. Apache MINA SSHD is a Java library for client-side and server-side SSH. A git server implemented with Apache MINA SSHD component sshd-git and running on Windows could allow an authenticated remote user access to git repositories outside of the configured server-side root directory. The path validation applied for CVE-2026-48827 in Apache MINA SSHD 2.18.0 and 3.0.0-M4 was partly ineffective for Servers running on Windows. Applications are affected if they use org.apache.sshd:sshd-git to implement a git server and run on Windows. Applications not using sshd-git or not running on Windows are not affected. Users are advised to upgrade affected applications to Apache MINA SSHD 2.19.0, which fixes the issue. The issue also is present in the pre-release milestones 3.0.0-M1 to 3.0.0-M4 for a new upcoming new major version 3.0.0. Again, applications are affected only if they use sshd-git and run on Windows. Upgrade affected applications to 3.0.0-M5.	7.1	More Details
CVE-2026-52584	Buffer Overflow vulnerability in libjxl v.0.11.2 and before allows a local attacker to obtain sensitive information via the DecodeImageAPNG function	7.1	More Details
CVE-2024-34268	EQ-3 Eqiva CC-RT-BLE Bluetooth Smart Radiator Thermostat Firmware up to the latest version 1.46 was discovered to allow unsecured bluetooth connections. This vulnerability allows attackers to gain full access to the device without authentication.	7.1	More Details
CVE-2026-9833	The Tag Groups is the Advanced Way to Display Your Taxonomy Terms WordPress plugin before 2.2.0 does not properly escape one of its AJAX parameters before reflecting it in the response body served with an HTML content type, allowing unauthenticated attackers to execute arbitrary JavaScript in the browser of a logged-in user with `edit_pages` capability (Editor or higher) who is tricked into following a crafted link.	7.1	More Details
CVE-2026-12970	The LearnPress WordPress plugin before 4.4.1 does not escape a search parameter before reflecting it into an HTML attribute, leading to Reflected Cross-Site Scripting that executes in the browser of a logged-in instructor or administrator who is tricked into opening a crafted link.	7.1	More Details
	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix posix_acl leak and ignored error in		

CVE-2026-53396	nfsd4_create_file nfsd4_create_file() has two bugs in its ACL handling: The return value of nfsd4_acl_to_attr() is silently discarded. When the NFSv4-to-POSIX ACL conversion fails (e.g., -EINVAL for unsupported ACE types), the file is created without any ACL and the client receives NFS4_OK. This violates RFC 7530/8881 which require the server to reject unsupported attributes on CREATE. When start_creating() fails after ACL attributes have been populated in attrs (either via nfsd4_acl_to_attr or via ownership transfer from open->op_dpacl/op_pacl), the function jumps to out_write which skips nfsd_attrs_free(). The posix_acl allocations are leaked. A client can trigger this repeatedly with OPEN(CREATE), ACL attributes, and an invalid filename (e.g., longer than NAME_MAX). Fix both by capturing the nfsd4_acl_to_attr() return value and by changing the early error paths to jump to out instead of out_write. Initialize child to ERR_PTR(-EINVAL) so that end_creating() is safe to call even if start_creating() was never reached.	7.1	More Details
CVE-2026-53402	In the Linux kernel, the following vulnerability has been resolved: fbdev: fbcon: fix out-of-bounds read in err_out of fbcon_do_set_font() When fbcon_do_set_font() fails (e.g., due to a memory allocation failure inside vc_resize() under heavy memory pressure), it jumps to the `err_out` label to roll back the console state. However, the current rollback logic forgets to restore the `hi_font` state, leading to a severe state machine corruption. Earlier in the function, `set_vc_hi_font()` might be called to change `vc->vc_hi_font_mask` and mutate the screen buffer. If `vc_resize()` subsequently fails, the `err_out` path restores `vc_font.charcount` but entirely skips rolling back the `vc_hi_font_mask` and the screen buffer. This mismatch leaves the terminal in a desynchronized state. Because `vc_hi_font_mask` remains set, the VT subsystem will still accept character indices greater than 255 from userspace and write them to the screen buffer. Subsequent rendering calls (e.g., `fbcon_putcs()`) will then use these inflated indices to access the reverted, 256-character font array, leading to a deterministic out-of-bounds read and potential kernel memory disclosure. Fix this by adding the missing rollback logic for the `hi_font` mask and screen buffer in the error path.	7.1	More Details
CVE-2026-64172	In the Linux kernel, the following vulnerability has been resolved: KVM: SVM: Disable AVIC IPI virtualization on Hygon Family 18h (erratum #1235) Hygon Family 18h CPUs are derived from AMD Family 17h (Zen1) silicon and share the same erratum #1235: hardware may read a stale IsRunning=1 bit during ICR write emulation and silently fail to generate an AVIC_IPI_FAILURE_TARGET_NOT_RUNNING VM-Exit on the sending vCPU. The absence of the VM-Exit causes KVM to miss the required wakeup of blocking target vCPUs, leading to hung vCPUs and unbounded delays in guest execution. Extend the existing AMD Family 17h erratum #1235 workaround to also cover Hygon Family 18h. With IPI virtualization disabled, KVM never sets IsRunning=1 in the Physical ID table, so every non-self IPI generates a VM-Exit and is correctly emulated.	7.1	More Details
CVE-2026-50144	ncnn is a high-performance neural network inference framework optimized for the mobile platform. In commit e54f7b1f88434e1d844ea0551b880a1cfb079ce1 and earlier, ncnn allows an out-of-bounds heap write in ncnn::ParamDict::load_param() when Net::load_param() loads a malicious .param model file because the parsed parameter id is checked only against id >= NCNN_MAX_PARAM_COUNT, allowing a negative id to index before the params[NCNN_MAX_PARAM_COUNT] array. This vulnerability is fixed by commit 5a0288f255daa6c3294f77109f67718e434ec020.	7.1	More Details
CVE-2026-53387	In the Linux kernel, the following vulnerability has been resolved: iio: light: veml6075: add bounds check to veml6075_it_ms index veml6075_it_ms has 5 elements but VEML6075_CONF_IT can yield values 0-7. If it returns a value >= 5, this causes an out-of-bounds array access. Add a bounds check and return -EINVAL if the index is out of range. The problem values are reserved so should never be read from the register. Hence this is hardening against fault device, missprogramming or bus corruption.	7.1	More Details
CVE-2026-52869	The MCP Python SDK, called mcp on PyPI, is a Python implementation of the Model Context Protocol (MCP). Prior to 1.27.2, the SSE and stateful Streamable HTTP transports mcp.server.sse.SseServerTransport and mcp.server.streamable_http_manager.StreamableHTTPSessionManager route requests to existing sessions using only the session_id query parameter or Mcp-Session-Id header without verifying the authenticated principal that created the session, allowing a different bearer-token-authenticated client with a known session ID to inject JSON-RPC messages into that session. This issue is fixed in version 1.27.2.	7.1	More Details
CVE-2026-16118	A flaw was found in xdgmime. A heap-based buffer overflow can be triggered in _xdg_mime_magic_parse_magic_line() in the xdgmimemagic.c file on little-endian systems when an attacker-controlled MIME magic file in a user-writable XDG data location (e.g., in the $XDG_DATA_HOME/mime/magic path) is parsed by an application performing MIME type detection (e.g., via g_content_type_guess()). When performing byte-swap, incorrect pointer arithmetic on the write side causes an out-of-bounds write of 2 bytes, resulting in an application crash or memory corruption.	7.1	More Details
CVE-2026-47130	NextCRM is open-source customer relationship management (CRM) software. Versions prior to 0.12.0 have a Broken Object Level Authorization (BOLA/IDOR) vulnerability exists in the CRM contact and target update endpoints. The application fails to verify if the authenticated user has ownership of the specific resource being modified. This allows any authenticated user (even with a standard `member` role) to arbitrarily modify sensitive CRM contacts and targets belonging to other users or organizations (cross-tenant data tampering). Version 0.12.0 fixes the issue.	7.1	More Details
CVE-2026-62205	OpenClaw versions 2026.4.12-beta.1 before 2026.6.6 contain a missing-authorization vulnerability in the MS Teams message actions feature. When the affected feature is enabled and reachable, a lower-trust caller or a configured input path can perform actions that should have required a stronger authorization or policy check. Practical impact depends on the operator's configuration and whether lower-trust input can reach that path.	7.1	More Details

	The issue is fixed in 2026.6.6.		
CVE-2026-56171	Exposure of private personal information to an unauthorized actor in Windows RDP allows an unauthorized attacker to disclose information over a network.	7.1	More Details
CVE-2026-53368	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix fsck inconsistency caused by incorrect nat_entry flag usage f2fs_need_dentry_mark() reads nat_entry flags without mutual exclusion with the checkpoint path, which can result in an incorrect inode block marking state. The scenario is as follows: create & write & fsync 'file A' write checkpoint - f2fs_do_sync_file // inline inode - f2fs_write_inode // inode folio is dirty - f2fs_write_checkpoint - f2fs_flush_merged_writes - f2fs_sync_node_pages - f2fs_fsync_node_pages // no dirty node - f2fs_need_inode_block_update // return true - f2fs_fsync_node_pages // inode dirtied - f2fs_need_dentry_mark //return true - f2fs_flush_nat_entries - f2fs_write_checkpoint end - __write_node_folio // inode with DENT_BIT_SHIFT set SPO, "fsck --dry-run" find inode has already checkpointed but still with DENT_BIT_SHIFT set The state observed by f2fs_need_dentry_mark() can differ from the state observed in __write_node_folio() after acquiring sbi->node_write. The root cause is that the semantics of IS_CHECKPOINTED and HAS_FSYNCED_INODE are only guaranteed after the checkpoint write has fully completed. This patch moves set_dentry_mark() into __write_node_folio() and protects it with the sbi->node_write lock.	7.1	More Details
CVE-2026-63771	Adminer before 5.4.3 contains a cookie injection vulnerability that allows attackers to manipulate cookie attributes by injecting arbitrary values through the unsanitized X-Forwarded-Prefix HTTP header used in Set-Cookie path attributes. Attackers can exploit a misconfigured reverse proxy to downgrade SameSite protection and enable cross-origin authenticated requests, bypassing cookie security controls.	7.1	More Details
CVE-2026-62206	OpenClaw versions before 2026.6.9 contain a missing authorization vulnerability in Discord moderation actions. In affected versions, a lower-trust caller or configured input path could perform moderation actions that should have required a stronger authorization or policy check. Practical impact depends on the operator's configuration and whether lower-trust input can reach the affected path.	7.1	More Details
CVE-2026-62219	OpenClaw 2026.2.12 before 2026.5.26 contain an authorization bypass vulnerability in the hooks allowedAgentIds validation. A lower-trust caller or configured input path can bypass agent ID restrictions by submitting blank agent IDs, allowing actions that should require stronger authorization or policy checks.	7.1	More Details
CVE-2026-59255	BloodHound through 9.4.0, fixed in commit 8f79035, contains a missing authorization vulnerability in the custom-nodes API endpoints that allows any authenticated user to modify the global graph schema. Attackers with valid session tokens can create, update, or delete custom node types affecting all users and tenants by invoking unprotected POST, PUT, and DELETE operations on the custom-nodes endpoints.	7.1	More Details
CVE-2026-47870	VMware Avi Load Balancer contains a privilege escalation vulnerability. A malicious authenticated user with network access may be able to execute remote code. Affected versions: 32.1.1 (fixed in 32.1.2) 31.1.1 through 31.2.2 (fixed in 31.2.2-2p3) 30.1.1 through 30.2.6 (fixed in 30.2.7) 22.1.1 through 22.1.7 (fixed in 30.2.7)	7.1	More Details
CVE-2026-39879	Due to a missing sanitization call in [`afsql_dd_run_query`](https://github.com/syslog-ng/syslog-ng/blob/649e6e18e3459fb4467000a88dfb12fa97f9719c/modules/afsql/afsql.c#L219), syslog-ng before 4.12 are vulnerable to SQL injection from an untrusted source. This is not part of the default configuration, the SQL driver has to be manually configured. Fixes are in syslog-ng 4.12, syslog-ng Premium Edition 8.2 and syslog-ng Store Box 7.8	7.1	More Details
CVE-2026-53515	Better Auth is an authentication and authorization library for TypeScript. From 1.2.10 until 1.6.11, the @better-auth/sso plugin's POST /sso/register endpoint lets any organization member attach a new SSO provider to that organization because registerSSOProvider checks only for a membership row and does not require an owner or admin role, allowing attacker-controlled OIDC or SAML providers to drive /sso/callback/{providerId} organization provisioning. This issue is fixed in version 1.6.11.	7.1	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: KVM: Replace guest-triggerable BUG_ON() in ioeventfd datamatch with get_unaligned() Drop a BUG_ON() that has been reachable since it was first added, way back in 2009, and instead use get_unaligned() to perform potentially-unaligned accesses. For a given store, KVM x86's emulator tracks the entire value in the destination operand, x86_emulate_ctxt.dst. If the destination is memory, and the target splits multiple pages and/or is emulated MMIO, then KVM handles each fragment independently. E.g. on a page split starting at page offset 0xffc, KVM writes 4 bytes to the first page, then the remaining bytes to the second page, using ctxt->dst as the source for both (with appropriate offsets). If the destination splits a page *and* hits emulated MMIO on the second page, then KVM will complete the write to the first page, then emulate the MMIO access to the second page. If there is a datamatch-enabled ioeventfd at offset 0 of the second page, then KVM will process the remainder of the store as a potential ioeventfd signal. Putting it all together, if the guest emits a store that splits a page starting at page offset N, and the second page has a datamatch-enabled ioeventfd at offset 0, then KVM will check for datamatch using &dst.valptr[N] as the source. Due to dst (and thus dst.valptr) being 32-byte aligned, if N is not aligned to @len, the BUG_ON() fires. E.g. with a 16-byte store at page offset 0xffc, to an ioeventfd of len 8, all initial checks in ioeventfd_in_range() will succeed, and the BUG_ON() fires due to @val being 4-byte aligned, but not 8-byte aligned. ------------[cut here]------------ kernel BUG at	7.1	More Details

63806	arch/x86/kvm/../../../virt/kvm/eventfd.c:783! Oops: invalid opcode: 0000 [#1] SMP CPU: 0 UID: 1000 PID: 615 Comm: repro Not tainted 7.1.0-rc2-ff238429d1ea #365 PREEMPT Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 0.0.0 02/06/2015 RIP: 0010:ioeventfd_write+0x6c/0x70 [kvm] Call Trace: <TASK> __kvm_io_bus_write+0x85/0xb0 [kvm] kvm_io_bus_write+0x53/0x80 [kvm] vcpu_mmio_write+0x66/0xf0 [kvm] emulator_read_write_onepage+0x12a/0x540 [kvm] emulator_read_write+0x109/0x2b0 [kvm] x86_emulate_insn+0x4f8/0xfb0 [kvm] x86_emulate_instruction+0x181/0x790 [kvm] kvm_mmu_page_fault+0x313/0x630 [kvm] vmx_handle_exit+0x18a/0x590 [kvm_intel] kvm_arch_vcpu_ioctl_run+0xc81/0x1c90 [kvm] kvm_vcpu_ioctl+0x2d5/0x970 [kvm] __x64_sys_ioctl+0x8a/0xd0 do_syscall_64+0xb7/0x890 entry_SYSCALL_64_after_hwframe+0x4b/0x53 RIP: 0033:0x7f19c931a9bf </TASK> Modules linked in: kvm_intel kvm irqbypass ---[end trace 0000000000000000]--- In a perfect world, the fix would be to simply delete the BUG_ON(), as KVM x86 doesn't perform alignment checks on "normal" memory accesses at CPL0. Sadly, C99 ruins all the fun; while the x86 architecture plays nice, dereferencing an unaligned pointer directly is undefined behavior in C, e.g. triggers splats when running with CONFIG_UBSAN_ALIGNMENT=y.		
CVE-2026-62212	OpenClaw before 2026.5.28 contains a race condition in the MS Teams safeFetch DNS rebinding check. When the affected feature is enabled and reachable, a lower-trust caller or configured input path could win a timing window between the DNS validation check and use, allowing actions that should have required a stronger authorization or policy check. Practical impact depends on the operator's configuration and whether lower-trust input can reach that path.	7.1	More Details
CVE-2026-47697	Shelf is a platform for tracking physical assets. Shelf is multi-tenant; data is isolated per organization (workspace). Prior to version 1.20.2, several endpoints accepted entity IDs from request input and `connect`-ed / read / updated them without verifying the IDs belonged to the caller's organization. An authenticated user in Org A who knew or obtained an ID belonging to Org B could act on Org B's data across organization boundaries (a cross-tenant IDOR). A loader-only restriction on personal-workspace bookings was also bypassable via a crafted POST. Version 1.20.2 patches the issue. No known workarounds are available.	7.1	More Details
CVE-2026-62387	The Grav API plugin (getgrav/grav-plugin-api) before 1.0.0-rc.16 shipped Access-Control-Allow-Origin: * as its default CORS configuration on all responses, including authenticated endpoints and preflight (OPTIONS) responses. Because the plugin accepts credentials via the Authorization and X-API-Token headers (set programmatically by JavaScript rather than via cookies), an attacker who obtains a valid access token (e.g., via log leakage, Referer headers, browser history, or network capture) can issue fully authenticated cross-origin requests from any malicious website to read sensitive data and perform write operations as the token's user. Fixed in 1.0.0-rc.16.	7.1	More Details
CVE-2026-59867	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.5, Kiota resolved OpenAPI $ref values by fetching remote http(s) URLs and reading local absolute or out-of-tree file paths, allowing `kiota generate` on an attacker-controlled or attacker-influenced description to perform build-time SSRF, remote file inclusion, and local file inclusion by inlining external schemas such as REMOTE_KIOTA_PROP or Leaked into generated clients. This issue is fixed in version 1.32.5 by AllowedExternalOriginsStreamLoader and the --allowed-external-origins option.	7.1	More Details
CVE-2026-50163	oras-go is a Go library for managing OCI artifacts. Prior to 2.6.2, ensureLinkPath in content/file/utils.go:262-275 validates a hardlink target relative to the extract base but returns the unresolved target, causing os.Link("victim.secret", "<extract_base>/payload.tar.gz/evil_cwd_link") to resolve header.Linkname against the process current working directory for a Typeflag=TypeLink entry such as Name=payload.tar.gz/evil_cwd_link and Linkname="victim.secret" with io.deis.oras.content.unpack: "true", which can expose or tamper with files such as .env, .git/config, .aws/credentials, and ~/.ssh/config. This issue is fixed in version 2.6.2.	7.1	More Details
CVE-2026-63920	In the Linux kernel, the following vulnerability has been resolved: ipv6: validate extension header length before copying to cmsg ip6_datagram_recv_specific_ctl() builds IPV6_{HOPOPTS,DSTOPTS,RTHDR} cmsgs (and their IPV6_2292* legacy counterparts) by trusting the on-wire hdrlen byte (ptr[1]) when computing the put_cmsg() length. The length was validated only at parse time (ipv6_parse_hopopts(), etc.). An nftables payload-write expression can rewrite hdrlen after parsing and before the skb reaches recvmsg; the write itself is in-bounds but put_cmsg() then reads up to ((hdrlen+1) << 3) = 2040 bytes from an 8-byte header. nftables is reachable from an unprivileged user namespace, so this is an unprivileged slab-out-of-bounds read: BUG: KASAN: slab-out-of-bounds in put_cmsg+0x3ac/0x540 put_cmsg+0x3ac/0x540 udpv6_recvmsg+0xca0/0x1250 sock_recvmsg+0xdf/0x190 ___sys_recvmsg+0x1b1/0x620 Add ipv6_get_exthdr_len() which validates that at least two bytes are accessible before reading the hdrlen field, then checks the computed length against skb_tail_pointer(skb), returning 0 on failure. Extension headers are kept in the linear skb area by pskb_may_pull() during input, so skb_tail_pointer() is the correct bound. Use ipv6_get_exthdr_len() at all non-AH call sites: the five standalone cmsg blocks (HbH, 2292HbH, 2292DSTOPTS x2, 2292RTHDR) and the three standard cases in the extension-header walk loop (DSTOPTS, ROUTING, default). AH retains an inline bounds check because its length formula differs ((ptr[1]+2)<<2). The walk loop also gets a pre-read bounds check at the top to validate ptr before any case accesses ptr[0] or ptr[1]. When the walk loop detects a corrupted header, return from the function instead of continuing to process later socket options.	7.1	More Details
	Cloudreve is a self-hosted file management and sharing system. Prior to 4.16.1, a Cloudreve WebDAV		

CVE-2026-54563	account rooted at a configured folder can send paths such as /dav/%2e%2e/outside.txt because stripPrefix in pkg/webdav/webdav.go joins the decoded request suffix to the account root with fs.URI.JoinRaw without checking containment, allowing the scoped credential to read and list files outside the configured folder and writable credentials to create, overwrite, move, or delete them. This issue is reported as fixed in version 4.16.1.	7.1	More Details
CVE-2026-58484	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.12.2, `EnvironmentManager.listBackups()` reads each backup's `_manifest.json` and trusts the manifest's `path` field. `EnvironmentManager.pruneBackups()` later passes that trusted `entry.path` directly to `rmSync(entry.path, { recursive: true, force: true })`. An attacker who can place or modify a manifest inside `data/<env>/.backups/<name>/_manifest.json` can cause `network-ai env backup prune --env <env> --keep <n>` or any code path invoking `pruneBackups()` to recursively delete an arbitrary path accessible to the Network-AI process user. This is fixed in v5.12.2. `pruneBackups()` no longer passes `entry.path` from the on-disk manifest to `rmSync`. The deletion path is recomputed from a format-validated `entry.backupId`, and a `dirname` containment check confines deletion to exactly one level under the backups directory. A poisoned manifest (e.g. `"path": "/"`) is now inert.	7.1	More Details
CVE-2026-46336	Manyfold is an open source, self-hosted web application for managing a collection of 3d models, particularly focused on 3d printing. From 0.96.0 until 0.140.0, authenticated users can rename uploaded files with path traversal sequences because app/models/model_file.rb uses the user-controlled filename in File.join(model.path, filename) without sufficient sanitization, allowing files to be moved or written outside the configured library directory. This issue is fixed in version 0.140.0.	7.1	More Details
CVE-2026-12978	The FunnelKit WordPress plugin before 3.15.0.6 does not escape a user-supplied parameter before reflecting it into the HTML response of one of its page-builder AJAX actions, allowing unauthenticated attackers to perform Reflected Cross-Site Scripting against logged-in users who open a crafted page. The affected action is only registered when the Divi /builder is active.	7.1	More Details
CVE-2026-49284	SimpleSAMLphp versions before 1.18.6 contain an information disclosure vulnerability. Prior to 2.4.7 and 2.5.2, SimpleSAMLphp's SAML SP ACS path does not enforce the IdP selected for an SP-initiated login when unsigned Response/InResponseTo is combined with a signed assertion lacking SubjectConfirmationData/InResponseTo, allowing a response issued by one trusted IdP to be bound to SP state created for another IdP and bypass flows that route users to a specific IdP, including deployments that set enable_unsolicited to false. This issue is fixed in versions 2.4.7 and 2.5.2.	7.1	More Details
CVE-2026-63833	In the Linux kernel, the following vulnerability has been resolved: ntfs3: reject direct userspace writes to reserved $LX* xattrs NTFS3 uses $LXUID, $LXGID, $LXMOD and $LXDEV as internal WSL permission metadata and reloads them into i_uid, i_gid and i_mode from ntfs_get_wsl_perm(). Because the empty-prefix xattr handler also lets file owners call setxattr() on these names directly, an unprivileged writer on a writable ntfs3 mount can plant root ownership and S_ISUID on their own file and gain euid 0 after inode reload. Reject direct userspace writes to the reserved $LX* names. Internal ntfs3 metadata updates are unchanged because ntfs_save_wsl_perm() writes them via ntfs_set_ea() directly. [almaz.alexandrovich@paragon-software.com: added an additional check for non privileged users]	7.1	More Details
CVE-2026-55550	NextCRM is open-source customer relationship management (CRM) software. The CRM product catalog is an organization-wide business object. Normal application server actions restrict product creation, update, and deletion to `manager` and `admin` roles. However, in version 0.12.1, the MCP product tools expose the same write operations through `/api/mcp/mcp` using user-generated Bearer tokens and do not enforce role checks. Any authenticated low-privileged user who can generate an MCP API token can create, modify, archive, or soft-delete products in the shared CRM product catalog. Version 0.12.3 contains a fix.	7.1	More Details
CVE-2026-64095	In the Linux kernel, the following vulnerability has been resolved: batman-adv: bla: avoid double decrement of bla.num_requests The bla.num_requests is increased when no request_sent was in progress. And it is decremented in various places (announcement was received, backbone is purged, periodic work). But the check if the request_sent is actually set to a specific state and the atomic_dec/_inc are not safe because they are not atomic (TOCTOU) and multiple such code portions can run concurrently. At the same time, it is necessary to modify request_sent (state) and bla.num_requests atomically. Otherwise batadv_bla_send_request() might set request_sent to 1 and is interrupted. batadv_handle_announce() can then set request_sent back to 0 and decrement num_requests before batadv_bla_send_request() incremented it. The two operations must therefore be locked. And since state (request_sent) and wait_periods are only accessed inside this lock, they can be converted to simpler datatypes. And to avoid that the bla.num_requests is touched by a parallel running context with a valid backbone_gw reference after batadv_bla_purge_backbone_gw() ran, a third state "stopped" is required to correctly signal that a backbone_gw is in the state of being cleaned up.	7.1	More Details
CVE-2026-64111	In the Linux kernel, the following vulnerability has been resolved: lsm: hold cred_guard_mutex for lsm_set_self_attr() Just as proc_pid_attr_write() already does before calling the LSM hook. This only matters for SELinux and AppArmor which check whether the process is being ptraced and if so, whether to allow the transition.	7.1	More Details
	Authorization Bypass Through User-Controlled Key (CWE-639) in Kibana can lead to unauthorized information		

CVE-2026-56147	disclosure and case attachment integrity compromise via Privilege Abuse (CAPEC-122). An inconsistency in Kibana's file access authorization logic allows a low-privileged authenticated user to retrieve, modify, and delete case attachments that belong to feature areas they are not authorized to access. Because the access control check and the resource retrieval use different resolution mechanisms, an authenticated attacker with limited file management permissions can obtain the contents of, modify, or delete protected case attachments — such as those associated with Security Solution cases — without holding the privileges required to access those features.	7.1	More Details
CVE-2026-52890	Wekan is open source kanban built with Meteor. Prior to 9.31, Wekan allows a logged-in board member to insert an attachment document through the /attachments/insert DDP method with attacker-controlled versions.original.path and versions.original.storage fields. The server/permissions/attachments.js insert rule checks only board write access, and FileStoreStrategyFilesystem.getReadStream() in models/lib/fileStoreStrategy.js streams the stored path without a storage-root containment check, allowing arbitrary file reads and denial of service through special files such as /dev/zero. This issue is fixed in version 9.31.	7.1	More Details
CVE-2026-58598	Concurrent execution using shared resource with improper synchronization ('race condition') in Windows Backup Engine allows an authorized attacker to elevate privileges locally.	7.0	More Details
CVE-2026-9046	A potential insecure permissions vulnerability was reported in Legion Zone and the Lenovo App Store Windows applications, distributed exclusively in the Chinese market, that when installed on a non-system partition, could allow a local user to execute arbitrary code.	7.0	More Details
CVE-2026-60063	An out-of-bounds write vulnerability in the Productivity Suite allows a local attacker to trigger kernel memory corruption via a crafted IOCTL request, potentially resulting in privilege escalation or system instability.	7.0	More Details
CVE-2026-61389	An out-of-bounds write vulnerability in the Productivity Suite allows a local attacker to trigger kernel memory corruption via a crafted IOCTL request, potentially resulting in privilege escalation or system instability.	7.0	More Details
CVE-2026-53410	A time-of-check to time-of-use (TOCTOU) race condition in the installation and uninstallation process of certain Zoom Clients for Windows could allow an authenticated local user to escalate privileges.	7.0	More Details
CVE-2026-51385	An issue in safishamsi Open-Source GRAPHIFY v.0.3.2 through v0.4.29 allows a remote attacker to execute arbitrary code via the validate_url, safe_fetch, _build_opener, _fetch_html and _download_binary functions.	6.9	More Details
CVE-2026-52888	NocoBase is an AI-powered no-code/low-code platform for building business applications and enterprise solutions. In 2.0.59 and earlier, NocoBase @nocobase/plugin-collection-sql used the checkSQL() function in packages/plugins/@nocobase/plugin-collection-sql/src/server/utils.ts with an incomplete keyword blacklist that did not restrict PostgreSQL system catalog tables such as pg_shadow, pg_roles, and pg_stat_activity, allowing an admin-role user to read password hashes and database metadata through the SQL Collection feature. This vulnerability is fixed in 2.1.0-alpha.46.	6.8	More Details
CVE-2026-48009	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, a low-privilege admin user with user_recovery:read ACL can take over any admin account by triggering POST /api/_action/user/user-recovery, reading the password recovery hash through POST /api/search/user-recovery, and using PATCH /api/_action/user/user-recovery/password; the root cause is that src/Core/System/User/Recovery/UserRecoveryDefinition.php exposes the hash field through the Admin API without ApiAware(false) or ReadProtection. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	6.8	More Details
CVE-2026-46404	BigBlueButton is an open-source virtual classroom. Prior to 3.0.23, the presentation URL validation did not properly restrict access to site local and link local addresses. The redirect following logic now pins resolved IPs. This issue is fixed in version 3.0.23.	6.8	More Details
CVE-2026-12283	Amazon Athena is a serverless, interactive query service that lets you analyze data directly in Amazon S3 using standard SQL. Athena Query Federation is a feature that allows you to connect to data sources outside of Amazon S3 like DynamoDB, Azure Synapse, and custom connectors using standard SQL syntax. Improper neutralization of special elements used in an SQL command in the Synapse connector in Amazon aws-athena-query-federation v2022.20.1 through v2026.19.1 might allow an authenticated remote user to execute injected read-only SQL queries that return unintended data from the connected database via a crafted table name. To remediate this issue, users should upgrade to version v2026.21.1 or later.	6.8	More Details
CVE-2026-54497	view_component is a framework for building reusable, testable, and encapsulated view components in Ruby on Rails. From 4.0.0 until 4.12.0, ViewComponent::Base instances retain render-scoped objects across calls to render_in; if the same component, collection, or spacer component instance is reused across requests, users, tenants, or threads, later renders can use stale helpers, controller, request, view_flow, format/variant details, and slot child context from an earlier render. This can cause authorization-aware components to render privileged UI for a lower-privileged user, generate links using a stale Host header, leak slot/helper state, and	6.8	More Details

	mix request context under concurrent rendering. This issue is fixed in version 4.12.0.		
CVE-2026-62843	File Browser is a file managing interface for uploading, deleting, previewing, renaming, and editing files within a specified directory. From 2.63.6 to 2.63.16, File Browser's archive builder uses strings.ReplaceAll(nameInArchive, "\", "/"), which turns a POSIX filename such as ..\..\evil.sh into the archive entry ../../evil.sh, allowing a user with upload permission to plant a backslash-named file that escapes the extraction directory when another user downloads and extracts the generated zip or tar archive. This issue is fixed in version 2.63.17.	6.8	More Details
CVE-2026-15927	A flaw was found in Red Hat Quay's repository-level mirror configuration feature. The POST and PUT handlers in endpoints/api/mirror.py accept an external_reference parameter without SSRF validation, unlike the organization-level mirror handlers which apply validate_external_registry_url(). A repository administrator can supply a crafted hostname that causes the Quay mirror worker to make requests via Skopeo to internal network services, cloud metadata endpoints, or other resources not intended to be reachable from the Quay application.	6.8	More Details
CVE-2026-49501	Dell PowerScale OneFS versions 9.5.0.0 through 9.10.1.7, and versions 9.11.0.0 through 9.13.0.2 contains an Improper Privilege Management vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to Elevation of privileges.	6.7	More Details
CVE-2026-55410	NocoBase is an AI-powered no-code/low-code platform for building business applications and enterprise solutions. Prior to 2.1.19, NocoBase @nocobase/plugin-backups restored PostgreSQL backups by interpolating the database.schema value from _metadata.json into shell command strings executed with Node.js child_process.exec(), allowing a backup-management user restoring a crafted backup to execute commands as the NocoBase server process. This vulnerability is fixed in 2.1.19.	6.7	More Details
CVE-2026-15370	A flaw was found in libssh. During SFTP server directory listing, the longname field is constructed with unsafe concatenation into a fixed-size stack buffer. When a client causes the server to list attacker-controlled filenames, sufficiently long names can overflow that stack buffer and may lead to crashes or possible code execution on the server.	6.7	More Details
CVE-2026-63729	The SyncTeX parser (synctex_parser.c) shipped with TeX Live and embedded by downstream consumers such as GNOME Evince contains a heap use-after-free vulnerability that allows attackers to crash applications or potentially execute arbitrary code by supplying a malformed .synctex or .synctex.gz file. A malformed SyncTeX file can construct a ref node with a NULL parent pointer, causing the replacement routine to fail to detach the node from its sibling chain, which triggers recursive freeing of live tree nodes and leaves dangling pointers that are later accessed by the parser during document load.	6.6	More Details
CVE-2026-58554	Permission control vulnerability in the Settings module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	6.6	More Details
CVE-2026-58555	Permission bypass vulnerability in the card module. Impact: Successful exploitation of this vulnerability may affect availability.	6.6	More Details
CVE-2026-62210	OpenClaw versions before 2026.6.1 contain a denial of service vulnerability where remote media URLs can trigger slow-read attacks that exhaust gateway worker resources. Attackers with access to configured input paths can supply remote media URLs that consume gateway resources and reduce availability.	6.5	More Details
CVE-2024-58359	SurrealDB versions before 2.1.0 contain a denial of service vulnerability in the sorting mechanism when using ORDER BY rand() clause. Authorized clients can execute queries with ORDER BY rand() to trigger a panic in the sorting function, crashing the server.	6.5	More Details
CVE-2026-62208	OpenClaw before 2026.6.5 could forward Authorization headers during MCP SSE redirects. When the affected feature is enabled and reachable, a lower-trust caller or configured input path could execute or persist actions beyond the caller's intended authorization. Impact depends on the operator's configuration and whether lower-trust input can reach the affected path.	6.5	More Details
CVE-2026-44251	Wazuh is a free and open source platform used for threat prevention, detection, and response. In versions 3.0.0 and above, prior to 4.14.5, a size_t integer underflow in os_crypto/shared/msgs.c:389 allows any enrolled Wazuh agent to crash the wazuh-remoted process on the manager, immediately disconnecting all agents from the manager. A second code path reached by the same underflow may allow heap memory corruption. This issue has been fixed in version 4.14.5.	6.5	More Details
CVE-2026-45295	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.219, the open tracking endpoint `GET /thread/read/{conversation_id}/{thread_id}` allows unauthenticated attackers to enumerate valid conversation and thread IDs, and modify thread state (`opened_at` timestamp) without any authentication. Version 1.8.219 patches the issue.	6.5	More Details
CVE-2024-58357	SurrealDB versions before 2.1.0 contain an uncaught exception vulnerability in the rand::time() function that panics when unwrap is called on a None result from timestamp_opt. Authorized clients can repeatedly invoke rand::time() to reliably trigger server panics and cause denial of service.	6.5	More Details

CVE-2026-62237	Grav before 2.0.4 contains a regular expression denial of service (ReDoS) vulnerability in the regex_replace filter and function, which are allowlisted in the Twig content sandbox. When Twig processing in page content is enabled (security.twig_content.process_enabled: true, disabled by default), an authenticated page editor can supply a catastrophically backtracking PCRE pattern that is passed directly to PHP's preg_replace(), causing unbounded CPU consumption and denial of service to the web server process.	6.5	More Details
CVE-2026-62213	OpenClaw versions before 2026.5.27 contain a token leakage vulnerability in MS Teams outbound requests that allows lower-trust callers to expose Bot Framework tokens. Attackers can access configured input paths to retrieve credentials that should remain within the trusted boundary.	6.5	More Details
CVE-2026-62214	OpenClaw versions before 2026.5.28 Bot Framework contains an improper input validation vulnerability that allows lower-trust callers to expose bot tokens and credentials by failing to properly validate serviceUrl parameters. Attackers can supply malicious serviceUrl values through configured input paths to retrieve sensitive authentication data outside the trusted boundary.	6.5	More Details
CVE-2025-32781	Apollo is a reliable configuration management system suitable for microservice configuration management scenarios. Prior to 2.5.0, Apollo Portal does not verify application and namespace permissions when an authenticated user requests a release by ID through GET /envs/{env}/releases/{releaseId} while configView.memberOnly.envs is enabled, allowing a low-privileged Portal user who obtains or guesses a valid releaseId to read configuration data from other applications and namespaces without calling UserPermissionValidator.shouldHideConfigToCurrentUser(...). This issue is fixed in version 2.5.0.	6.5	More Details
CVE-2026-48022	@hapi/wreck is an HTTP client utility. Prior to 18.1.2, Wreck strips credential headers including Authorization, Cookie, and Proxy-Authorization before following a cross-origin redirect, but the origin check compares hostnames only and ignores scheme and port, so credentials are forwarded intact across same-host port changes and HTTPS-to-HTTP downgrades, allowing a co-tenant on an adjacent port or a network-position attacker capable of forging a redirect to capture bearer tokens, session cookies, and proxy credentials and impersonate the victim against the upstream service. This issue is fixed in version 18.1.2.	6.5	More Details
CVE-2024-58364	SurrealDB versions before 1.2.1 contain an uncaught exception handling vulnerability in span rendering when parsing queries with errors on line terminator characters. Authorized clients can submit malformed queries that trigger a panic in the span rendering code, crashing the server and causing denial of service.	6.5	More Details
CVE-2026-56434	NGINX Plus and NGINX Open Source have a vulnerability in the ngx_http_ssi_module module. This vulnerability may exist when the Server-Side Includes (SSI), proxy_pass, and proxy_buffering off directives are configured. With this configuration, an unauthenticated attacker with man-in-the-middle (MITM) ability to control responses from an upstream server may be able to cause a use-after-free in the NGINX worker process. This issue may lead to limited modification of memory or a restart of the NGINX worker process. Impact: This vulnerability may allow remote attackers to have limited control to modify memory contents or restart the NGINX worker process. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	6.5	More Details
CVE-2026-14503	The pCloud WP Backup plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 2.0.3 via the wp2pcl_ajax_process_request_inner. This makes it possible for authenticated attackers, with subscriber-level access and above, to extract force generation of a full-site backup archive written to a publicly accessible directory, exposing wp-config.php database credentials, WordPress secret salts, and the complete PHP source tree. The resulting archive is deposited in the plugin's unprotected tmp/ directory at a predictable URL, making the extracted data accessible to unauthenticated visitors once the backup is triggered.	6.5	More Details
CVE-2026-54562	Cloudreve is a self-hosted file management and sharing system. Prior to 4.16.1, Cloudreve's remote download workflow accepts user-supplied URLs at POST /api/v4/workflow/download and passes them to the configured downloader without blocking loopback, localhost, IPv6 localhost, or redirect-to-loopback targets, allowing a non-admin user with remote download permission to fetch internal-only URLs and read the response after it is imported into the user's own files. This issue is fixed in version 4.16.1.	6.5	More Details
CVE-2024-58361	SurrealDB versions before 2.0.4 contain an uncaught exception handling vulnerability in the parser error rendering code when processing empty strings. Authorized clients can execute malformed queries with empty string conversions to record, duration, or datetime types that cause a panic in error rendering, crashing the server.	6.5	More Details
CVE-2026-63759	SurrealDB before 3.1.0 fails to enforce recursion depth limits in the type/kind parser when processing nested type annotations. Authenticated attackers can send queries with deeply nested type annotations to exhaust server memory and crash the process.	6.5	More Details
CVE-2024-58365	SurrealDB versions before 1.2.0 contain an uncaught exception vulnerability in the query executor when processing calls to nonexistent built-in functions. Authorized clients can craft pre-parsed queries invoking nonexistent functions to trigger a panic that crashes the server.	6.5	More Details
CVE-2026-12395	The WP Job Portal WordPress plugin before 2.5.5 does not properly sanitize and escape a parameter before using it in a SQL query, allowing authenticated users with a subscriber-level (self-registerable) account to perform SQL injection attacks.	6.5	More Details

CVE-2026-36425	An issue in OPSWAT AppRemover Driver (ardrv.sys) v2017.10.02.1551 and earlier in IOCTL handler 0x2420031. Any local user can open the device and send process termination requests without privilege validation.	6.5	More Details
CVE-2026-33731	WWBN AVideo is an open source video platform. In versions prior to 29.0, the Authorize.Net webhook handler at plugin/AuthorizeNet/webhook.php contains a signature verification bypass that allows an attacker to forge webhook requests with arbitrary payment amounts and target user IDs. By supplying a valid transaction ID from a small legitimate purchase, the attacker bypasses signature validation and credits arbitrary wallet balances to any user account via attacker-controlled payload fields. Three flaws combine into an exploit chain: signature bypass via OR logic (webhook.php:33), payload values override API-fetched values (AuthorizeNet.php:169-171, webhook.php:44-48) and a missing approval check (webhook.php:61-75). By forging payment metadata, an attacker can credit arbitrary amounts to any user's wallet without a corresponding payment and include a plans_id to activate premium subscriptions (webhook.php:86-134), enabling free access to all paid and premium content and causing direct revenue loss to the platform owner. This issue has been fixed in version 29.0.	6.5	More Details
CVE-2026-11889	SALTO ProAccess Space software using the tenancy feature / logical partition is vulnerable to a privilege escalation attack that could allow an authorized attacker to access any space managed by the affected product.	6.5	More Details
CVE-2026-12898	The All-in-One WP Migration and Backup WordPress plugin before 7.106 does not properly sanitise a user-supplied value before using it to build a file path, allowing unauthenticated attackers to create or append a log file in arbitrary locations outside its intended storage directory.	6.5	More Details
CVE-2026-52892	Wekan is open source kanban built with Meteor. Prior to 9.32, Wekan REST handlers in server/models/customFields.js use read-level Authentication.checkBoardAccess instead of write-level Authentication.checkBoardWriteAccess for mutating custom-field routes. A read-only board member can call POST, PUT, and DELETE handlers for /api/boards/:boardId/custom-fields and custom-field dropdown items to create, update, or delete board custom fields. This issue is fixed in version 9.32.	6.5	More Details
CVE-2026-53447	Wekan is open source kanban built with Meteor. Prior to 9.35, the Wekan cloneBoard Meteor method in models/import.js uses caller-supplied sourceBoardId to build a board export through models/exporter.js without invoking canExport() or checking source-board membership. Any authenticated user who knows a private board ID can clone the board into their own account and read its cards, comments, attachments, member information, and activities. This issue is fixed in version 9.35.	6.5	More Details
CVE-2026-12941	The MultiVendorX – WooCommerce Multivendor Marketplace AI Powered Solutions plugin for WordPress is vulnerable to generic SQL Injection via the 'order_by' parameter in all versions up to, and including, 5.0.9 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with subscriber-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. This vulnerability is exploitable by any authenticated subscriber-level user when the plugin's store approval setting is configured to automatically approve store owners (described as the default), as this allows any logged-in user to self-register as a store_owner via the public Stores REST endpoint, thereby obtaining the edit_stores capability required to reach the vulnerable transactions endpoint.	6.5	More Details
CVE-2026-47084	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. The LOCALDELETE command bypassed ACL checks. An authenticated but non-admin user could invoke the admin-only LOCALDELETE IMAP command and delete mailboxes for which they had no permissions.	6.5	More Details
CVE-2026-46514	Frogman provides headless PBX control through MCP and HTTP API. Prior to 1.6.2, fm_reset_password in Tools/ResetPassword.php:48-53 returned a plaintext password and fm_add_extension in Tools/AddExtension.php:172 returned a plaintext secret; Frogman.class.php:2207-2211 used auditOutcome to JSON-encode those responses into oc_audit_log.detail, allowing any PERM_READ caller with access to fm_audit_search to recover the stored credentials. This issue is fixed in version 1.6.2.	6.5	More Details
CVE-2026-33754	Wazuh is a free and open source platform used for threat prevention, detection, and response. In versions 3.9.0 and above, prior to 4.14.5, a remote attacker can trigger memory exhaustion in the cluster protocol parser by sending a crafted message header with an arbitrarily large payload length. The length is trusted before authentication/decryption and used directly to allocate memory, allowing unauthenticated denial of service of the cluster service. This issue has been fixed in version 4.14.5.	6.5	More Details
CVE-2026-12684	The Customer Reviews for WooCommerce WordPress plugin before 5.113.0 does not perform authentication, capability, or nonce checks on one of its media upload AJAX actions when the review media attachment feature is enabled, allowing unauthenticated users to upload media files (bounded to an image and video allowlist) to the Media Library and create attachment posts, leading to media library pollution and disk space exhaustion.	6.5	More Details
CVE-2026-13754	The Tickera – Sell Tickets & Manage Events plugin for WordPress is vulnerable to generic SQL Injection via the 's' parameter in all versions up to, and including, 3.6.0.0 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with custom-level access and above, to append additional SQL queries into already	6.5	More Details

	existing queries that can be used to extract sensitive information from the database.		
CVE-2026-13767	The Quiz Master Next plugin for WordPress is vulnerable to SQL Injection via stored quiz page data in versions up to, and including, 11.2.0. This is due to insufficient escaping on the user-supplied 'pages' parameter persisted by the qsm_ajax_save_pages() AJAX handler (sanitize_text_field only) and lack of sufficient preparation on the existing SQL query built in qsm_options_questions_tab_content() at line 143, where the stored page IDs are interpolated into an IN() clause via implode() with no $wpdb->prepare() and no integer casting. This makes it possible for authenticated attackers, with Author-level access and above (who can own a quiz they are entitled to edit), to plant a SQL payload that is executed second-order whenever any user (including an administrator) views the quiz's Questions tab, allowing them to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-15022	The Tutor LMS – eLearning and online course solution plugin for WordPress is vulnerable to generic SQL Injection via Stored Quiz Answer Array in all versions up to, and including, 4.0.0 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with custom-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. The payload is stored at quiz-attempt time via the wp_ajax_tutor_quiz_abandon handler, but the injected SQL executes only when a privileged user or Tutor REST API key holder requests the /wp-json/tutor/v1/quiz-attempt-details/{id} endpoint, making this a second-order (stored) injection chain.	6.5	More Details
CVE-2024-58360	stoatchat versions before 0.7.8 fail to enforce account creation restrictions including invite-only mode, email verification, captcha, and shield verification. Attackers can create unlimited accounts with unverified email addresses, increasing denial-of-service risk and compromising service integrity.	6.5	More Details
CVE-2026-47729	Squid is a caching proxy for the Web. Prior to 7.6, due to an improper validation of syntactic correctness of input in the FTP gateway (src/clients/FtpGateway.cc), Squid is vulnerable to an out-of-bounds read: when a listing entry date in the TypeA or TypeB directory-listing formats is not followed by a filename, parsing was not restricted to the input buffer, so a trusted client accessing a misbehaving FTP server through Squid's gateway feature could read memory from random unrelated transactions. This issue is fixed in version 7.6.	6.5	More Details
CVE-2026-44596	Yamcs is a mission control framework. Prior to 5.12.7, the authentication endpoint POST /auth/token in yamcs-core, handled by yamcs-core/src/main/java/org/yamcs/http/auth/AuthHandler.java, lacked any rate limiting, account lockout, or failed-attempt throttling, so an unauthenticated remote attacker could perform unlimited password-guessing attempts against any user account, significantly increasing the risk of successful brute-force attacks. This issue is fixed in versions 5.12.7 and 5.13.0.	6.5	More Details
CVE-2026-55440	Microsoft UFO open-source framework for intelligent automation across devices and platforms. Prior to 3.0.7, the COMMAND_RESULTS handler in ufo/server/ws/handler.py called get_or_create_session in ufo/server/services/session_manager.py without owner_client_id, allowing an authenticated client to create an unowned attacker-chosen session_id such as constellation_task_id = f"{task_name}@{task_id}" and deny the legitimate owner or exhaust memory with phantom sessions. This issue is fixed in version 3.0.7.	6.5	More Details
CVE-2026-12973	The PayPlus Payment Gateway WordPress plugin before 8.2.2 does not perform authorization or order-ownership validation in one of its AJAX actions available to unauthenticated users, allowing them to disclose the secret order key of arbitrary WooCommerce orders and, under some configurations, to modify order statuses.	6.5	More Details
CVE-2026-15813	A vulnerability was found in the network packet de-fragmentation engine of kronosnet (Version affected <= 1.34). The internal reassembly code does not properly validate sequence numbers of incoming payload fragments. An attacker can exploit this lack of verification by transmitting malformed packets with corrupted sequence parameters. Under specific conditions, this forces the packet processing layer to parse data outside the designated bounds of the internal memory structures, causing an out-of-bounds memory access or heap corruption. This behavior can result in sudden application crashes or system instability.	6.5	More Details
CVE-2026-63737	SurrealDB versions before 3.1.5 contain a denial of service vulnerability where authenticated users can crash the server with queries containing long chains of operators. Attackers can submit queries with tens of thousands of chained operators that create unbounded expression trees, causing stack overflow during query processing and aborting the entire process.	6.5	More Details
CVE-2026-63740	SurrealDB versions before 3.1.4 fail to properly enforce SELECT permissions on array elements (field.*) for record users, leaking denied array elements instead of hiding them. Attackers with record scope access can read array elements that element-level permissions should deny by exploiting incorrect index handling during permission filtering.	6.5	More Details
CVE-2024-58369	SurrealDB versions before 1.1.1 fail to properly validate invocation of custom parameters and functions at root or namespace levels, causing server panic. Authorized clients can invoke these entities at unsupported levels to crash the SurrealDB server, resulting in denial of service.	6.5	More Details
CVE-2024-58370	SurrealDB versions before 1.1.0 fail to enforce recursion depth limits when parsing nested SurrealQL statements including IF, RELATE, and attribute access idioms. Authorized attackers can submit queries with excessive nesting depth to cause stack overflow and crash the server.	6.5	More Details

CVE-2026-58481	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.12.2, `AgentRuntime` promises scoped file access under a configured sandbox `basePath`, but its path containment checks use raw string prefix tests. A sandbox base such as `/tmp/network-ai-sandbox` also matches a sibling path such as `/tmp/network-ai-sandbox_evil/secret.txt`. An agent/user that can call `AgentRuntime.readFile()` or `AgentRuntime.listDir()` can read or list files outside the intended sandbox when the target path is in a sibling directory sharing the base path prefix. This breaks the documented sandbox boundary. The issue is fixed in v5.12.2. `SandboxPolicy.resolvePath()` and `isPathAllowed()` now use separator-anchored prefix checks (`resolved === base || resolved.startsWith(base + path.sep)`) for both the allow-list and block-list. A sibling directory that merely shares a name prefix (e.g. `/srv/app-evil` vs base `/srv/app`) is no longer treated as in-scope.	6.5	More Details
CVE-2026-55645	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a vulnerability concerning the processing of Client Control PDUs. During the RDP connection sequence, the parser does not perform sufficient length validation before reading specific data fields from the network stream. A remote, unauthenticated attacker could potentially exploit this flaw by sending a specially crafted, truncated Client Control PDU. Due to missing bounds checks, the xrdp process may perform out-of-bounds memory reads, which can result in the termination of the service (Denial of Service). However, since xrdp forks a new process for each connection by default, an out-of-bounds read causing a process crash is unlikely to bring down the entire xrdp service.This issue has been fixed in version 0.10.6.1.	6.5	More Details
CVE-2026-47276	In nanomq versions 0.24.11 and earlier, a NULL pointer dereference in `properties_parse()` allows an authenticated attacker to crash the NanoMQ broker by sending a POST request to `/api/v4/mqtt/publish` with `user_properties` as a JSON array instead of a JSON object. The crash occurs because `strlen()` is called on a NULL `item->string` pointer when iterating over array elements. An authenticated attacker can exploit this to crash the NanoMQ broker process. This is patched in version 0.24.14.	6.5	More Details
CVE-2026-35217	NanoMQ contains a protocol-semantics flaw in its MQTT v5 `SUBSCRIBE` handling: if a subscription entry is missing the final 1-byte `Subscription Options` field, the broker may still accept the malformed packet and install the subscription into internal broker state. Under a specific packet-length construction, the same parser flaw also causes a 1-byte out-of-bounds read that crosses the real heap allocation boundary and is detected by ASAN as a `heap-buffer-overflow`. If the consumed byte happens to look acceptable, NanoMQ may continue and append the malformed subscription entry into its internal `subinfol` state. In that case, a `SUBSCRIBE` packet that should be rejected by MQTT rules is instead treated as a successful subscription. Whether ASAN reports the bug does not depend on MQTT's logical `remain` boundary; it depends on whether the read crosses the real heap allocation boundary of the underlying message buffer. In other words, these are not two unrelated issues. They are two manifestations of the same parsing defect: by default, it appears as a semantic vulnerability, and under suitable input conditions, it also becomes a verifiable out-of-bounds read vulnerability.	6.5	More Details
CVE-2026-51026	Directory Traversal vulnerability in FileThingie v.2.5.7 allows a remote attacker to obtain sensitive information via a crafted request.	6.5	More Details
CVE-2026-63091	ProFTPD before 1.3.9c and 1.3.10rc3 contains a signed integer overflow vulnerability in the mod_sftp module's SCP size-record parser that allows authenticated low-privilege attackers to bypass ASLR by sending a crafted file size value of UINT64_MAX, which results in a negative off_t value. Attackers can exploit the subsequent conversion to uint32_t, causing an approximately 4 GB requested read length and forcing the server to read beyond the end of the SSH channel data and write overread process memory into the uploaded file. In tested configurations, the disclosed data contains libc, libcrypto, and PIE pointers sufficient to derive their randomized base addresses, thereby bypassing ASLR and enabling reliable exploitation of memory corruption vulnerabilities in the same process.	6.5	More Details
CVE-2026-45139	CI4MS is a CodeIgniter 4-based content management system skeleton. Prior to version 0.31.9.0, the Fileeditor module enforces an extension allowlist (`['css','js','html','txt','json','sql','md']`) on content-write operations (`saveFile`, `createFile`), but two destructive endpoints — `deleteFileOrFolder` and `renameFile` — never validate the extension of the *source* path. A backend user with file-editor permissions can therefore unlink or rename any file inside the project root that is not explicitly listed in the small `$hiddenItems` blocklist. Critical framework files such as `app/Config/Routes.php`, `app/Config/App.php`, `app/Config/Database.php`, `app/Config/Filters.php`, `public/index.php`, and `public/.htaccess` all live outside that blocklist and can be destroyed, producing a persistent denial of service that requires filesystem-level redeployment to recover. Version 0.31.9.0 patches the issue.	6.5	More Details
CVE-2026-16277	A stack-based buffer overflow was found in rpcbind's rpcinfo utility. When querying a remote rpcbind service with `rpcinfo -l`, address information returned by the server is copied into a fixed-size buffer without sufficient bounds checking. A malicious or compromised rpcbind server could use this flaw to crash the rpcinfo client, resulting in a denial of service. The highest threat from this vulnerability is to system availability.	6.5	More Details
	Strands Agents is an open-source Python SDK for building and running AI agents. The strands-agents-tools package provides pre-built tools for use with the SDK, including the elasticsearch_memory tool for agent memory storage. We identified CVE-2026-15746, a server-side request forgery (SSRF) issue in the		

CVE-2026-15746	elasticsearch_memory tool. The tool exposed its connection parameters (es_url, cloud_id, api_key) as fields the large language model (LLM) could control through the tool schema. When a caller omitted the api_key parameter, the tool fell back to the operator's ELASTICSEARCH_API_KEY environment variable and sent it to whichever host the LLM specified. A crafted prompt could cause the tool to connect to a threat-actor-controlled server and disclose the operator's Elasticsearch API key in the Authorization header. We recommend you upgrade to strands-agents-tools version 0.7.0 or later. As a precautionary measure, we recommend all operators rotate their ELASTICSEARCH_API_KEY, even if there is no indication the credential was exposed.	6.5	More Details
CVE-2026-54171	Excon is usable, fast, simple HTTP 1.1 for Ruby. Prior to 1.5.0, Excon's RedirectFollower middleware failed to strip additional sensitive headers when following redirects and did not provide a custom list of headers to strip. This could cause inadvertent leakage of sensitive data when the initial request includes header information that is not intended for the new target. This issue is fixed in version 1.5.0.	6.5	More Details
CVE-2026-16215	A security flaw has been discovered in geex-arts django-jet up to 1.0.8. This impacts an unknown function of the component OAuth Credential Revoke Handler. Performing a manipulation results in missing authorization. The attack is possible to be carried out remotely. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	6.5	More Details
CVE-2026-63755	SurrealDB before 3.1.0 evaluates user-supplied WHERE clauses in SELECT statements (and SET/MERGE/CONTENT/PATCH clauses in UPDATE, UPSERT, INSERT ON DUPLICATE KEY UPDATE, and RELATE update-variant statements) against full record data before enforcing PERMISSIONS FOR SELECT WHERE restrictions. An authenticated user — including Record and Scope users — can exploit this ordering flaw to read the full contents of any table in the database they are authenticated against, bypassing table-level permission checks. Exfiltration is most direct when scripting functions are enabled (--allow-scripting), but is also possible via SurrealQL's THROW statement and timing-based side channels without scripting. The vulnerability is confined to the attacker's current database and does not cross namespace or database isolation boundaries.	6.5	More Details
CVE-2026-63754	SurrealDB versions before 3.1.0 contain a denial of service vulnerability where malicious LIVE queries with WHERE clauses that evaluate to errors cause all CREATE, UPDATE, and DELETE operations on the watched table to fail. An authenticated user with only select permission can prevent write operations on a table for any user, including root, by registering a LIVE query that triggers evaluation errors until the query is killed or the session ends.	6.5	More Details
CVE-2026-63746	SurrealDB versions before 3.1.0 fail to enforce table SELECT permissions when traversing graph edges or back-references. Authenticated users can read records from any table reachable through graph edges regardless of the target table's PERMISSIONS FOR select clause.	6.5	More Details
CVE-2026-63741	SurrealDB versions before 3.1.0 fail to validate DEFINE NAMESPACE or DEFINE DATABASE permissions when processing USE NS and USE DB statements. Unauthenticated attackers can create arbitrary namespaces and databases by issuing USE commands, bypassing authorization checks in the RPC use method and SurrealQL executor.	6.5	More Details
CVE-2026-52865	When NGINX Ingress Controller processes Ingress or TransportServer resources, an authenticated, remote attacker with permission to create or modify Ingress or TransportServer resources can cause the NGINX Ingress Controller process to terminate. Impact: The NGINX Ingress Controller control plane process terminates and enters a persistent crash loop while the malformed Ingress or TransportServer resource remains in the cluster. This vulnerability allows a remote, authenticated attacker with at least Ingress or TransportServer resource write access to cause a denial-of-service (DoS) on the NGINX Ingress Controller system. There is no data plane exposure; this is a control plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	6.5	More Details
CVE-2025-45870	LogicalDOC Enterprise up to and for v9.1.1 is vulnerable to Local File Inclusion (LFI) in the OnlyOfficeEditor servlet class, allowing authenticated user to exploit path traversal flaws in the fileExt parameter, enabling unauthorized access to sensitive files outside the designated directories.	6.5	More Details
CVE-2026-47183	Zeroconf is a pure Python implementation of multicast DNS service discovery. Prior to 0.149.6, DNSIncoming._log_exception_debug and the four QuietLogger exception-dedup methods stored an unbounded _seen_logs dictionary keyed by attacker-influenced IncomingDecodeError messages, retaining sys.exc_info() tracebacks whose frame locals kept raw packet self.data buffers and allowing unauthenticated hosts on the local link over UDP/5353 (224.0.0.251 / ff02::fb) to drive memory growth until mDNS-dependent features degrade or the process is OOM-killed. This issue is fixed in version 0.149.6.	6.5	More Details
CVE-2026-63139	Uncontrolled Resource Consumption (CWE-400) in Kibana can lead to denial of service via Excessive Allocation (CAPEC-130). An authenticated low-privileged user can exploit an uncontrolled resource consumption vulnerability in Kibana's Canvas functionality by sending a specially crafted request, causing the Kibana server process to terminate and resulting in a denial of service for all users of the affected Kibana instance.	6.5	More Details

CVE-2026-48045	Zeroconf is a pure Python implementation of multicast DNS service discovery. Prior to 0.149.12, AsyncListener.handle_query_or_defer retained every truncated TC-bit incoming query, each up to _MAX_MSG_ABSOLUTE = 8966 bytes, in self._deferred[addr] and armed a per-address timer in self._timers[addr] without capping the per-address list or distinct addr keys, allowing unauthenticated hosts on the local link over UDP/5353 (224.0.0.251 / ff02::fb) to spoof sources, grow _deferred and _timers, and cause memory exhaustion and quadratic CPU burn. This issue is fixed in version 0.149.12.	6.5	More Details
CVE-2026-63080	Aptabase through commit 5a89368 contains a SQL injection vulnerability in the ClickHouse query backend that allows authenticated attackers to read event data across all tenants by injecting unsanitized filter parameters into Liquid SQL templates. Attackers can supply malicious values through EventName, CountryCode, OsName, DeviceModel, AppVersion, or SessionId parameters to inject a UNION ALL statement that bypasses the app_id tenant isolation filter across thirteen of the fifteen stats API endpoints.	6.5	More Details
CVE-2026-63136	Uncontrolled Resource Consumption (CWE-400) in Elasticsearch can lead to denial of service via Excessive Allocation (CAPEC-130). A user with search privileges can submit a specially crafted search request that causes a data node to exhaust available heap memory, resulting in node unavailability and cluster degradation. An attacker could leverage this vulnerability to cause cluster downtime requiring manual intervention to restore service.	6.5	More Details
CVE-2026-47408	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an Insecure Direct Object Reference. The `GET /workspaces/{workspace_id}/issues/{issue_id}/activity` endpoint is gated by `require_workspace_member(workspace_id)` and dispatches to `ActivityService.list_for_issue(issue_id)`, which executes `SELECT * FROM activity WHERE issue_id = :issue_id` with no workspace constraint. A user who is a member of any workspace can read the full activity log of any issue across the entire multi-tenant deployment. PraisonAI Platform version 0.1.4 patches the issue.	6.5	More Details
CVE-2026-6792	Missing Authorization vulnerability in Universal Software Inc. FlexCity allows Exploiting Incorrectly Configured Access Control Security Levels. This issue affects FlexCity: from 5.536.0 through 11052026.	6.5	More Details
CVE-2026-16397	Clickjacking issue in the WebExtensions component in Firefox for Android. This vulnerability was fixed in Firefox 153.	6.5	More Details
CVE-2026-63140	Reachable Assertion (CWE-617) in Elasticsearch can lead to denial of service via Input Data Manipulation (CAPEC-153). A specially crafted search request containing a null value in a specific query clause causes an internal assertion to be raised during query parsing. Because Elasticsearch treats assertion failures as fatal errors, this terminates the affected node process. A low-privileged authenticated user with read access to at least one index can exploit this condition with a single request to cause a node to terminate, disrupting search availability. In a single-node deployment this fully stops Elasticsearch; in a multi-node cluster it reduces cluster capacity for each affected node.	6.5	More Details
CVE-2024-23566	HCL Aftermarket EPC is vulnerable to brute force attacks since application doesn’t have captcha implemented. It can lead to various security issues like brute force , automated attacks & account enumeration	6.5	More Details
CVE-2026-51083	Incorrect access control in Proxmox Virtual Environment (PVE) 9.x qemu-server before 9.1.8 and 8.x before 8.4.8 allows users within limited privileges to obtain hashed passwords via the cloudinit/dump API.	6.5	More Details
CVE-2026-47411	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Versions prior to 0.1.4 have an authorization bypass enabling workspace metadata + settings tampering. The `PATCH /workspaces/{workspace_id}` endpoint is gated only by `require_workspace_member(workspace_id)` (default `min_role="member"`). Any member can rewrite the workspace's `name`, `description`, and the `settings` JSON blob. The settings field is a free-form JSON object — depending on which downstream code reads it, this becomes a configuration-injection primitive for any setting the platform exposes there. PraisonAI Platform version 0.1.4 patches the issue.	6.5	More Details
CVE-2026-59844	A flaw was found in libssh. A remote authenticated client can issue SSH_FXP_READ requests with an arbitrarily large length, causing a libssh SFTP server to allocate excessive memory and potentially exhaust it through repeated requests.	6.5	More Details
CVE-2026-59843	A flaw was found in libssh. A remote authenticated peer can advertise a zero maximum packet size in SSH_MSG_CHANNEL_OPEN, causing later channel writes to loop indefinitely and consume CPU, leading to denial of service.	6.5	More Details
CVE-2026-16461	A stack-based buffer overflow was found in rpcbind's rpcinfo utility. In rpcbdump() short mode (used by `rpcinfo -s`), version numbers from a remote RPCBPROC_DUMP reply are written into a fixed-size stack buffer without bounds checking. A user or administrator who runs `rpcinfo -s` against a malicious or compromised rpcbind endpoint could experience a crash or denial of service of the rpcinfo client.	6.5	More Details
CVE-			

2026-58559	DoS vulnerability in the vibration service. Impact: Successful exploitation of this vulnerability may affect availability.	6.5	More Details
CVE-2026-47184	Zeroconf is a pure Python implementation of multicast DNS service discovery. Prior to 0.149.7, DNSCache._async_add inserted every response record into cache, _expirations, _expire_heap, and service_cache without a cap, allowing unauthenticated hosts on the local link over UDP/5353 (224.0.0.251 / ff02::fb) to multicast valid mDNS responses with unique names and cause memory exhaustion, slower cache lookups, slower async_expire passes, and broken discovery, registration, and ServiceBrowser callbacks. This issue is fixed in version 0.149.7.	6.5	More Details
CVE-2026-61440	PraisonAI Platform before 0.1.9 fails to properly authorize label and issue-label mutations, allowing workspace members to rename and recolor shared labels and add or remove labels on owner-created issues. Attackers with workspace member privileges can exploit PATCH and POST/DELETE endpoints to alter shared label taxonomy and manipulate issue-label associations without owner or admin authorization.	6.5	More Details
CVE-2026-47180	Zeroconf is a pure Python implementation of multicast DNS service discovery. Prior to 0.149.5, DNSIncoming._decode_labels_at_offset recurses once per DNS-name compression pointer, and a single mDNS packet carrying chained pointers can trigger a RecursionError that escapes DNSIncoming.__init__, causing sustained CPU burn, log flooding, and degraded mDNS-dependent features for unauthenticated hosts on the local link over UDP/5353 (224.0.0.251 / ff02::fb). This issue is fixed in version 0.149.5.	6.5	More Details
CVE-2026-9602	Mattermost Desktop App versions <=6.2 6.0.2 5.6.13.0 fail to validate payloads sent from the Mattermost Web App to the Desktop App which allows a malicious server owner to crash the Mattermost Desktop App via changing the payload of a method to a malformed one. Mattermost Advisory ID: MMSA-2026-00678	6.5	More Details
CVE-2026-65050	Ninja Forms WordPress plugin version 3.14.8 and prior contains a missing authorization vulnerability in the render callback of the `ninja-forms/submissions-table` Gutenberg block that allows authenticated attackers with Author-level privileges to expose stored form submissions to unauthenticated visitors by embedding the block with an arbitrary formID on a published post. Attackers can retrieve the signed bearer token injected into every page visitor's browser via `wp_localize_script` and use it against the REST API submissions endpoint to access all saved form submission field values, including sensitive personally identifiable information such as names, email addresses, and phone numbers.	6.5	More Details
CVE-2026-63095	Dendrite through 0.13.8 contains an improper authorization vulnerability in the Matrix Client-Server API that allows any authenticated local user to delete third-party identifier bindings belonging to other users by submitting an arbitrary address and medium to the account deletion endpoint without ownership verification. Attackers can exploit the unverified Forget3PID handler to remove a victim's email or MSISDN binding and subsequently rebind the address through an identity server to hijack the victim's password reset flow.	6.5	More Details
CVE-2026-65051	Ninja Forms WordPress plugin version 3.14.8 contains a client-side enforcement of server-side security vulnerability that allows unauthenticated attackers to bypass all form validation by merging attacker-controlled field metadata over server-loaded form definitions before validation runs. Attackers can craft a malicious AJAX submission overriding field types, removing required flags, and disabling CAPTCHA checks through the nopriv AJAX endpoint to trigger form actions such as email notifications and database storage with unverified, attacker-controlled content.	6.5	More Details
CVE-2026-46556	FlaskBB is a Forum Software written in Python using the micro framework Flask. Prior to version 2.2.1, a Server-Side Request Forgery (SSRF) vulnerability in get_image_info() allows any authenticated user to force the server to send HTTP requests to arbitrary internal endpoints, including cloud metadata services. This is a blind SSRF with confirmed internal port scanning and internal API triggering capabilities. Version 2.2.1 patches the issue.	6.5	More Details
CVE-2026-63099	TheHive through 4.1.24 contains a broken object-level authorization vulnerability in the attachment download endpoints that allows any authenticated user to access attachments belonging to other organizations by supplying a content-hash identifier. Attackers can exploit the missing organization-scoped authorization check in AttachmentSrv.visible, which is implemented as a pass-through traversal, to download arbitrary attachments.	6.5	More Details
CVE-2026-63100	Maybe through 0.6.0 contains a missing authorization vulnerability that allows authenticated low-privilege member-role users to access and modify global hosting settings by exploiting unprotected show and update actions in the Settings::HostingsController, where the before_action ensure_admin filter is applied only to the clear_cache action. Attackers can read the operator's Synth API key rendered in plaintext via a form field value attribute, overwrite it with an attacker-controlled value, toggle public registration settings, and disable email confirmation requirements to disrupt the entire instance.	6.5	More Details
CVE-2026-48014	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, the order state transition features /api/_action/order/{orderId}/state/{transition} and similar transaction and delivery transition routes in src/Core/Checkout/Order/Api/OrderActionController.php do not declare PlatformRequest::ATTRIBUTE_ACL or perform an explicit privilege check, so AclAnnotationValidator exits when route ACL metadata is absent and low-privileged users without order:update, order_transaction:update, or order_delivery:update can trigger StateMachineRegistry::transition() writes in SYSTEM_SCOPE. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	6.5	More Details

CVE-2026-11763	Authorization bypass through User-Controlled key vulnerability in Gis Informatics Engineering Consulting Laboratory R&D and Software Services Inc. GisLab Laboratory Management System allows Exploitation of Trusted Identifiers. This issue affects GisLab Laboratory Management System: from 1.4.03 through 08072026.	6.5	More Details
CVE-2026-8075	Mattermost Desktop App versions <=6.2 5.5.13 6.0.2.0 fail to properly null check when checking for headers in the Mattermost Desktop App which allows any user to crash another channel members Desktop App via posting a malicious link with an embedded image that misses one of those headers. Mattermost Advisory ID: MMSA-2026-00668	6.5	More Details
CVE-2026-16317	Missing validation of the outer content_type byte on TLS 1.3 encrypted records in s2n-tls allows an active man-in-the-middle to silently discard individual application data records without either endpoint detecting the modification. RFC 8446 Section 5.2 requires that the outer content_type of all encrypted TLS 1.3 records must be application_data (0x17). The s2n-tls AEAD implementation hardcodes this value in the additional authenticated data rather than using the actual wire byte, so the outer content_type is not covered by the authentication tag. This enables selective suppression of application data. In HTTP pipelining scenarios, dropping a TLS record containing an HTTP request can cause request/response desynchronization, where subsequent responses are delivered to the wrong requests. In write-heavy workloads, a dropped record containing a write request can result in undetectable data loss when the client interprets a subsequent success response as confirmation of the dropped write. All TLS 1.3 connections are affected. Both TLS clients and servers are affected. TLS 1.2 and QUIC connections are not affected. We recommend you upgrade s2n-tls to version v1.7.6	6.5	More Details
CVE-2026-56145	Uncontrolled Resource Consumption (CWE-400) in Elasticsearch can lead to denial of service via Excessive Allocation (CAPEC-130). A low-privileged authenticated user with permission to execute EQL sequence queries against an index they control can send a specially crafted query that triggers excessive memory consumption, causing the Elasticsearch node to crash.	6.5	More Details
CVE-2026-42397	Allocation of Resources Without Limits or Throttling (CWE-770) in Kibana can lead to a denial of service via Excessive Allocation (CAPEC-130). An authenticated user can submit a specially crafted request to affected Entity Analytics endpoints containing an oversized input value that causes excessive resource consumption, which may render Kibana unavailable.	6.5	More Details
CVE-2026-48010	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, UserController::upsertUser() in src/Core/Framework/Api/Controller/UserController.php writes raw user data in SYSTEM_SCOPE without filtering the admin field, so a non-admin API user with user:create or user:update ACL permission can set admin: true on new or existing users; IntegrationController::upsertIntegration() contains an isAdmin() check for the same field, but UserController was missing this check. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	6.5	More Details
CVE-2026-63307	Chat2DB before 5.3.0 contains an insecure direct object reference vulnerability in the GET /api/connection/datasource/{id} endpoint. The handler calls dataSourceService.queryExistent(id, ...) without an ownership check and returns the decrypted password field, allowing any authenticated non-admin user to enumerate datasource IDs and read the plaintext database credentials of datasources owned by other users.	6.5	More Details
CVE-2026-48008	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, a non-admin API user with integration:create ACL privilege can escalate to full administrator by creating an integration with admin: true through the Sync API POST /api/_action/sync; the regular integration endpoint POST /api/integration blocks this, but SyncController::sync() routes writes through SyncService to EntityWriter::upsert(), and src/Core/Framework/Integration/IntegrationDefinition.php lacks WriteProtection on the admin field. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	6.5	More Details
CVE-2026-59259	n8n before versions 1.123.61, 2.27.4, and 2.28.1 contains a permission bypass vulnerability in external secrets handling caused by a mismatch between the static validation check and the runtime expression engine. An authenticated user with credential create or update permissions but without the externalSecret:list scope can embed external secret references into credentials in forms the static validation does not detect; these references resolve at workflow execution time, exposing secret values the user is not authorized to access. This issue only affects instances where an external secrets provider is configured and Advanced Permissions are in use.	6.5	More Details
CVE-2026-13694	The Bit Form WordPress plugin before 3.1.0 does not properly validate its workflow-trigger token once the associated transient has expired, allowing unauthenticated attackers to re-trigger a form's configured workflow actions such as notification emails and integrations.	6.5	More Details
CVE-2026-49209	Symfony UX is a JavaScript ecosystem for Symfony. From 2.5.0 until 2.36.0 and 3.1.0, Symfony\UX\LiveComponent\Controller\BatchActionController::__invoke() iterates over the client-supplied actions array and issues a full HttpKernel sub-request for each entry; because the array size is never bounded, an authenticated client can submit a single _batch request containing thousands of actions and exhaust CPU, memory, and database connections on the application server. This issue is fixed in versions 2.36.0 and 3.1.0.	6.5	More Details

CVE-2026-4938	IBM Verify Identity Access 11.0 through 11.0.2 and IBM Security Verify Access 10.0 through 10.0.9.1 and IBM Verify Identity Access Container 11.0 through 11.0.2 and IBM Security Verify Access Container 10.0 through 10.0.9.1 could allow an attacker with read-only privileges to make unauthorized modifications and deployments outside of their assigned permissions.	6.5	More Details
CVE-2026-21770	HCL Traveler for Microsoft Outlook (HTMO) is susceptible to a DLL hijacking vulnerability which could allow an attacker to modify or replace the application with malicious content.	6.5	More Details
CVE-2026-10677	The CONFIG_USERSPACE syscall verifier z_vrfy_k_poll() in kernel/poll.c allocates a kernel-side copy of the user-supplied k_poll_event[] via z_thread_malloc() and then validates each event's object handle. Before this fix, validation used K_OOPS(K_SYSCALL_OBJ(...)) inline inside the loop, which kills the calling thread without freeing events_copy. A user thread can pass num_events >= 1 with a forged object handle to leak the allocation; because newly spawned user threads inherit the parent's resource_pool (kernel/thread.c), an attacker spawns sacrificial threads to repeat the leak until the shared kernel heap is exhausted. Once depleted, legitimate kernel allocations from that pool (k_queue alloc nodes, k_msgq buffers, future k_poll calls, etc.) fail, causing a system-level denial of service. The fix replaces each inline K_OOPS with a conditional goto oops_free so the buffer is freed before the thread is killed. Affects Zephyr releases from v1.12.0 (when k_poll was first exposed to user mode) through v4.4.1.	6.5	More Details
CVE-2026-61449	Grav 2.0.1 contains a decompression-bomb size-cap bypass in ZipArchiver and GPM\Installer. The size bound introduced in 2.0.1 sums the uncompressed size declared in each entry's ZIP central-directory header (ZipArchive::statIndex()['size']) and rejects archives exceeding system.gpm.archive.max_uncompressed_size before extraction. Because this declared size is attacker-forgeable and is not cross-checked against the actual inflated stream, a crafted archive declaring tiny per-entry sizes passes the cap while extractTo() writes the real, much larger content, filling disk or exhausting inodes. The archive must be supplied by a package source or admin upload (admin/operator trust). Fixed in 2.0.2. This is an incomplete fix for GHSA-928x-9mpw-8h56.	6.5	More Details
CVE-2026-60062	The NGINX Agent config_dirs directive allows a low-privileged attacker to gain limited read and write access to files outside of the designated secure directory. The config_dirs directive required for this issue can also be configured through NGINX Instance Manager. A successful exploit may allow an attacker to cross a security boundary. Impact: A remotely authenticated low-privileged attacker could gain limited read and write access outside of the list of directories specified in the NGINX Agent configuration. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	6.4	More Details
CVE-2026-56678	9Router is an AI router & token saver. Prior to 0.5.6, the Kiro API-key validation endpoint POST /api/oauth/kiro/api-key builds an upstream URL using a user-controlled region value, allowing an authenticated attacker to supply a crafted region such as kiro-canary.local:8443# and cause 9Router to send the Kiro validation request to an attacker-controlled host while forwarding the submitted Kiro API key as an Authorization header. This issue is fixed in version 0.5.6.	6.4	More Details
CVE-2026-56352	n8n before 2.19.3 contains a file path restriction bypass in the legacy ExecuteWorkflow node's localFile source option, which reads workflow files from disk without the file-access checks enforced by other file-reading nodes. Although hidden from the UI since v1.2, it remains reachable via the REST API. An authenticated user with permission to create or modify workflows can supply an arbitrary file path to bypass the N8N_RESTRICT_FILE_ACCESS_TO restriction and determine whether arbitrary files exist on the host; where the targeted path contains a valid workflow JSON file, that file can additionally be loaded and executed.	6.4	More Details
CVE-2026-15021	The wpForo Forum plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'location' Profile Field in all versions up to, and including, 3.1.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The sanitize_text_field() function applied at input does not encode double quotes, allowing attribute breakout via a payload that escapes the href attribute context and injects event handler attributes.	6.4	More Details
CVE-2026-2594	The Smart Custom Fields plugin for WordPress is vulnerable to Stored Cross-Site Scripting in versions up to, and including, 5.0.7. This is due to insufficient input sanitization and output escaping of uploaded image attachment titles. This makes it possible for authenticated attackers, with Author-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. NOTE: This vulnerability was partially patched in 5.0.7.	6.4	More Details
CVE-2026-12900	The Spectra Gutenberg Blocks – Website Builder for the Block Editor plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the `uagb/image` block in all versions up to, and including, 2.19.28 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with Contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
	The Ninja Forms - Excel Export plugin for WordPress is vulnerable to Stored Cross-Site Scripting in versions up to, and including, 3.3.6. This is due to the save_filter() AJAX handler storing the raw $_POST['filter'] array into		

CVE-2026-15161	a WordPress option via update_option() without any capability check, nonce verification, or input sanitization, combined with the get_filter_row() method on the admin Excel Export screen concatenating the stored filter values (field_key, condition, value) directly into HTML attributes without esc_attr(). This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-14987	The GiveWP – Donation Plugin and Fundraising Platform plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'twitter_message' Sequoia Template Setting in all versions up to, and including, 4.16.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with give worker-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The injected script executes specifically when a donor clicks the Share on Twitter button on the Sequoia donation confirmation view, as that is when the unescaped twitter_message value is evaluated inside the JavaScript template literal.	6.4	More Details
CVE-2026-15099	The Delicious Recipes plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'steps' block attribute in versions up to, and including, 1.10.2. This is due to insufficient input sanitization and output escaping in the wrap_direction_text() function, which interpolates the user-supplied href value from nested link nodes ($node['props']['href']) directly into an anchor tag via sprintf() at line 1627 without esc_url() or any URL scheme validation. This makes it possible for authenticated attackers, with Contributor-level access and above, to inject arbitrary web scripts (including javascript: URIs) in pages that will execute whenever a user (such as an editor or administrator previewing the pending post) accesses an injected page and clicks the malicious link.	6.4	More Details
CVE-2026-15652	The Easy Accordion – AI-Powered FAQ & Accordion Blocks, Product FAQ plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'align' Block Attribute in all versions up to, and including, 3.1.6 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-63397	remorses/genql before version 6.3.4 allows an authenticated attacker with control of the GraphQL schema that is passed to genql to inject arbitrary JavaScript or TypeScript. The malicious code is injected into the generated schema.ts file and executes when the genql client is bundled and imported.	6.4	More Details
CVE-2026-13755	The Tickera – Sell Tickets & Manage Events plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'price_wrapper' Shortcode Attribute in all versions up to, and including, 3.6.0.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Successful execution of the injected script is limited to victims who have the referenced ticket ID present in their cart cookie, meaning the payload only fires for users who have previously added that ticket to their cart.	6.4	More Details
CVE-2026-64626	AVideo versions from commit 0dbadbca through latest master contain a server-side request forgery vulnerability in the encoder download-by-URL flow due to an unpinned retry fallback that bypasses DNS pinning validation. An authenticated attacker can supply a downloadURL that redirects to an internal address, causing the unpinned retry to follow the redirect and reach internal targets for blind SSRF attacks.	6.4	More Details
CVE-2026-15759	The ChatHelp – Click to Chat Button, WooCommerce Chat to Order & Floating Chat Form plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'number' and 'group' Shortcode Attributes in all versions up to, and including, 3.5.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-12705	Missing support for integrity check vulnerability in ABB KNX Update Tool (ABB), ABB KNX Update Tool (BJE). This issue affects KNX Update Tool (ABB): through 2.0.175; KNX Update Tool (BJE): through 2.0.175.	6.4	More Details
CVE-2026-15156	The Essential Addons for Elementor – Popular Elementor Templates & Widgets plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Reading Progress Global Color Settings in all versions up to, and including, 6.6.11 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15145	The Essential Addons for Elementor – Popular Elementor Templates & Widgets plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Fancy Text Widget in all versions up to, and including, 6.6.11 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-63743	SurrealDB before 3.1.0 contains a capability bypass vulnerability in HTTP redirect handling that allows authenticated users to circumvent port-scoped --deny-net rules. Attackers can chain an HTTP redirect from an allowed hostname to a denied host:port combination, and the redirect is followed because the port information is dropped during redirect policy evaluation.	6.4	More Details

CVE-2026-45703	Pimcore is an Open Source Data & Experience Management Platform. Prior to 11.5.17 (LTS) and 12.3.7, the WordExport export flow in bundles/WordExportBundle/src/Controller/TranslationController.php only checks the word_export feature permission and directly resolves attacker-controlled type/id input without enforcing view permission on page, snippet, email, or object elements, allowing a low-privileged backend user to export document content the user is not allowed to view. This issue is fixed in versions 11.5.17 (LTS) and 12.3.7.	6.4	More Details
CVE-2026-16196	A weakness has been identified in Sipeed PicoClaw up to 0.2.9. Impacted is the function isPrivateOrRestrictedIP of the file pkg/tools/integration/web.go of the component web_fetch. This manipulation causes server-side request forgery. The attack can be initiated remotely. The exploit has been made available to the public and could be used for attacks. Patch name: 2efbe5d560e7ed9bc5209c203dc4aa6ecdbc7405. To fix this issue, it is recommended to deploy a patch.	6.3	More Details
CVE-2026-35146	HCL DFXServer is affected by an Unencrypted Communication vulnerability. The application permits users to establish connections over unencrypted channels via the HTTP protocol, which could allow a remote attacker to intercept network traffic and expose sensitive data transmitted between the user and the application.	6.3	More Details
CVE-2026-16195	A security flaw has been discovered in Sipeed PicoClaw up to 0.2.9. This issue affects the function dispatchIncoming of the file pkg/channels/wecom/wecom.go of the component Group Message Handler. The manipulation results in incorrect authorization. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks. The reported GitHub issue was closed automatically due to inactivity.	6.3	More Details
CVE-2026-16194	A vulnerability was determined in zhayujie CowAgent up to 2.1.1. This affects the function WebFetch.execute of the file agent/tools/web_fetch/web_fetch.py. Executing a manipulation of the argument url can lead to server-side request forgery. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized. Upgrading to version 2.1.2 is able to mitigate this issue. This patch is called ea47f3097eed4f8295c4cb3d76ecb97e0f43d632. It is recommended to upgrade the affected component. The vendor was contacted early, responded in a very professional manner and quickly released a fixed version of the affected product.	6.3	More Details
CVE-2026-35148	HCL DFXServer is affected by a Missing Access Control vulnerability. This vulnerability states that certain endpoints are accessible without any form of authentication in another browser. This allows any network user to invoke these APIs and interact with the application without verification of their identity or authorization level.	6.3	More Details
CVE-2026-16124	A security vulnerability has been detected in nextlevelbuilder GoClaw up to 3.15.0-beta.32. This affects the function CheckSSRF/isPrivateIP of the file internal/tools/web_shared.go of the component web_fetch. Such manipulation leads to server-side request forgery. The attack can be launched remotely. The exploit has been disclosed publicly and may be used. Upgrading to version 3.15.0-beta.33 is able to mitigate this issue. The name of the patch is 12a0168271827650ddb0026d6277fbadf3dcf3ea. Upgrading the affected component is recommended.	6.3	More Details
CVE-2026-16119	A vulnerability was found in nextlevelbuilder GoClaw up to 3.13.2. This affects the function RequestApproval of the file internal/tools/exec_approval.go of the component WebSocket Approval Endpoint. Performing a manipulation results in incorrect authorization. The attack is possible to be carried out remotely. The exploit has been made public and could be used.	6.3	More Details
CVE-2026-16120	A vulnerability was determined in nextlevelbuilder GoClaw up to 3.13.3-beta.3. This impacts the function matchesAllowlist/extractBin of the file internal/tools/exec_approval.go. Executing a manipulation can lead to incorrectly-resolved name. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized.	6.3	More Details
CVE-2026-16449	A vulnerability was determined in zsadmin2025 ZS-Admin up to b52e14536d59fda11e56e2536a1c32e82a38cead. The impacted element is the function OrderItem.asc/OrderItem.desc of the file /api/system/sys/dept/page of the component com.zs.sys.dept.controller.SysDeptController. This manipulation of the argument orderField causes sql injection. The attack is possible to be carried out remotely. The exploit has been publicly disclosed and may be utilized. This product adopts a rolling release strategy to maintain continuous delivery. Therefore, version details for affected or updated releases cannot be specified. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16121	A vulnerability was identified in nextlevelbuilder GoClaw up to 3.13.2. Affected is the function isSafeBin of the file internal/tools/exec_approval.go. The manipulation leads to improper authorization. It is possible to initiate the attack remotely. The exploit is publicly available and might be used.	6.3	More Details
CVE-2026-16123	A weakness has been identified in nextlevelbuilder GoClaw up to 3.13.2. Affected by this issue is the function ToolsInvokeHandler.ServeHTTP of the file internal/http/tools_invoke.go of the component Invoke Endpoint. This manipulation causes missing authorization. The attack can be initiated remotely. The exploit has been made available to the public and could be used for attacks.	6.3	More Details
	A vulnerability was found in D-Link DNS-120, DNR-202L, DNS-315L, DNS-320, DNS-320L, DNS-320LW, DNS-321, DNR-322L, DNS-323, DNS-325, DNS-326, DNS-327L, DNR-326, DNS-340L, DNS-343, DNS-345, DNS-726-		

CVE-2026-16448	4, DNS-1100-4, DNS-1200-05 and DNS-1550-04 up to 20260205. The affected element is the function cgi_check_rsync_rw of the file /cgi-bin/remote_backup.cgi. The manipulation of the argument ip results in command injection. The attack can be executed remotely. The exploit has been made public and could be used.	6.3	More Details
CVE-2026-45737	Argo CD is a declarative, GitOps continuous delivery tool for Kubernetes. From 3.2.0 until 3.2.12, 3.3.10, and 3.4.2, Argo CD ServerSideDiff can expose Kubernetes Secret values embedded in the kubectl.kubernetes.io/last-applied-configuration annotation because HideSecretData(target, live, ...) does not fully sanitize ResourceDiff.TargetState and LiveState predicted live Secret objects, allowing sensitive data, stringData, and annotations to appear in UI or CLI diffs. This issue is fixed in versions 3.2.12, 3.3.10, and 3.4.2.	6.3	More Details
CVE-2026-16225	A security flaw has been discovered in davenardella snap7 up to 1.4.3. The impacted element is the function TSnap7Peer::NegotiatePDULength of the file src/core/s7_peer.cpp. The manipulation of the argument PDULength results in out-of-bounds write. The attack can be executed remotely. The exploit has been released to the public and may be used for attacks.	6.3	More Details
CVE-2026-16131	A weakness has been identified in itsourcecode Hospital Management System 1.0. This affects an unknown function of the file /prescriptionrecord.php. This manipulation of the argument delid causes sql injection. It is possible to initiate the attack remotely. The exploit has been made available to the public and could be used for attacks.	6.3	More Details
CVE-2026-16150	A vulnerability was found in RobinHerbots Inputmask up to 5.0.9. Affected by this issue is the function extendDefaults/extendDefinitions/extendAliases in the library lib/dependencyLibs/extend.js of the component Internal Deep Merge Helper. The manipulation results in improperly controlled modification of object prototype attributes. The attack may be performed from remote. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16151	A vulnerability has been found in CartoDB carto-api-client 0.5.29. This impacts the function addFilter of the file src/filters.ts. Such manipulation of the argument column leads to improperly controlled modification of object prototype attributes. The attack can be executed remotely. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16334	A vulnerability was identified in itsourcecode Hospital Management System 1.0. This vulnerability affects unknown code of the file /prescriptionorder.php. Such manipulation of the argument editid leads to sql injection. The attack can be launched remotely. The exploit is publicly available and might be used.	6.3	More Details
CVE-2026-63728	Gitleaks prior to 8.30.1 contains a template injection vulnerability that allows attackers who can supply or influence report templates to read arbitrary environment variables and exfiltrate sensitive data by leveraging non-hermetic Sprig template functions. Attackers can craft malicious report templates using the env, expandenv, and getHostByName functions to extract credentials, tokens, and API keys from the host process and exfiltrate them through DNS queries, including secrets discovered during the scan itself.	6.3	More Details
CVE-2026-16223	A vulnerability was determined in 1Panel-dev CordysCRM up to 1.4.1. Impacted is the function getSqlBotSrc of the file backend/crm/src/main/java/cn/cordys/crm/system/service/IntegrationConfigService.java of the component Third Party Edit Endpoint. Executing a manipulation of the argument appSecret can lead to server-side request forgery. The attack may be launched remotely. The exploit has been publicly disclosed and may be utilized.	6.3	More Details
CVE-2026-15909	A vulnerability has been found in RafyMrX TOKO-ONLINE-ROTI up to ddfe1cd587be0a0b5135d8b6e85cce2ec3aece99. Affected is an unknown function of the file proses/add.php. The manipulation of the argument kd_cs leads to authorization bypass. The attack is possible to be carried out remotely. This product adopts a rolling release strategy to maintain continuous delivery. Therefore, version details for affected or updated releases cannot be specified. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-16197	A security vulnerability has been detected in Sipeed PicoClaw up to 0.2.9. The affected element is the function handleMessageReceive of the file pkg/channels/feishu/feishu_64.go of the component Group Message Handler. Such manipulation leads to missing authorization. The attack can be launched remotely. The exploit has been disclosed publicly and may be used. The reported GitHub issue was closed automatically due to inactivity.	6.3	More Details
CVE-2026-16219	A flaw has been found in Croogo CMS up to 4.0.7. This affects the function FileManager::isEditable of the file FileManager/src/Utility/FileManager.php of the component Admin File Manager. This manipulation causes path traversal. The attack can be initiated remotely. The exploit has been published and may be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16008	A security vulnerability has been detected in sagold json-schema-library 11.5.0/11.5.1. This impacts the function parsePropertyDependencies of the file src/keywords/propertyDependencies.ts. The manipulation leads to improperly controlled modification of object prototype attributes. The attack can be initiated remotely. Upgrading to version 11.6.0 will fix this issue. The identifier of the patch is 432287ee6f68a02ce6f015354618486ec427a32d. It is advisable to upgrade the affected component.	6.3	More Details

CVE-2026-62235	Grav Flex-Objects before version 1.4.3 contains a broken access control vulnerability in the admin-next REST API that allows authenticated users with only api.access permission to perform unauthorized CRUD operations on permission-less directories. Attackers with api.access credentials can create, read, update, delete, and export objects from any directory lacking an explicit permissions configuration, bypassing intended authorization controls.	6.3	More Details
CVE-2026-16214	A vulnerability was identified in geex-arts django-jet up to 1.0.8. This affects an unknown function of the file jet/dashboard/views.py of the component Dashboard Module. Such manipulation leads to authorization bypass. The attack can be executed remotely. The exploit is publicly available and might be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16017	A security flaw has been discovered in mosaxiv clawlet up to 0.2.10. Impacted is the function list/remove of the file tools/tool_cron.go of the component cron Chat Tool. The manipulation results in missing authorization. The attack may be performed from remote. The exploit has been released to the public and may be used for attacks. The reported GitHub issue was closed with the label "not planned".	6.3	More Details
CVE-2026-16076	A vulnerability has been found in AstrBotDevs AstrBot up to 4.25.5. This issue affects the function OpenApiRoute.chat_send of the file astrbot/dashboard/routes/open_api.py of the component API. Such manipulation of the argument Username leads to authentication bypass by spoofing. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-16244	A security vulnerability has been detected in itsourcecode Hospital Management System 1.0. Affected by this vulnerability is an unknown functionality of the file /prescriptionorderreport.php. Such manipulation of the argument delid leads to sql injection. The attack may be launched remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2026-16217	A security vulnerability has been detected in guohongze adminset up to 0.61. Affected by this vulnerability is an unknown functionality of the file delivery/deli.py of the component Delivery Deployment Endpoint. The manipulation of the argument project_id leads to authorization bypass. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16015	A vulnerability was determined in poco-ai poco-claw up to 0.5.4. This vulnerability affects the function create_task of the file executor_manager/app/api/v1/tasks.py of the component executor_manager API. Executing a manipulation can lead to missing authentication. The exploit has been publicly disclosed and may be utilized. Upgrading to version 0.5.7 is able to resolve this issue. This patch is called 67fcc88505c57f77d3fcf04eb5b89425b10cbf48. It is recommended to upgrade the affected component.	6.3	More Details
CVE-2026-16206	A security vulnerability has been detected in django-oauth django-oauth-toolkit 3.3.0. This issue affects the function _load_id_token of the file oauth2_provider/oauth2_validators.py. The manipulation leads to session expiration. The attack can be initiated remotely. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-46403	Klever-Go is the Go implementation of the Klever blockchain protocol. Prior to 1.7.17, KVM exposes `ExecuteReadOnlyWithTypedArguments` as a read-only execution mechanism. The hook saves the previous read-only state, sets `runtime.SetReadOnly(true)`, executes the destination context, and then restores the previous read-only state. However, the indirect contract delete and upgrade paths do not reject execution when `runtime.ReadOnly()` is true. As a result, a contract reached through read-only execution can call the production delete hook for a target contract it owns. The delete path appends the target address to `vmOutput.DeletedAccounts`, the output context merges `DeletedAccounts` into the caller output, and the smart contract processor later processes the VM output by deleting accounts listed in that field. The root cause is that read-only mode is applied as runtime state, but not enforced by the state-changing delete and upgrade host-core paths. This breaks the expected isolation boundary for workflows that rely on read-only calls to inspect another contract without allowing that callee to produce state-changing VM output. The issue is fixed in v1.7.17. Contract delete and upgrade host-core paths now reject execution when `runtime.ReadOnly()` is true. The invariant is regression-tested for delete, upgrade, storage writes, value transfers, and any VM output field that can later mutate chain state.	6.3	More Details
CVE-2026-16009	A vulnerability was detected in itsourcecode Hospital Management System 1.0. Affected is an unknown function of the file /prescriptionorderdetail.php. The manipulation of the argument delid results in sql injection. The attack can be launched remotely. The exploit is now public and may be used.	6.3	More Details
CVE-2026-44968	dbt-mcp is a Model Context Protocol server for interacting with dbt. Prior to 1.17.1, _run_dbt_command() in src/dbt_mcp/dbt_cli/tools.py appended unsanitized node_selection and resource_type values to the dbt subprocess argument list, allowing an MCP client to inject dbt global flags such as --profiles-dir, --project-dir, and --target into subprocess.Popen even though shell=False prevents shell metacharacter injection. This issue is fixed in version 1.17.1.	6.3	More Details
CVE-2026-	A security flaw has been discovered in zevorn rt-claw up to 0.2.0. This affects the function tool_run_script_execute of the file claw/services/tools/script.c of the component Telegram-to-AI Tool Execution Flow. Performing a manipulation results in code injection. It is possible to initiate the attack	6.3	More

16204	remotely. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.		Details
CVE-2026-16451	A security flaw has been discovered in zsadmin2025 ZS-Admin up to b52e14536d59fda11e56e2536a1c32e82a38cead. This impacts an unknown function of the file /api/system/file/upload of the component com.zs.file.controller.SysFileController. Performing a manipulation of the argument File results in unrestricted upload. It is possible to initiate the attack remotely. The exploit has been released to the public and may be used for attacks. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-16199	A flaw has been found in nextlevelbuilder GoClaw up to 3.13.3-beta.3. This affects the function ExecTool.Execute of the file goclaw/internal/tools/credentialed_exec.go. Executing a manipulation can lead to improper authorization. The attack may be launched remotely. The exploit has been published and may be used.	6.3	More Details
CVE-2026-16222	A vulnerability was found in 1Panel-dev CordysCRM up to 1.4.1. This issue affects some unknown processing of the file backend/crm/src/main/java/cn/cordys/crm/integration/sso/service/TokenService.java of the component Third Party Endpoint. Performing a manipulation of the argument mkAddress results in server-side request forgery. The attack may be initiated remotely. The exploit has been made public and could be used. The project closed the issue report, stating that this is not the official way to report a security vulnerability.	6.3	More Details
CVE-2026-16074	A vulnerability was detected in AstrBotDevs AstrBot up to 4.25.2. This affects the function update_plugin/update_all_plugins of the file astrbot/dashboard/routes/plugin.py of the component Plugin Update Handler. The manipulation of the argument download_url/download_urls/proxy results in server-side request forgery. The attack may be performed from remote. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2024-58363	SurrealDB before 1.5.4 fails to properly validate authentication when a scope user switches databases using the USE clause or use method. Attackers with an authenticated session can impersonate an unrelated user in a different database if a user record with an identical identifier exists, allowing unauthorized actions if permissions rely solely on the $auth parameter.	6.3	More Details
CVE-2026-46377	Dasel is a command-line tool and library for querying, modifying, and transforming data structures. From 3.0.0 until 3.10.1, the escape sequence handler in (*Tokenizer).parseCurRune in selector/lexer/tokenize.go increments past a trailing backslash in a quoted string such as "\ or '\ and then reads p.src[pos] without a bounds check, allowing attacker-controlled selector strings to trigger a Go index-out-of-range panic. This issue is fixed in version 3.10.1.	6.2	More Details
CVE-2026-13713	YAML::Syck versions before 1.47 for Perl allow a use-after-free and double-free via an anchor node freed while still on the parser value stack. In the bundled libsyck, when an anchor name is redefined or removed, syck_hdlr_add_anchor and syck_hdlr_remove_anchor free the node stored under that name with syck_free_node. That node can still be live on the parser's value stack, so syck_hdlr_add_node reaches it again and frees it a second time. On a normal build the 48-byte node chunk is freed twice and the interpreter aborts. Anchors need no special flags, so this is reached on the default Load path, and a 7-byte document that redefines an anchor triggers it. Any caller that runs Load or LoadFile on an untrusted document that redefines an anchor mid-parse crashes the interpreter, a denial of service.	6.2	More Details
CVE-2026-14961	Pegatron `Tdelo64.sys` exposes a privileged device interface, `\\.\TdeIo`, that fails to properly restrict access to sensitive IOCTL functionality. The driver's IOCTL dispatcher does not validate caller privileges or verify user-supplied kernel memory addresses before performing memory operations. By sending crafted requests to IOCTL, a local attacker can achieve arbitrary kernel memory read and write operations, leading to privilege escalation to `NT AUTHORITY\SYSTEM`, security product bypass, credential theft, or complete system compromise.	6.2	More Details
CVE-2026-44722	pyzipper is a replacement for Python's zipfile that can read and write AES encrypted zip files. Prior to 0.4.0, a Python operator precedence bug in pyzipper/zipfile_aes.py caused the AE-2 format to never be automatically selected during encryption, causing encrypted entries to be written in AE-1 format and exposing the plaintext CRC32 checksum in the ZIP header and, for unseekable zip archives, in the datadescripter section, allowing an attacker who possesses the archive to brute-force candidate plaintexts for small or low-entropy files by comparing CRC32 values. This issue is fixed in version 0.4.0.	6.2	More Details
CVE-2026-28315	SolarWinds Serv-U was found to be affected by a stored cross-site scripting vulnerability that could lead to session hijacking or information disclosure from an administrator account.	6.2	More Details
CVE-2026-46378	Dasel is a command-line tool and library for querying, modifying, and transforming data structures. From 3.0.0 until 3.10.1, the selector lexer matchRegexPattern closure in (*Tokenizer).parseCurRune in selector/lexer/tokenize.go loops while tokenizing an unterminated regex literal such as r/ because peekRuneEqual returns false after the end of input, allowing attacker-controlled selector strings to consume CPU indefinitely. This issue is fixed in version 3.10.1.	6.2	More Details

CVE-2026-45785	OpenMcdf is a fully .NET / C# library to manipulate Compound File Binary File Format files, also known as Structured Storage. In 3.1.3 and earlier, the BST name-lookup loop in DirectoryTree.TryGetDirectoryEntry (OpenMcdf/DirectoryTree.cs:35-46) walks directory entries by repeatedly calling directories.TryGetSibling(child, siblingType, validateColor). A crafted CFB file with cyclic Left/Right sibling links among directory entries, constructed so the per-step BST-order check in TryGetSibling (DirectoryEntries.cs:84-85) is satisfied at every step, drives this while (child is not null) loop forever. There is no cycle detection in TryGetDirectoryEntry, and the bug is reachable from RootStorage.OpenStorage(name), TryOpenStorage(name), OpenStream(name), and TryOpenStream(name), causing an unrecoverable denial of service. This issue is fixed in version 3.1.4.	6.2	More Details
CVE-2026-56087	Dell ThinOS 10, versions prior to 2605_10.2100 contain a Protection Mechanism Failure vulnerability. An attacker with physical access could potentially exploit this vulnerability, leading to unauthorized access to encrypted data.	6.1	More Details
CVE-2026-2445	The affected product accepts user-supplied input within a URL parameter without enforcing expected sanitization or encoding before rendering it within the response. This condition allows for the injection of malicious JavaScript payloads. An attacker can leverage this vulnerability to cause the user's browser to redirect to a malicious website, modify the user interface of the webpage, or retrieve sensitive information from the browser. However, the impact is mitigated for session hijacking as all session-related sensitive cookies are protected by the httpOnly flag.	6.1	More Details
CVE-2026-49210	Symfony UX is a JavaScript ecosystem for Symfony. From 2.8.0 until 2.36.0 and 3.1.0, Symfony\UX\LiveComponent\Util\ChildComponentPartialRenderer::createHtml() interpolates the client-controlled children[id].tag value from LiveComponentSubscriber and InterceptChildComponentRenderSubscriber directly into HTML as a tag name without escaping or validation, allowing arbitrary HTML, including <script> tags, on any Live Component re-render that contains at least one child component. This issue is fixed in versions 2.36.0 and 3.1.0.	6.1	More Details
CVE-2026-58643	Improper neutralization of input during web page generation ('cross-site scripting') in Windows Admin Center allows an unauthorized attacker to perform spoofing over a network.	6.1	More Details
CVE-2026-57896	An out-of-bounds read vulnerability in the Productivity Suite allows a local attacker to trigger kernel memory corruption by sending a crafted IOCTL request. This could lead to limited information disclosure or disruption of the affected product.	6.1	More Details
CVE-2026-61453	Grav v2.0.0 contains a cross-site scripting vulnerability (fixed in 2.0.1). The XSS blueprint validator (Security::detectXss()) runs on raw page content before Twig processing. When Twig content processing is enabled (twig_content.process_enabled: true), an attacker with page-write API permission can use Twig's string concatenation operator (~) to dynamically construct event handler names, dangerous tag names, or dangerous protocols at render time (e.g. {% set x = "on" ~ "error" %}). The validator sees only the harmless Twig expression and allows the content, but after Twig rendering the output (rendered via {{ page.content|raw }}) contains an active payload such as <img src=1 onerror=alert(1)>, executing arbitrary JavaScript in visitors' browsers.	6.1	More Details
CVE-2026-60087	PraisonAI before 1.6.78 caches tool approval decisions by tool name only, allowing attackers to reuse initial approvals for subsequent calls with arbitrary arguments. Attackers can exploit this by obtaining approval for a benign operation and then executing dangerous file write operations with unreviewed parameters in the same session.	6.1	More Details
CVE-2026-32822	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, any unauthenticated attacker can place arbitrary HTML into flash notifications on public routes and rely on the frontend toast component to inject that content into the DOM with `innerHTML`. This creates a reflected DOM XSS that can be delivered with a crafted link to a public page such as `/docs`. Because the vulnerable JavaScript is loaded by the normal application layout, the issue is not limited to a special debug page or an isolated admin-only view.	6.1	More Details
CVE-2026-51081	A cross-site scripting (XSS) vulnerability in Proxmox Virtual Environment (PVE) 9.x 5.1.8 and Proxmox Virtual Environment (PVE) 8.x 4.3.16 allows attackers to execute arbitrary web scripts or HTML via injecting a crafted payload.	6.1	More Details
CVE-2026-33213	Redash is a package for data visualization and sharing. From 5.0.2 to 26.3.0, the get_next_path() function in Redash's authentication module stripped the scheme and netloc from user-supplied next parameters but did not normalize multiple leading slashes, allowing a crafted login URL such as /login?next=////evil.com to redirect users to an external attacker-controlled site after authentication.	6.1	More Details
CVE-2026-	Vulnerability in the Oracle Commerce Service Center product of Oracle Commerce (component: Commerce Service Center). The supported version that is affected is 11.4.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Commerce Service Center. Successful attacks require human interaction from a person other than the attacker and while the vulnerability is in Oracle Commerce Service Center, attacks may significantly impact additional products	6.1	More

34316	(scope change). Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle Commerce Service Center accessible data as well as unauthorized read access to a subset of Oracle Commerce Service Center accessible data. CVSS 3.1 Base Score 6.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:C/C:L/I:L/A:N).		Details
CVE-2026-51025	Cross Site Scripting vulnerability in fuint Member Marketing System <=v1.0 allows a remote attacker to execute arbitrary code via the ClientMessageController.java file	6.1	More Details
CVE-2026-11324	The WooCommerce Placetopay Gateway and PlacetoPay/AvalPay gateway plugins for WordPress are vulnerable to Reflected Cross-Site Scripting via the 'redirect-url' parameter in versions up to, and including, 3.2.2 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that execute if they can successfully trick a user into performing an action such as clicking on a link.	6.1	More Details
CVE-2026-47128	nono is software that allows users to run AI agents in a zero-latency sandbox. Prior to version 0.55.0, the nono Landlock/seccomp policies allow access to local Unix domain sockets (concrete and abstract). This allows an easy sandbox escape by talking to the per-user systemd dbus socket. Version 0.55.0 patches the issue.	6.1	More Details
CVE-2026-54243	Statamic is a Laravel and Git powered content management system (CMS). Prior to 5.73.24 and 6.20.1, form submission values in src/Forms/Exporters/CsvExporter.php were not neutralized for spreadsheet formula characters when exported to CSV. A submission containing a value beginning with a formula trigger character, such as =, +, -, or @, could be interpreted as a live formula when a Control Panel user opens the export in a spreadsheet application. Form submissions can come from unauthenticated front-end visitors, so the malicious value can be supplied by an anonymous user and is later triggered by an editor opening the export. This issue is fixed in versions 5.73.24 and 6.20.1.	6.1	More Details
CVE-2026-46341	The Apify MCP server enables AI agents to extract data from websites using ready-made scrapers, crawlers, and automation tools available on the Apify Store. Prior to 0.9.21, the fetch-apify-docs tool in src/tools/common/fetch_apify_docs.ts validates allowlisted documentation domains with String.startsWith() rather than URL hostname comparison, allowing attacker-controlled URLs such as `https://docs.apify.com.evil.com/` and `https://docs.apify.com@evil.com/` to pass the ALLOWED_DOC_DOMAINS check and return arbitrary fetched content to the LLM. This issue is fixed in version 0.9.21.	6.1	More Details
CVE-2026-41580	Stirling-PDF is a locally hosted web application that facilitates various operations on PDF files. Prior to 2.0.0, Stirling-PDF's /get-info-on-pdf endpoint rendered PDF Title and Author metadata fields without proper HTML encoding or sanitization, allowing a crafted PDF to execute attacker-controlled JavaScript in the browser of a user who views the resulting page. This issue is fixed in version 2.0.0.	6.1	More Details
CVE-2026-12869	The Header Footer Builder for Elementor WordPress plugin before 1.2.1 does not require an administrative capability for its dashboard template-import action (it allows any edit_posts user), so a Contributor can import a template containing an Elementor HTML widget configured to display site-wide, injecting JavaScript that executes in the session of any visitor or administrator who loads the site.	6.1	More Details
CVE-2026-15094	The WP Hotel Booking plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the 'check_in_date' parameter in all versions up to, and including, 2.3.2 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that execute if they can successfully trick a user into performing an action such as clicking on a link.	6.1	More Details
CVE-2026-11371	The BetterDocs WordPress plugin before 4.5.5 does not sanitise an AI-generated documentation summary before storing and outputting it, and the feature that generates it is exposed to unauthenticated users, allowing them to store a malicious payload via prompt injection that executes in the browser of any visitor who views the affected page, including administrators.	6.1	More Details
CVE-2026-60140	An out-of-bounds read vulnerability in the Productivity Suite allows a local attacker to trigger kernel memory corruption by sending a crafted IOCTL request. This can lead to exposing sensitive information or causing the affected product to become unstable or unavailable.	6.1	More Details
CVE-2026-15306	The Product Feed Manager For WooCommerce – Sell on 200+ Online Marketplaces plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via 's' Search Parameter in all versions up to, and including, 7.6.1 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that execute if they can successfully trick a user into performing an action such as clicking on a link.	6.1	More Details
CVE-2026-44230	RT is an open source, enterprise-grade issue and ticket tracking system. Versions 5.0.4 up to (but not including) 5.0.10, and 6.0.0 up to (but not including) 6.0.3 contain a reflected Cross-Site Scripting (XSS) vulnerability where an attacker who can induce an authenticated RT user to visit a crafted URL can execute arbitrary JavaScript in that user's browser session. This issue has been fixed in versions 5.0.10 and 6.0.3.	6.1	More Details
CVE-			

2026-61901	The Joomla extension Hikashop is vulnerable to an open redirect.	6.1	More Details
CVE-2026-26719	Cross Site Scripting vulnerability in xxl-job-admin v.3.0.0 allows a remote attacker to execute arbitrary code via a crafted HTTP GET request containing a malicious script	6.1	More Details
CVE-2026-15779	A flaw was found in samba's pam_winbind. When mkhomedir is enabled, pam_winbind chowns the target account's home directory without validating the path is not a critical system directory such as /. On affected systems, accounts with / as their home directory (a common default for system accounts) can have this triggered not only by root, but by a non-root user holding a narrow sudo delegation to run commands as that account, causing ownership of / to change and resulting in severe denial of service (SSH, sudo, and package-manager failures). The change does not grant write access to / (which ships with restrictive 0555 permissions on RHEL), so the impact is availability loss rather than further privilege escalation.	6.1	More Details
CVE-2026-58413	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.12.2, `EnvironmentManager.restore(env, backupId)` computes the backup path with `join(envDir, '.backups', backupId)` and only checks that this path exists. It does not resolve the result or verify that it remains under `data/<env>/.backups`. A caller can pass a traversal backup ID such as `../../../outside/source-dir` to restore files from an arbitrary directory into the target environment data directory. The issue is fixed in v5.12.2. `restore()` now validates `backupId` against `/^[\w\-]+$/` and asserts `dirname(resolve(join(backupsDir, backupId))) === resolve(backupsDir)` before touching the filesystem. Backup IDs containing path separators or `..` are rejected, so a crafted ID can no longer copy directories from outside `.backups/` into the environment.	6.1	More Details
CVE-2026-47714	libheif is a HEIF and AVIF file format decoder and encoder. In versions 1.21.2 and prior, the inline mask parsing code in `libheif/region.cc` contains an integer overflow. Both `width` and `height` are `unsigned int` (32-bit) values parsed from the HEIF file. Their product can exceed `UINT32_MAX`, wrapping to a small value before the division by 8. This causes an undersized buffer allocation, leading to out-of-bounds memory access when the mask data is later interpreted as a `width x height` bitmap. Version 1.22.0 patches the issue.	6.1	More Details
CVE-2026-26483	Mettle SendPortal 3.0.1 and earlier contains a stored cross-site scripting (XSS) vulnerability in the template management functionality. The application fails to properly sanitize user-supplied input in the content parameter of the /templates endpoint, allowing an attacker to persistently inject malicious JavaScript code that is executed in the browsers of users who access the affected template.	6.1	More Details
CVE-2026-10525	The NEX-Forms WordPress plugin before 9.2.3 does not sanitise and escape some submitted form data before storing it and outputting it back in the admin dashboard, leading to a Stored Cross-Site Scripting vulnerability which could allow unauthenticated users to perform Stored Cross-Site Scripting attacks against high privilege users such as administrators when they view the submitted entries.	6.1	More Details
CVE-2026-44227	RT is an open source, enterprise-grade issue and ticket tracking system. Versions 6.0.0 and above, prior to 6.0.3 contain a reflected Cross-Site Scripting (XSS) vulnerability. An attacker who can induce an authenticated RT user to visit a crafted URL can execute arbitrary JavaScript in that user's browser session. There are no effective workarounds. Avoid following untrusted RT URLs. This issue has been fixed in version 6.0.3.	6.1	More Details
CVE-2026-47121	Sparkle is a software update framework for macOS. Prior to version 2.9.2, `Autoupdate/SUBinaryDeltaApply.m` enforces `relativePath.pathComponents containsObject:@".."` and rejects writes whose immediate parent directory IS itself a symbolic link, but does not detect symlinks deeper in the relative path. `Autoupdate/SPUSparkleDeltaArchive.m`'s `extractItem:` will create symlinks in the destination tree from archive content (no `..` check on the symlink target), and a subsequent `Extract` item targeting `<symlink>/foo/bar` then escapes the destination tree via `fopen(path, "wb")` because the kernel resolves the intermediate symlink during the open call. This is a defense-in-depth issue: exploitation requires a maliciously-crafted `.delta` that passes EdDSA signature verification, i.e. EdDSA private-key compromise. With the AppInstaller running as root for system-domain installs, it gives the holder of a stolen signing key arbitrary file write at root level via the delta-apply path, which is a strictly broader primitive than the "drop-in replacement bundle" install they would otherwise have. Version 2.9.2 contains a patch for the issue.	6.1	More Details
CVE-2026-8284	URL redirection to untrusted site ('open redirect') vulnerability in Universal Software Inc. FlexCity allows Input Data Manipulation. This issue affects FlexCity: from 5.536.0 through 11052026.	6.1	More Details
CVE-2026-50182	WWBN AVideo is an open source video platform. Versions prior to 29.0 contain an unauthenticated Reflected XSS vulnerability through AVideo YouTubeAPI Gallery Pagination. The $_GET['search'] query parameter is concatenated directly into the href attribute of two pagination links in plugin/YouTubeAPI/gallerySection.php (lines 67 and 74) with no htmlspecialchars, no urlencode, and no allow-list check. An injected <script> element is then extracted by the AVideo Layout plugin and concatenated into a single trailing inline script block at the bottom of the page, where the browser executes it. Any unauthenticated attacker can lure a victim into following a crafted URL to execute arbitrary JavaScript under the AVideo origin, which can read non-HttpOnly cookies and issue authenticated AJAX requests as the victim, and when the victim is an	6.1	More Details

	administrator, it can perform any cookie-authenticated admin action (create user, promote to admin, change configuration, install plugin), escalating a single click into full administrative takeover. This issue has been patched by this commit: https://github.com/WWBN/AVideo/commit/f50fc033b7adb36f1ffd6640e7826468bdafdec3.		
CVE-2026-10587	A potential out-of-bounds write vulnerability could allow a local privileged attacker to modify power management settings in System Management Mode.	6.0	More Details
CVE-2026-10589	A potential out of bounds write vulnerability could allow a local privileged attacker to execute code in System Management Mode.	6.0	More Details
CVE-2026-60073	An out-of-bounds read in the Productivity Suite allows a physical attacker to control the length of data sent to a USB device. This can lead to a system crash or disclosure of kernel memory.	5.9	More Details
CVE-2026-12510	The AI Engine WordPress plugin before 3.5.5 does not verify that a user owns the chatbot conversation referenced by a client-supplied identifier, allowing users with subscriber-level access to read other users' private conversations and take over their conversation records when the discussions feature is enabled.	5.9	More Details
CVE-2026-33443	CVE-2026-33443 is a memory management error in Secure Access servers prior to 14.55. Attackers with an intimate knowledge of and total control over the tunnel protocol can create a persistent DoS against the server.	5.9	More Details
CVE-2026-56742	Cilium is a networking, observability, and security solution. Prior to 1.17.17, 1.18.11, and 1.19.5, Cilium clusters using Gateway API allow users with permissions to create or update namespaced HTTPRoutes to mirror HTTP traffic to any Service in any namespace, bypassing the ReferenceGrant authorization mechanism. Gateway API functionality is disabled by default. This issue is fixed in versions 1.17.17, 1.18.11, and 1.19.5.	5.9	More Details
CVE-2026-46968	Vulnerability in Oracle Java SE (component: JSSE). Supported versions that are affected are Oracle Java SE: 8u491, 8u491-perf, 11.0.31, 17.0.19, 21.0.11, 25.0.3, 26.0.1; Oracle GraalVM for JDK: 17.0.19 and 21.0.11; Oracle GraalVM Enterprise Edition: 21.3.18. Difficult to exploit vulnerability allows unauthenticated attacker with network access via TLS to compromise Oracle Java SE. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Java SE accessible data. Note: This vulnerability can only be exploited by supplying data to APIs in the specified Component without using Untrusted Java Web Start applications or Untrusted Java applets, such as through a web service. CVSS 3.1 Base Score 5.9 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:H/A:N).	5.9	More Details
CVE-2026-56454	HCL DFXAnalytics is affected by a Deprecated Protocol vulnerability due to the use of TLS 1.0 and TLS 1.1. These legacy protocols contain numerous cryptographic design flaws that expose data to interception and decryption. To remediate this risk, the application must disable all support for TLS 1.0 and TLS 1.1, and exclusively enable support for secure protocols, specifically TLS 1.2 and TLS 1.3.	5.9	More Details
CVE-2026-59838	A improper neutralization of script-related html tags in a web page (basic xss) vulnerability in Fortinet FortiSIEM 7.4.0, FortiSIEM 7.3.0 through 7.3.4, FortiSIEM 7.2.0 through 7.2.6, FortiSIEM 7.1 all versions, FortiSIEM 7.0 all versions, FortiSIEM 6.7 all versions, FortiSIEM 6.6 all versions, FortiSIEM 6.5 all versions, FortiSIEM 6.4 all versions may allow attacker to execute unauthorized code or commands via <insert attack vector here>	5.9	More Details
CVE-2026-61643	FastGPT is a knowledge-based AI application platform. From 4.14.17 until 4.15.0-beta5, an authenticated FastGPT user can save a workflow node that points to another user's private HTTP toolset by using a crafted saved tool id such as http-<victim_toolset_app_id>/<tool_name>. The normal toolset routes deny access, but the workflow save and runtime path did not apply the same authorization check to the referenced toolset, allowing /api/v2/chat/completions to resolve the saved reference and execute the victim-owned HTTP tool. This issue is fixed in version 4.15.0-beta5.	5.9	More Details
CVE-2026-33445	CVE-2026-33445 is a memory management vulnerability in Secure Access servers prior to 14.55. Attackers with an intimate knowledge of and total control over the tunnel protocol can create a persistent DoS against the server.	5.9	More Details
CVE-2026-49834	sigstore-go is a Go library for Sigstore signing and verification. Prior to 1.2.0, a verifier configured with WithTransparencyLog(N>1) or WithSignedCertificateTimestamps(N>1) counts verified witnesses per entry or per validation path rather than per log authority, allowing a single compromised transparency log or CT log to satisfy multi-log threshold requirements and defeat the multi-log policy. This issue is fixed in version 1.2.0.	5.9	More Details
CVE-2026-59847	A flaw was found in libssh. Incorrect AES-GCM finalization checks in builds using the OpenSSL backend can effectively remove integrity protection, allowing an in-path attacker to modify plaintext on the wire without detection.	5.9	More Details
CVE-2026-	WordPress 6.8.x before 6.8.6, 6.9.x before 6.9.5, and 7.0.x before 7.0.2 does not properly sanitise the author__not_in parameter of WP_Query, which could allow SQL Injection when a plugin or theme passes	5.9	More

60137	untrusted input to the parameter.		Details
CVE-2026-4942	IBM i 7.6, 7.5, 7.4, and 7.3 could allow a remote attacker to send a specifically crafted message and downgrade the Transport Layer Security (TLS) protocol to a version disabled in the server configuration.	5.9	More Details
CVE-2026-49835	Sigstore Timestamp Authority is a service for issuing RFC 3161 timestamps. Prior to 2.1.0, the global wrapMetrics middleware records raw HTTP request path r.URL.Path and raw HTTP request method r.Method as Prometheus labels for latency and request count metric vectors before routing, allowing an unauthenticated remote attacker to issue requests with random paths such as /api/v1/timestamp/<uuid> or random HTTP methods and create unbounded permanent time-series entries that exhaust memory. This issue is fixed in version 2.1.0.	5.9	More Details
CVE-2026-58482	Network-AI, a TypeScript/Node.js multi-agent orchestrator, has a shipped, exported, documented feature called `ApprovalInbox` (`lib/approval-inbox.ts`). It is the network surface of the human-in-the-loop Approval Gate, which `ApprovalGate` uses to require explicit human approval for high-risk operations. The HTTP server it exposes has no authentication of any kind and sets `Access-Control-Allow-Origin: *` on every route, including the state-changing `POST /approvals/:id/approve` and `/deny`. As a result, in versions 5.0.0 through 5.12.1, any party who can send an HTTP request to the inbox port — a co-located process, a container/SSRF on the same host, a remote client when the operator binds a non-loopback address, or any website the operator visits in a browser (via the wildcard CORS) — can enumerate pending approvals and approve them, defeating the entire human-in-the-loop control and causing the gated high-risk action (e.g. a shell command the agent was holding for review) to execute without consent. This issue is fixed in v5.12.2. `ApprovalInbox` now accepts a `secret` option. When set, the mutating endpoints `POST /:id/approve` and `POST /:id/deny` require an `Authorization: Bearer <secret>` header, validated in constant time with `crypto.timingSafeEqual`. `startServer()` already binds to `127.0.0.1` by default; operators exposing the inbox on a network must set a secret.	5.9	More Details
CVE-2026-44452	h2o is an HTTP server with support for HTTP/1.x, HTTP/2 and HTTP/3. Prior to commit 8dc37cb, when h2o receives a ClientHello message over TLS or QUIC and it contains a zero-length SNI extension, the h2o server runs over the zero-length hostname while trying to copy the hostname, assuming that it is NULL-terminated. This is a potential denial-of-service attack vector in sense that it might trigger segmentation violation. This issue has been fixed by commit 8dc37cb.	5.9	More Details
CVE-2026-13693	The Bit Form WordPress plugin before 3.1.0 does not restrict a form file-field value to a safe path before reading the file and attaching it to a notification email, allowing unauthenticated attackers to read arbitrary server files such as the WordPress configuration file.	5.9	More Details
CVE-2026-45712	Mailpit is an email testing tool and API for developers. Prior to version 1.30.0, the screenshot/print proxy (/proxy?data=...) maintains a package-level assets map[string]MessageAssets cache, but reads the map without holding assetsMutex while a long-running cleanup goroutine and (re-entrant) CSS-rewriting code path concurrently write to it under the lock. When the unsynchronized read coincides with a synchronized write, Go's runtime raises fatal error: concurrent map read and map write — a runtime.throw that is not recoverable by http.Server's handler-panic recover. The whole Mailpit process exits, taking the SMTP, POP3 and HTTP listeners down with it. Version 1.30.0 contains a patch.	5.9	More Details
CVE-2026-45711	Mailpit is an email testing tool and API for developers. Prior to version 1.30.0, the mailpit dump --http <base-url> <out-dir> sub-command downloads every message from a remote Mailpit instance and writes each one as <id>.eml inside the user-supplied output directory. The message ID field is taken verbatim from the JSON response of the remote server and concatenated into the output path with path.Join, which silently normalizes `..` segments. A malicious HTTP server impersonating Mailpit can therefore make mailpit dump write attacker-controlled bytes to any path the running user can write, fully outside the intended output directory. Version 1.30.0 contains a patch.	5.9	More Details
CVE-2026-47160	Vaultwarden is a Bitwarden-compatible server written in Rust. Prior to 1.36.0, Vaultwarden's /icons/{domain}/icon.png endpoint used src/http_client.rs checks including should_block_address() and post_resolve() that missed decimal, hexadecimal, and octal IP representations, allowing SSRF through the icon-fetching HTTP client for blind internal network or port discovery. This issue is fixed in version 1.36.0.	5.8	More Details
CVE-2026-63096	Dendrite through 0.13.8 contains a server-side request forgery vulnerability that allows unauthenticated attackers to cause the server to open outbound TLS connections to arbitrary hosts and ports by supplying an unvalidated serverName parameter to the legacy media download endpoint. Attackers can exploit distinguishable error response classes and leaked internal IP addresses in error messages to perform blind port scanning and enumerate internal network topology.	5.8	More Details
CVE-2026-63428	HeyForm is an open-source form builder. Prior to version 3.0.0-rc.9, `completeSubmission` accepts a `hiddenFields: [{id, name, value}]` array from the submitter and stores it verbatim in `submission.hiddenFields`, without validating the supplied `id`/`name` against the form's declared `form.hiddenFields` schema. An anonymous form submitter can therefore inject arbitrary key/value pairs (including XSS payloads, fake authorization metadata, integration-relevant values) into the stored submission. These fields are subsequently forwarded as-is to every webhook integration registered on the form. Version 3.0.0-rc.9 contains a patch for the issue.	5.8	More Details

CVE-2026-45709	Mailpit is an email testing tool and API for developers. The fix for GHSA-6jxm-fv7w-rw5j (CVE-2026-23845, "Server-Side Request Forgery (SSRF) via HTML Check API"), shipped in mailpit `v1.28.3`, hardened `internal/htmlcheck/css.go::downloadCSSToBytes` with a 5MB size cap, a `text/css` content-type check, login-info stripping in `isValidURL`, and an opt-in `--block-remote-css-and-fonts` config flag — but did not add the IP-filtering dialer that the same codebase already uses on the two sister SSRF endpoints (the proxy handler and link-check). Prior to version 1.30.0, `internal/htmlcheck/css.go::newSafeHTTPClient` is mis-named — it builds an `http.Client` whose `Transport.DialContext` calls `net.Dialer.DialContext` directly with no IP allowlisting. As a result, the SSRF originally reported by Bao Anh Phan still permits the server to dial loopback, private, link-local, and any other reserved/multicast range, provided the target replies with `HTTP/200` and a content-type beginning with `text/css`. With redirect-following (`CheckRedirect` allows redirects to any `isValidURL` URL with no IP filter), an attacker-controlled public site can redirect mailpit's request into the private network without ever appearing in the email's HTML. In the default mailpit deploy (no UI auth, no SMTP auth, port 1025/8025 exposed), this is an unauthenticated, network-reachable SSRF triggered by sending an HTML email and then issuing one HTTP `GET` to `/api/v1/message/{id}/html-check`. Version 1.30.0 contains an updated fix.	5.8	More Details
CVE-2026-15811	A vulnerability was found in kronosnet's (version <=1.34) cryptographic configuration management. The framework does not correctly zero-out or wipe sensitive memory segments after executing changes to its cryptographic configuration. This omission leaves raw encryption keys resident in memory after the associated structures are freed. A local attacker capable of leveraging memory disclosure techniques could exploit this flaw to retrieve the active encryption key, allowing them to decrypt cluster network communications or inject malicious packets to cause severe high-availability cluster instability.	5.8	More Details
CVE-2026-62314	Anubis is a Web AI Firewall Utility that challenges users' connections in order to protect upstream resources from scraper bots. From 1.22.0 until 1.26.0-pre1, lib/policy/checker.go PathChecker.Check() trusted the client-controlled X-Original-URI header before matching r.URL.Path, allowing an HTTP client to match default data/common/keep-internet-working.yaml ALLOW rules such as ^/\.well-known/.*$ and bypass the Anubis challenge. This issue is fixed in version 1.26.0-pre1.	5.8	More Details
CVE-2026-46975	Vulnerability in the RDBMS component of Oracle Database Server. Supported versions that are affected are 19.3-19.31, 21.3-21.22 and 23.4.0-23.26.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via Oracle Net to compromise RDBMS. While the vulnerability is in RDBMS, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of RDBMS accessible data. CVSS 3.1 Base Score 5.8 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:N/I:L/A:N).	5.8	More Details
CVE-2026-15737	AWS Bedrock AgentCore Python SDK is an open-source Python library that provides client tools for building AI agents on the Amazon Bedrock AgentCore platform. Unintended logging of sensitive user content in the OpenTelemetry instrumentation in AWS Bedrock AgentCore Python SDK versions 1.4.8 and 1.5.0 might allow a local authenticated user with access to CloudWatch Logs to access raw user prompts and agent responses containing sensitive data via span attributes. The SDK wrote raw user prompts and complete agent responses into OpenTelemetry span attributes on every invocation without filtering or masking. These spans flow into the customer's aws/spans CloudWatch log group, exposing sensitive content to any principal with log read access. We recommend you upgrade to version 1.5.1 or later. Users who ran affected versions should also review and purge sensitive content from their aws/spans CloudWatch log groups.	5.7	More Details
CVE-2024-32387	An issue in Kerlink Kerlink Wirnet iStation 868 KerOS v.4.3.3_20200803132042 allows a remote attacker to obtain sensitive information via the community string component.	5.7	More Details
CVE-2026-16198	A vulnerability was detected in Sipeed PicoClaw up to 0.2.9. The impacted element is an unknown function of the file web/backend/middleware/access_control.go of the component First Run Setup. Performing a manipulation of the argument allowed_cidrs results in authentication bypass using alternate channel. The attack may be initiated remotely. The attack's complexity is rated as high. The exploitability is regarded as difficult. The exploit is now public and may be used. The patch is named 017601354be38cb027ff3ffb01aed79bd5d12610. Applying a patch is the recommended action to fix this issue.	5.6	More Details
CVE-2026-57852	Grav CMS scheduler-webhook plugin contains an authentication bypass vulnerability that allows unauthenticated remote attackers to trigger configured scheduled jobs by exploiting a short-circuit logic flaw in the webhook token validation. Attackers can send a single unauthenticated POST request to the scheduler webhook endpoint to execute all configured scheduled jobs or target a specific job, causing unintended execution of operator-defined commands under the web server process user.	5.6	More Details
CVE-2024-	An access control bypass and information disclosure vulnerability exists in the base AppArmor security profile configuration of Canonical snapd. The abstraction rules located in /etc/apparmor.d/abstractions/nss-systemd (inherited via) inadvertently permit strictly confined snap applications, which lack the privileged account-control interface, to interact directly with the io.systemd.Multiplexer and io.systemd.NameServiceSwitch UNIX domain sockets under /run/systemd/userdb/. On systems where the systemd-userdbd service is installed and operational, the service fails to distinguish between an unconfined root user on the host system and a	5.6	More

5300	restricted root user running within a snap application's sandbox (such as a daemon or configuration hook). Because systemd-userdbd returns "complete" user records—including sensitive hashed user passwords from /etc/shadow—when queried by a process running as root, a compromised or malicious strictly confined snap executing code as root can successfully query the Varlink interface to retrieve all system password hashes, bypassing intended snap sandbox restrictions. This issue is mitigated by the fact that systemd-userdbd is not installed by default on standard Ubuntu deployments.		Details
CVE-2026-46420	setup-php is a GitHub action to set up PHP with extensions, php.ini configuration, coverage drivers, and tools. From 2.25.0 prior to 2.37.1, shivammathur/setup-php resolves the PHP version from repository-controlled files such as .php-version, composer.lock through platform-overrides.php, and composer.json through config.platform.php, and insufficiently constrains those values before incorporating them into generated shell or PowerShell setup scripts, allowing command injection on a GitHub Actions runner when workflows such as pull_request_target check out attacker-controlled contents before invoking setup-php. This issue is fixed in version 2.37.1.	5.6	More Details
CVE-2026-62361	listmonk is a standalone, self-hosted, newsletter and mailing list manager. Prior to 6.2.0, listmonk's GET /api/subscribers/export endpoint injects the user-controlled query parameter into QuerySubscribersForExport in internal/core/subscribers.go without calling validateQueryTables, unlike GET /api/subscribers, allowing an authenticated user with subscribers:sql_query and subscribers:get_all to read arbitrary database tables such as users and settings and execute data-modifying PostgreSQL CTEs. This issue is fixed in version 6.2.0.	5.5	More Details
CVE-2026-15943	A flaw was found in the Keycloak keycloak-services component, which handles the management of identity providers. The issue occurs when a delegated administrator updates an OIDC identity provider using a masked client secret sentinel value. Due to improper validation, Keycloak reuses the existing real secret even if security-sensitive fields like the token URL have been changed, allowing an attacker to redirect and capture the secret.	5.5	More Details
CVE-2026-9494	An information disclosure vulnerability exists in Canonical ubuntu-pro-client (formerly ubuntu-advantage-tools). The client validates Ubuntu Pro APT credentials by executing /usr/lib/apt/apt-helper using the download-file command. During this process, the secret bearer token is embedded directly in the cleartext URL component passed via the command-line arguments (argv), resulting in a URL format such as https://bearer:<token>@esm.ubuntu.com/.../. On systems utilizing a default-mounted /proc file system where process-hiding mitigations (such as hidepid) are disabled, an unprivileged local attacker can monitor system processes and read the sensitive bearer token directly from /proc/cmdline while the helper process is actively running. This leaked token can subsequently be used to gain unauthorized access to the victim's Ubuntu Pro or Expanded Security Maintenance (ESM) repositories.	5.5	More Details
CVE-2026-20146	A vulnerability in Cisco Identity Services Engine (ISE) and Cisco ISE Passive Identity Connector (ISE-PIC) could allow an authenticated, remote attacker to perform path traversal attacks on the underlying operating system to either read or delete arbitrary files. To exploit this vulnerability, the attacker must have valid administrative credentials. This vulnerability is due to improper validation of user-supplied input. An attacker could exploit this vulnerability by sending a crafted HTTP request to an affected system. A successful exploit could allow the attacker to access sensitive files or delete arbitrary files on the affected system.	5.5	More Details
CVE-2026-45612	rz-libdemangle is a Rizin library for demangling symbols. Prior to 6bf56d3, the Rust demangler in src/rust/rust_v0.c can perform an out-of-bounds read when the demangler structure is not yet initialized. This issue is fixed in commit 6bf56d3.	5.5	More Details
CVE-2026-58414	Network-AI is a TypeScript/Node.js multi-agent orchestrator. Prior to version 5.12.2, `EnvironmentManager.backup()` recursively collects files using `_collectBackupFiles()`. `_collectBackupFiles()` uses `statSync(full)`, which follows symlinks. If `data/<env>` contains a symlink to a directory outside the environment root, backup recursion follows the symlink and copies external files into `data/<env>/.backups/<backupId>/`. An attacker who can place a symlink under the environment data directory can cause backup operations to disclose files outside the environment root into backup artifacts. The issue is fixed in v5.12.2. `_collectBackupFiles()` now uses `lstatSync` instead of `statSync` and skips any entry where `isSymbolicLink()` is true. Symlinks are never traversed, so `backup()` can no longer follow a link out of the environment root and copy external files into a backup artifact.	5.5	More Details
CVE-2026-56453	HCL DFXAnalytics is affected by an Account Takeover via Response Manipulation vulnerability. A remote attacker can intercept and alter the contents of the server's HTTP responses before they reach the client application, allowing them to manipulate the authentication or authorization logic to bypass controls and gain unauthorized access to targeted user accounts.	5.5	More Details
CVE-2026-50012	Squid is a caching proxy for the Web. Prior to 7.6, due to an improper input validation bug in cache digest reply handling (peerDigestSwapInMask in src/peer_digest.cc), Squid is vulnerable to a heap-based buffer overflow: a cache digest's on-the-wire size may be larger than the mask_size declared within the digest, so a trusted peer sending a maliciously crafted reply to a cache_digest request message can trigger the overflow. This attack is limited to Squid instances compiled with the --enable-cache-digests option and configured with cache_peer entries. This issue is fixed in version 7.6.	5.5	More Details
	The NXP LPUART serial driver (drivers/serial/uart_mcux_lpuart.c), when		

CVE-2026-10674	CONFIG_UART_USE_RUNTIME_CONFIGURE is enabled, called LPUART_Deinit() at the start of mcux_lpuart_configure(), which disables the LPUART peripheral clocks. The requested configuration is validated only afterwards (in mcux_lpuart_configure_basic), and unsupported parity/data-bit/stop-bit/flow-control values return -ENOTSUP before the clock is re-enabled. As a result, a uart_configure() request with an unsupported configuration left the LPUART in a clock-disabled state; any subsequent access to LPUART registers (poll_out/poll_in, interrupt handling, or a later reconfigure) faults on the gated peripheral and escalates to a hard fault, crashing the system. uart_configure() is a Zephyr syscall whose verifier (z_vrfy_uart_configure) only checks that cfg is readable user memory and forwards the caller-supplied configuration unchanged, so an unprivileged userspace thread with access to an LPUART device can deterministically trigger the fault, a persistent system-wide denial of service. Introduced in v2.5.0 and present in all subsequent releases until this fix, which removes the LPUART_Deinit() call and instead only disables the transmitter/receiver, leaving the clock running.	5.5	More Details
CVE-2026-15415	AWS HealthOmics is a HIPAA-eligible service that fully manages the compute, storage, and workflow engine infrastructure required to run bioinformatics analyses at scale for clinical diagnostics, drug discovery, and agricultural research. Improper limitation of a pathname to a restricted directory in the linting tools of the AWS HealthOmics MCP Server (aws-healthomics-mcp-server) before version 0.0.36 might allow an actor who can influence the MCP agent to write an actor-controlled content to arbitrary locations outside the intended workflow bundle directory, via directory traversal sequences in the workflow_files input. To remediate this issue, users should upgrade to version 0.0.36 or later.	5.5	More Details
CVE-2026-7771	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to a trap when compiling a specially crafted statements containing subqueries could lead to a denial of service.	5.5	More Details
CVE-2026-47144	Shamefile is a linter for undocumented linter warnings. Prior to version 0.1.7, a path traversal vulnerability in `shame next` allows an attacker-controlled `shamefile.yaml` to disclose contents of files outside the repository, one line at a time, to the terminal of a user who runs the command. See patch commit for technical details. The issue is fixed in 0.1.7. Upgrade to either 0.1.7 or later versions to incorporate the patch. As a workaround, do not run `shame next` against untrusted `shamefile.yaml`. Use `shame me --dry-run` for CI validation.	5.5	More Details
CVE-2026-6511	During an internal security assessment, a potential improper access control vulnerability was discovered in Lenovo Smart Connect for Windows that could allow a local authenticated user to access files owned by a different user on the same system.	5.5	More Details
CVE-2026-57848	Stoat for Android exports the chat.stoat.activities.ShareTargetActivity component (reachable to any process on the device via the android.intent.action.SEND intent) and accepts the file to share as a URI supplied through the android.intent.extra.STREAM extra. The activity does not validate or filter the incoming URI before using it as the outgoing attachment, so a caller can pass a file:// URI pointing at the application's own internal storage (for example /data/data/chat.revolt/databases/revolt.db, cached authentication token files, or preferences) and have the app treat that internal file as a user-selected attachment. An attacker who can invoke intents on the victim's device (via ADB access, a co-installed malicious application, or any other route that reaches Android's intent dispatch) can launch ShareTargetActivity with such a URI and cause the victim, on a single channel-selection interaction, to send the internal file to any Stoat channel or user of the attacker's choosing. The composer displays the attachment as \"attachment\" with no filename indication, so the victim has no visible signal that the file being sent is their own internal application data. Consequences include disclosure of the local Stoat database (message history, contact list, cached content), disclosure of authentication tokens permitting full account takeover, and disclosure of any other file readable by the app process.	5.5	More Details
CVE-2026-47390	PraisonAI is a multi-agent teams system. Prior to version 4.6.40 of PraisonAI, corresponding to version 1.6.40 of praisonaiagents, `spider_tools` URL validation can be bypassed using alternate loopback host encodings. The tool contains a URL validation function intended to block local or unsafe targets before fetching attacker-controlled URLs. However, the validation only blocks a small set of exact host strings such as `localhost` and `127.0.0.1`. It does not normalize hostnames, resolve DNS, parse numeric IPv4 variants, or validate the final resolved IP address before making the request. As a result, certain URLs may bypass the protection and still reach loopback services. After the weak validation passes, `scrape_page()` calls `requests.Session.get()` on the attacker-controlled URL. This allows an attacker who can influence URLs passed to `scrape_page`, `crawl`, or `extract_text` to induce SSRF requests against loopback-only services. This is a server-side request forgery protection bypass. PraisonAI version 4.6.40 and praisonaiagents version 1.6.40 contain a patch.	5.5	More Details
CVE-2026-47395	PraisonAI is a multi-agent teams system. Prior to version 4.6.40 of PraisonAI, corresponding to version 1.6.40 of praisonaiagents, PraisonAI's direct-prompt CLI automatically expands `@url:` mentions in raw prompt text before agent execution begins. If a prompt contains `@url:<http-or-https-url>`, the CLI calls `MentionsParser.process(...)`. The `@url:` handler then performs a direct `urllib.request.urlopen()` request to the attacker-controlled URL and returns the response body. That response body is prepended to the final model prompt context. There is no loopback/private-address restriction, no metadata-service restriction, and no approval gate before the fetch. As a result, attacker-influenced prompt text can cause the operator's machine to fetch localhost-only HTTP resources and inject the response into model context. PraisonAI version	5.5	More Details

	4.6.40 and praisonaiagents version 1.6.40 contain a fix.		
CVE-2026-61378	A divide-by-zero vulnerability in the Productivity Suite allows a local attacker to cause a division by zero leading to a system crash.	5.5	More Details
CVE-2026-12979	The FunnelKit WordPress plugin before 3.15.0.6 does not validate a user-supplied path before deleting a file during a template-import operation, allowing users with administrator privileges to delete arbitrary .json files outside the intended directory through path traversal, which can disable other FunnelKit WordPress plugin before 3.15.0.6 or (denial of service).	5.5	More Details
CVE-2026-11580	The Kali Forms — Contact Form & Drag-and-Drop Builder WordPress plugin before 2.4.17 does not perform a per-object capability check in its post-duplication AJAX action, allowing users with Contributor-level access or above to duplicate any post (regardless of owner, post type, or status) into a published post they own and read its private post metadata, including secrets stored by other Kali Forms — Contact Form & Drag-and-Drop Builder WordPress plugin before 2.4.17.	5.5	More Details
CVE-2026-44229	RT is an open source, enterprise-grade issue and ticket tracking system. Versions 5.0.0 and 6.0.0 and above, prior to both 5.0.10 and 6.0.3 contain a Cross-Site Scripting (XSS) vulnerability where uploaded content is served inline rather than as an attachment. An authenticated user who can upload content can include JavaScript in the upload that will execute in the browser session of any RT user who later views or downloads it. This issue has been fixed in versions 5.0.10 and 6.0.3.	5.4	More Details
CVE-2026-62221	OpenClaw 2026.5.12 before 2026.5.26 contain an incorrect authorization vulnerability in the ClickClack allowFrom feature. When the affected feature is enabled and reachable, a lower-trust caller or configured input path could execute or persist actions beyond the caller's intended authorization, including running non-allowlisted commands.	5.4	More Details
CVE-2026-15069	IBM Engineering AI Hub 1.0.0, 1.1.0, and 1.2.0 could allow a remote attacker to execute arbitrary script code due to improper neutralization of input during web page generation.	5.4	More Details
CVE-2026-63082	Perfect Support Ticketing & Document Management System through 1.7 contains a broken access control vulnerability that allows authenticated attackers with Agent-level privileges to manipulate the Support Agent assignment field of tickets by bypassing intended authorization checks. Attackers can add or remove any user, including Superadmin accounts, from the Support Agent field of any ticket to which they are assigned, circumventing role-based access controls.	5.4	More Details
CVE-2026-47082	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. The vacation "fcc" feature skips the destination-mailbox ACL. A user whose vacation Sieve script used :fcc (to save a copy of the sent message) could deliver vacation auto-reply copies into any mailbox the script could name, regardless of whether the script owner had insert permissions on the destination mailbox.	5.4	More Details
CVE-2026-63081	Perfect Support Ticketing & Document Management System through 1.7 contains a stored cross-site scripting vulnerability that allows authenticated attackers with Agent-level privileges to inject malicious payloads into the Notes field of assigned support tickets. Attackers can store malicious scripts that execute in the browser context of any user who views the affected ticket notes, including Superadmin users, enabling session hijacking or unauthorized actions on behalf of the victim.	5.4	More Details
CVE-2026-15995	IBM Cognos Analytics 12.1.3 GA Version with build number through 12.1.3-2606251736 could allow an attacker to obtain incorrect report summary results or cause report-processing failures due to a race condition in the Agentic AI assistant's concurrent request-handling logic when multiple authenticated users submit report-related tasks simultaneously.	5.4	More Details
CVE-2026-16089	A flaw was found in the keycloak-services component of Red Hat Build of Keycloak. The issue occurs because OAuth 2.0 authorization codes are not properly bound to the client that originally requested them. An attacker who can intercept an authorization code can modify it to be redeemed by their own client, potentially allowing them to obtain access tokens for a victim's identity.	5.4	More Details
CVE-2026-11866	The Appointment Booking Plugin WordPress plugin before 5.6.3 does not validate a CSRF nonce on several state-changing actions handled by its central request dispatcher, allowing attackers to perform privileged actions, such as overwriting the booking-form configuration or disconnecting the connected payment gateway, via Cross-Site Request Forgery against a logged-in administrator.	5.4	More Details
CVE-2026-57980	Authentication bypass using an alternate path or channel in Microsoft Edge (Chromium-based) allows an unauthorized attacker to perform tampering over a network.	5.4	More Details
CVE-2026-62224	OpenClaw MS Teams before 2026.5.12 contain an authorization bypass vulnerability where the allowFrom feature binds to mutable display names. Attackers with lower-trust access can perform actions requiring stronger authorization by exploiting the mutable display name binding in the affected feature.	5.4	More Details
	Grav contains a stored cross-site scripting vulnerability in shortcode-core attribute handlers where the XSS		

CVE-2026-64628	detection scan only matches payloads containing literal angle brackets, allowing shortcode parameters to bypass validation. Attackers with admin.pages permission can inject malicious JavaScript through shortcode attributes that execute in any viewer's browser, including administrators, enabling session hijacking via admin nonce theft.	5.4	More Details
CVE-2026-62225	OpenClaw versions before 2026.5.18 contain an authorization bypass vulnerability in skill command dispatch that allows lower-trust callers to execute or persist actions beyond their intended authorization. Attackers can bypass tool policy restrictions through configured input paths to perform unauthorized actions when the affected feature is enabled and reachable.	5.4	More Details
CVE-2026-44585	Paymenter is a free and open-source webshop solution for management of hosting services. In versions prior to 1.5.0, the ticket creation endpoint accepts a user-supplied service identifier without enforcing ownership validation, allowing authenticated users to create support tickets referencing services belonging to other accounts by modifying the service ID in the request. An attacker could modify the service ID value in the client-side request and successfully create a ticket associated with another user's service. The vulnerability requires authentication and does not provide direct access to service contents or customer data. However, referenced service information could become visible to support personnel handling the ticket. Successful exploitation could allow an authenticated user to: create support tickets referencing services belonging to other users, potentially cause support staff to interact with or review unrelated customer services. The vulnerability did not allow direct access to another user's service, modification of another user's service or retrieval of confidential service data through the vulnerable endpoint itself. This issue has been fixed in version 1.5.0.	5.4	More Details
CVE-2026-14184	The Academy LMS WordPress plugin before 3.8.1 does not verify ownership of a user-supplied user identifier in several of its lesson AJAX handlers, allowing authenticated users with subscriber-level access to read and modify other users' lesson notes and mark other users' lesson content as completed.	5.4	More Details
CVE-2026-62236	grav-plugin-login before 3.8.11 contains a cross-site request forgery (CSRF) vulnerability in the login.regenerate2FASecret frontend task, which regenerates and persists a new TOTP secret for the authenticated session user without any anti-CSRF nonce or Origin/Referer check. Because Grav core dispatches the task from the GET 'task:' URI parameter and the default session cookie is SameSite=Lax, an attacker can lure a logged-in victim to an off-site page that performs a top-level GET navigation, rotating the victim's TOTP secret so their enrolled authenticator no longer matches the server, effectively forcing 2FA re-enrollment. Sites configured with session.samesite: Strict are not affected.	5.4	More Details
CVE-2026-58624	Improper input validation in sshd-git in Apache MINA SSHD. Apache MINA SSHD is a Java library for client-side and server-side SSH. Component org.apache.sshd:sshd-git provides though its GitPgmCommandFactory a way to configure an Apache MINA SSHD server such that SSH clients can remotely execute git commands via the JGit library on git repositories stored on the server. This GitPgmCommandFactory allowed a user authenticated via SSH to run any JGit command available, including commands that could write files at arbitrary places such as git archive with the --output option. Affected are SSH servers implemented with Apache MINA SSHD and using the GitPgmCommandFactory. If the GitPgmCommandFactory is not configured on the server, the server is not affected. It is recommended to upgrade affected servers to Apache MINA SSHD 2.19.0 or 3.0.0-M5, which fix this issue. The issue is fixed by restricting the available commands to a small whitelist of uncritical commands (such as git log). git archive is also allowed, but its --output argument is ignored and the archive is always sent through the SSH channel to the client.	5.4	More Details
CVE-2026-49215	Symfony UX is a JavaScript ecosystem for Symfony. From 2.22.0 until 2.36.0 and 3.1.0, Symfony\UX\LiveComponent\EventListener\LiveComponentSubscriber::isLiveComponentRequest() gates #[LiveAction] invocations on Accept: application/vnd.live-component+html, but the Accept header is CORS-safelisted and cross-origin fetch() can set it without preflight, allowing forged cross-origin #[LiveAction] requests against a victim session when applications use SameSite=None, credentials: 'include', a permissive cookie policy, or a same-origin pivot. This issue is fixed in versions 2.36.0 and 3.1.0.	5.4	More Details
CVE-2026-16093	Keycloak provides a mechanism called Client Policies to enforce security requirements on clients, such as requiring them to use signed JWTs for authentication. A flaw was discovered where this enforcement can be bypassed. An attacker with valid client credentials can provide a fake, unsigned assertion header that tricks the system into thinking the policy requirements have been met. This allows the attacker to authenticate using simpler methods like a client secret even when the administrator has mandated more secure, signed assertions.	5.4	More Details
CVE-2026-62355	TDengine is an open source, time-series database optimized for Internet of Things devices. Prior to 3.4.1.15, a Data Reader admin_user on a TDengine Cloud DB instance could run create udf even though standard users should have read-only permissions for non-database objects and show dnodes and create user were denied. This issue is fixed in version 3.4.1.15.	5.4	More Details
CVE-2026-56146	Improper Access Control (CWE-284) in Kibana can lead to unauthorized modification of Entity Analytics Watchlist configuration and potential information disclosure. A low-privileged authenticated user with read-only Security Solution access could perform write operations on watchlist data that should require elevated privileges. Under specific deployment conditions, this could also allow such a user to access data beyond their authorized scope.	5.4	More Details

CVE-2026-13156	The MailerSend WordPress plugin before 1.0.8 does not perform a nonce check on its configuration-delete action (it verifies the manage_options capability but ignores the nonce), so an attacker can trick a logged-in administrator into visiting a crafted page that wipes the MailerSend WordPress plugin before 1.0.8's SMTP configuration and deactivates the MailerSend WordPress plugin before 1.0.8, breaking the site's email delivery.	5.4	More Details
CVE-2026-62348	TDengine is a time-series database optimized for Internet of Things devices. Prior to 3.4.1.15, TDengine Enterprise allowed an authenticated low-privilege SQL user to run KILL SSMIGRATE <id> against an active shared-storage migration because mndProcessKillSsMigrateReq called mndKillSsMigrate while the intended MND_OPER_SSMIGRATE_DB privilege check was commented out. This issue is fixed in version 3.4.1.15.	5.4	More Details
CVE-2026-13432	The ThumbPress WordPress plugin before 6.2.2 does not perform a capability check on one of its AJAX actions, allowing authenticated users with subscriber-level access or higher to deactivate the ThumbPress WordPress plugin before 6.2.2, disrupting the site's image-handling functionality.	5.4	More Details
CVE-2026-26032	The PackagerResolver of Apache Ivy is able to download online artifacts and to (re)package them in a format defined by a packager.xml file. This repackaging is done by an Ant script, which is stored in a subdirectory of the configured "buildRoot" directory. This subdirectory is calculated based on modules coordinates, like the organisation, name or version. If one of the coordinates contains "../" sequences - which are valid characters for Ivy coordinates in general- it is possible to break out of the configured "buildRoot" directory where other files can be overwritten. In order to exploit this vulnerability an attacker needs to have access to a packager repository and add or modify the coordinates in ivy.xml files to have such "../" sequences. Users of Apache Ivy 2.0.0 to 2.5.3 (inclusive) should upgrade to Ivy 2.6.0.	5.4	More Details
CVE-2026-63758	SurrealDB versions before 3.1.0 contain an authorization bypass vulnerability in the KILL statement that allows authenticated database users to terminate other users' LIVE SELECT subscriptions. Attackers can issue KILL statements with target live query UUIDs to disrupt real-time data subscriptions of other users without ownership verification.	5.4	More Details
CVE-2026-45138	CI4MS is a CodeIgniter 4-based content management system skeleton. Prior to version 0.31.9.0, the custom `html_purify` validation rule used to sanitize blog post bodies relies on by-reference mutation (`?string &$str`), but CodeIgniter 4's validator passes a local copy of the value, so the sanitized text is silently discarded. The Blog controller writes `$lanData['content']` directly into `blog_langs.content`, and the public template echoes it without escaping — yielding stored XSS executable in any visitor's browser, including the superadmin when previewing or editing posts. Version 0.31.9.0 patches the issue.	5.4	More Details
CVE-2026-63102	rConfig Core before 8.2.8 contains a privilege escalation vulnerability that allows authenticated users to assign arbitrary roles to any account by submitting an unvalidated role field through the Users API during user creation or profile updates. Attackers can exploit the missing allowlist validation and absent admin-level authorization check in StoreUserRequest to mass-assign the Admin role directly to the User model, granting access to privileged features. rConfig Pro and Enterprise are not affected.	5.4	More Details
CVE-2026-6793	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Bifra Engineering Consulting Ltd. Q-smart NexT Poll allows Stored XSS. This issue affects Q-smart NexT Poll: before 1.8.7.	5.4	More Details
CVE-2026-61718	bunkerweb is an Open-source and next-generation Web Application Firewall (WAF). From 1.6.2 until 1.6.12, the BunkerWeb web UI BiscuitMiddleware authorization bypass list included the /cache/ URL prefix, so routes in src/ui/app/routes/cache.py protected only by @login_required, including POST /cache/delete, allowed low-privilege read-only reader accounts to permanently delete job cache files containing blacklist, greylist, DNSBL, CrowdSec, GeoIP, ModSecurity CRS, Let's Encrypt, ACME, and custom configuration data. This issue is fixed in version 1.6.12.	5.4	More Details
CVE-2026-47671	Nhost is an open source Firebase alternative with GraphQL. In versions of Nhost CLI prior to 1.46.0, the hidden `nhost configserver` used by `nhost dev` exposes the Mimir GraphQL API with dummy authorization directives and permissive CORS. When a developer is running the local development environment, any process that can reach the developer's localhost service, including a web page loaded from an arbitrary origin, can query the configserver for local Nhost configuration and secrets and can mutate the local `.secrets` file. This impacts developers using `nhost dev`: project admin secrets, JWT signing keys, webhook secrets, Grafana credentials, and custom environment variables can be read, and attacker-controlled secrets can be written to the local development project. Version 1.46.0 of Nhost CLI contains a fix.	5.4	More Details
CVE-2026-44228	RT is an open source, enterprise-grade issue and ticket tracking system. Versions 6.0.0 and above, prior to 6.0.3, contain a stored Cross-Site Scripting (XSS) vulnerability, where user-controlled data is rendered without proper HTML escaping. An authenticated user with permission to set the relevant data can inject JavaScript that executes when another RT user views the affected page. This issue has been fixed in version 6.0.3.	5.4	More Details
CVE-2026-12393	The WPS Bookings for WooCommerce WordPress plugin before 3.11.7 does not verify that a booking order belongs to the requesting user before cancelling it, allowing any authenticated user, such as a Subscriber or Customer, to cancel and void other customers' booking orders.	5.4	More Details
	Cilium is a networking, observability, and security solution. From 1.19.0 to 1.19.4, standard Kubernetes		

CVE-2026-56743	NetworkPolicy specifications using CIDR-based ipBlock rules without pod or namespace selectors erroneously generate a wildcard namespace allow rule when Cilium is configured with a custom clusterName rather than the default any value. The parser incorrectly instantiates a pod selector on selectorless peer definitions, allowing traffic from other workloads in the same namespace as the subject of the policy. This issue is fixed in version 1.19.5.	5.4	More Details
CVE-2026-63745	SurrealDB versions before 3.1.0 contain an authorization bypass vulnerability where authenticated users can spoof composite record-id field values by writing to editable body fields. Attackers can bypass permission rules that gate access on id components like tenant isolation by setting same-named body fields to spoofed values that permission checks incorrectly read instead of the immutable id key.	5.4	More Details
CVE-2026-49216	Symfony UX is a JavaScript ecosystem for Symfony. From 2.2.0 until 2.36.0 and 3.1.0, the Stimulus controller in symfony/ux-autocomplete renders AJAX response items in _createAutocompleteWithRemoteData() by interpolating the text field into HTML template literals (${item[labelField]}) rather than text, allowing attacker-controlled markup from user-supplied dropdown values to execute in the browser of any user who opens an autocomplete widget backed by the same data. This issue is fixed in versions 2.36.0 and 3.1.0.	5.4	More Details
CVE-2026-49997	SurrealDB is a scalable, distributed, collaborative, document-graph database for the realtime web. Prior to 3.1.0, Document::purge_edges in surrealdb/core/src/doc/delete.rs automatically removed graph edge records with permissions disabled through opt.clone().with_perms(false) when a connected node was deleted, bypassing the edge table's PERMISSIONS FOR delete and PERMISSIONS FOR select clauses. This issue is fixed in version 3.1.0.	5.4	More Details
CVE-2026-62353	TDengine is a time-series database optimized for Internet of Things devices. Prior to 3.4.1.14, source/libs/parser/src/parTokenizer.c tGetToken() incremented past a trailing backslash in a SQL string literal such as 'abc\ and read one byte beyond the null terminator, allowing an authenticated user who can submit SQL queries to crash the server and possibly leak adjacent memory. This issue is fixed in version 3.4.1.14.	5.4	More Details
CVE-2026-20298	In Splunk Enterprise versions below 10.4.1, 10.2.5, 10.0.8, and 9.4.13, and Splunk Cloud Platform versions below 10.5.2605.0, 10.4.2604.6, 10.3.2512.15, 10.2.2510.18, and 10.1.2507.24, a low-privileged user that does not hold the 'admin' or 'power' Splunk roles could view stored credential hashes when they access the `/servicesNS/-/-/storage/passwords` REST endpoint through the `|rest` Search Processing Language (SPL) command. The exposure happens because the `|rest` SPL command returns the `encr_password` field in the results of the `/servicesNS/-/-/storage/passwords` REST endpoint.	5.3	More Details
CVE-2026-54685	FileBrowser Quantum is a free, self-hosted, web-based file manager. Prior to version 1.3.2-beta, the `/api/auth/login` authentication endpoint does not execute in constant time. When a non-existent username is supplied, the server returns a `401`/`403` response almost immediately. When a valid username is provided, the server performs a bcrypt password comparison, causing a measurable delay in the response time. Version 1.3.2-beta patches the issue.	5.3	More Details
CVE-2026-46917	Vulnerability in the Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: JSSE). Supported versions that are affected are Oracle Java SE: 11.0.31, 17.0.19, 21.0.11, 25.0.3, 26.0.1; Oracle GraalVM for JDK: 17.0.19 and 21.0.11; Oracle GraalVM Enterprise Edition: 21.3.18. Easily exploitable vulnerability allows unauthenticated attacker with network access via TLS to compromise Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Successful attacks of this vulnerability can result in unauthorized ability to cause a partial denial of service (partial DOS) of Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Note: This vulnerability can only be exploited by supplying data to APIs in the specified Component without using Untrusted Java Web Start applications or Untrusted Java applets, such as through a web service. CVSS 3.1 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L).	5.3	More Details
CVE-2026-56456	HCL DFXAnalytics is affected by an Internal File Path Disclosure vulnerability. The application dashboard inadvertently leaks sensitive information regarding its internal file structure and directory paths through unhandled error messages, system logs, or debugging output, which could allow a remote attacker to map the underlying server environment and identify targets for further exploitation.	5.3	More Details
CVE-2026-56455	HCL DFXAnalytics is affected by a Buffer Overflow vulnerability that can lead to a Denial of Service (DoS). The application fails to properly validate input sizes, allowing an attacker to pass an excessive amount of information into a memory container, which can cause the system to crash or become unresponsive. To mitigate this flaw, comprehensive input length checks must be implemented and enforced on both the client and server sides.	5.3	More Details
CVE-2026-15588	A denial-of-service and resource exhaustion vulnerability exists within the `GDBus` component of GLib. The `gdbusauth` authentication mechanism fails to enforce proper length limitations on data lines read from a client. An unauthenticated local or remote attacker can exploit this lack of input validation by sending excessively long streams of data, causing the application to consume massive amounts of system memory and CPU, potentially leading to a crash or system hang.	5.3	More Details
CVE-	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: Agent Next Gen). Supported versions that are affected are 13.5 and 24.1. Easily exploitable		

2026-46986	vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Enterprise Manager Base Platform accessible data. CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N).	5.3	More Details
CVE-2026-46985	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: Agent Next Gen). Supported versions that are affected are 13.5 and 24.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Enterprise Manager Base Platform accessible data. CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N).	5.3	More Details
CVE-2026-15106	The WPBot – AI ChatBot for Live Support, Lead Generation, AI Services plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 8.5.6. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for unauthenticated attackers to delete arbitrary chat session records from the wpbot_user and wpbot_conversation tables, including chat history and conversation logs, by supplying a crafted userid value.	5.3	More Details
CVE-2026-63750	SurrealDB versions before 3.1.0 fail to apply the SURREAL_WEBSOCKET_MAX_MESSAGE_SIZE limit to anonymous /sql WebSocket connections, allowing attackers to buffer unbounded frames in the per-connection read buffer. Attackers can stream WebSocket frames larger than the configured limit across multiple concurrent connections to consume excessive memory and degrade /sql availability.	5.3	More Details
CVE-2026-46984	Vulnerability in the Oracle Enterprise Manager Base Platform product of Oracle Enterprise Manager (component: Agent Next Gen). Supported versions that are affected are 13.5 and 24.1. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle Enterprise Manager Base Platform. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Enterprise Manager Base Platform accessible data. CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N).	5.3	More Details
CVE-2026-61452	The Grav API plugin (getgrav/grav-plugin-api) before 2.0.4 contains an improper session invalidation vulnerability where JWT access tokens are issued without a jti (JWT ID) claim and therefore cannot be revoked server-side. Unlike refresh tokens, access tokens remain valid for their full lifetime (default 1 hour) regardless of logout, password change, new token issuance, or account disablement. An attacker who has stolen an access token retains full API access until the token naturally expires.	5.3	More Details
CVE-2026-59845	A flaw was found in libssh. When ProxyCommand is used, an unchecked fork() failure can be stored as process ID -1; during cleanup, signals may then be sent across the caller's accessible process tree, leading to local denial of service.	5.3	More Details
CVE-2026-56144	Incorrect Authorization (CWE-863) in Elasticsearch can allow an authenticated user with limited index privileges to exploit insufficient authorization controls in the ingest simulation feature. By targeting indices they are not authorized to access directly, the user can cause those indices' configured ingest pipelines to execute and return their output, potentially disclosing data processed or enriched by those pipelines. Additionally, the same feature can be used to retrieve index mapping metadata for indices the user are not authorized to access directly.	5.3	More Details
CVE-2026-11579	The Kali Forms — Contact Form & Drag-and-Drop Builder WordPress plugin before 2.4.17 does not verify that a file upload is made against an existing form configured with a file-upload field, accepting uploads regardless of whether any such form exists, which allows unauthenticated users to upload files to the WordPress Media Library; the uploads are limited to WordPress's default-allowed MIME types, so this does not lead to code execution.	5.3	More Details
CVE-2026-64822	djangoSIGE through 1.10 (commit a6fe7e8) contains a user enumeration vulnerability in ForgotPasswordView within djangosige/apps/login/views.py that allows unauthenticated attackers to identify valid accounts by observing distinct error messages returned by the password reset endpoint. Attackers can submit arbitrary usernames or email addresses to the POST login/esqueceu/ endpoint and distinguish between existing and non-existing accounts based on observable discrepancies in the application's responses.	5.3	More Details
CVE-2026-38974	Dulwich through 1.1.0 was found to be missing SSH host key verification in contrib/paramiko_vendor.py.	5.3	More Details
CVE-2026-44583	Paymenter is a free and open-source webshop solution for management of hosting services. In versions prior to 1.5.0, the PayPal webhook endpoint /extensions/paypal/webhook processes the PAYPAL-CERT-URL HTTP header without validation, allowing attackers to control server-side HTTP request destinations. This value is passed directly into a server-side HTTP request via file_get_contents, allowing attackers to control the destination of the request. No allowlist, validation, or signature verification is applied to the header before usage. As a result, the application can be coerced into performing HTTP requests to attacker-controlled or internal network destinations. This vulnerability allows remote unauthenticated attackers to induce server-side HTTP GET requests to arbitrary external or internal endpoints. Depending on network configuration, this may lead to: blind SSRF to external attacker-controlled systems, and potential access to internal network	5.3	More Details

	services No direct response data is returned to the attacker (blind SSRF), but the issue may still enable sensitive network probing or data exfiltration via side channels. This issue has been fixed in version 1.5.0.		
CVE-2026-59848	A flaw was found in libssh. A malicious SFTP server can send responses for unknown request IDs that libssh clients keep queued indefinitely, causing unbounded memory growth and client-side denial of service.	5.3	More Details
CVE-2026-33684	WWBN AVideo is an open source video platform. Prior to version 29.0, Privilege Escalation is possible through unguarded permission parameters in signUp API, which allows any user who can solve a CAPTCHA to self-grant elevated permissions during account registration. The set_api_signUp method in the API plugin accepts emailVerified, canUpload, canStream, and canCreateMeet parameters from user-supplied input and applies them to newly created accounts without verifying that the request was authenticated with a valid APISecret. By self-granting account attributes, attackers can mark their own accounts as email-verified without owning the address (bypassing email-gated functionality) and award themselves upload, streaming, and meeting-creation permissions, circumventing administrator access controls that intentionally restrict these capabilities for new users. This issue has been fixed in version 29.0	5.3	More Details
CVE-2026-55238	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a vulnerability concerning the processing of RDP Confirm Active PDU, where during the capability negotiation phase, the parser did not perform sufficient length validation for specific capability sets. A remote, unauthenticated attacker could potentially exploit this flaw by sending a specially crafted RDP packet containing malformed capability data. Due to missing bounds checks, the xrdp process may perform out-of-bounds memory reads, which can result in the termination of the service (Denial of Service). However, since xrdp forks a new process for each connection by default, an out-of-bounds read causing a process crash is unlikely to bring down the entire xrdp service. This issue has been fixed in version 0.10.6.1.	5.3	More Details
CVE-2026-44978	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a heap out-of-bounds read vulnerability within the FIPS-specific receive paths. This vulnerability does not affect the default configuration of xrdp. The vulnerability is only exploitable when the security layer is set to security_layer=negotiate or security_layer=rdp, and the crypto level is changed to crypt_level=fips in xrdp.ini. In this specific non-default mode, the server fails to validate the FIPS padding length field, leading to a pointer underflow and a subsequent negative length calculation. An unauthenticated remote attacker can exploit this by sending a crafted FIPS-protected PDU, causing a heap out-of-bounds read that results in a process crash and denial of service (DoS). However, since xrdp forks a new process for each connection by default, an out-of-bounds read causing a process crash is unlikely to bring down the entire xrdp service. This issue has been fixed in version 0.10.6.1.	5.3	More Details
CVE-2026-65055	Taiga 6.10.1 contains a missing authorization vulnerability that allows unauthenticated attackers to disclose the full member roster and internal workflow configuration of any private project by supplying a project ID to the filters_data API endpoints on UserStory, Task, Issue, and Epic viewsets. Attackers can send unauthenticated GET requests to the filters_data endpoints with sequential integer project IDs to enumerate private project membership details including user IDs, full names, and gravatar hashes, bypassing the access controls that correctly restrict other project API endpoints.	5.3	More Details
CVE-2026-42218	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a timing side-channel vulnerability in the login interface. Due to a discrepancy in response processing times, a remote attacker can infer the existence of a username on the system, leading to unauthorized information disclosure via username enumeration. This issue has been fixed in version 0.10.6.1.	5.3	More Details
CVE-2026-65058	Trezor Safe 3, Safe 5, and Safe 7 firmware contains a confirmation-binding flaw in the Ethereum sign_tx / sign_tx_eip1559 flow. For contract interactions, the device confirms only the initial calldata chunk while the signature commits to the full streamed calldata. An attacker could present calldata to a victim then supply a different tail that changes the signed transaction. Fixed in 70c9b0c.	5.3	More Details
CVE-2026-16318	The QUIC transport parameters extension handler in s2n-tls incorrectly uses s2n_alloc instead of s2n_realloc to store the peer's transport parameters. When a TLS 1.3 connection goes through a HelloRetryRequest, the handler is called twice on the same connection. On the second call, s2n_alloc zeroes the existing pointer before allocating new memory, causing the first allocation to be leaked. This can occur during normal QUIC traffic when a client offers a key share group the server does not prefer. An unauthenticated user can amplify the issue by deliberately forcing HelloRetryRequests, causing up to approximately 64 KB of unreachable memory per handshake. Over time, this can lead to increased memory consumption on long-running server processes. The unreachable memory is only reclaimed when the process is restarted. Only server-side QUIC-enabled deployments are affected. Non-QUIC TLS connections are not affected. We recommend you upgrade s2n-tls to version v1.7.6	5.3	More Details
CVE-2026-53596	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. Prior to version 1.8.224, the FreeScout helpdesk application does not enforce rate limiting on the file upload endpoint. Any user can flood the server with upload requests, leading to database overload and potential denial of service for all users. Version 1.8.224 contains a fix.	5.3	More Details
	Paymenter is a free and open-source webshop solution for management of hosting services. In versions prior to 1.5.5, the credit payment implementation in app/Livewire/Invoices/Show.php executes a pessimistic row		

CVE-2026-55219	lock (lockForUpdate()) outside of an active database transaction. Because MySQL/MariaDB requires an enclosing transaction to enforce row-level locks, the guard is ineffective. Concurrent payment requests can exploit this race condition to read the same credit balance simultaneously, allowing users to pay multiple invoices using the same credit balance. In database systems like MySQL, a row lock only works inside a formal transaction; without one, the lock is completely ignored. Because there is no active lock, two payment requests sent at the exact same millisecond can look at the database at the same time. Both requests see the original credit balance, decide it is sufficient, and approve the payment. Because the payment processes successfully through ExtensionHelper::addPayment(), the application provisions the corresponding services or digital goods, resulting in direct financial or resource loss to the platform. This issue has been fixed in version 1.5.5.	5.3	More Details
CVE-2026-11868	The WP Travel WordPress plugin before 11.7.1 does not perform capability or ownership checks on its booking cancellation action, which is also exposed to unauthenticated users, allowing them to cancel arbitrary bookings on the site.	5.3	More Details
CVE-2026-12723	The Kirki WordPress plugin before 6.0.12 does not perform any authorisation check on one of its REST routes, allowing unauthenticated users to overwrite the content of arbitrary existing comments and to create pre-approved comments under a spoofed identity, bypassing comment moderation.	5.3	More Details
CVE-2026-48824	Mailpit is an email testing tool and API for developers. Prior to version 1.30.1, the fix for GHSA-fpxj-m5q8-fphw (CVE-2026-45710, "Mailpit: Set a default 50MB p/m limit to prevent DoS via unlimited SMTP DATA and /api/v1/send body sizes") wrapped only `POST /api/v1/send` with `http.MaxBytesReader`. The four other Mailpit JSON-body API endpoints `PUT /api/v1/messages` (SetReadStatus), `DELETE /api/v1/messages` (DeleteMessages), `PUT /api/v1/tags` (SetMessageTags), and `POST /api/v1/message/{id}/release` (ReleaseMessage) still call `json.NewDecoder(r.Body)` directly with no body-size cap and remain reachable unauthenticated in the default `docker run axllent/mailpit:latest` deploy. An unauthenticated remote attacker can post a multi-million-element `IDs` slice and drive RSS from ~25 MiB baseline to ~450 MiB per 16 MB request body. Repeating across multiple connections accumulates the same per-request amplification per process. Version 1.30.1 contains a patch.	5.3	More Details
CVE-2026-12972	The PayPlus Payment Gateway WordPress plugin before 8.2.2 does not perform authorization or order-ownership validation in one of its AJAX actions available to unauthenticated users, allowing them to tamper with the payment-related metadata of arbitrary WooCommerce orders.	5.3	More Details
CVE-2026-55639	xrdp is an open source RDP server. Versions 0.10.6 and prior contain a vulnerability concerning the parsing of Client Security Data within the Client MCS Connect Initial PDU with GCC Conference Create Request during the connection sequence. During the initial capability and security negotiation phase, the parser fails to perform sufficient length validation for the incoming data block. A remote, unauthenticated attacker could potentially exploit this flaw by sending a specially crafted RDP packet containing malformed data. Due to missing bounds checks, the xrdp process may read a small number of bytes beyond the declared data block boundary, potentially disclosing process memory contents that could be combined with other vulnerabilities. This issue has been fixed in version 0.10.6.1.	5.3	More Details
CVE-2024-23568	HCL Aftermarket EPC is vulnerable to attacks since the server software version used by the application is revealed by the web server. Displaying version information of software could allow an attacker to determine which vulnerabilities are present in the software, particularly if an outdated software version is in use with published vulnerabilities.	5.3	More Details
CVE-2026-13082	GD::SecurityImage versions through 1.75 for Perl use rand to generate secrets. The random method creates the challenge text used for the CAPTCHA by sampling characters from an array using Perl's built-in rand function, and generates a (by default) six-character string. The built-in rand function is unsuitable for security applications because it is predictable and reversible.	5.3	More Details
CVE-2026-58149	The Joomla extension Events Booking is vulnerable to an unauthenticated user enumeration that allows to retrieve account usernames and email addresses.	5.3	More Details
CVE-2026-16082	A vulnerability was identified in Sipeed PicoClaw up to 0.2.9. The impacted element is the function ExecTool.executeRun of the file pkg/agent/pipeline_execute.go. The manipulation of the argument cwe leads to time-of-check time-of-use. The attack must be carried out locally. The exploit is publicly available and might be used. The reported GitHub issue was closed automatically with the label "not planned" by a bot.	5.3	More Details
CVE-2026-48504	OpenTelemetry Rust is the Rust OpenTelemetry implementation. In 0.32.0 and earlier, BaggagePropagator::extract_with_context in opentelemetry_sdk did not enforce W3C Baggage size limits before parsing an inbound baggage header, so a large attacker-controlled header could cause unnecessary CPU work and short-lived heap allocations while parsing entries later discarded by the SDK's baggage storage limits. Services that accept untrusted inbound propagation headers may experience increased per-request resource usage when processing oversized baggage headers. This issue is fixed in version 0.32.1.	5.3	More Details
CVE-2026-	A security flaw has been discovered in Sipeed PicoClaw up to 0.2.9. This affects the function webhook.ParseRequest of the file pkg/channels/line/line.go of the component LINE Webhook. The manipulation results in authentication bypass by capture-replay. The attack may be launched remotely. The	5.3	More

16083	exploit has been released to the public and may be used for attacks. The reported GitHub issue was closed automatically with the label "not planned" by a bot.		Details
CVE-2026-16013	A vulnerability has been found in liftoff-sr CIPster up to 632336d414ef708a542377c1aa8d6fdb7c70a760. Affected by this issue is the function CipAppPath::deserialize_symbolic of the file source/src/cip/cipepath.cc. Such manipulation leads to out-of-bounds read. The attack may be launched remotely. The exploit has been disclosed to the public and may be used. This product operates on a rolling release basis, ensuring continuous delivery. Consequently, there are no version details for either affected or updated releases. The name of the patch is 886a4d090e1c5b0475f0b1c2fe0606a8f0d6a519. A patch should be applied to remediate this issue.	5.3	More Details
CVE-2026-63098	TheHive through 4.1.24 contains an unauthenticated information disclosure vulnerability that allows unauthenticated attackers to retrieve sensitive configuration data by sending a GET request to the /api/status endpoint, which lacks authentication enforcement in the StatusCtrl.scala handler. Attackers can obtain the datastore attachment protection password, configured authentication providers, SSO settings, MFA capabilities, and clustered node addresses and roles without any credentials.	5.3	More Details
CVE-2026-16077	A vulnerability was found in AstrBotDevs AstrBot up to 4.25.5. Impacted is the function _normalize_rw_path of the file astrbot/core/tools/computer_tools/fs.py of the component Filesystem Computer-Use Tool. Performing a manipulation results in link following. The attack is only possible with local access. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2026-49208	Symfony UX is a JavaScript ecosystem for Symfony. From 2.8.0 until 2.36.0 and 3.1.0, when a #[LiveProp] is typed as DateTimeInterface and no explicit format is configured, Symfony\UX\LiveComponent\LiveComponentHydrator::hydrateObjectValue() falls back to new $className($value), allowing client-supplied relative strings such as now, tomorrow, or +10 years to move a writable, format-less date prop past time-based business logic checks. This issue is fixed in versions 2.36.0 and 3.1.0.	5.3	More Details
CVE-2026-9537	Mojo::JWT versions before 1.02 for Perl verify HMAC signatures with a non-constant-time string comparison. The decode() method compares the supplied signature to the recomputed HMAC with Perl's eq operator, which stops at the first differing byte, so the comparison time varies with the number of matching leading bytes. A caller that decodes attacker supplied tokens leaks the expected signature through this timing variation, which can be aggregated over many requests to recover the signature and forge a token.	5.3	More Details
CVE-2024-23565	HCL Aftermarket EPC is vulnerable to email flooding as the application does not have a proper mail limitation mechanism at Forget Password functionality. The actor could b e a human or an automated process such as a virus or bot. This could be used to cause a denial of service, compromise program logic or other consequences.	5.3	More Details
CVE-2024-23575	HCL Aftermarket EPC is vulnerable to attack since the application returns detailed error messages that leak information about the processing on the server. An attacker may use the contents of error messages to help launch another ,more focused attack.	5.3	More Details
CVE-2026-13402	The Royal Addons for Elementor WordPress plugin before 1.7.1063 does not check the post status of menu items or the templates they reference in one of its REST endpoints, allowing unauthenticated users to retrieve the rendered HTML content of private or draft Elementor templates linked from non-public navigation menu items.	5.3	More Details
CVE-2026-44434	Quicly is an IETF QUIC protocol implementation intended primarily for use within the H2O HTTP server. Prior to commit dccf5d4, Quicly was vulnerable to stateless reset injection through lack of packet entry validation. The QUIC protocol is designed to withstand packet injection attacks, once the handshake is complete. Only packets that carry some secret patterns are considered as stateless resets. Quicly allows the peer to share up to 4 such patterns per connection. However, until now, it failed to determine which of the 4 slots that it uses to retain the secret patterns contains a valid entry. As the slots are zero-initialized, the failure meant that, unless the peer advertised 4 of such patterns, an all-zero pattern was treated as a stateless reset.In effect, this allowed an on-path attacker to reset QUIC connections governed by Quicly. This issue has been fixed by commit dccf5d4.	5.3	More Details
CVE-2026-16085	A security vulnerability has been detected in Sipeed PicoClaw up to 0.2.9. Affected is the function NewContextBuilder of the file pkg/agent/context.go. Such manipulation leads to inclusion of functionality from untrusted control sphere. The attack needs to be performed locally. The exploit has been disclosed publicly and may be used. The reported GitHub issue was closed automatically with the label "not planned" by a bot.	5.3	More Details
CVE-2026-14979	IBM Engineering Lifecycle Management 7.0.3 (Interim Fix 001 through) Interim Fix 021, 7.1.0 (Interim Fix 001 through) Interim Fix 009, and 7.2.0 and 7.2.0 Interim Fix 001 DOORS could allow a remote attacker to cause a denial of service due to improper handling of XML entity expansion.	5.3	More Details
CVE-	The Janssen Project is an open-source identity and access management (IAM) platform. Prior to 2.0.0, jans-auth-server accepts unsigned JWE request objects because JwtAuthorizationRequest skips inner signature		

2026-45795	validation when jwe.getSignedJWTPayload() returns null, and AuthzRequestService.processRequestObject() does not reject the unrecognized RSA-OAEP algorithm when forceSignedRequestObject=true. This issue is fixed in version 2.0.0.	5.3	More Details
CVE-2026-48049	@hapi/inert provides static file and directory handlers for hapi.js. From 4.0.0 to 7.1.0, @hapi/inert serves static files from a directory configured with path in the directory or file handlers or relativeTo for h.file(), with confinement enforced by the confine option, but the confinement check compared the resolved absolute path against the confine directory using a raw string-prefix test, so a sibling directory such as /app/static-secret next to /app/static was incorrectly accepted and could allow an unauthenticated remote attacker to read files via /..%2fstatic-secret/secret.txt. This issue is fixed in version 7.1.1.	5.3	More Details
CVE-2024-23574	HCL Aftermarket EPC is vulnerable to attack since It was found that a malicious actor can use brute-force techniques to either guess or confirm valid users in the system. Use renumeration is when a malicious actor can use brute-force techniques to either guess or confirm valid users in a system	5.3	More Details
CVE-2026-62220	OpenClaw 2026.2.25 before 2026.5.26 allow a lower-trust caller or configured input path to bypass non-browser rate limits on WebSocket authentication attempts. When the affected feature is enabled and reachable by lower-trust input, this can consume gateway resources and reduce service availability.	5.3	More Details
CVE-2026-44433	Quicly is an IETF QUIC protocol implementation intended primarily for use within the H2O HTTP server. Prior to commit 8b178e6, an adversarial peer could send a STREAM frame carrying just one byte at the largest offset being permitted to obtain additional flow control credit, which under certain circumstances could lead to a Denial of Service. Assuming the application prepares a receive buffer for storing all data that arrive out-of-order, up to the largest offset being received, this behavior could lead to the application allocating large amount of memory with the peer sending only a handful of packets, resulting in memory exhaustion. In addition to the receive buffer allocation strategy, the severity of this vulnerability depends on how the application controls the stream concurrency. In case of the H2O HTTP server, under its default setting, this bug increases the maximum amount of memory allocated per connection by about 4 times. This issue has been fixed by commit 8b178e6.	5.3	More Details
CVE-2026-11966	The User Registration & Membership WordPress plugin before 5.2.3 does not perform a capability check for unauthenticated callers on one of its membership payment actions and acts on a caller-supplied user identifier, allowing unauthenticated attackers to delete recently-registered, payment-pending user accounts.	5.3	More Details
CVE-2026-62299	CoreDNS is a DNS server written in Go. Prior to 1.14.5, the CoreDNS rewrite plugin supports edns0 rewrite rules with an optional revert flag, and two response rules, edns0SetResponseRule and edns0ReplaceResponseRule[T] in plugin/rewrite/edns0.go, call res.IsEdns0() and immediately dereference the returned *dns.OPT without a nil check when a downstream plugin returns a response with no OPT record. A remote, unauthenticated client can send a single ordinary DNS query matching a rewrite edns0 <local|nsid|subnet> <set|append|replace> ... revert rule, causing ResponseReverter in plugin/rewrite/reverter.go to panic, return SERVFAIL, and degrade availability, or crash the CoreDNS process if the debug directive disables recovery. This issue is fixed in version 1.14.5.	5.3	More Details
CVE-2026-8616	The Fense Proxy & VPN Blocker plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check and missing nonce validation on the fense_bpvt_save_settings() function in versions up to, and including, 3.0.1. The callback is registered to both wp_ajax_* and wp_ajax_nopriv_* hooks and unconditionally calls delete_option() on four plugin options and delete_transient() on three transients tied to the plugin's API key cache and settings. This makes it possible for unauthenticated attackers to delete plugin options and transients, effectively resetting the plugin's API key/data cache and forcing the plugin to refetch state.	5.3	More Details
CVE-2026-16129	A vulnerability has been found in princezuda SafestClaw up to 4.2.4. This vulnerability affects the function ShellAction._validate_command of the file src/safestclaw/actions/shell.py of the component Built-in Web Interface. Such manipulation leads to incomplete blacklist. An attack has to be approached locally. The exploit has been disclosed to the public and may be used. The presence of this vulnerability remains uncertain at this time. The project maintainer explains: "On paper you're correct, this is a vulnerability. In practice, nothing your AI generated shows how it makes users vulnerable. It's open source. Someone can mod the shell allow list or remove that system. Present an actual poc that shows a threat to users."	5.3	More Details
CVE-2026-16201	A vulnerability was found in zevorn rt-claw up to 0.2.0. Affected is the function claw_net_get/claw_net_post of the file claw/services/tools/net.c of the component http_request. The manipulation results in information disclosure. The attack can be executed remotely. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-8861	IBM Security Verify could allow a remote attacker to obtain sensitive information when a detailed technical error message is returned in the browser. This information could be used in further attacks against the system.	5.3	More Details
CVE-2024-42214	HCL Aftermarket EPC is vulnerable to attack since HTTP OPTIONS method is enabled on this web server. The OPTIONS method provides a list of the methods that are supported by the Web server which allows an attacker to narrow and intensify their efforts.	5.3	More Details

CVE-2026-58556	Permission control vulnerability in the Bluetooth module. Impact: Successful exploitation of this vulnerability may affect availability.	5.1	More Details
CVE-2026-38754	A heap overflow in the ifsbreakup() function (shell/ash.c) of Busybox v1.38.0 allows attackers to cause a Denial of Service (DoS) via supplying a crafted input.	5.1	More Details
CVE-2026-58552	Out-of-bounds read vulnerability in the image codec module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	5.1	More Details
CVE-2026-47143	Capstone is a disassembly framework. Versions prior to 6.0.0-Alpha8 and 5.0.8 have a NULL pointer dereference in `modRMRequired()` and `decode()` when disassembling 3DNow! opcodes (`0F 0F`) in builds compiled with `-DCAPSTONE_X86_REDUCE`, allowing a remote attacker to crash any application using the reduced X86 Capstone library by supplying a crafted input containing the 4-byte sequence `0F 0F <modrm> <imm8>`. Versions 6.0.0-Alpha8 and 5.0.8 patch the issue.	5.1	More Details
CVE-2026-58551	Out-of-bounds read vulnerability in the image codec module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	5.1	More Details
CVE-2026-63730	HyperDX before 2.31.0 contains a server-side request forgery vulnerability that allows authenticated team members to direct the server to make requests to arbitrary internal network destinations by supplying a caller-controlled URL to the webhook test endpoint. Attackers can bypass the insufficient hostname blacklist validation in the webhook handler to enumerate internal services, interact with internal containers, or access cloud instance metadata services including provider metadata endpoints.	5.0	More Details
CVE-2026-16133	A flaw has been found in LiuMengxuan04 MiniCode 0.1.0. Affected by this vulnerability is the function child_process.spawn of the file mcp.ts. Executing a manipulation can lead to command injection. The attack can be launched remotely. The attack requires a high level of complexity. The exploitation appears to be difficult. The exploit has been published and may be used. The pull request to fix this issue awaits acceptance.	5.0	More Details
CVE-2026-16208	A flaw has been found in django-tastypie up to 0.15.1. The affected element is the function CacheThrottle/CacheDBThrottle of the file tastypie/throttle.py. This manipulation causes race condition. The attack may be initiated remotely. The complexity of an attack is rather high. The exploitability is described as difficult. The project was informed of the problem early through an issue report but has not responded yet.	5.0	More Details
CVE-2026-62216	OpenClaw 2026.4.20 before 2026.5.28 contain a policy bypass in the QQBot media upload feature. A lower-trust caller or configured input path could cause the media upload to reach network destinations that should have been blocked by OpenClaw policy (server-side request forgery). The practical impact depends on the operator's configuration and whether lower-trust input can reach that path.	5.0	More Details
CVE-2026-62211	OpenClaw versions before 2026.6.1 contain a credential redaction bypass vulnerability in the trajectory export feature that allows lower-trust callers to access data that should remain within trusted boundaries. Attackers can exploit misconfigured input paths or feature accessibility to expose sensitive credentials and data through the export mechanism.	5.0	More Details
CVE-2026-12391	An insecure symlink following vulnerability exists in Canonical ubuntu-pro-client (formerly ubuntu-advantage-tools) within the pro collect-logs command framework. The utility creates or utilizes predictable temporary file paths or user-accessible log directories when gathering diagnostic information without verifying the file type or ownership. An unprivileged local attacker can exploit this behavior by creating a symbolic link (symlink) at a predictable destination path pointing to an arbitrary, root-readable file (such as /etc/shadow or private files within /root). When a root administrator or operator subsequently executes the pro collect-logs command, the tool follows the user-controlled symlink, reads the target file, and compresses its contents into the resulting diagnostic support archive. Because the output archive remains readable by the unprivileged user, the attacker can extract and read the sensitive root-owned files, leading to a complete information disclosure of system secrets.	5.0	More Details
CVE-2026-63734	SurrealDB versions before 3.2.0 contain a denial of service vulnerability in the SurrealML header parser that allows authenticated Owner-role users to crash the server by uploading a malformed .surml file to the /ml/import endpoint. Attackers can supply non-numeric input-dimensions or other malformed header fields that trigger unchecked unwrap calls, causing a panic that aborts the entire server process and denies service to all databases.	4.9	More Details
CVE-2026-15651	The WP TripAdvisor Review Slider plugin for WordPress is vulnerable to generic SQL Injection via the 'filtersource' parameter in all versions up to, and including, 14.6 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
	The WP Bulk Delete plugin for WordPress is vulnerable to generic SQL Injection via the 'delete_user_roles'		

CVE-2026-15727	parameter in all versions up to, and including, 1.4.2 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. wp_unslash() is applied to the raw POST body before parse_str() decomposes it, stripping WordPress magic-quotes protection and leaving attacker-controlled values fully unescaped prior to reaching the SQL sink.	4.9	More Details
CVE-2026-15445	The SEO Booster plugin for WordPress is vulnerable to time-based SQL Injection via the 'orderby' parameter in all versions up to, and including, 7.3.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. Although esc_sql() and sanitize_text_field() are applied, neither neutralizes SQL keywords, commas, parentheses, or subquery syntax in an unquoted ORDER BY context, leaving the clause fully attacker-controlled.	4.9	More Details
CVE-2026-15458	The SEO Booster plugin for WordPress is vulnerable to generic SQL Injection via the 'sort_field' parameter in all versions up to, and including, 7.3.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2024-58358	SurrealDB versions before 2.1.0 contain a denial of service vulnerability in role conversion that allows privileged owner users to define users with nonexistent roles. Attackers can trigger an uncaught panic by signing in with a user assigned an invalid role, crashing the server.	4.9	More Details
CVE-2026-15457	The Kirki – Freeform Page Builder, Website Builder & Customizer plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 6.0.13 via the 'family' parameter. This makes it possible for authenticated attackers, with editor-level access and above, to delete arbitrary directories on the server, which can result in loss of data and availability.	4.9	More Details
CVE-2026-15782	The WPForms – AI Form Builder for WordPress – Contact Forms, Payment Forms, Survey Form, Quiz & More plugin for WordPress is vulnerable to Stored Cross-Site Scripting via OptinMonster Integration data-sitekey Attribute in Post Content in all versions up to, and including, 2.0.0.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Exploitation requires the OptinMonster plugin to be installed and configured with an active inline campaign that outputs matching #om-{id} markup on the target page, as the WPForms handler only fires when OptinMonster emits its 'om.Campaign.load' event.	4.9	More Details
CVE-2026-16106	A flaw was found in the admin REST API of Keycloak, a solution for identity and access management. The issue occurs when a delegated administrator attempts to remove a child role from a composite role. Due to missing authorization checks, an attacker with limited administrative permissions can remove privileged roles they are not authorized to manage, leading to a loss of access for other users and administrators.	4.9	More Details
CVE-2026-14782	The Booking for Appointments and Events Calendar – Amelia plugin for WordPress is vulnerable to SQL Injection via the Customer Import in all versions up to, and including, 2.4.3 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with wpamelia-manager role, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-53594	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. FreeScout's `Manage -> Logs -> App Logs` feature uses the bundled `rap2hpoutre/laravel-log-viewer` override to decrypt a user-supplied file identifier and then pass the resolved path to Laravel's download response. Prior to version 1.8.224, the path resolution logic accepts any existing absolute path before applying the intended `storage/logs` restriction. As a result, an attacker who can access the App Logs route and forge a valid Laravel-encrypted `dl` parameter can download arbitrary server-local files readable by the PHP process, not just log files. Version 1.8.224 contains a fix.	4.9	More Details
CVE-2026-62947	OpenWrt is a Linux operating system targeting embedded devices. Prior to 25.12.5, the cgi-download handler in cgi-io authorizes the requested path against the caller's ubus session file ACL before canonicalization, and rpcd session.c uses fnmatch() without FNM_PATHNAME, allowing traversal such as an allowed wildcard prefix followed by ../ to read root-readable files including /etc/shadow. This vulnerability is fixed in 25.12.5.	4.9	More Details
CVE-2026-48015	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, SVG files are in the allowed_extensions whitelist in src/Core/Framework/Resources/config/packages/shopware.yaml and can be uploaded via the media manager without SVG content sanitization in the upload pipeline from MediaUploadController to FileSaver to TypeDetector, allowing malicious SVG JavaScript such as onload, <script>, and <foreignObject> to execute in the Shopware domain when the uploaded SVG is viewed. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	4.9	More Details
	Statamic is a Laravel and Git powered content management system (CMS). Prior to 5.73.24 and 6.20.1, the Glide image proxy's URL validation in src/Imaging/RemoteUrlValidator.php and		

CVE-2026-54242	src/Imaging/GuzzleAdapter.php could be bypassed using DNS rebinding. The remote hostname was validated as publicly routable, but resolved again when the image was actually fetched, so an attacker controlling the hostname's DNS could rebind it to an internal address after validation and cause the server to make HTTP requests to internal addresses, including loopback, private network, and cloud metadata endpoints. This affects sites that pass user-supplied URLs to Glide. This issue is fixed in versions 5.73.24 and 6.20.1.	4.9	More Details
CVE-2026-38753	A use-after-free in the awk_sub() function (editors/awk.c) of Busybox v1.38.0 allows attackers to cause a Denial of Service (DoS) via supplying a crafted AWK script.	4.9	More Details
CVE-2026-16072	A flaw was found in the organization management component of Keycloak. A delegated administrator with permission to manage organizations can create an invitation for a non-existent email address and then retrieve the secret registration link directly through the application programming interface. By using this link, the administrator can create new user accounts and add them to the organization without having the required user management permissions or access to the invited email account. This allows an administrator to bypass security boundaries and add unauthorized members to an organization.	4.9	More Details
CVE-2026-8825	The Elementor Website Builder WordPress plugin before 4.1.4 does not properly check user permissions before returning post data through one of its REST endpoints, allowing authenticated users with Contributor-level access and above to retrieve the title, body and metadata of private posts, private pages and drafts authored by other users (including administrators).	4.9	More Details
CVE-2026-26081	HAProxy Community Edition 3.0 through 3.3 before 3.3.3 lacks a length check for the NEW_TOKEN format. HAProxy Enterprise and ALOHA are also affected.	4.8	More Details
CVE-2026-58557	Design defect vulnerability in Expedition mode. Impact: Successful exploitation of this vulnerability may affect availability.	4.8	More Details
CVE-2026-56353	n8n contains an authentication bypass in the Chat Trigger node when configured with n8n User Auth (a non-default configuration). In affected releases — before 1.123.22, the 2.0.0 through 2.9.2 line, and 2.10.0 — the authentication check on the Chat Trigger webhook endpoint can be circumvented, allowing access without valid credentials. Fixed in 1.123.22, 2.9.3, and 2.10.1.	4.8	More Details
CVE-2026-15812	A vulnerability was found in the internal Access Control List (ACL) subsystem of kronosnet (Version affected: <= 1.34). When the framework is explicitly configured to manage dynamic links (accepting network traffic from any IP address) without network payload encryption, the validation architecture implicitly trusts the link ID provided within incoming data packets. A remote, unauthenticated attacker can exploit this lack of validation by spoofing a legitimate link ID inside crafted network frames. This allows the attacker to fully bypass the ACL framework and inject arbitrary data packets into the application layer, potentially leading to data corruption or service instabilities.	4.8	More Details
CVE-2026-48819	Hey API is an ecosystem for turning API specifications into production-ready code. Prior to 0.97.3, dist/clients/core/params.ts ships a runtime template copied into generated SDKs as params.gen.ts, and buildClientParams writes unknown slot-prefixed keys such as $body_, $headers_, $path_, and $query_ directly to the corresponding slot, allowing $query___proto__ alongside a legitimate q field to set params.query through params["query"]["__proto__"] = value, call Object.setPrototypeOf(params.query, value), and expose inherited attacker-controlled keys during for..in iteration. This issue is fixed in version 0.97.3.	4.8	More Details
CVE-2026-1562	Pega Platform versions 8.1.0 through 25.1.2 are affected by an Stored Cross-site scripting (XSS) vulnerability in a user interface component. Requires a high privileged user with a developer role.	4.8	More Details
CVE-2026-10724	The Reviews Feed WordPress plugin before 2.6.5 does not neutralize WordPress shortcodes contained in third-party review content before rendering it through its dynamic block, allowing unauthenticated attackers to execute arbitrary shortcodes on pages that display the feed by planting a shortcode in a review on the connected source.	4.8	More Details
CVE-2026-1563	Pega Platform versions 8.1.0 through 25.1.2 are affected by an Reflected Cross-site scripting (XSS) vulnerability in a user interface component. Requires a high privileged user with a developer role.	4.8	More Details
CVE-2026-55254	NCalc is a fast, lightweight expression evaluator for .NET. Prior to 6.1.1, the factorial operator implementation in src/NCalc.Core/Helpers/MathHelper.cs permits specially crafted expressions with extremely large factorial operands, causing excessive CPU consumption or a non-terminating loop due to integer overflow in the factorial calculation logic when applications evaluate untrusted expressions. This issue is fixed in version 6.1.1.	4.8	More Details
	WWBN AVideo is an open source video platform. Versions 29.0 and below contain a stored Cross-Site Scripting vulnerability in the YouTubeAPI plugin. The plugin renders the snippet.title field returned by the		

CVE-2026-50183	YouTube Data API into the homepage gallery markup with no HTML encoding. The title is set by the YouTube video uploader (anyone in the world) and is treated by AVideo as trusted content. A YouTube uploader who controls a video matching the operator's configured query injects HTML into the AVideo homepage by setting their video's title to a JavaScript-bearing string; the payload then executes in the browser of every visitor who loads any page that renders the gallery. When the visitor is an AVideo administrator, the injected JavaScript performs any admin action (create user, promote to admin, change configuration, install plugin) that uses cookie-based authentication without an additional CSRF token, escalating the bug into full administrative takeover. The payload persists for the duration of cacheTimeout (default 3600 seconds) after the malicious title is set on YouTube and survives YouTube removing the hostile video for the same window. This issue has been addressed by commit https://github.com/WWBN/AVideo/commit/7292129eaee5f609beae103b5cb387d55f17b877.	4.7	More Details
CVE-2026-16088	A vulnerability was detected in halo-dev halo up to 2.24.2. Affected by this vulnerability is the function Download of the file MigrationEndpoint.java of the component Files Backup Endpoint. Performing a manipulation results in path traversal. The attack is possible to be carried out remotely. The exploit is now public and may be used.	4.7	More Details
CVE-2026-64823	Home Assistant Core before 2026.5.4 contains a cross-site scripting vulnerability in the Shelly integration's async_get_media_image() method that allows attackers controlling a Shelly device's thumb field to serve arbitrary HTML content by supplying a data URI with a text/html content type without validation against an image-only allowlist. Attackers can cause the media player proxy endpoint to serve attacker-controlled bytes with Content-Type text/html in the Home Assistant web origin, enabling theft of session tokens from local storage and authenticated calls to sensitive service endpoints including lock, alarm, and cover controls.	4.7	More Details
CVE-2026-54163	secure_headers manages application of security headers with many safe defaults. Prior to 7.3.0, secure_headers builds the Content-Security-Policy value by stitching directives with ; separators, and build_sandbox_list_directive, build_media_type_list_directive, and build_report_to_directive interpolate caller-supplied strings without scrubbing ;, \r, or \n. When untrusted input reaches SecureHeaders.override_content_security_policy_directives or append APIs for :sandbox, :plugin_types, or :report_to, an attacker can inject a CSP directive such as script-src 'unsafe-inline' * before the legitimate script-src, enabling XSS reachability through these sinks or CSP report exfiltration. This issue is fixed in version 7.3.0.	4.7	More Details
CVE-2026-40106	Wazuh is a free and open source platform used for threat prevention, detection, and response. Versions 4.6.0 and above prior to 4.14.5 contain a heap-based buffer overflow vulnerability in the syscheck component of the Wazuh agent for Windows. When expanding registry paths containing wildcards (* or ?), the agent allocates a fixed-size heap buffer of 256 bytes (OS_SIZE_256). By creating a registry subkey with a maximum allowed length (255 characters) inside a monitored path, a low-privileged local attacker can force an out-of-bounds write during string concatenation. Since wazuh-agent.exe runs as NT AUTHORITY\SYSTEM, this can lead to a silent Denial of Service (blinding the agent) or potentially Local Privilege Escalation (LPE). This issue has been fixed in version 4.14.5.	4.7	More Details
CVE-2026-16226	A weakness has been identified in SourceCodester Pizzafy Ecommerce System 1.0. This affects the function save_settings of the file /admin/admin_class_novo.php. This manipulation of the argument img causes unrestricted upload. The attack is possible to be carried out remotely.	4.7	More Details
CVE-2026-46948	Vulnerability in the Oracle Utilities Network Management System product of Oracle Utilities Applications (component: Security). Supported versions that are affected are 2.4.0.1.0-2.4.0.1.32, 2.5.0.1.0-2.5.0.1.17, 2.5.0.2.0-2.5.0.2.11, 2.6.0.2.0-2.6.0.2.7 and 25.12.0.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Utilities Network Management System. Successful attacks require human interaction from a person other than the attacker. Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle Utilities Network Management System accessible data as well as unauthorized read access to a subset of Oracle Utilities Network Management System accessible data. CVSS 3.1 Base Score 4.6 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:R/S:U/C:L/I:L/A:N).	4.6	More Details
CVE-2026-62826	Improper neutralization of input during web page generation ('cross-site scripting') in Microsoft Office SharePoint allows an authorized attacker to perform spoofing over a network.	4.6	More Details
CVE-2026-21760	HCL DevOps Loop is affected by an Unauthorized Access to Admin Functionality (Forced Browsing) vulnerability. Improper authorization checks may allow unauthorized users to access restricted administrative functionality by directly accessing protected application endpoints.	4.6	More Details
CVE-2026-53592	FreeScout is a free help desk and shared inbox built with PHP's Laravel framework. A Prototype Pollution condition in the `getQueryParam` function `/public/js/main.js` and was addressed in version 1.8.139 by blocking URL query keys matching the pattern `__proto__`. However, this mitigation is incomplete: it only filters top-level `__proto__` keys and fails to sanitize nested forms such as `b[__proto__][polluted]=PWNED`. As a result, an attacker-controlled URL query string can still write into `Object.prototype` on any page that loads `main.js`. Version 1.8.223 contains a updated fix.	4.6	More Details

CVE-2026-47689	FOG is a free open-source cloning/imaging/rescue suite/inventory management system. Prior to versions 1.5.10.1832 and 1.6.0-beta.2313, the `buildRow()` method in `fogpage.class.php` substitutes data values into HTML table cell templates using `str_replace()` without any HTML escaping. An unauthenticated attacker who knows any registered host's MAC address can POST malicious inventory values (e.g. `sysproduct`, `sysserial`) to `/service/inventory.php`, which stores them in the database. When an administrator opens the Group Inventory tab, the payload renders as executable HTML/JavaScript in the admin's browser. Versions 1.5.10.1832 and 1.6.0-beta.2313 fix the issue.	4.6	More Details
CVE-2026-46671	Rust OneNote File Parser is a parser for Microsoft OneNote files implemented in Rust. Prior to version 1.1.1, a maliciously crafted `.onetoc2` table-of-contents file can cause `Parser::parse_notebook` to open arbitrary files on the host filesystem outside the notebook's directory. The parser reads entry names listed inside the `.onetoc2` and joins them against the notebook's base directory without validating that they are relative paths confined to that directory. The parser will bail out when the target file fails to parse as a OneNote section, so direct content exfiltration through the parser's return value is not practical, though file-existence probing and denial-of-service via large or special files remain possible. Anyone using `onenote_parser` to parse .onetoc2 files received from untrusted sources is affected. Users who only ever parse their own notebooks are not at meaningful risk. The issue is fixed in onenote_parser 1.1.1. The fix rejects absolute paths, parent-directory components, and other invalid path characters in entry names, and additionally canonicalises the resolved path to confirm it stays inside the notebook's base directory. For users who cannot upgrade to 1.1.1, only call `Parser::parse_notebook` on `.onetoc2` files from trusted sources. Alternatively, use `Parser::parse_section` / `Parser::parse_section_buffer` on individual .one files, which do not perform the directory walk.	4.4	More Details
CVE-2026-46936	Vulnerability in the MySQL Server, MySQL Cluster product of Oracle MySQL (component: Server: DDL). Supported versions that are affected are MySQL Server: 8.4.0-8.4.10, 9.7.0-9.7.1; MySQL Cluster: 8.0.0-8.0.47, 8.4.0-8.4.10 and 9.7.0-9.7.1. Difficult to exploit vulnerability allows high privileged attacker with network access via multiple protocols to compromise MySQL Server, MySQL Cluster. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of MySQL Server, MySQL Cluster. CVSS 3.1 Base Score 4.4 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:H/UI:N/S:U/C:N/I:N/A:H).	4.4	More Details
CVE-2026-12139	Tanium addressed an information disclosure vulnerability in Connect.	4.4	More Details
CVE-2026-15324	The SysBasics Customize My Account for WooCommerce – Live My Account Customizer plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'row_type' parameter in all versions up to, and including, 4.4.14 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with shop manager-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	4.4	More Details
CVE-2026-10588	A potential vulnerability could allow a local privileged attacker to disclose the address of protected System Management Mode memory.	4.4	More Details
CVE-2026-16130	A vulnerability was identified in nearai ironclaw up to 0.29.1. The affected element is the function validate_path of the file src/tools/builtin/path_utils.rs of the component write_file. The manipulation leads to link following. Local access is required to approach this attack. The exploit is publicly available and might be used. The identifier of the patch is 369ff3d240cf3c0787b50e1e9f182e1a06c71255. It is recommended to apply a patch to fix this issue.	4.4	More Details
CVE-2026-10590	A potential missing authentication vulnerability could allow a local privileged attacker to use WMI commands to arbitrarily trigger a System Management Interrupt handler.	4.4	More Details
CVE-2026-13005	The MxChat – AI Chatbot & Content Generation for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via admin settings in all versions up to, and including, 3.2.10 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with administrator-level permissions and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This only affects multi-site installations and installations where unfiltered_html has been disabled.	4.4	More Details
CVE-2026-40953	CVE-2026-40953 is a heap overflow in the certificate parsing function of Secure Access clients prior to 14.55. Attackers with local access and administrator permissions can create a denial of service attack against the client over which they have control.	4.4	More Details
CVE-2026-41993	Improper Access Control vulnerability in the Removable Media Validation function of TXOne Networks products allows a local attacker with administrator privileges to bypass the file lockdown mechanism, resulting in unauthorized file transfer to the victim device. The attacker needs to deploy unauthorized file on the removable media in advance. This issue affects SafePortAgent: before 3.2.5024; StellarProtect: from 3.2.4011 before 5.0.1083.	4.4	More Details

CVE-2024-32385	An issue in Kerlink Kerlink Wirnet iStation 868 KerOS v.4.3.3_20200803132042 allows a remote attacker to obtain sensitive information via a boardID and revisionID components	4.3	More Details
CVE-2026-24232	NVIDIA Tranformers4Rec contains a vulnerability where an attacker could cause improper deserialization of untrusted data. A successful exploit of this vulnerability might lead to code execution, data tampering, and information disclosure.	4.3	More Details
CVE-2026-16254	A flaw was found in claircore's apk package scanner. Malformed package-database data in a container layer can cause an out-of-bounds access that panics the scanner. If that panic is not recovered, the Clair indexer process can crash, leading to a denial of service.	4.3	More Details
CVE-2026-49092	Unintended Proxy or Intermediary ('Confused Deputy') (CWE-441) in Kibana can lead to unauthorized information exposure via Accessing Functionality Not Properly Constrained by ACLs (CAPEC-1). Under certain conditions, a lower-privileged user can cause data from sources they are not authorized to access to be processed using another user's privileges.	4.3	More Details
CVE-2026-57857	The Flow Payment plugin for WordPress (flow.cl) version 3.0.8 is vulnerable to reflected cross-site scripting on the WooCommerce checkout page. When the plugin handles an order cancellation, the error_message GET parameter is passed directly to wc_add_notice() in flowpayment-fl.php (lines 57-58) without input sanitization (for example sanitize_text_field()) or output escaping (for example esc_html()) before being rendered in the checkout notice HTML. An unauthenticated attacker can craft a URL containing a JavaScript payload in the error_message parameter (for example /checkout/?add-to-cart={product-id}&cancel_order=true&error_message={payload}); when a victim with an active WooCommerce checkout session follows the link, the payload executes in the victim's browser in the origin of the WordPress site.	4.3	More Details
CVE-2026-16108	A flaw was found in the default-groups REST endpoint and realm representation of Keycloak. This component is responsible for managing groups that are automatically assigned to new users within a realm. The issue allows a delegated administrator with realm-viewing permissions to see the names and identifiers of hidden default groups, even if they lack the specific permissions to view those groups. This can lead to the exposure of sensitive organizational structures or internal group names.	4.3	More Details
CVE-2026-16229	A flaw has been found in itsourcecode Courier Management System up to 1.0. Affected by this vulnerability is an unknown functionality of the file /index.php. Executing a manipulation of the argument page can lead to cross site scripting. It is possible to launch the attack remotely. The exploit has been published and may be used.	4.3	More Details
CVE-2026-12724	The Kirki WordPress plugin before 6.0.12 does not sanitise or escape the email subject and body values supplied in a request before including them in the password-reset email it sends as HTML, allowing unauthenticated users to inject arbitrary HTML into the message delivered to a registered user, which can be used for phishing.	4.3	More Details
CVE-2026-55399	CVE-2026-55399 is a resource exhaustion vulnerability in the Secure Access publisher prior to 14.55. Attackers with valid credentials to the Secure Access tunnel can create a non-persistent DoS against the publisher.	4.3	More Details
CVE-2026-63733	SurrealDB versions before 3.2.0 contain a permissions bypass vulnerability where data-modifying statements within PERMISSIONS clauses execute with enforcement disabled. Attackers with permission to perform a guarded operation can write to tables they lack permission for by embedding CREATE, UPDATE, DELETE, or UPSERT statements in the PERMISSIONS clause, causing unintended writes and data corruption.	4.3	More Details
CVE-2026-47089	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. LISTRIGHTS os not limited to users with admin access. An authenticated user could call IMAP LISTRIGHTS against any mailbox they could name and learn what principals had what access to it. (This action should have been restricted to users with admin access on the target mailbox.)	4.3	More Details
CVE-2026-63761	SurrealDB before 3.1.0 silently substitutes the ES384 algorithm when a JWT access method is configured with ALGORITHM ES512 (DEFINE ACCESS ... TYPE JWT ALGORITHM ES512), because the underlying jsonwebtoken crate (v10.x) has no ES512 variant and the mapping defaults to ES384 without any error, warning, or log message. Users who supply the correct P-521 key for ES512 experience authentication handshake failures due to the curve mismatch with ES384 (which expects P-384), and tokens are rejected by external systems expecting genuine ES512 signatures. The flaw cannot be used to forge tokens or compromise data confidentiality or integrity, as ES384 remains cryptographically strong.	4.3	More Details
CVE-2026-47083	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. There is an ESEARCH cross-user content oracle. By using the ESEARCH command, an authenticated IMAP user could enumerate folder names under any account they could name. Search would return UIDs of messages matching the search, creating a content oracle (without allowing arbitrary reads of the target's content).	4.3	More Details
CVE-2026-63742	SurrealDB versions before 3.1.0 contain a field-level SELECT permission bypass vulnerability in indexed COUNT fast paths. Attackers can execute COUNT queries on indexed fields with field-level SELECT restrictions to confirm or recover restricted field values through repeated guesses.	4.3	More Details

CVE-2026-16450	A vulnerability was identified in zsadmin2025 ZS-Admin up to b52e14536d59fda11e56e2536a1c32e82a38cead. This affects the function getTenantId of the file /api/system/sys/dept/page of the component MyBatis-Plus Tenant Plugin. Such manipulation of the argument X-Tenant-Id leads to authorization bypass. The attack may be performed from remote. The exploit is publicly available and might be used. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-15336	The Catch Themes Demo Import plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 3.3. This is due to the catch_themes_demo_import_activate_plugin() function, hooked on admin_init when the activate_plugin GET parameter is present, calling Plugin_Upgrader::install() to download and install a plugin from WordPress.org before performing the current_user_can('activate_plugins') capability check. This makes it possible for authenticated attackers, with subscriber-level access and above, to install the hardcoded 'essential-content-types' plugin from the WordPress.	4.3	More Details
CVE-2026-63308	Helm through 4.2.3, fixed in commit ba6c9a2, contains a denial of service vulnerability in the Files.Lines template helper in pkg/engine/files.go that allows attackers to trigger an index out of range panic by including zero-length byte slices in chart files. Attackers can include empty files in Helm charts to cause deterministic render failures across template, install, upgrade, lint, and SDK Engine.Render operations.	4.3	More Details
CVE-2026-16224	A vulnerability was identified in jxxghp MoviePilot up to 2.13.5. The affected element is an unknown function of the file /jxxghp/MoviePilot of the component Application API. The manipulation leads to improper authorization. Remote exploitation of the attack is possible. The identifier of the patch is dc2b6910a423b3bfadeffaa303e1ba75cfb33900. Applying a patch is the recommended action to fix this issue.	4.3	More Details
CVE-2026-8285	Improper restriction of excessive authentication attempts vulnerability in Universal Software Inc. FlexCity allows Excessive Allocation. This issue affects FlexCity: from 5.536.0 through 11052026.	4.3	More Details
CVE-2026-63309	SurrealDB before 3.1.5 fail to apply field-level SELECT permissions to ORDER BY clauses, allowing authenticated users to leak the relative ordering of restricted field values. Attackers can issue ORDER BY queries on indexed restricted fields to recover the hidden values' sort order across records, even though the field itself returns null as intended.	4.3	More Details
CVE-2026-63748	SurrealDB versions before 3.1.0 contain an information disclosure vulnerability where authenticated users with UPDATE access can read field values hidden by field-level SELECT permissions through error messages. Attackers can trigger arithmetic or extend operations on hidden fields to embed raw operand values in error responses, bypassing field-level access controls.	4.3	More Details
CVE-2026-9656	The HubSpot All-In-One Marketing – Forms, Popups, Live Chat plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 11.3.62 via the wp_localize_script() / window.leadinConfig JavaScript object. This makes it possible for authenticated attackers, with contributor-level access and above, to extract the site's plaintext HubSpot OAuth refresh token exposed via the window.leadinConfig JavaScript object, which can then be used to access or modify data in the connected HubSpot tenant. Although the refresh token is stored at rest with AES-256-CTR encryption, decryption occurs server-side before the plaintext value is passed to wp_localize_script(), rendering the at-rest encryption ineffective against this exposure path.	4.3	More Details
CVE-2026-63749	SurrealDB versions before 3.1.0 contain an authentication bypass vulnerability in LIVE SELECT subscriptions where permission expressions referencing $value, $before, $after, or $event are evaluated against attacker-controlled bindings instead of actual documents. Authenticated subscribers can bind chosen values to these parameter names and register LIVE SELECT queries to receive notifications for records that SELECT permission expressions should have hidden.	4.3	More Details
CVE-2026-48016	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, the Store API endpoint /store-api/handle-payment in src/Core/Checkout/Payment/SalesChannel/HandlePaymentMethodRoute.php accepts a user-controlled orderId and forwards it to src/Core/Checkout/Payment/PaymentProcessor.php without verifying order ownership or guest-order authentication, allowing a normal customer or guest context to trigger the payment flow for another user's order while /store-api/order enforces the expected ownership model. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	4.3	More Details
CVE-2026-16220	A vulnerability has been found in code-projects Online Examination System 1.0. This vulnerability affects unknown code of the file /account.php?q=quiz. Such manipulation of the argument eid/n/t leads to cross site scripting. The attack can be launched remotely. The exploit has been disclosed to the public and may be used.	4.3	More Details
CVE-2026-63751	SurrealDB versions before 3.1.0 contain a field-level permission bypass vulnerability in JSON Patch operations that allows authenticated users to read protected fields. Attackers can use UPDATE PATCH with an empty from pointer in copy or move operations to duplicate all record fields, including those restricted by field-level SELECT permissions, into attacker-chosen destination fields.	4.3	More Details
CVE-	SurrealDB before 3.1.0 contains an authorization bypass vulnerability in the RELATE statement that allows		

2026-63752	authenticated users with CREATE permission to overwrite existing edge records without UPDATE permission. Attackers can issue a RELATE statement with a SET id clause pointing to an existing edge id, causing the storage layer to silently overwrite the target record instead of rejecting the operation.	4.3	More Details
CVE-2026-12409	The Landing Page Builder – Coming Soon page, Maintenance Mode, Lead Page, WordPress Landing Pages plugin for WordPress is vulnerable to Cross-Site Request Forgery in all versions up to, and including, 1.5.3.6. This is due to missing or incorrect nonce validation on the ulpb_admin_ajax function. This makes it possible for unauthenticated attackers to create, update, retitle, or change the post status, slug, and type of arbitrary posts and write ULPB_DATA post meta via a forged request granted they can trick a site administrator into performing an action such as clicking on a link. This attack requires the victim to hold an editor-level or administrator session, as the wp_ajax_ulpb_admin_data action enforces a capability check that the forged request must satisfy by inheriting the logged-in user's session cookies.	4.3	More Details
CVE-2026-12434	The List category posts plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 0.95.0 via the sanitize_status. This makes it possible for authenticated attackers, with contributor-level access and above, to extract titles, full content, excerpts, dates, authors, and custom-field metadata of other users' pending-review, scheduled, and trashed posts by embedding a crafted [catlist] shortcode in their own draft and previewing it. This vulnerability is a bypass of the incomplete fix introduced for CVE-2025-11377 in version 0.93.0.	4.3	More Details
CVE-2026-63753	SurrealDB before 3.1.0 fails to refresh authentication state in LIVE SELECT subscriptions when session state changes. Attackers can continue receiving real-time notifications under revoked or expired session credentials until the connection closes.	4.3	More Details
CVE-2026-16216	A weakness has been identified in geex-arts django-jet up to 1.0.8. Affected is an unknown function of the component OAuth Handler. Executing a manipulation can lead to cross-site request forgery. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-63738	SurrealDB versions 3.1.0 before 3.1.5 fail to enforce field-level SELECT permissions when records are accessed through graph-edge or back-reference traversals. Attackers with table-level SELECT access can read field values hidden by field-level permissions by materializing records through graph traversals instead of direct table scans.	4.3	More Details
CVE-2026-59850	A flaw was found in libssh. If data packets are processed after a channel is closed, channel data callbacks can be invoked after the associated data has already been freed, leading to crashes or possible use-after-free conditions.	4.3	More Details
CVE-2026-14501	IBM Db2 Genius Hub 1.1, 1.1.1, 1.1.2 and IBM Agentics 1.0 could allow an attacker to execute arbitrary code or obtain sensitive information due to the use of dangerous functions without sufficient restrictions.	4.3	More Details
CVE-2026-15160	The Ninja Forms - Excel Export plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 3.3.6 via the 'spreadsheet_export_tmp_name' parameter. This makes it possible for authenticated attackers, with subscriber-level access and above, to write .xls/.xlsx files to arbitrary locations on the server, which can be used to stage further attacks.	4.3	More Details
CVE-2026-16103	A flaw was found in the keycloak-services component of Keycloak. This issue is an incomplete fix for CVE-2026-9798, where brute-force protection checks were added to the Client-Initiated Backchannel Authentication (CIBA) initiation handler but were omitted from the token redemption handler. This allows an attacker with valid client credentials to obtain access and refresh tokens for a user account that has been locked due to brute-force protection, provided the authentication request was started before the lockout occurred and was approved by the user.	4.3	More Details
CVE-2026-14183	The Classified Listing WordPress plugin before 5.3.9 does not verify that the order targeted by its payment-receipt handler belongs to the requesting user, allowing authenticated users with subscriber-level access to read the payment receipt details of any other user's order.	4.3	More Details
CVE-2026-63097	Dendrite through 0.13.8 contains an improper access control vulnerability in the syncapi /context endpoint (syncapi/routing/context.go) that allows authenticated local users to access post-leave room state events by exploiting a flawed membership check that evaluates only the RoomExists field while ignoring IsInRoom, HasBeenInRoom, and Membership fields. Attackers who have left a room can call the rooms context API endpoint for a previously permitted event and receive unfiltered current room state that the /messages and /sync endpoints correctly withhold.	4.3	More Details
CVE-2026-3182	Zohocorp ManageEngine Endpoint Central versions before 11.4.2528.34 are affected by cleartext transmission of sensitive information vulnerability.	4.3	More Details
CVE-	Wazuh is a free and open source platform used for threat prevention, detection, and response. In versions 4.6.0 and above, prior to 4.14.5, a logic error in CheckRateLimitsMiddleware.dispatch() causes the /events endpoint rate check to unconditionally overwrite the general rate limit result. When the global		More

2026-33434	max_request_per_minute is exceeded, requests to /events still succeed if the events-specific counter (hardcoded 30/min) has not been reached. This allows event injection into analysisd beyond the admin-configured global rate limit. This issue has been fixed in version 4.14.5.	4.3	Details
CVE-2026-16336	A vulnerability was found in trinodb trino 481. Affected is an unknown function of the file core/trino-main/src/main/java/io/trino/server/ExternalUriInfo.java of the component OAuth2/OIDC. Performing a manipulation of the argument redirect_uri results in open redirect. It is possible to initiate the attack remotely. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-63092	kirby-modules through 5.5.7, fixed in commit 315417e, contains an information disclosure vulnerability that allows any authenticated Kirby Panel user to retrieve the full plaintext commercial license key by sending a GET request to the modules/activate dialog endpoint. The plugin's activate dialog handler in lib/areas.php returns the complete key via ModulesLicense::readKey() without performing an administrator check, as the dialog is gated only by the access.system permission which defaults to true for all non-admin roles, enabling attackers to use the disclosed key to activate the plugin on arbitrary third-party installations.	4.3	More Details
CVE-2026-15349	The ERP: Complete HR, Accounting & CRM Suite Built for WooCommerce plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 1.17.6. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for authenticated attackers, with subscriber-level access and above, to create arbitrary company locations in the ERP database.	4.3	More Details
CVE-2026-64821	djangoSIGE through 1.10 (commit a6fe7e8) contains a cross-site request forgery vulnerability that allows unauthenticated attackers to cancel sales or purchase orders on behalf of authenticated users by exploiting order-cancellation logic implemented inside HTTP GET method handlers in CancelarOrcamentoVendaView, CancelarPedidoVendaView, CancelarOrcamentoCompraView, and CancelarPedidoCompraView. Attackers can lure an authenticated victim with change_orcamentovenda or equivalent permissions to a page containing a cross-origin reference such as an img tag pointing to the cancellation endpoint, bypassing CSRF token validation entirely since Django's CsrfViewMiddleware only enforces CSRF checks on unsafe HTTP methods.	4.3	More Details
CVE-2026-15159	The Ninja Forms - Excel Export plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 3.3.6 via the 'spreadsheet_export_form_id' parameter due to missing validation on a user controlled key. This makes it possible for authenticated attackers, with subscriber-level access and above, to enumerate any Ninja Forms form ID and download all stored submission data — including names, email addresses, phone numbers, physical addresses, and any other PII collected by site forms — as a downloadable XLSX file.	4.3	More Details
CVE-2026-49977	tarteaucitron.js is a compliant and accessible cookie banner. Prior to 1.33.0, tarteaucitron.cookie.purge() is called on any element with the purgeBtn class and does not check whether the element is a legitimate tarteaucitron button or whether the cookie corresponds to a service handled by tarteaucitron. If an attacker can write HTML with data attributes, an element with data-cookie can silently delete a non-HttpOnly cookie with a known name when clicked by a user. This issue is fixed in version 1.33.0.	4.3	More Details
CVE-2026-44595	Yamcs is a mission control framework. Prior to 5.12.7, the IAM API endpoints listUsers, getUser, listGroups, and getGroup in yamcs-core did not enforce the required SystemPrivilege.ControlAccess check in yamcs-core/src/main/java/org/yamcs/http/api/IamApi.java, so any authenticated user, even one with low or no privileges, could enumerate all user accounts in the system including their usernames, superuser status, and group memberships. This issue is fixed in versions 5.12.7 and 5.13.0.	4.3	More Details
CVE-2026-63768	cal.diy through 6.2.0 contains an open redirect vulnerability in the conferencing OAuth callback endpoint that allows attackers to redirect users to arbitrary URLs by crafting malicious state parameters. Attackers can exploit the unsigned state parameter and onErrorReturnTo field to silently redirect visitors from the trusted domain to attacker-controlled URLs for phishing attacks.	4.3	More Details
CVE-2026-10675	In Zephyr's Bluetooth Mesh PB-ADV provisioning bearer (subsys/bluetooth/mesh/pb_adv.c), prov_msg_recv() rescheduled the provisioning protocol watchdog timer unconditionally at the top of the function, before the FCS check and before the ADV_LINK_INVALID check. Once a provisioning attempt fails, prov_failed() sets ADV_LINK_INVALID and the only recovery path is the protocol timer firing (protocol_timeout -> prov_link_close -> close_link -> reset_adv_link and re-enabling of scanning and the unprovisioned device beacon). A remote, unauthenticated attacker on the BLE advertising channel can first induce a provisioning failure (e.g. with a malformed generic-provisioning PDU) and then transmit any FCS-valid PB-ADV transaction PDU on the same link ID more often than once per protocol timeout (60 s, or 120 s for OOB input/output). Because each such packet reset the timer even on an invalidated link, protocol_timeout never fired, the dead link was never torn down, and the device remained pinned in an un-provisionable state with its unprovisioned beacon disabled and new Link Open requests rejected. PB-ADV PDUs are processed without authentication and the FCS is a keyless CRC, so no pairing or prior trust is required and the attacker chooses the link ID itself. The impact is a persistent denial of provisioning/re-provisioning service; there is no memory-safety, confidentiality, or integrity impact. The vulnerable code shipped in releases through v4.4.1. The fix moves the timer reschedule to after the ADV_LINK_INVALID check (and the FCS check before the reset) so an invalidated link can no longer be kept alive by incoming packets.	4.3	More Details
	HCL Aftermarket EPC is affected by Sensitive Information in GET method & in URL which allows application to		

CVE-2024-23567	pass sensitive data via URL parameters during normal usage. Data passed in this manner can be exposed because it may end up stored in unintended locations, including server logs, local browser history and proxy logs.	4.3	More Details
CVE-2026-54568	Microsoft UFO open-source framework for intelligent automation across devices and platforms. From 3.0.0 until 3.0.6, a client connected to the UFO WebSocket server as a DEVICE could call DEVICE_INFO_REQUEST with another device's target_id and receive that device's server-side system_info through ufo/server/ws/handler.py, because handle_device_info_request and get_device_info did not enforce the constellation-only role or object-level authorization boundary. This issue is fixed in version 3.0.6.	4.3	More Details
CVE-2026-16081	A vulnerability was determined in Sipeed PicoClaw up to 0.2.9. The affected element is an unknown function of the file web/backend/api/auth.go. Executing a manipulation can lead to cross-site request forgery. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized. This patch is called 4b0229351678f479429b8d8b19207757266f246b. Applying a patch is advised to resolve this issue.	4.3	More Details
CVE-2026-57205	SimpleChat is a secure AI conversation application with personal and group workspaces for document-grounded interactions. Prior to 0.241.203, the authenticated GET /api/user/info/<user_id> and GET /api/user/profile-image/<user_id> endpoints in application/single_app/route_backend_users.py accepted a caller-supplied user_id and read the matching Cosmos DB user-settings document without object-level authorization, allowing a low-privilege authenticated user to retrieve another user's email address, display name, and profile image. This issue is fixed in version 0.241.203.	4.3	More Details
CVE-2026-21954	Vulnerability in the Oracle Retail Xstore Point of Service product of Oracle Retail Applications (component: Xstore Mobile). The supported version that is affected is 21.0.3. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Retail Xstore Point of Service. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Retail Xstore Point of Service accessible data. CVSS 3.1 Base Score 4.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N).	4.3	More Details
CVE-2024-23577	HCL Aftermarket EPC is vulnerable since the application does not have a validation for HOST header and accepts arbitrary hosts when requested in http protocol. When an application doesn’t adequately validate or sanitize this header, it can lead to several security risks, including Host header poisoning, server misconfigurations.	4.3	More Details
CVE-2026-9734	The W3SC Elementor to Zoho CRM plugin for WordPress is vulnerable to Cross-Site Request Forgery in all versions up to, and including, 2.2.0. This is due to missing or incorrect nonce validation on the storeInfo function. This makes it possible for unauthenticated attackers to modify the plugin's Zoho CRM integration settings, replacing the configured data center, client ID, client secret, and user email credentials with attacker-controlled values via a forged request granted they can trick a site administrator into performing an action such as clicking on a link.	4.3	More Details
CVE-2026-16075	A flaw has been found in AstrBotDevs AstrBot up to 4.25.5. This vulnerability affects the function OpenApiRoute.get_chat_sessions of the file astrbot/dashboard/routes/open_api.py of the component session-listing Endpoint. This manipulation of the argument Username causes authorization bypass. It is possible to initiate the attack remotely. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-44584	Paymenter is a free and open-source webshop solution for management of hosting services. In versions prior to 1.5.0, the email update functionality fails to invalidate the existing verification state when a user changes their email address, allowing a verified account to retain its verified status after switching to an unverified or unowned email address. When a user updated their email address, the system did not reset or revalidate the associated email verification status. As a result, the verification column remained set to “true” even after the email address was changed. Exploitation could potentially result in: misrepresentation of email ownership, bypass of verification-based trust assumptions, and abuse of features gated behind verified status. No direct unauthorized access to other users accounts or data is possible through this issue alone. This issue has been fixed in version 1.5.0.	4.3	More Details
CVE-2026-46980	Vulnerability in the Oracle Utilities Network Management System product of Oracle Utilities Applications (component: Mobile). Supported versions that are affected are 2.5.0.1.0-2.5.0.1.17, 2.5.0.2.0-2.5.0.2.11, 2.6.0.1.0-2.6.0.1.12, 2.6.0.2.0-2.6.0.2.8 and 25.12.0.0.0-25.12.0.0.2. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Utilities Network Management System. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Utilities Network Management System accessible data. CVSS 3.1 Base Score 4.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N).	4.3	More Details
CVE-2024-23571	HCL Aftermarket EPC is vulnerable to attack since the application does not have an appropriate caching policy specifying the extent to which the page and its form fields should be cached. If sensitive information in application responses is stored in the local cache, then this may be retrieved by other users who have access to the same computer at a future time.	4.3	More Details
	HCL Aftermarket EPC is affected by clickjacking vulnerability Cross-Frame Scripting is an attack technique		

CVE-2024-23570	where an attacker loads a vulnerable application in an iFrame on his malicious site. The attacker can then launch a Clickjacking attack, which may lead to Phishing, Cross-Site Request Forgery, sensitive information leakage and more.	4.3	More Details
CVE-2026-14185	The WPBot WordPress plugin before 8.2.0 does not perform a capability or nonce check in one of its retrieval-augmented-generation settings handlers, allowing authenticated users with subscriber-level access to modify the WPBot WordPress plugin before 8.2.0's configuration.	4.3	More Details
CVE-2024-23569	HCL Aftermarket EPC is vulnerable to attack since the server is not configured with “X-XSS-Protection" header	4.3	More Details
CVE-2026-32823	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, the application exposes server-side state changes through `GET` routes. Because browsers automatically send cookies on same-site top-level navigation and Rails does not apply CSRF protections to `GET`, an attacker can force a logged-in victim to modify application state by embedding a link, image, iframe, or redirect to one of these endpoints. This was confirmed on the target with a normal `Standard` account: a cross-site-style `GET` to `watch_lists/:id/add_item?thing_id=...` inserted content into a watch list with no CSRF token. Additional `GET` mutation routes exist in the codebase, including user impersonation for authorized admins and cache or translation state changes. This is patched in version 26.06.08.	4.3	More Details
CVE-2026-15093	IBM Engineering AI Hub 1.0.0, 1.1.0, and 1.2.0 could allow a remote attacker to redirect users to malicious websites due to improper validation of user-supplied URLs.	4.3	More Details
CVE-2026-46338	PyMdown Extensions is a set of extensions for the Python-Markdown markdown project. From 10.0.1 until 10.21.3, pymdownx.snippets uses a string-prefix containment check in SnippetPreprocessor.get_snippet_path() in pymdownx/snippets.py when `restrict_base_path: True`, allowing markdown snippet directives to read files from sibling paths that share the same base_path prefix, such as docs and docs_internal. This is a regression of CVE-2023-32309. This issue is fixed in version 10.21.3.	4.3	More Details
CVE-2026-15610	The WPBot – AI ChatBot for Live Support, Lead Generation, AI Services plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 8.5.6. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for authenticated attackers, with subscriber-level access and above, to trigger arbitrary re-embedding of stored RAG documents, modifying the rag_documents table and consuming the site owner's paid third-party AI API credits (OpenAI, Gemini, OpenRouter, or xAI).	4.3	More Details
CVE-2026-55548	Yamcs is a mission control framework. Prior to 5.12.8 and 5.13.2, the PacketsApi.exportPackets endpoint in yamcs-core/src/main/java/org/yamcs/http/api/PacketsApi.java failed to enforce object-level ReadPacket privileges when a request omitted specific packet names: with an empty name list the ctx.checkObjectPrivileges(ObjectPrivilegeType.ReadPacket, nameSet) call passed over an empty set, no WHERE pname IN filter was applied to the resulting SELECT * FROM tm query, and the onTuple handler streamed every retrieved packet without any per-row authorization check, so a low-privileged or zero-privilege authenticated user could dump the entire raw telemetry packet archive and bypass the role-based access control model. This issue is fixed in versions 5.12.8 and 5.13.2, which enforce per-packet ReadPacket checks in exportPackets.	4.3	More Details
CVE-2026-1372	The Tutor LMS Elementor Addons plugin for WordPress is vulnerable to Missing Authorization in all versions up to, and including, 4.0.0 This is due to missing capability checks on the `activate_tutor_free()` and `activate_elementor_free()` functions registered as `admin_action_*` handlers. This makes it possible for authenticated attackers, with Subscriber-level access and above, to activate the Tutor LMS and Elementor plugins without proper authorization.	4.3	More Details
CVE-2026-16122	A security flaw has been discovered in nextlevelbuilder GoClaw up to 3.13.2. Affected by this vulnerability is the function extractBin/RequestApproval/matchesAllowlist of the file internal/tools/exec_approval.go. The manipulation results in incorrect authorization. The exploit has been released to the public and may be used for attacks.	4.3	More Details
CVE-2026-16454	In Eclipse hawkBit versions 1.0.3 and prior, a privilege escalation vulnerability (CWE-284 / CWE-862) has been identified in the Direct Device Integration (DDI) Controller. This vulnerability allows an authenticated device to escalate its permissions and bypass the strict boundaries of its assigned updates. Under normal operation, a device should be restricted strictly to the specific firmware artifacts explicitly assigned to it. However, this flaw enables any authenticated device to bypass this restriction and download any firmware artifact within the same tenant. This is not an authentication bypass; the requesting device must possess valid credentials for its respective tenant. Instead, the issue stems from a flaw in object-level authorization validation. A related, lower-severity helper issue exists in the listing software modules artifacts metadata endpoint. This endpoint does not enforce assignment checks, enabling an authenticated device to list and enumerate available firmware artifacts, which can facilitate targeted exfiltration using the main download authorization bypass.	4.3	More Details

CVE-2026-32819	dataCycle is a data management system for centrally storing, managing, searching, finding, and distributing data. In dataCycle-CORE, the module handling core processing and framework rules, before and including version 25.07.3, a Standard user can enumerate other users' names and email addresses through `/users/search`, even though direct access to those user profiles is denied. This leaks internal staff addresses, full names, and existence of guest and external test accounts.	4.3	More Details
CVE-2026-16104	A flaw was found in the authentication configuration endpoint of the keycloak-services component, which is the core engine for Red Hat Build of Keycloak identity and access management. The issue occurs because the system fails to mask sensitive configuration values, such as reCAPTCHA secret keys, when they are requested by administrators with view-only permissions. This can lead to the exposure of third-party service credentials to unauthorized personnel or through administrative logs.	4.3	More Details
CVE-2026-15350	The The Cache Purger plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.3.20. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for authenticated attackers, with subscriber-level access and above, to permanently truncate the plugin's cache-purge audit log (wp-content/purge.log), destroying the entire cache-purge audit history. The tcp_log_purge nonce is rendered in the admin bar on frontend pages accessible to all authenticated users including subscribers, meaning any authenticated user possesses the nonce required to trigger the deletion.	4.3	More Details
CVE-2026-65009	OpenRemote versions before 1.26.2 contain an information disclosure vulnerability in the SyslogResource REST endpoint that fails to filter operational logs by realm. Attackers with the read:rules role can access the GET /api/{realm}/syslog/event endpoint to retrieve operational logs from all tenants, exposing asset IDs, agent connection details, rule names, and protocol errors across the multi-tenant deployment.	4.3	More Details
CVE-2026-15407	The Themify Builder plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 7.7.7. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for authenticated attackers, with subscriber-level access and above, to overwrite or delete the generated CSS stylesheet file of arbitrary posts, including private and draft posts owned by other users, and modify plugin-scoped font options. The required CSRF nonce (tf_nonce) is emitted on public front-end builder pages via wp_localize_script, making it trivially obtainable by any authenticated user visiting such a page.	4.3	More Details
CVE-2026-13724	Client-Side Enforcement of Server-Side Security vulnerability in Gobito Informatics Technologies Engineering Industry and Trade Ltd. Co. Corporate Training Management System allows Input Data Manipulation. This issue affects Corporate Training Management System: before dd1a9df64.	4.3	More Details
CVE-2026-15945	A flaw was found in the group search functionality of the Keycloak server's administrative API. When Fine-Grained Admin Permissions (FGAP) v2 is enabled, a delegated administrator can bypass access restrictions to view parent groups they are not authorized to see. By searching for a child group they have permission to view, the system incorrectly returns the full details of the parent group in the response, leading to the disclosure of sensitive group attributes and configuration.	4.3	More Details
CVE-2026-55608	n8n-MCP is an MCP server that provides AI assistants access to n8n node documentation, properties, and operations. Prior to 2.57.4, multi-tenant HTTP mode with ENABLE_MULTI_TENANT=true could allow an authenticated tenant to access default-scope workflow_versions backups instead of being confined to the tenant scope, exposing or deleting workflow-version backups from prior single-tenant deployments or migrations. This issue is fixed in version 2.57.4.	4.2	More Details
CVE-2026-21761	HCL DevOps Loop is affected by a Cross-Origin Resource Sharing (CORS) misconfiguration. Improper CORS configuration may allow unauthorized cross-origin requests, potentially exposing application resources to untrusted domains.	4.2	More Details
CVE-2024-23578	HCL Aftermarket EPC is vulnerable to attack as the application implements an HTML5 cross-origin resource sharing (CORS) policy for this request that allows access from any domain (*-Wildcard).	4.2	More Details
CVE-2026-47122	Sparkle is a software update framework for macOS. In versions up to and including 2.9.1, `Autoupdate/AppInstaller.m`'s `shouldAcceptNewConnection:` only enforces `SUCodeSigningVerifier validateConnection:` before stage 1 completes. After `_performedStage1Installation = YES`, new connections to the registered Mach service `<bundleId>-spki` are accepted from any local process without team-ID or code-signing checks. As of time of publication, no known patched versions are available.	4.2	More Details
CVE-2026-12548	A heap out-of-bounds read flaw was found in libsoup. When parsing multipart HTTP messages, an integer type mismatch between the caller and soup_headers_parse() can cause the length parameter to be incorrectly truncated, leading to a heap buffer over-read. A remote attacker could use this flaw to crash an application using libsoup or potentially disclose heap memory contents.	4.2	More Details
CVE-2024-23572	HCL Aftermarket EPC is vulnerable to attack as cookie appears to contain a session token, which may increase the risk associated with this issue. You should review the contents of the cookie to determine its function.	4.2	More Details

CVE-2026-16212	A vulnerability was identified in awesto django-shop up to 1.2.4. Affected is an unknown function of the file shop/models/inventory.py of the component Purchase Stock Handler. The manipulation leads to race condition. The attack is possible to be carried out remotely. The attack is considered to have high complexity. The exploitability is told to be difficult. The exploit is publicly available and might be used. The project was informed of the problem early through an issue report but has not responded yet.	4.2	More Details
CVE-2026-63736	SurrealDB before 3.2.0 contains a server-side request forgery vulnerability in the JWKS fetcher that validates only the URL hostname string against allow-lists without checking resolved IP addresses. An Owner role attacker can point an access method at an allow-listed hostname resolving to private or loopback addresses, causing the server to issue GET requests to internal addresses that would be blocked by direct URL.	4.1	More Details
CVE-2026-63744	SurrealDB before 3.1.5 contains a server-side request forgery vulnerability in the JWKS fetcher that follows HTTP redirects without re-validating redirect targets against network capabilities. Attackers with Owner role can configure a JWKS URL pointing to an allowlisted host that redirects to blocked internal addresses, bypassing network access controls.	4.1	More Details
CVE-2026-58549	Out-of-bounds read vulnerability in the image codec module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	4.0	More Details
CVE-2026-58550	Out-of-bounds read vulnerability in the image codec module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	4.0	More Details
CVE-2026-16266	Versions of the package mongo-object before 3.0.3 are vulnerable to Prototype Pollution via the expandKey() function in util.js. An attacker can modify the JavaScript prototype chain by supplying a crafted property path containing special keys such as __proto__.	4.0	More Details
CVE-2026-58553	Out-of-bounds read vulnerability in the image codec module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	4.0	More Details
CVE-2026-47085	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. URLAUTH token forgery can occur via a missing mboxkey. If an attacker knew a folder name on the victim's account for which the victim had never issued an auth URL, they could forge a working URLAUTH token by computing an HMAC-SHA1 value with a predictable key, giving them read access to the mailbox. (URLAUTH is an obscure feature, meaning that the odds of any user actually being susceptible to this attack are very low. Perhaps no public clients use URLAUTH.)	4.0	More Details
CVE-2026-14971	IBM PowerVM Novalink 2.2.02.2.12.2.1.1, and 2.3.02.3.0.12.3.12.3.2 IBM NovaLink APIs misconfiguration may increase attack surface and enable unintended or unauthorized operations under non-default conditions.	3.9	More Details
CVE-2026-59846	A flaw was found in libssh. A malicious username expanded through %r in ProxyCommand handling can inject shell metacharacters, exposing environment variables and causing unintended shell behavior.	3.9	More Details
CVE-2026-61868	ImageMagick before 7.1.2-26 and 6.9.x before 6.9.13-51 contains a memory leak in the YUV decoder that occurs when opening of the blob fails. Repeated triggering can lead to resource exhaustion (denial of service).	3.7	More Details
CVE-2026-56584	HCL IEM was affected with the Information disclosure nginx server. It may enable attackers to identify outdated software versions and target known vulnerabilities or publicly available exploits.	3.7	More Details
CVE-2026-26080	HAProxy Community Edition 3.2.x through 3.3.x before 3.3.3 can enter a loop or crash because varint is mishandled. HAProxy Enterprise and ALOHA are also affected.	3.7	More Details
CVE-2026-33444	CVE-2026-33444 is a memory management vulnerability in Secure Access servers prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can create a non-persistent DoS against the server.	3.7	More Details
CVE-2024-23573	HCL Aftermarket EPC is vulnerable to attack since the Application is vulnerable to Lucky 13. that makes the SS LLUCKY13 possible affects the TLS1.1and 1.2 and DTLS1.0 or 1.2 implementations . It also affects previous versions such as SSL3.0 and TLS1.0. This can also be considered a type of man-in-the-middle attack.	3.7	More Details
CVE-2026-54335	Feathersjs is a framework for creating web APIs and real-time applications with TypeScript or JavaScript. In 5.0.44 and earlier, the _.merge(target, source) utility exported by @feathersjs/commons recursively merges source into target by iterating Object.keys(source). When source was produced by JSON.parse and contains a __proto__, constructor, or prototype key, that key is returned as an own-enumerable property; the recursive merge then resolves target['__proto__'] to Object.prototype and writes attacker-supplied properties onto it, polluting the prototype for all plain objects in the process for the lifetime of the Node process. This issue is	3.7	More Details

	fixed in version 5.0.45.		
CVE-2026-55398	CVE-2026-55398 is a memory management vulnerability in Secure Access clients and servers prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can create a non-persistent DoS against the server.	3.7	More Details
CVE-2026-56587	HCL IEM was affected with Strict transport security not enforced. It may enable attackers to perform SSL stripping or man-in-the-middle attacks and compromise secure communications.	3.7	More Details
CVE-2026-61871	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a memory leak in the ICON decoder that occurs when a memory allocation fails. Processing a crafted ICON file that triggers an allocation failure leaks memory, which may lead to a denial of service.	3.7	More Details
CVE-2026-40956	CVE-2026-40956 is a memory disclosure vulnerability in Secure Access client versions prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can cause a small amount of random memory to leak.	3.7	More Details
CVE-2026-16207	A vulnerability was detected in django-tastypie up to 0.15.1. Impacted is the function ApiKeyAuthentication of the file tastypie/authentication.py. The manipulation results in use of get request method with sensitive query strings. The attack can be launched remotely. This attack is characterized by high complexity. The exploitability is considered difficult. The project was informed of the problem early through an issue report but has not responded yet.	3.7	More Details
CVE-2026-61860	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a use-after-free vulnerability that occurs when freetype initialization fails: the method does not exit and continues to use memory that was already freed. This can be triggered during image processing and may lead to a denial of service.	3.7	More Details
CVE-2026-40958	CVE-2026-40958 is a input validation error in Secure Access clients prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can create a non-persistent DoS against their client.	3.7	More Details
CVE-2026-40955	CVE-2026-40955 is an integer underflow vulnerability in the traffic parsing function of Secure Access clients prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can create a non-persistent DoS against their client.	3.7	More Details
CVE-2026-62994	CoreDNS is a DNS server written in Go. From 1.9.4 until 1.14.5, a network DNS client allowed to request AXFR for a CoreDNS zone can trigger a panic when CoreDNS is configured with k8s_external headless-service zone transfers and Kubernetes contains a headless service endpoint with no declared ports; plugin/kubernetes/object/endpoint.go creates Port: -1, plugin/k8s_external/msg_to_dns.go skips that service, plugin/k8s_external/transfer.go sends an empty []dns.RR batch, and plugin/transfer/transfer.go indexes records[0] without checking the batch is non-empty. This issue is fixed in version 1.14.5.	3.7	More Details
CVE-2026-60065	When NGINX Plus is configured to use the Message Queuing Telemetry Transport (MQTT) filter module (ngx_stream_mqtt_filter_module), unauthenticated attackers can send requests with conditions beyond the attacker's control to cause a heap buffer over-read in the NGINX worker process, leading to a restart. Impact: This vulnerability may allow remote unauthenticated attackers to have limited control to restart the NGINX worker process. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	3.7	More Details
CVE-2026-40954	CVE-2026-40954 is an integer underflow vulnerability in the traffic parsing function of Secure Access clients prior to 14.55. Attackers with intimate knowledge of and total control over the tunnel protocol can create a non-persistent DoS against their client	3.7	More Details
CVE-2026-59842	A flaw was found in libssh. During server-side GSSAPI key exchange, a client-supplied Curve25519 public key shorter than the expected length is copied without proper length validation, leading to an out-of-bounds heap read. This could allow a remote unauthenticated attacker to disclose small amounts of server memory.	3.7	More Details
CVE-2026-21762	HCL DevOps Loop is affected by missing HTTP security headers. Missing security headers may reduce browser protections against common web-based attacks such as clickjacking, MIME-type sniffing, and cross-site scripting.	3.7	More Details
CVE-2026-56764	Hono before 4.11.10 contains a timing attack vulnerability in the basicAuth and bearerAuth middlewares due to non-constant-time string comparison in the timingSafeEqual function. Attackers can exploit early termination of string equality checks to infer valid credentials through precise timing measurements.	3.7	More Details
CVE-2026-16073	A security vulnerability has been detected in AstrBotDevs AstrBot up to 4.25.2. Affected by this issue is the function Star.text_to_image/NetworkRenderStrategy.render of the file astrbot/core/star/base.py of the component T2I Feature. The manipulation leads to cross site scripting. The attack is possible to be carried out remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.5	More Details
CVE-	A vulnerability was identified in SourceCodester Class and Exam Timetabling System 1.0. Affected by this		

2026-16203	issue is some unknown functionality of the file /forCYS.php. Such manipulation of the argument course leads to cross site scripting. The attack may be performed from remote. The exploit is publicly available and might be used.	3.5	More Details
CVE-2026-16202	A vulnerability was determined in SourceCodester Class and Exam Timetabling System 1.0. Affected by this vulnerability is an unknown functionality of the file /CYS.php. This manipulation of the argument course causes cross site scripting. The attack is possible to be carried out remotely. The exploit has been publicly disclosed and may be utilized.	3.5	More Details
CVE-2024-32389	Buffer Overflow vulnerability in Kerlink Kerlink Wirnet iStation 868 KerOS v.4.3.3_20200803132042 allows a remote attacker to obtain sensitive information via the update URLs component.	3.5	More Details
CVE-2026-16155	A vulnerability was identified in SourceCodester Class and Exam Timetabling System 1.0. Affected by this issue is some unknown functionality of the file /schoolyr.php. The manipulation of the argument sy leads to cross site scripting. It is possible to initiate the attack remotely. The exploit is publicly available and might be used.	3.5	More Details
CVE-2026-16156	A security flaw has been discovered in SourceCodester Class and Exam Timetabling System 1.0. This affects an unknown part of the file /forexam.php. The manipulation of the argument day results in cross site scripting. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks.	3.5	More Details
CVE-2026-54244	Statamic is a Laravel and Git powered content management system (CMS). Prior to 5.74.0 and 6.20.3, the Live Preview endpoint for existing entries and terms in src/Http/Controllers/CP/PreviewController.php only checked view authorization, but it accepts and renders caller-supplied field values. A Control Panel user with view but not edit permission could therefore submit content they were not authorized to author and generate a shareable Live Preview URL rendering it. This issue is fixed in versions 5.74.0 and 6.20.3.	3.5	More Details
CVE-2026-47087	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. URLAUTH does not honor revoked authorizer access. A URLAUTH URL minted while the authorizer had access continued to work after that access was revoked.	3.5	More Details
CVE-2026-47086	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. GENURLAUTH-issued tokens can bypass ACLs. Any authenticated user could mint a URLAUTH token (via the GENURLAUTH command) for any mailbox they could name, even without read access on it. This would allow reading mail from mailboxes despite having no granted permissions.	3.5	More Details
CVE-2026-12547	SoupAuthManager caches proxy authentication credentials without scoping them to the proxy authority (host:port). When the proxy configuration changes (e.g., via system settings or WPAD), cached Proxy-Authorization headers from the previous proxy are sent to the new proxy, leaking credentials.	3.4	More Details
CVE-2026-61859	ImageMagick before 7.1.2-26 and 6.9.13-x before 6.9.13-51 contains a policy bypass vulnerability in the -script operation due to missing security policy checks. This allows reading files from paths that are otherwise disallowed by the configured security policy.	3.3	More Details
CVE-2026-16213	A security flaw has been discovered in Fantomas42 django-blog-zinnia up to 0.20. Affected by this vulnerability is an unknown functionality of the file zinnia/views/mixins/entry_protection.py of the component Protected Entry Password Handler. The manipulation results in cleartext storage of sensitive information. The attack needs to be approached locally. The project was informed of the problem early through an issue report but has not responded yet.	3.3	More Details
CVE-2026-21953	Vulnerability in the Oracle Retail Xstore Point of Service product of Oracle Retail Applications (component: Xstore Mobile). The supported version that is affected is 21.0.3. Easily exploitable vulnerability allows low privileged attacker with logon to the infrastructure where Oracle Retail Xstore Point of Service executes to compromise Oracle Retail Xstore Point of Service. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Retail Xstore Point of Service accessible data. CVSS 3.1 Base Score 3.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:L/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N).	3.3	More Details
CVE-2026-10679	The DesignWare SPI driver (drivers/spi/spi_dw.c) computed the SPI BAUDR clock divider as info->clock_frequency / config->frequency without validating config->frequency. spi_transceive is a Zephyr __syscall and its verify handler (drivers/spi/spi_handlers.c) copies the caller-supplied spi_config from userspace without checking the frequency field, so a userspace thread that has been granted access to a DesignWare SPI device kernel object can pass frequency = 0 and trigger an unsigned integer divide-by-zero in spi_dw_configure(). On Cortex-M Mainline (SCB->CCR.DIV_0_TRP is set in z_arm_fault_init()) and on ARC (a dedicated __ev_div_zero vector) this raises a CPU exception, resulting in a kernel fault and local denial of service. The fix rejects zero frequency and frequencies above clock_frequency / 2 (the DesignWare SSI databook minimum SCKDIV of 2) with -EINVAL. The defect affects all Zephyr releases up to and including v4.4.0; exploitation requires CONFIG_USERSPACE=y and an unprivileged thread already granted SPI driver permission. There is no memory-corruption or information-disclosure impact.	3.3	More Details
CVE-			

2026-56375	ImageMagick through 7.1.2-18 contains a memory leak vulnerability in the ASHLAR coder when an action fails. Attackers can trigger failed actions to exhaust memory resources and cause denial of service.	3.3	More Details
CVE-2025-59866	The HCL DFMPro, DFXAnalytics and DFXServer installers are affected by 'Insecure file permissions Leading to Privilege Escalation' vulnerability, which enables any logged-in non-administrative user to overwrite or replace the executable file with a malicious binary.	3.3	More Details
CVE-2026-56586	HCL IEM was affected with X-Content-Type-Options Header Missing. It may enable attackers to perform SSL stripping or man-in-the-middle attacks and intercept sensitive data.	3.1	More Details
CVE-2026-35145	HCL DFXAnalytics is affected by a Missing HTTP Strict-Transport-Security Header vulnerability. The application fails to implement the HTTP Strict Transport Security (HSTS) policy within its responses, which could allow a remote attacker to downgrade the communication channel to an unencrypted connection (HTTP) and conduct man-in-the-middle (MitM) attacks. To remediate this, the application must include the "Strict-Transport-Security" header in all web application responses.	3.1	More Details
CVE-2026-56583	HCL MyCloud was affected with Concurrent Login Vulnerability. It may increase the risk of unauthorized access, session hijacking, and account misuse.	3.1	More Details
CVE-2026-15921	Node Version Manager (nvm) is a POSIX-compliant shell function for managing multiple node.js versions. In versions 0.32.1 through 0.40.5, `nvm ls-remote` (and other commands that refresh remote LTS aliases, such as `nvm install --lts`) parse the node.js mirror's `index.tab` and use each release's LTS codename field as an alias filename without validating it. A malicious, compromised, or man-in-the-middled mirror can return an LTS codename containing path-traversal sequences such as `../../../.bashrc`, causing nvm to write the associated version string to a path outside `$NVM_DIR/alias`. With the default layout (`$NVM_DIR` is `~/.nvm`), this can create or overwrite files in the user's home directory, including shell startup files, which can lead to code execution in a later shell session. Exploitation requires the victim to use a hostile mirror -- via a compromised mirror or CDN, a network man-in-the-middle, or a maliciously configured `NVM_NODEJS_ORG_MIRROR`/`NVM_IOJS_ORG_MIRROR` -- and to run an affected command. Version 0.40.6 validates remote LTS codenames as safe alias filenames and rejects `..` path components when writing alias files.	3.1	More Details
CVE-2026-62683	File Browser is a file managing interface for uploading, deleting, previewing, renaming, and editing files within a specified directory. Prior to 2.63.17, File Browser can leave a public directory share behind when the shared directory is deleted through a path with a trailing slash because the share cleanup path calls DeleteWithPathPrefix(file.Path, userID) and the Bolt backend performs the database prefix query with the unnormalized path before trimming the slash for boundary checks, so deleting /a/ does not delete the stored /a share and the stale public share exposes future content if the same path is recreated. This issue is fixed in version 2.63.17.	3.1	More Details
CVE-2026-47081	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. There is an XAPPLEPUSHSERVICE folder existence oracle and push hijack. An authenticated IMAP user could probe for the existence of arbitrary mailboxes on other users' accounts via the XAPPLEPUSHSERVICE command and then create Apple Push Notification Service notifications for new mail in those mailboxes to their own APNS device. This did not leak any data about the content of mailboxes. Instead, a "mailbox has changed" notice would be pushed when the mailbox modseq changed.	3.1	More Details
CVE-2026-56579	HCL MyCloud was affected with License Key Revealed in HTTP Response. It may enable attackers to misuse the exposed information and compromise the application's security.	3.1	More Details
CVE-2026-59849	A flaw was found in libssh. Logic errors in automatic certificate-based public key authentication can cause libssh clients to loop indefinitely when configured certificates are missing or repeatedly rejected by a server, leading to denial of service.	3.1	More Details
CVE-2026-56585	HCL IEM was affected with the Anti Clickjacking XFrame Options Header Missing. It may allow attackers to embed the application in malicious pages and induce unauthorized user actions.	3.1	More Details
CVE-2026-56582	HCL MyCloud was affected by the SSL/TLS LUCKY13 Vulnerability. An attacker may exploit this vulnerability to decrypt sensitive information through a TLS/SSL padding oracle attack.	3.1	More Details
CVE-2026-65054	MediaCMS 8.2.0 contains an information disclosure vulnerability that allows authenticated users to expose private media metadata belonging to other users by adding arbitrary media tokens to their own playlist without access control checks. Attackers can issue a PUT request to the playlist API endpoint with a known media token to bypass state and ownership validation, then retrieve the playlist to read private media fields including title, description, view count, like count, file size, author username, and encoding status through the unfiltered playlist owner branch in the playlist detail view.	3.1	More Details

CVE-2026-44970	dbt-mcp is a Model Context Protocol server for interacting with dbt. Prior to 1.17.1, DefaultUsageTracker.emit_tool_called_event() in src/dbt_mcp/tracking/tracking.py serialized every MCP tool call's complete arguments dictionary and sent it through dbtlabs_vortex.producer.log_proto without redaction, including sql_query from show, vars from run, build, and test, and node_selection from compile, while usage_tracking_enabled in settings.py enabled telemetry by default unless DBT_SEND_ANONYMOUS_USAGE_STATS=false or DO_NOT_TRACK=1 was set. This issue is fixed in version 1.17.1.	3.1	More Details
CVE-2026-21764	HCL DevOps Loop is affected by insufficient input validation that allows special characters where they should be restricted. This may result in unintended application behavior under certain conditions.	3.1	More Details
CVE-2026-56577	HCL MyCloud was affected with Weak Password Policy. It may increase the risk of account compromise through brute-force or credential-based attacks.	3.1	More Details
CVE-2026-47088	An issue was discovered in cyrus-imapd in Cyrus IMAP through 3.12.2. There is heap exposure in nested MIME comment parsing. An authenticated IMAP user could craft an email message containing an RFC 822 comment ending with a backslash. When parsing the message, the server would read past the message's end in memory, and read into the heap, returning the read content to the user.	3.1	More Details
CVE-2026-7364	IBM Verify Identity Access 11.0 through 11.0.2 and IBM Security Verify Access 10.0 through 10.0.9.1 and IBM Verify Identity Access Container 11.0 through 11.0.2 and IBM Security Verify Access Container 10.0 through 10.0.9.1 could allow a remote attacker to conduct phishing attacks, caused by an open redirect vulnerability. An attacker could exploit this vulnerability using a specially crafted request to redirect a victim to arbitrary Web sites.	3.1	More Details
CVE-2026-35143	HCL DFXAnalytics is affected by a Missing SameSite Attribute vulnerability. The application fails to set the "SameSite" attribute on session cookies generated during authentication, which could allow a remote attacker to execute Cross-Site Request Forgery (CSRF) attacks if additional mitigations, such as Anti-CSRF tokens, are not implemented.	3.0	More Details
CVE-2026-35140	HCL DFXAnalytics is affected by a Missing Secure Attribute in Encrypted Session (SSL) Cookie vulnerability. The application fails to set the "secure" attribute on session cookies generated during authentication, which could allow a remote attacker to intercept network traffic and capture sensitive cookies, session tokens, or credentials sent in cleartext over unencrypted channels.	3.0	More Details
CVE-2026-38755	A heap overflow in the evalcommand() function (shell/ash.c) of Busybox v1.38.0 allows attackers to cause a Denial of Service (DoS) via supplying a crafted input.	2.9	More Details
CVE-2026-38752	A stack overflow in the evaluate() function (editors/awk.c) of BusyBox commit 371fe9 allows attackers to cause a Denial of Service (DoS) via supplying a crafted AWK script.	2.9	More Details
CVE-2026-61862	ImageMagick before 7.1.2-26 and 6.9.13-51 contains an information disclosure vulnerability: when a profile is displayed with the identify command and the profile value is not printable, a single byte at the end of the profile can be printed (read past the profile boundary). This behavior occurs when debug output is enabled.	2.9	More Details
CVE-2026-61867	ImageMagick before 7.1.2-26 contains a memory leak vulnerability in the TIFF encoder when memory allocation fails. Attackers can trigger allocation failures during TIFF image processing to cause memory exhaustion and denial of service.	2.9	More Details
CVE-2026-61864	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a memory leak in color transformation to the log colorspace: when the operation fails, a small amount of memory is not released.	2.9	More Details
CVE-2026-61863	ImageMagick before 7.1.2-26 (and 6.x before 6.9.13-51) contains a memory leak in the TIFF encoder that occurs when a temporary file cannot be created, resulting in a small memory leak.	2.9	More Details
CVE-2026-61865	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a memory leak in the hough lines operation: when a specific operation fails, a small memory leak occurs.	2.9	More Details
CVE-2026-61866	ImageMagick before 7.1.2-26 contains a memory leak vulnerability in the JNG encoder when a blob cannot be opened. Attackers can trigger the memory leak by providing malformed JNG files that fail blob operations, causing resource exhaustion.	2.9	More Details
CVE-2026-61869	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a memory leak in the MIFF encoder that occurs when a memory allocation fails during MIFF image processing, which can lead to denial of service.	2.9	More Details

CVE-2026-11925	Tanium addressed a User Interface (UI) Misrepresentation of Critical Information vulnerability in Tanium Server.	2.7	More Details
CVE-2026-12907	The RTMKit WordPress plugin before 2.0.9 does not perform a proper capability check on one of its -builder AJAX actions, allowing users with at least the Author role to create and activate a site-wide template that overrides the header, footer or other global areas displayed to all visitors, which is normally restricted to administrators.	2.7	More Details
CVE-2026-10755	The All in One SEO WordPress plugin before 4.9.9 does not correctly restrict access to some of its AI integration REST API endpoints, allowing users with low-level privileges such as Contributors to overwrite or reset the site-wide AI integration state.	2.7	More Details
CVE-2026-12906	The RTMKit WordPress plugin before 2.0.9 does not perform a capability check in one of its AJAX actions and resolves a request-supplied post identifier directly, allowing users with at least the Contributor role to read the titles of other users' private, draft, pending, scheduled and trashed posts.	2.7	More Details
CVE-2026-16211	A vulnerability was determined in allegro up to bcf65b994ef29fb3fc2e10b660e6288723d5209e. This impacts the function AssetLastHostname.increment_hostname of the file src/ralph/assets/models/assets.py of the component Hostname Allocation Handler. Executing a manipulation of the argument counter can lead to race condition. Attacks of this nature are highly complex. The exploitability is said to be difficult. The exploit has been publicly disclosed and may be utilized. The project was informed of the problem early through an issue report but has not responded yet.	2.6	More Details
CVE-2026-35142	HCL DFXAnalytics is affected by an Internal IP Address Disclosure vulnerability. The application includes internal IP address details within its generated server responses, which could allow a remote attacker to gather sensitive network topology information and use it to map the internal infrastructure for further targeted attacks.	2.6	More Details
CVE-2026-47275	In nanomq versions 0.24.11 and earlier, a NULL pointer dereference in `nni_mqttv5_msg_decode_connect()` allows a malicious MQTT broker to crash any connecting NanoMQ MQTTv5 client (including bridge mode) with a single packet, causing remote denial of service via SIGSEGV. In `nni_mqttv5_msg_decode_connect()` (`mqtt_codec.c:1863`), the code iterates over CONNECT properties using variable `prop` when it should use `will_prop`. When a CONNECT packet has no connect-level properties (`prop == NULL`) but has will properties (`will_prop != NULL`), dereferencing `prop->next` causes SIGSEGV at address `0x38` (NULL + `offsetof(property, next)`). This affects both `nanomq_cli` and **NanoMQ bridge mode** (Core component), as both use the same `mqtt_client.c` receive path. This can lead to remote DoS if a malicious MQTT broker can crash the client process with a single 35-byte packet and persistent DoS if auto-reconnect causes infinite crash loop.	2.6	More Details
CVE-2026-56581	HCL MyCloud was affected with Cookie Attribute Path Not Set. It may increase the risk of unauthorized access to session data or authentication tokens.	2.6	More Details
CVE-2026-35141	HCL DFXAnalytics is affected by a Login Replay Attack vulnerability. The application allows a remote attacker to intercept, delay, or fraudulently retransmit valid authentication data to achieve unauthorized access. To mitigate this risk, the application must implement a mechanism to include timestamps with every message, ensuring that messages exceeding a specific age threshold are automatically rejected by the recipient system.	2.6	More Details
CVE-2026-16218	A vulnerability was detected in hunvreus devpush up to 0.4.6. Affected by this issue is the function reset_storage of the file app/workers/tasks/storage.py of the component Storage Reset Failure Handler. The manipulation results in improper check or handling of exceptional conditions. A high complexity level is associated with this attack. The exploitation is known to be difficult. The project was informed of the problem early through an issue report but has not responded yet.	2.6	More Details
CVE-2026-61872	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a memory leak in the TIFF encoder when an invalid tiff:tile-geometry is specified. Supplying malformed tile geometry parameters causes allocated memory not to be released, which can lead to increased memory consumption.	2.5	More Details
CVE-2026-44969	dbt-mcp is a Model Context Protocol server for interacting with dbt. Prior to 1.17.1, DbtMCP.call_tool() in src/dbt_mcp/mcp/server.py logged the raw arguments dictionary at INFO level before each tool call and at ERROR level on exceptions, and configure_file_logging() wrote those records to dbt-mcp.log when DBT_MCP_SERVER_FILE_LOGGING=true, preserving sensitive sql_query, vars, and node_selection values in plaintext without automatic rotation or deletion. This issue is fixed in version 1.17.1.	2.5	More Details
CVE-2026-16205	A weakness has been identified in Pluck CMS up to 4.7.21. This vulnerability affects the function htmlspecialchars_decode of the file data/modules/albums/albums.admin.php of the component Albums Module. Executing a manipulation of the argument Info can lead to cross site scripting. It is possible to launch the attack remotely. The exploit has been made available to the public and could be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	2.4	More Details
CVE-			

2026-56578	HCL MyCloud was affected by Server Version Disclosure. It may help attackers identify and exploit known vulnerabilities affecting the disclosed software versions.	2.2	More Details
CVE-2026-56580	HCL MyCloud was affected by Using Components with Known Vulnerability (IIS Server). It may allow attackers to exploit publicly disclosed weaknesses and compromise the system.	2.2	More Details
CVE-2026-61464	ImageMagick before 7.1.2-26 and 6.9.13-51 contains a heap-based buffer over-write vulnerability that occurs when running an X11 import with a crafted window title, which can result in heap memory corruption and denial of service.	1.8	More Details
CVE-2026-59235	Missing Authorization (CWE-862) in BankAccountListController (app/Http/Controllers/Api/BankAccount/BankAccountListController.php), exposed at GET /api/bank-account, in Prospero Flow CRM <5.5.3, which allows a remote, authenticated attacker holding a low-privileged role (e.g. the "User"/"Usuario" role) to read arbitrary bank account records belonging to their company by sending an authenticated request to the endpoint with a valid bearer token, because the API route is protected only by the auth:api middleware and carries no permission gate, unlike the equivalent web route, which enforces can('read bank'), and the handler resolves records with Account::where('company_id', Auth::user()->company_id)->get(), performing only company scoping and no role or permission check before returning the data. This results in the unauthorized disclosure of sensitive banking information (e.g. IBAN, SWIFT/BIC, account identifiers) to users who should not have access to it.	N/A	More Details
CVE-2026-47695	CC: Tweaked is a mod for Minecraft which adds programmable computers, turtles, and more to the game. Prior to version 1.119.0, CC-Tweaked's HTTP API (`http.request`, `http.websocket`) blocks requests to private network ranges to prevent server-side request forgery (SSRF). This protection can be bypassed on IPv6-capable servers using NAT64 well-known prefix addresses (`64:ff9b::/96`). An attacker who can execute Lua code can reach any internal IPv4 service that the filter is intended to block, by addressing it as `http://[64:ff9b::<ipv4-as-hex>]/` instead of its direct IPv4 address. This affects any CC-Tweaked deployment on a network with NAT64 routing — a configuration that is standard on AWS, GCP, and other cloud platforms when using IPv6-only subnets. Version 1.119.0 fixes the issue.	N/A	More Details
CVE-2023-37507	HCL DevOps Plan is susceptible to an information disclosure that can allow an attacker to focus their attacks based upon the information revealed.	N/A	More Details
CVE-2026-56349	n8n before version 2.10.0 contains an input validation vulnerability in the Guardrail node that allows attackers to bypass default guardrail instructions. End users can craft malicious inputs to circumvent guardrail protections and compromise workflow integrity.	N/A	More Details
CVE-2026-48978	oras-go is a Go library for managing OCI artifacts. Prior to 2.6.1, auth.Client follows the realm URL from a registry's WWW-Authenticate: Bearer challenge without validating the scheme or host, allowing a malicious or compromised registry to cause SSRF to internal networks such as http://169.254.169.254/, http://10.0.0.x/, and http://127.0.0.1/, or to downgrade a registry contacted over https:// to an http:// token endpoint in registry/remote/auth/client.go through Client.Do(), Client.fetchBearerToken(), fetchDistributionToken, and fetchOAuth2Token. This issue is fixed in version 2.6.1.	N/A	More Details
CVE-2023-37508	HCL DevOps Plan is potentially susceptible to Cross-Site Scripting (XSS) which could allow an attacker to exploit this vulnerability if certain browser weaknesses are present.	N/A	More Details
CVE-2026-57833	The Joomla extension 4Analytics is vulnerable to an unauthenticated stored XSS in relation to the AI analysis feature.	N/A	More Details
CVE-2026-47708	MCP-for-Stata is an MCP server for Stata to integrate Stata into an agent. Prior to version 1.17.3, the `log_file_name` parameter in the `stata_do` API and CLI is directly interpolated into a Stata command string without sanitization. The security guard (`GuardValidator`) only scans the do-file content but does not validate this parameter. An attacker can inject arbitrary Stata commands (including `shell`, `python`, `erase`, etc.) by crafting a malicious `log_file_name` containing quotes, newlines, or Stata command separators. Version 1.17.3 contains a patch for the issue.	N/A	More Details
CVE-2026-58077	The Joomla extension 4Analytics is vulnerable to an unauthenticated stored XSS. A specially crafted unauthenticated request may result in website takeover under some circumstances.	N/A	More Details
CVE-2026-8281	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-	An authenticated local file inclusion vulnerability exists in Sangoma Switchvox SMB Edition 8.3 (104997). The play_file functionality accepts user-controlled input through the sound_path parameter and fails to properly validate file paths before accessing the underlying filesystem. By supplying absolute paths, an authenticated	N/A	More Details

9587	attacker can retrieve files outside the intended directory scope.		
CVE-2026-45383	libde265 is an open source implementation of the h.265 video codec. Versions prior to 1.0.19 have a heap buffer overflow (out-of-bounds READ) exists in `decoder_context::decode_slice_unit_WPP()` in `libde265/decctx.cc`. When decoding a WPP (Wavefront Parallel Processing) HEVC slice, `ctbAddrRS` is computed as `ctbRow * ctbsWidth` inside the entry-point loop. If the PPS/SPS headers are crafted so that this value exceeds `pps.CtbAddrRStoTS.size()`, the subsequent array access `pps.CtbAddrRStoTS[ctbAddrRS]` reads past the end of the allocated vector, triggering a heap-buffer-overflow confirmed by AddressSanitizer. Version 1.0.19 patches the issue.	N/A	More Details
CVE-2026-65064	Data::HashMap::Shared versions before 0.14 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in shm_generic.h with open(path, O_RDWR | O_CREAT | O_CLOEXEC, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-16351	Sandbox escape due to use-after-free in the DOM: Navigation component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16350	Incorrect boundary conditions in the Audio/Video: cubeb component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-64614	Data::Deque::Shared versions before 0.06 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in deque.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-64615	Data::Graph::Shared versions before 0.04 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in graph.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-16349	Same-origin policy bypass in the DOM: Navigation component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-64616	Data::NDArray::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in ndarray.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-64627	Parse Server versions >= 9.0.0 before 9.10.0-alpha.4 and versions before 8.6.85 contain a schema disclosure vulnerability. When the GraphQL API is mounted with public introspection disabled (graphQLPublicIntrospection: false, the default), schema-derived 'Did you mean ...?' suggestions were still returned in GraphQL error messages produced during variable coercion, which were not covered by the introspection-hardening control (that only handled validation errors). An unauthenticated caller possessing only the public application id can iteratively recover hidden schema identifiers — including registered Cloud Code function names and Parse class and field names — by submitting queries or mutations whose variables contain near-miss enum values or input-object field names. This is a follow-up bypass of GHSA-8cph-rgr4-g5vj. The issue is fixed in 9.10.0-alpha.4 and 8.6.85.	N/A	More Details
CVE-2026-64617	Data::PubSub::Shared versions before 0.07 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in pubsub.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details

CVE-2026-60080	Use After Free vulnerability in the Rust deserialization logic of Apache Fory. This issue affects Apache Fory from 0.13.0 through 1.3.0. A crafted Fory payload could cause undefined behavior, process crash, or potential memory disclosure. Users are recommended to upgrade to version 1.4.0, which fixes the issue.	N/A	More Details
CVE-2019-25764	**UNSUPPORTED WHEN ASSIGNED** Exposed IOCTL with Insufficient Access Control in the ASUS AURA SYNC driver allows a local user to bypass the driver's verification and invoke arbitrary IOCTLs, resulting in privilege escalation. Refer to the 'End-of-Life Notice and Driver Update for Legacy ASUS Drivers ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-65061	Data::ReqRep::Shared versions before 0.05 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in reqrep.h with open(path, O_RDWR | O_CREAT, 0666), for both the request-reply and the integer-variant segments. The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-65062	Data::SortedSet::Shared versions before 0.03 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in sortedset.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-65063	Data::RadixTree::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in radix.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-65065	Data::RoaringBitmap::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in roaring.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-45704	Pimcore is an Open Source Data & Experience Management Platform. Prior to 11.5.17 (LTS) and 12.3.6, CustomReports uses inconsistent authorization between the report listing endpoint and the report detail endpoint in bundles/CustomReportsBundle/src/Controller/Reports/CustomReportController.php and bundles/CustomReportsBundle/src/Tool/Config/Listing/Dao.php, allowing a low-privileged backend user with the reports permission to directly request an unshared report such as poc-secret-report by name and read report name, grouping information, display and icon metadata, data source configuration, column configuration, and sharing settings even when shareGlobally is false. This issue is fixed in versions 11.5.17 (LTS) and 12.3.6.	N/A	More Details
CVE-2026-65066	Data::RingBuffer::Shared versions before 0.04 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in ring.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-65067	Data::Intern::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in intern.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
	Data::SpatialHash::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in sphash.h with open(path, O_RDWR|O_CREAT,		

CVE-2026-65068	0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-65069	Data::DisjointSet::Shared versions before 0.02 for Perl create a world-readable mmap backing file and open it without O_EXCL or O_NOFOLLOW. The segment is created in dsu.h with open(path, O_RDWR|O_CREAT, 0666). The mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable). O_NOFOLLOW is absent, so a symlink planted at the path is followed, and O_EXCL is absent, so the open silently uses a pre-planted file instead of failing. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted file or symlink at the path lets a local attacker win a pre-creation race or redirect the open.	N/A	More Details
CVE-2026-43636	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-8593	Improper permission enforcement in Checkmk versions 2.5.0 before 2.5.0p9, 2.4.0 before 2.4.0p34, 2.3.0 before 2.3.0p49, and 2.2.0 (EOL) allows users without permissions to view and modify BI packs and rules	N/A	More Details
CVE-2026-58148	The Joomla extension ChronoForms is vulnerable to an unauthenticated stored XSS vulnerability.	N/A	More Details
CVE-2026-60060	Improper Handling of Length Parameter Inconsistency (CWE-130) vulnerability exists in TTSSH2 plugin of Tera Term provided by TeraTerm Project. When Tera Term attempts to establish an SSH connection to a server set up by an attacker, out-of-bounds read/write may occur. As a result, the contents of adjacent memory regions may be transmitted to the server, and Tera Term may behave unexpected or terminate abnormally.	N/A	More Details
CVE-2026-30631	An issue was discovered in bytebot-ai in commit 3d37894ce07ef8d8b40adc7fd309ad96c2a71313 (2025-09-11) allowing attackers to execute arbitrary code via crafted path to `computer_write_file`.	N/A	More Details
CVE-2026-30633	Directory traversal vulnerability in knowns-dev/knowns 0.11.4 via crafted path value to the get_doc and update_doc tools.	N/A	More Details
CVE-2026-44974	@hapi/content provided HTTP Content-* headers parsing. Prior to 6.0.2, Content.disposition() retained the last occurrence of each duplicate parameter while Content.type() retained the first occurrence of duplicate charset and boundary parameters, creating a parameter-smuggling primitive when another component in the request-processing chain resolves duplicates the opposite way. This can allow an upload filename allowlist bypass in headers such as Content-Disposition: form-data; name="file"; filename="safe.txt"; filename="shell.php". This issue is fixed in version 6.0.2.	N/A	More Details
CVE-2026-45382	libde265 is an open source implementation of the h.265 video codec. Prior to version 1.0.19, `decoder_context::decode_slice_unit_tiles` (libde265/decctx.cc:920) reads `pps.CtbAddrRStoTS[ctbAddrRS]` at line 966 where `ctbAddrRS = ctbY * ctbsWidth + ctbX` is computed from PPS-supplied `colBd[]`/`rowBd[]` arrays without validating the result against `CtbAddrRStoTS.size() == sps->PicSizeInCtbsY`. A malformed PPS that passes `set_derived_values` but encodes geometry inconsistent with the SPS produces a `ctbAddrRS` past the allocation, causing a 4-byte heap-buffer-overflow READ. Version 1.0.19 fixes the issue.	N/A	More Details
CVE-2026-58317	Unsigned to Signed Conversion Error (CWE-196) vulnerability exists in TTSSH2 plugin of Tera Term provided by TeraTerm Project. When Tera Term attempts to establish an SSH connection to a server set up by an attacker, out-of-bounds read/write may occur. As a result, the contents of adjacent memory regions may be transmitted to the server, and Tera Term may behave unexpected or terminate abnormally.	N/A	More Details
CVE-2026-59776	Missing Cryptographic Step (CWE-325) vulnerability exists in certain FeliCa IC chips shipped in or before 2017. If the vulnerability is exploited, information stored in the IC chip may be read or tampered with.	N/A	More Details
CVE-2026-52476	SQL Injection vulnerability in aiflowy <= 2.1.2 allows a remote attacker to obtain sensitive information via the getPageData method in the DatacenterQuery.java file	N/A	More Details
CVE-2026-49852	joserfc is a Python library that provides an implementation of several JSON Object Signing and Encryption (JOSE) standards. Prior to 1.6.8, joserfc.jwt.decode accepts attacker-forged HMAC-signed tokens when the caller-supplied verification key is the empty string or None, because HMACAlgorithm.sign and HMACAlgorithm.verify in src/joserfc/_rfc7518/jws_algs.py pass the output of OctKey.get_op_key(...) to hmac.new(...) and OctKey.import_key in src/joserfc/_rfc7518/oct_key.py only emits a SecurityWarning for keys shorter than 14 bytes without rejecting zero-length input. This issue is fixed in version 1.6.8.	N/A	More Details

CVE-2026-44510	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-43620. Reason: This candidate is a duplicate of CVE-2026-43620. Notes: All CVE users should reference CVE-2026-43620 instead of this candidate.	N/A	More Details
CVE-2026-57831	The Joomla extension DP Calendar is vulnerable to an unauthenticated SQL injection.	N/A	More Details
CVE-2026-8982	Two undocumented privileged accounts exist in Autel Maxi Charger Single firmware through V1.03.51. The accounts use vendor-defined password derivation mechanisms based on device-specific values, allowing an attacker with knowledge of the algorithm and required inputs to authenticate to the web management interface with administrative privileges.	N/A	More Details
CVE-2026-8983	Autel Maxi Charger Single firmware through V1.03.51 contains a hard-coded authentication token that bypasses authorization checks for multiple management endpoints. An attacker can supply the special token value to invoke privileged functionality without valid authentication.	N/A	More Details
CVE-2026-42936	The installer of HYPER SBI 2 insecurely loads Dynamic Link Libraries. If there is a crafted DLL at the same directory when invoking the affected installer, arbitrary code may be executed with the privilege of the user invoking the installer.	N/A	More Details
CVE-2026-62238	OpenRemote before 1.26.0 contain an authenticated SQL injection vulnerability in the datapoint crosstab export endpoint that constructs PostgreSQL queries by concatenating asset display names into raw SQL. An authenticated attacker with asset creation or rename permissions can inject SQL through the asset name parameter and receive query results in the exported CSV response, enabling database data exfiltration.	N/A	More Details
CVE-2026-47134	ClearanceKit intercepts file-system access events on macOS and enforces per-process access policies. The ECDSA private key used to sign the on-disk policy database (`/Library/Application Support/clearancekit/store.db`) is stored in the macOS System Keychain. The key was created via the two-step pattern `SecKeyCreateRandomKey` (in-memory) followed by `SecItemAdd(kSecValueRef:, kSecAttrAccess:)` (persist). Prior to version 5.0.10, for `kSecClassKey` items in the legacy System Keychain, `kSecAttrAccess` passed to `SecItemAdd` is silently ignored — the persisted key inherits no ACL restriction. The same access builder applied to `kSecClassGenericPassword` items correctly binds the ACL, making this bug specific to the EC key. The result is that any process running as root can use the key to produce valid signatures over arbitrary policy content. Version 5.0.10 fixes the issue. No known workarounds are available. Disabling the system extension and manually removing the System Keychain item labelled `clearancekit policy signing key` would prevent the forged-signature path but also disables policy enforcement.	N/A	More Details
CVE-2026-47133	ClearanceKit intercepts file-system access events on macOS and enforces per-process access policies. Prior to version 5.0.10, each table in the on-disk SQLite policy store (`/Library/Application Support/clearancekit/store.db`) is verified using an ECDSA signature stored in the `data_signatures` table. The signed payload contains only the canonical row content, with no version counter or freshness binding. An attacker who can write `store.db` and the matching `data_signatures` row — feasible during the opfilter-update window when the Endpoint Security filter is offline, or via offline-boot / decrypted-backup scenarios — can substitute a previously-captured legitimately-signed snapshot. opfilter accepts the older snapshot as fully valid on next boot because the existing signatures still verify. Version 5.0.10 patches the issue.	N/A	More Details
CVE-2026-8920	Improper Restriction of Communication Channel to Intended Endpoints and External Control of File Name or Path in Aura Wallpaper Service allow a local user to perform file operations by sending crafted commands containing an arbitrary file path and bypassing the service’s path restrictions . On specific models , this can also cause a single feature to become unavailable . Refer to the ' Security Update for Aura Wallpaper Service ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-43945	FUXA is a web-based Process Visualization (SCADA/HMI/Dashboard) software. Versions 1.2.11 until 1.3.1 allow an unauthenticated remote attacker to achieve Full Remote Code Execution (RCE) as root. The exploit succeeds even when the platform is configured in its most secure state (Secure Mode Enabled and Node-RED Secure Auth Enabled). Version 1.3.1 fixes the issue.	N/A	More Details
CVE-2026-53727	css_parser is a Ruby CSS parser. From 2.2.0 until 3.0.0, CssParser::Parser#read_remote_file in lib/css_parser/parser.rb, and therefore load_uri! and the @import-following branch of add_block!, issued HTTP and HTTPS requests against any host, port, and URI without a scheme allowlist, host or IP filtering, or protection against link-local, loopback, or RFC-1918 addresses. Location: redirects were followed recursively back into the same function, which also serviced file:// URIs, so a single attacker-controlled HTTP redirect could upgrade the bug from SSRF to arbitrary local file disclosure. Any consumer of css_parser that hands it attacker-influenced CSS together with a base_uri: option is exposed. This issue is fixed in version 3.0.0.	N/A	More Details
CVE-2026-43946	FUXA is a web-based Process Visualization (SCADA/HMI/Dashboard) software. Version 1.3.0 has an authorization bypass in the /api/getTagValue endpoint allows unauthenticated access to tag values when the referenced script does not exist. Version 1.3.1 patches the issue.	N/A	More Details
CVE-	Cursor is a code editor built for programming with AI. Prior to the Cloud Agent fix on 03/31/2026, browser-enabled Cursor Cloud Agent sessions allowed attacker-controlled web content to connect from inside the		

2026-61613	agent container to an unauthenticated local agent endpoint, enabling code execution within the affected Cloud Agent sandbox or session and access to files, repository contents, environment variables, credentials, and GitHub App access tokens available to that session. This issue was fixed on 03/31/2026 by requiring authentication for the relevant agent endpoint.	N/A	More Details
CVE-2026-45784	rust-openssl provides OpenSSL bindings for the Rust programming language. From 0.10.50 until 0.10.80, CipherCtxRef::cipher_update_inplace in openssl/src/cipher_ctx.rs incorrectly sized output buffers when used with AES key-wrap-with-padding ciphers EVP_aes_{128,192,256}_wrap_pad. For a non-multiple-of-8 input, OpenSSL writes up to 7 bytes past the end of the caller's buffer or Vec, producing attacker-controllable heap corruption when the plaintext length is attacker-influenced. This issue is fixed in version 0.10.80.	N/A	More Details
CVE-2026-43947	FUXA is a web-based Process Visualization (SCADA/HMI/Dashboard) software. Version 1.3.0 has an unauthenticated Remote Code Execution vulnerability when `secureEnabled` is set to `true`. The `POST /api/runscript` endpoint checks authorization against the stored script's permission by ID, but when `test: true` is set in the request, it compiles and executes attacker-supplied code instead of the stored script's code. An unauthenticated attacker who knows a valid script ID and name may execute arbitrary code via test mode if at least one server-side script exists and is accessible without restrictive permissions. Script IDs and names can be obtained through the unauthenticated information disclosure in `GET /api/project` (reported separately). The only prerequisite is that at least one server-side script exists in the project. Version 1.3.1 fixes the issue.	N/A	More Details
CVE-2026-61646	FastGPT is a knowledge-based AI application platform. Prior to 4.15.0-beta5, FastGPT's shared SSRF guard validates only the initial request URL before handing the request to axios, and axios follows redirects by default. An authenticated workflow user can configure an HTTP request node to call an attacker-controlled public URL that redirects to cloud metadata, loopback, or internal services that the guard would block on direct request, and the HTTP node returns the response body to the workflow caller. This issue is fixed in version 4.15.0-beta5.	N/A	More Details
CVE-2026-64651	The `@ai-sdk/harness-opencode` tool connects HarnessAgent to OpenCode through a sandboxed bridge. Prior to version 1.0.28, the tool relay authorizes requests from any process whose command line contains an allowed helper script path (`host-tool-mcp.mjs`). This allows untrusted code executing in the sandbox to invoke arbitrary host-exposed tools including secret lookups, deployment operations, and cloud API calls without a corresponding model-authorized tool-call event. Exploitation requires a Linux environment (the vulnerable fallback checks `process.platform === 'linux'` and reads `/proc`); an active harness session with one or more host-provided tools; and untrusted code executing in the sandbox (e.g. a malicious dependency, build script, or lifecycle hook) The fix in version 1.0.28 removes the process-path authorization fallback entirely. Relay requests are now only accepted after exact, short-lived, one-time authorization matching the tool name and input from a bridge-observed model event. Some workarounds are available. Do not run the OpenCode harness on untrusted repositories or with untrusted dependencies, and/or limit host-exposed tools to non-sensitive operations when working with untrusted code.	N/A	More Details
CVE-2026-64650	The `@ai-sdk/harness-opencode` tool is an HarnessV1 adapter backed by @openai/codex-sdk, which drives the codex command line interface. Prior to version 1.0.29, the tool relay authorizes requests from any process whose command line contains an allowed helper script path (the Codex CLI shim). This allows untrusted code executing in the sandbox to invoke arbitrary host-exposed tools, including secret lookups, deployment operations, and cloud API calls without a corresponding model-authorized tool-call event. Exploitation requires a Linux environment (the vulnerable fallback checks `process.platform === 'linux'` and reads `/proc`); an active harness session with one or more host-provided tools; and untrusted code executing in the sandbox (e.g. a malicious dependency, build script, or lifecycle hook) The fix in version 1.0.29 removes the process-path authorization fallback entirely. Relay requests are now only accepted after exact, short-lived, one-time authorization matching the tool name and input from a bridge-observed model event. Some workarounds are available. Do not run the Codex harness on untrusted repositories or with untrusted dependencies, and/or limit host-exposed tools to non-sensitive operations when working with untrusted code.	N/A	More Details
CVE-2026-8919	Permissive Cross-domain Security Policy with Untrusted Domains in ASUS GameSDK allows a remote user to obtain a local user’s NTLM hash by convincing the user to visit a crafted web page that sends a request containing a UNC path to the application’s local service endpoint. This can result in information disclosure or data tampering, may cause GameSDK to become unavailable, and may also enable access to the victim’s information on other services. Refer to the ' Security Update for ASUS GameSDK ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-15030	Out-of-bounds Read in ASUS System Control Interface v3, ASUS System Control Interface, and ASUS Business Manager allows a local administrator to read memory regions beyond the intended firmware boundary by supplying a crafted IOCTL request that bypasses the validation. Refer to the ' Security Update for ASUS System Control Interface ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-15029	Untrusted Pointer Dereference in ASUS System Control Interface v3, ASUS System Control Interface, and ASUS Business Manager allows a local administrator to perform arbitrary physical memory read and write operations via crafted IOCTL requests to the driver, bypassing OS-enforced memory protections. Refer to the ' Security Update for ASUS System Control Interface ' section on the ASUS Security Advisory for more information.	N/A	More Details

CVE-2026-13585	Allocation of Resources Without Limits and Throttling and Sensitive Information in Resource Not Removed Before Reuse in the ASUS System Control Interface driver and ASUS Business Manager allow a local administrator to disclose sensitive information via crafted IOCTL requests, which, in severe cases, may lead to a Denial of Service (DoS) on the system. Refer to the ' Security Update for ASUS System Control Interface ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-13385	An Improper Validation of Integrity Check Value and Improper Certificate Validation in certain ASUS router models allows a remote man-in-the-middle(MITM) user to make the router download and execute arbitrary command via a spoofed server. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-61684	FastGPT is a knowledge-based AI application platform. In 4.15.0-beta4, FastGPT plugin invoke reverse-call endpoints under /api/invoke/* authenticate only by verifying a JWT signed with INVOKE_TOKEN_SECRET, which defaults to the constant string token and was not set in official deployment templates. An unauthenticated attacker can self-sign an HS256 JWT and reach /api/invoke/userInfo to disclose cross-tenant user PII by attacker-supplied tmbId values, or /api/invoke/fileUpload to write attacker-controlled content into chat files. This issue is fixed in version 4.15.0-beta5.	N/A	More Details
CVE-2026-11851	Improper Neutralization of Special Elements used in an SQL Command ("SQL Injection") in the web management interface of certain ASUS router models allows a remote authenticated user to disclose confidential information via a crafted request that bypasses existing input validation Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-9592	SEPPmail Secure Email Gateway & SEPPmail Cloud before version 15.0.4.2 allows an attacker to replay & hijack a user session in the GINA web portal, as the session token is disclosed inside the URL and a HTTP header.	N/A	More Details
CVE-2026-57832	The Joomla extension EDocman is vulnerable to an unauthenticated SQL injection.	N/A	More Details
CVE-2026-50162	oras-go is a Go library for managing OCI artifacts. Prior to 2.6.1, resolveWritePath() in content/file/file.go uses a lexical filepath.Rel check for workingDir and does not account for symlink traversal, so when AllowPathTraversalOnWrite=false an attacker-controlled blob title through ocispec.AnnotationTitle such as out/pwn.txt can follow a workingDir symlink out -> /some/outside/dir and cause pushFile() to create /some/outside/dir/pwn.txt outside workingDir. This issue is fixed in version 2.6.1.	N/A	More Details
CVE-2026-16353	Invalid pointer in the DOM: Bindings (WebIDL) component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-50197	Skipper is an HTTP router and reverse proxy for service composition. Prior to 0.26.10, zalando/skipper's OpenPolicyAgent integration silently bypasses request-body inspection on HTTP/1.1 Transfer-Encoding: chunked and HTTP/2 requests that omit the content-length pseudo-header, because the opaAuthorizeRequestWithBody filter and OpenPolicyAgentInstance.ExtractHttpBodyOptionally in filters/openpolicyagent/openpolicyagent.go produce an empty raw_body and input.parsed_body while the upstream service receives the full attacker-controlled body. This issue is fixed in version 0.26.10.	N/A	More Details
CVE-2026-52469	SQL injection vulnerability in Crocus v.1.3.44 allows a remote attacker to escalate privileges via the DeviceInfoMapper.xml file	N/A	More Details
CVE-2026-52470	SQL injection vulnerability in Crocus v.1.3.44 allows a remote attacker to escalate privileges via the RecordStateMapper.xml file	N/A	More Details
CVE-2026-50289	systeminformation is a System and OS information library for node.js. Prior to 5.31.7, networkInterfaces() on Linux is vulnerable to OS command injection through the Debian/Ubuntu interfaces(5) source directive because lib/network.js checkLinuxDCHPInterfaces() reads /etc/network/interfaces, extracts a source <path> token from file content, and interpolates it unquoted into cat ${file} 2> /dev/null | grep 'iface|source' executed by execSync(cmd, util.execOptsLinux), allowing a path containing shell metacharacters to execute commands in any process that calls networkInterfaces(), including via getStaticData() and getAllData(). This issue is fixed in version 5.31.7.	N/A	More Details
CVE-2026-52472	SQL injection vulnerability in Wgcloud 3.6.4 allows a remote attacker to escalate privileges via the PortInfoMapper.xml file	N/A	More Details
CVE-2026-54463	websocket-driver is a WebSocket protocol handler with pluggable I/O. Prior to 0.8.1, draft versions of the WebSocket protocol in websocket-driver include a length header that allows an arbitrarily large integer to be encoded as bytes with the high bit set, and a server or client can send an indefinite sequence of 0x80 or higher bytes that the peer parses into an ever-growing Ruby integer. This can make a WebSocket connection	N/A	More Details

	consume an unbounded amount of memory and lead to the host process running out of memory. This issue is fixed in version 0.8.1.		
CVE-2026-62175	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-60091. Reason: This candidate is a duplicate of CVE-2026-60091. Notes: All CVE users should reference CVE-2026-60091 instead of this candidate.	N/A	More Details
CVE-2026-52474	An issue in aiflowy <= 2.1.2 allows a remote attacker to obtain sensitive information via the JobUtil.java file.	N/A	More Details
CVE-2026-54464	### Impact If this library is used in tandem with the `permessage-deflate` extension, a WebSocket server or client can be made to accept messages that are larger than the configured maximum message size. This is because this limit is checked against the message frames' length headers, which give the size of the compressed data, not the size after decompression. This can lead to applications accepting larger messages than expected and exceeding their intended resource usage. ### Patches The issue has been patched in version 0.8.1, by checking the length of messages after they are processed by incoming extensions. All users should upgrade to this version. ### Workarounds No known workarounds exist. ### Acknowledgements This issue was discovered and reported by Pranjali Thakur, DepthFirst Security Research Team.	N/A	More Details
CVE-2026-52475	Cross Site Scripting vulnerability in aiflowy <= 2.1.2 allows a remote attacker to obtain sensitive information via the UploadController.java file	N/A	More Details
CVE-2026-54465	websocket-driver is a WebSocket protocol handler with pluggable I/O. Prior to 0.8.1, when websocket-driver is used to implement a WebSocket server on top of a TCP server using WebSocket::Driver.server() or to complement a WebSocket client, a peer can make a single connection consume an unbounded amount of memory by sending an HTTP request or response with a never-ending list of headers. This can lead to the receiving process running out of memory. This issue is fixed in version 0.8.1.	N/A	More Details
CVE-2026-15783	A missing authorization vulnerability was identified in GitHub Enterprise Server that allowed an authenticated user with write access to any repository to read metadata from private repositories they did not have access to, including private repository owners and names, branch names, commit SHAs, commit messages, and the pushing actor. The delegated bypass endpoint resolved a rule suite directly from an attacker-supplied, encoded identifier without verifying that the requesting user could read the rule suite's repository, and because these identifiers are sequential an attacker could enumerate them across the instance. This vulnerability affected all versions of GitHub Enterprise Server prior to 3.22 and was fixed in versions 3.17.18, 3.18.12, 3.19.9, 3.20.5, and 3.21.3. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-12715	Missing Authorization in Google Cloud Firebase Studio versions prior to 2026-04-15 on Google Cloud Platform allows an attacker to download other users' deployed source code and access sensitive data via unauthorized GCS URL signing requests. This vulnerability was patched on 15 April 2026, and no customer action is needed.	N/A	More Details
CVE-2026-15343	A path traversal vulnerability was identified in GitHub Enterprise Server that allowed an attacker who had code execution inside the Dependabot updater container to write files to arbitrary repository paths, including GitHub Actions workflow files under .github/workflows/ as the path validation did not check the effective path which the attacker could control through the dependency file's directory and symlink target. If the repository used a pull_request_target workflow or had auto-merge enabled, an injected workflow could execute with access to the repository's GitHub Actions secrets. This vulnerability affected all versions of GitHub Enterprise Server prior to 3.22 and was fixed in versions 3.21.3, 3.20.5, 3.19.9, 3.18.12, 3.17.18.	N/A	More Details
CVE-2026-15904	Use after free in Ozone in Google Chrome on Linux prior to 150.0.7871.128 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	N/A	More Details
CVE-2026-15903	Out of bounds read and write in V8 in Google Chrome prior to 150.0.7871.128 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	N/A	More Details
CVE-2026-15902	Use after free in Cast in Google Chrome prior to 150.0.7871.128 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	N/A	More Details
CVE-2026-15901	Use after free in Network in Google Chrome prior to 150.0.7871.128 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Critical)	N/A	More Details
CVE-2026-15900	Use after free in GPU in Google Chrome on Android prior to 150.0.7871.128 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	N/A	More Details

CVE-2026-15899	Use after free in CameraCapture in Google Chrome on Mac prior to 150.0.7871.128 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	N/A	More Details
CVE-2026-15007	A denial of service vulnerability was identified in GitHub Enterprise Server that allowed an authenticated user to cause service disruption by supplying a repository release notes configuration file containing deeply nested YAML. When release notes were generated, the configuration file was parsed without a nesting depth limit, causing excessive resource consumption that could render the instance unresponsive. This vulnerability affected all versions of GitHub Enterprise Server prior to 3.22 and was fixed in versions 3.17.18, 3.18.12, 3.19.9, 3.20.5, and 3.21.3. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-14871	osTicket versions v1.18.3 and v1.17.7 contain a Broken Object Level Authorization (BOLA) leading to Insecure Direct Object Reference (IDOR) in the AJAX ticket-management subsystem.	N/A	More Details
CVE-2026-57495	AgenticMail gives AI agents real email addresses and phone numbers. In @agenticmail/claudecode prior to version 0.2.39, @agenticmail/codex prior to version 0.1.33, @agenticmail/core prior to version 0.9.43, and @agenticmail/openclaw prior to version 0.5.71, two inbound-mail handlers act on a privileged effect without verifying that the sender is the operator, while a sibling handler in the same repo does. The higher-impact one: any external email routed to the bridge inbox causes the dispatcher to resume the operator's Claude Code session with `permissionMode: 'bypassPermissions'`, embedding the attacker-controlled `from`/`subject`/`preview` verbatim into the prompt the resumed agent reads — an indirect prompt injection into a fully-privileged agent (Bash/Write/Edit/WebFetch + the agenticmail MCP toolbelt) running as the operator's OAuth identity. The sibling operator-query email-reply hook gates the same untrusted-From provenance with `isOperatorReplySender(replyFrom, config.operatorEmail)` (fail-closed); the bridge-wake path — a strictly higher-privilege effect — has no equivalent. @agenticmail/claudecode 0.2.39, @agenticmail/codex 0.1.33, @agenticmail/core 0.9.43, and @agenticmail/openclaw 0.5.71 contain a fix.	N/A	More Details
CVE-2026-57494	AgenticMail gives AI agents real email addresses and phone numbers. In @agenticmail/api prior to version 0.9.64, a low-privileged authenticated AgenticMail agent can enumerate another agent's pending/claimed tasks by supplying the target agent name to `GET /api/agenticmail/tasks/pending?assignee=<name>`. The returned task objects include the task IDs and payloads. The same task IDs can then be used with the capability-style task mutation endpoints (`/tasks/:id/claim`, `/tasks/:id/result`, `/tasks/:id/complete`, `/tasks/:id/fail`) to claim, complete, or fail tasks assigned to a different agent. Because ordinary authenticated agents can discover agent names through `GET /api/agenticmail/accounts/directory`, the task ID effectively stops being a secret capability. This turns the intended capability model into a cross-agent authorization bypass. Version 0.9.64 contains a fix.	N/A	More Details
CVE-2026-16352	Sandbox escape due to use-after-free in the Disability Access APIs component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-59143	Data::RoaringBitmap::Shared versions before 0.02 for Perl allow an out-of-bounds read via an unvalidated container offset and cardinality in rb_contains_locked. The attach-time validator rb_validate_header checks the header scalars and region layout against the file size, but does not validate the bucket contents it then trusts. rb_contains_locked forms a container pointer as pool + container_off * 8192 from a raw file-stored offset and then searches over a file-stored cardinality, neither bounded against the container pool capacity or the fixed 8192-byte slot size. A local peer that can write the backing file can leave the header valid while poisoning a bucket, so the next membership query dereferences a file-controlled wild pointer and scans a file-controlled count, reading adjacent memory or crashing the process.	N/A	More Details
CVE-2026-64613	Data::Buffer::Shared versions before 0.05 for Perl create a world-readable mmap backing file and open it without O_NOFOLLOW. The segment is created in buf_generic.h with open(path, O_RDWR|O_CREAT|O_EXCL, 0666). O_EXCL blocks a pre-seeded file on create, but the mode is 0666, so under the default umask 022 the file is created mode 0644 (world-readable), and O_NOFOLLOW is absent, so a symlink planted at the path is followed when the segment is attached. A "Shared" segment naturally lives in a shared directory such as /tmp or /dev/shm, where any local user can read the IPC payloads stored in the world-readable segment, and a pre-planted symlink at the path redirects the open to another file.	N/A	More Details
CVE-2026-9585	An unauthenticated reflected cross-site scripting (XSS) vulnerability exists in Sangoma Switchvox SMB Edition version 8.3 (104997). The application fails to properly sanitize the portal parameter supplied to the invalid_browser and invalid_browser_login handlers. User-supplied data is reflected into JavaScript generated by the application, allowing attacker-controlled script execution within a victim's browser.	N/A	More Details
CVE-2026-21577	This High severity DoS (Denial of Service) vulnerability was introduced in versions 9.0.1, 9.1.0, 9.2.0, 9.3.1, 9.4.0, 9.5.1, 10.0.2, 10.1.0 and 10.2.0 of Confluence Data Center. This DoS (Denial of Service) vulnerability, with a CVSS Score of 7.1, allows an authenticated attacker to cause a resource to be unavailable for its intended users by temporarily or indefinitely disrupting services of a host connected to a network. Atlassian recommends that Confluence Data Center customers upgrade to latest version, if you are unable to do so, upgrade your instance to one of the specified supported fixed versions: Confluence Data Center 9.2: Upgrade to a release greater than or equal to 9.2.17 Confluence Data Center 10.2: Upgrade to a release greater than	N/A	More Details

	or equal to 10.2.7 See the release notes ([https://confluence.atlassian.com/doc/confluence-release-notes-327.html]). You can download the latest version of Confluence Data Center from the download center ([https://www.atlassian.com/software/confluence/download-archives]). This vulnerability was reported via our Penetration Testing program.		
CVE-2026-53712	SCRAM (Salted Challenge Response Authentication Mechanism) is part of the family of Simple Authentication and Security Layer (SASL, RFC 4422) authentication mechanisms. Prior to 3.3, a flaw in com.ongres.scram:scram-client and com.ongres.scram:scram-common allows an attacker capable of a TLS man-in-the-middle attack to silently downgrade a connection from SCRAM-SHA-256-PLUS with channel binding to standard SCRAM-SHA-256 without channel binding when TlsServerEndpoint processes an X.509 certificate using a modern signature algorithm such as Ed25519; getChannelBindingData() can return an empty byte array after NoSuchAlgorithmException, and the ScramClient builder treats that as absent channel-binding data. This issue is fixed in version 3.3.	N/A	More Details
CVE-2026-15379	The Altiris WMI provider exposes a class (AltirisAgent_Stream) that allows any local standard user to read the contents of any file accessible to the SYSTEM account, bypassing filesystem ACLs. No admin privileges required. The provider reverts to the LocalSystem context when servicing WMI queries without re-impersonating the caller. Any local standard user can therefore read SYSTEM-readable files — including configuration files, service logs, and secrets stored with SYSTEM/Administrator-only ACLs — by querying the provider directly.	N/A	More Details
CVE-2026-21579	This High severity Information Disclosure vulnerability was introduced in versions 7.17.0, 7.19.0, 8.5.0, 8.9.0, 9.0.1, 9.1.0, 9.2.0, 10.0.2, 10.1.0, and 10.2.0 of Confluence Data Center. This Information Disclosure vulnerability, with a CVSS Score of 8.2, allows an unauthenticated attacker to view sensitive information via an Information Disclosure vulnerability. Atlassian recommends that Confluence Data Center customers upgrade to latest version, if you are unable to do so, upgrade your instance to one of the specified supported fixed versions: Confluence Data Center 9.2: Upgrade to a release greater than or equal to 9.2.22 Confluence Data Center 10.2: Upgrade to a release greater than or equal to 10.2.14 See the release notes ([https://confluence.atlassian.com/doc/confluence-release-notes-327.html]). You can download the latest version of Confluence Data Center from the download center ([https://www.atlassian.com/software/confluence/download-archives]). This vulnerability was reported via our Atlassian (Internal) program.	N/A	More Details
CVE-2026-16410	JIT miscompilation in the JavaScript Engine: JIT component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16409	Invalid pointer in the Security: PSM component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16407	Mitigation bypass in the DOM: Service Workers component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16404	Spoofing issue in Firefox for Android. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16403	Spoofing issue in the Address Bar component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16402	Integer overflow in the Graphics: ImageLib component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16401	Privilege escalation in the Data Loss Prevention component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16400	Information disclosure in the DOM: Security component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-21575	This High severity RCE (Remote Code Execution) vulnerability was introduced in version 3.4.11 of Sourcetree for Mac and Sourcetree for Windows. This RCE (Remote Code Execution) vulnerability, with a CVSS Score of 7.1, allows an authenticated attacker to execute arbitrary code which has high impact to confidentiality, high impact to integrity, high impact to availability, and requires user interaction. Atlassian recommends that Sourcetree for Mac and Sourcetree for Windows customers upgrade to latest version, if you are unable to do so, upgrade your instance to one of the specified supported fixed versions: * Sourcetree for Mac and Sourcetree for Windows 3.4: Upgrade to a release greater than or equal to 3.4.13 See the release notes (https://www.sourcetreeapp.com/download-archives). You can download the latest version of Sourcetree for	N/A	More Details

	Mac and Sourcetree for Windows from the download center (https://www.sourcetreeapp.com/download-archives). This vulnerability was reported via our Bug Bounty program.		
CVE-2026-59252	Improper Validation of Specified Quantity in Input in ZenHive mpp allows an unauthenticated remote client to drain the fee-payer wallet, resulting in denial of service for legitimate clients. When the mpp Elixir library is configured as fee payer (fee_payer: true), the MPP.Methods.Tempo payment method co-signs and broadcasts a client-supplied EVM transaction without first validating that the client-supplied gas_limit is sufficient to complete the intended call. A malicious client can submit a signed transferWithMemo transaction with gas_limit deliberately set just below the amount required for successful execution. The server co-signs the transaction and broadcasts it via rpc_broadcast_sync. The transaction runs out of gas during EVM execution and reverts, but the fee-payer wallet is still charged for the burned gas while the client pays nothing and receives no resource. Repeated requests from one or more malicious clients drain the fee-payer wallet at near-zero cost to the attacker, ultimately preventing the server from sponsoring gas for legitimate payment requests. The wait_for_confirmation = false (optimistic) path is also affected: it invokes simulate_payment_call via eth_call, but that simulation omits the gas parameter and therefore does not catch out-of-gas conditions. This issue affects mpp: from 0.2.0 before 0.6.0.	N/A	More Details
CVE-2026-59694	Improper Validation of Specified Quantity in Input in ZenHive mpp allows an unauthenticated remote client to inflate the fee-payer's gas cost per payment by a large multiplier, degrading the sponsor's operating margin. When the mpp Elixir library is configured as fee payer (fee_payer: true), MPP.Tempo.Transaction.cosign_fee_payer/3 re-signs the client-supplied base fields of the 0x76 AASigned envelope verbatim, including the EIP-2930 access list, without validating its length or contents. EIP-2930 access list entries incur intrinsic gas (~2,400 gas per address, plus 1,900 gas per storage key) charged before any opcode executes, regardless of whether the listed addresses are ever touched. A malicious client submits a valid transferWithMemo call alongside a large number of fabricated access-list entries. The server co-signs and broadcasts the transaction. The intended transfer executes normally, but the fee-payer wallet pays a large multiple of the expected gas cost with no corresponding on-chain work. At the maintainer's default of 137 access-list entries (fitting within Bandit's 10,000-byte per-header-field limit) and 100 Gwei max_fee_per_gas, per-payment gas cost rises from ~51,287 to ~380,087 gas, a 7.4x multiplier. Sustained abuse destroys the sponsor's operating margin on low-cost payments and, over time, drains the fee-payer wallet. This issue affects mpp: from 0.2.0 before 0.6.0.	N/A	More Details
CVE-2026-47425	Rattler is a library that provides common functionality used within the conda ecosystem. Prior to version 0.43.2, `EntryPoint::FromStr` in `rattler_conda_types` performs only `.trim()` on the `command` field before the linker joins it onto the install prefix and writes an executable Python script. A malicious `noarch:python` package can ship an `info/link.json` with an entry-point name containing `..`, `/`, `\`, or an absolute path; the resulting file is written outside the prefix (or clobbers an existing in-prefix entry-point such as `bin/pip`) with mode `0o775` on Unix and a copied launcher `.exe` on Windows. This affects the default install path of `pixi install`, `mamba install` via py-rattler, `rattler-build`, and any other consumer of the `rattler` install crate; no flag or post-link-script opt-in is involved. Version 0.43.2 contains a fix for the issue.	N/A	More Details
CVE-2026-47657	HumHub is an Open Source Enterprise Social Network. In versions 1.13.0 through 1.18.2, a missing authorization check in the Space member management controller allowed any authenticated user to trigger the removal of all members from any Space, regardless of their own role or membership in that Space. Versions 1.13.0 through 1.18.2 are affected. The vulnerability has been patched in version 1.18.3, and all users are encouraged to upgrade to this version or later immediately. No known workaround is available.	N/A	More Details
CVE-2026-16399	Site isolation issue in the DOM: Navigation component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16398	Site isolation issue in the Graphics component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16396	Privilege escalation in WebExtensions. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16393	Incorrect boundary conditions in the Graphics: WebGPU component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16392	JIT miscompilation in the JavaScript Engine: JIT component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16391	Information disclosure in the Storage: IndexedDB component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-	Mitigation bypass in the Enterprise Policies component. This vulnerability was fixed in Firefox 153 and Firefox		More

2026-16390	ESR 140.13.	N/A	Details
CVE-2025-66390	In Microsoft Azure API Management through 2025-10-17, when self-service signup (username/password Basic Authentication) is enabled in Tenant A, an attacker can reuse the registration flow by changing the hostname or tenant identifier to Tenant B, even when Tenant B has signup disabled at the UI level. In other words, disabling signup in the UI does not disable the underlying API endpoint (which still accepts cross-tenant requests based on the Host header). NOTE: The supplier states that they evaluated the report and determined it did not cross a security boundary (i.e., the observed behavior was a configuration/state issue rather than an exploitable product vulnerability affecting tenant isolation). NOTE: The supplier evaluated this report and determined that it did not cross a security boundary (i.e., the observed behavior was a configuration/state issue rather than an exploitable product vulnerability affecting tenant isolation).	N/A	More Details
CVE-2026-15380	A non-administrator interactive user can obtain full SYSTEM code execution through a DCOM/task scheduler logic chain — no network access, no memory corruption required (ITMS 8.7.3)	N/A	More Details
CVE-2026-16441	In Eclipse OpenJ9 versions up to 0.60, when executing class files where a previously concrete superclass method has been recompiled as abstract, execution is incorrectly delegated to an interface default method.	N/A	More Details
CVE-2026-9586	An unauthenticated SQL injection vulnerability exists in Sangoma Switchvox SMB Edition 8.3 (104997). The /pa endpoint processes XML content beginning with <PolycomIPPhone> and directly concatenates the user-controlled PhoneIP value into PostgreSQL queries without sanitization or parameterization. An unauthenticated remote attacker can execute arbitrary SQL statements against the backend PostgreSQL database using a single crafted request, including database operations and remote code execution.	N/A	More Details
CVE-2026-15793	BuildKit custom frontends or clients using the raw low-level API can set git.checkoutbundle=true when checking out Git sources. If the Git source is malicious, this could lead to a crafted command invocation on the host.	N/A	More Details
CVE-2026-15792	A malicious BuildKit client or frontend could craft a request that could lead to BuildKit daemon crashing with a panic.	N/A	More Details
CVE-2026-47397	PraisonAI is a multi-agent teams system. Prior to version 4.6.40, hidden metadata in a webpage causes PraisonAI agents to write attacker-controlled content to arbitrary paths. `write_file` skips path validation when `workspace=None` (always `None` in production). Version 4.6.40 fixes the issue.	N/A	More Details
CVE-2026-15791	A crafted message in the BuildKit low-level build API can be used to remove the contents of the /tmp directory. The action that can normally be used to delete files inside the build container rootfs can escape into the real host temp directory.	N/A	More Details
CVE-2026-15789	A custom client can produce such an upload request to the BuildKit daemon that files can escape from the BuildKit-controlled state directory. The client needs to have valid permissions to access the BuildKit control API to issue builds, e.g., bypass authentication, etc.	N/A	More Details
CVE-2026-15432	When verifying a mac with a ChunkedMacVerification object, Tink compares the resulting tag with non constant time comparison. This potentially allows an attacker to use timinig information as a side channel in order to get information how many bytes of a given tag match the correct tag. This in turn could allow to find a correct tag bytewise.	N/A	More Details
CVE-2026-15342	Plane contains a multi-tenant authorization flaw in its asset-management API that allows authenticated users from one workspace to access, delete, or duplicate assets belonging to another workspace by providing only the victim workspace slug and asset ID. The affected endpoints return presigned file URLs and enable destructive or duplicative actions without verifying that the requester is a member of the targeted workspace. This enables cross-tenant data exposure, data deletion, and persistent exfiltration of files into an attacker-controlled workspace.	N/A	More Details
CVE-2025-68640	The Apple Find My backend service through 2025-12-17 allows an attacker in possession of a valid PET (Private Endpoint Token) to enumerate devices and remove offline devices from an Apple ID account without triggering two-factor authentication or ownership verification. This may result in unauthorized removal of devices associated with the account.	N/A	More Details
CVE-2026-9588	A stored cross-site scripting (XSS) vulnerability exists in Sangoma Switchvox SMB Edition 8.3 (104997) within the voicemail notification template functionality. The submit_modify_voicemail_template endpoint fails to properly sanitize HTML content supplied by authenticated users, allowing malicious JavaScript supplied through the template_text parameter to be stored server-side and subsequently rendered to other users.	N/A	More Details
	PraisonAI is a multi-agent teams system. Prior to version 4.6.40, the fix for GHSA-9mqq-jqxf-grvw / CVE-2026-44336 is incomplete. The original advisory description named four vulnerable handlers in `mcp_server/adapters/cli_tools.py`. Commit `68cc9427` ("fix(security): harden MCP rules path handling…")		

CVE-2026-47394	added a `_resolve_rule_path()` helper and applied it to `rules.create`, `rules.show`, and `rules.delete`. `workflow.show` was left unchanged. Two adjacent handlers in the same file have the same pattern, `workflow.validate` and `deploy.validate`. Neither was mentioned in the original advisory. Both remained unchanged. The original advisory also identified the dispatcher (`server.py:281-298`) as a root cause. It accepts unvalidated `**kwargs` from `params["arguments"]` with no enforcement against the tool's declared `input_schema`. That code is unchanged prior to version 4.6.40. A single unauthenticated MCP `tools/call` to `praisonai.workflow.show` returns the contents of any file the host user can read: `/etc/passwd`, `~/.ssh/id_rsa`, `~/.aws/credentials`, or any project `.env`. Version 4.6.40 contains an updated fix.	N/A	More Details
CVE-2026-47407	PraisonAI Platform is the platform layer for the PraisonAI multi-agent teams system. Prior to version 0.1.4, the Platform server exposes resources under `/api/v1/workspaces/{workspace_id}/...` and protects them with a `require_workspace_member(workspace_id)` FastAPI dependency. The dependency only checks that the caller is a member of the workspace_id in the URL prefix. The route handlers then look up the inner resource (`agent_id`, `issue_id`, `project_id`, `label_id`, `comment_id`, `dependency_id`) by primary key alone. The resource's own `workspace_id` is never compared to the URL's `workspace_id`. A user can therefore put their own workspace in the URL prefix and any other workspace's resource ID in the path. The auth check passes, since they really are a member of the prefix workspace. The service then returns the cross-tenant resource for read, update, or delete. There is a second bug in the member-management routes (`add_member`, `update_member_role`, `remove_member`, `update_workspace`, `delete_workspace`). Each one inherits the default `min_role="member"` from `require_workspace_member`. Any basic member can therefore promote themselves to admin or owner, demote or remove other members, and delete the workspace. The role hierarchy exists in the schema but is not enforced. Registration is open at `/api/v1/auth/register` with no email verification. The default server bind is `0.0.0.0:8000` (`python -m praisonai_platform`). One curl from any unauthenticated network position is enough to bootstrap into the system. PraisonAI Platform version 0.1.4 patches the issue.	N/A	More Details
CVE-2026-22104	Improper access control in Hashtopolis server web-interface chunk activity component for versions prior to 0.14.8 allows any created account to read all cracked hashes of a Hashtopolis server instance.	N/A	More Details
CVE-2026-50185	RustCrypto CMOV provides conditional move CPU intrinsics which are guaranteed on major platforms to execute in constant-time and not be rewritten as branches by the compiler. From 0.1.1 until 0.5.4, the aarch64 implementations of Cmov and CmovEq in cmov/src/backends/aarch64.rs assume high bits are zero-extended when loading values smaller than a register, so set high bits such as [8..] in a Cmov selector or [16..] of self or other in the u16 and i16 CmovEq implementations can cause left.cmovz(&right, condition) to produce incorrect output. This issue is fixed in version 0.5.4.	N/A	More Details
CVE-2026-62764	Improper Handling of Insufficient Privileges vulnerability in Apache Accumulo. An authenticated, but low-privileged user without system permissions may issue a remote command to gracefully shutdown system components (compaction-coordinator, compactor, gc, manager, monitor, tserver, or sserver), leading to a denial of service. This issue affects Apache Accumulo 2.1.4 and 2.1.5. Users are recommended to upgrade to version 2.1.6, which fixes the issue.	N/A	More Details
CVE-2026-46458	ICU Scandinavia Boomerang is vulnerable to an information disclosure flaw where sensitive credential files are exposed via static HTTP. This allows an unauthenticated remote attacker to retrieve plaintext service account and SMTP credentials by requesting specific XML files from the webroot. This issue has been fixed in version 2.4.18.029	N/A	More Details
CVE-2026-61740	LightRAG provides simple and fast retrieval-augmented generation. Prior to 1.5.4, when LightRAG is deployed with LIGHTRAG_API_KEY set but AUTH_ACCOUNTS unset, X-API-Key protection can be bypassed because lightrag/api/auth.py falls back to a hardcoded DEFAULT_TOKEN_SECRET, /auth-status and /login can mint guest JWTs, and combined_dependency in lightrag/api/utils_api.py accepts a valid guest token before checking the API key. A remote unauthenticated attacker can call endpoints guarded by combined_auth, including document read, upload, deletion, graph mutation, and query endpoints. This vulnerability is fixed in 1.5.4.	N/A	More Details
CVE-2026-16243	In Eclipse OMR versions up to 0.11, the arraycmp SIMD implementation for Z and P does not check if the number of bytes to compare is zero.	N/A	More Details
CVE-2026-16439	In Eclipse OpenJ9 versions up to 0.60, using -Xtrace to trace method arguments can lead to buffer underflow.	N/A	More Details
CVE-2026-46681	@nevware21/ts-utils is a comprehensive TypeScript/JavaScript utility library. Prior to version 0.14.0, the _copyProps function in lib/src/object/copy.ts uses for...in to iterate over source object properties without an Object.hasOwnProperty check, and does not filter dangerous keys (__proto__, constructor, prototype). This allows an attacker to pollute the prototype chain of all objects in the application. Version 0.14.0 patches the issue.	N/A	More Details
CVE-2026-	ICU Scandinavia Boomerang is vulnerable to a missing authentication flaw in its device receiver endpoints. This allows an unauthenticated remote attacker to read full facility configurations and write unauthorized	N/A	More

46459	data to the sensor database. This issue has been fixed in version 2.4.18.029		Details
CVE-2026-11876	In zenml-io/zenml version 0.94.2, the `GET /api/v1/stack-deployment/stack` endpoint (`get_deployed_stack`) lacks proper RBAC authorization checks, allowing any authenticated user to enumerate all deployed stacks across all users and tenants. This includes stack component details, service connector information, and user IDs of stack owners. The vulnerability arises from two issues: missing endpoint-level RBAC checks and the use of a server-side `Client()` that bypasses the RBAC enforcement layer by directly accessing the database through `SqlZenStore`. This exposes sensitive information such as infrastructure topology, service connector details, stack ownership, and deployment metadata, potentially enabling cross-tenant reconnaissance and further attacks in multi-tenant ZenML Pro/Cloud deployments.	N/A	More Details
CVE-2026-45309	AsyncSSH is a Python package which provides an asynchronous client and server implementation of the SSHv2 protocol on top of the Python asyncio framework. Prior to 2.23.0, AsyncSSH expands the OpenSSH-compatible AuthorizedKeysFile %u token in asyncssh/config.py, asyncssh/connection.py, asyncssh/auth_keys.py, and asyncssh/misc.py with the raw SSH username during pre-authentication server config reload, allowing a server configured with AuthorizedKeysFile authorized_keys/%u to read an authorized-keys file outside the intended directory when the SSH username contains /, \, or .. path traversal segments and authenticate with an attacker-selected key file. This issue is fixed in version 2.23.0.	N/A	More Details
CVE-2026-9499	An out-of-bounds read (buffer over-read) vulnerability exists in QTextCodec::codecForName() in Qt. When the function is called with a QByteArray that is not NUL-terminated (for example, one created with QByteArray::fromRawData()), the codec-name matching routine reads past the end of the supplied buffer. In most cases this results in an incorrect text codec being selected; in the worst case, if the over-read reaches unmapped memory, the process crashes (denial of service). The over-read is bounded by the length of the longest codec-name candidate, and the out-of-bounds bytes are only compared internally against Qt's fixed list of codec names, so no data is disclosed to an attacker. Applications that do not pass non-NUL-terminated QByteArrays to QTextCodec::codecForName() are not exposed. The affected code resides in the Qt5Compat module from Qt 6.0.0 onward, and in Qt Core (qtbase) in Qt 4.x and Qt 5.x.	N/A	More Details
CVE-2026-48487	Zeroconf is a pure Python implementation of multicast DNS service discovery. Prior to 0.149.16, _read_character_string and _read_string in src/zeroconf/_protocol/incoming.py advanced self.offset by attacker-declared RDLENGTH without checking it against self._data_len, allowing unauthenticated hosts on the local link over UDP/5353 (224.0.0.251 / ff02::fb) to send a TXT, HINFO, or A/AAAA record with rdlength=65535 and seed DNSCache and ServiceInfo.properties with truncated, attacker-shaped key/value or address records. This issue is fixed in version 0.149.16.	N/A	More Details
CVE-2026-16389	Incorrect boundary conditions, integer overflow in the Libraries component in NSS. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16388	Sandbox escape in the DOM: Networking component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-59147	Data::DisjointSet::Shared versions before 0.02 for Perl allow out-of-bounds reads and writes via an unvalidated parent index in dsu_find. The attach-time validator dsu_validate_header checks the header scalars and region layout against the file size, but does not validate the array contents it then trusts. dsu_find walks and path-compresses parent[x] with x a raw file-stored index never bounded against the node count, so both the read and the compression write-back land out of bounds. A local peer that can write the backing file can leave the header valid while poisoning the parent array, so the next find or union both reads and writes through an out-of-bounds parent index, corrupting memory or crashing the process.	N/A	More Details
CVE-2026-59144	Data::RingBuffer::Shared versions before 0.04 for Perl allow a stack buffer overflow via an unvalidated elem_size in ring_read_seq. The attach-time validator ring_validate_header checks the capacity-overflow and total_size consistency of the header but never caps elem_size against the destination size. ring_read_seq does memcpy(out, ring_slot(h, seq), elem_size) with elem_size read raw from the mmap'd segment, copying into a fixed 8-byte destination scalar. An elem_size larger than 8 bytes writes past the destination. A local peer that can write the backing file can leave the header valid while setting a large elem_size, so the next read copies a file-controlled length into the fixed 8-byte stack buffer, corrupting adjacent stack frames.	N/A	More Details
CVE-2026-16371	Privilege escalation in the DOM: Navigation component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-50757	Directory Traversal vulnerability in DayuanJiang next-ai-draw-io 0.4.13 allowsa remote attacker to execute arbitrary code via the nex-ai-draw-io/mcp-server	N/A	More Details
CVE-2026-50758	Cross Site Scripting vulnerability in DayuanJiang next-ai-draw-io 0.4.13 allows a remote attacker to execute arbitrary code via the mcp parameter	N/A	More Details

CVE-2026-50759	An issue in exo-explore exo 1.0.69 allows a remote attacker to escalate privileges via the GET /state and DELETE /instance/{instance_id} endpoints with no authentication.	N/A	More Details
CVE-2026-16370	Mitigation bypass in the DOM: Networking component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16369	Integer overflow in the JavaScript: WebAssembly component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16368	Incorrect boundary conditions in the JavaScript: WebAssembly component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16367	Sandbox escape due to invalid pointer in the Disability Access APIs component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-56852	A norm.Iter can enter an infinite loop when handling input containing invalid UTF-8 bytes.	N/A	More Details
CVE-2026-15829	A SQL injection (CWE-89) and security boundary bypass (CWE-863) vulnerability exists in the prebuilt BigQuery forecasting tool (bigquery-forecast) of googleapis/mcp-toolbox. The tool accepts client-controlled parameters (data_col, timestamp_col, and id_cols) as plain strings and interpolates them unescaped via fmt.Sprintf directly into a generated AI.FORECAST table-valued SELECT statement. While MCP Toolbox utilizes an allowedDatasets mechanism to restrict queries, this defense only validates the history_data parameter; the final assembled query is executed without re-validation. An attacker can break out of the string literal fields (such as timestamp_col) to inject a valid multi-statement or cross-dataset query block. This allows an unauthorized user to bypass the operator-configured allowedDatasets boundary and read arbitrary BigQuery tables.	N/A	More Details
CVE-2026-16366	Privilege escalation in the DOM: Navigation component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-59145	Data::Intern::Shared versions before 0.02 for Perl allow an out-of-bounds read via unvalidated slot, reverse and arena indices in si_idx_find. The attach-time validator si_validate_header is thorough about the header and layout (magic, version, section offsets, total_size, count and arena_used) but does not validate the three arrays it then trusts. Every lookup in si_idx_find walks a triple indirection read straight from the mmap'd segment, arena[reverse[slots[i].id]], with no bound on slots[i].id against count, on the reverse[id] arena offset against arena_used, or on the arena record's length prefix. A local peer that can write the backing file can leave the header valid while poisoning a slot id, a reverse offset or an arena length prefix, so an id_of, intern or string lookup dereferences the chain out of bounds; because string() returns a file-controlled length of bytes from the arena, adjacent process memory can be disclosed.	N/A	More Details
CVE-2026-16373	Information disclosure in the Privacy component in Firefox for Android. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16365	Privilege escalation in the DOM: Workers component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-59146	Data::SpatialHash::Shared versions before 0.02 for Perl allow out-of-bounds reads and writes via unvalidated bucket, link and free-list indices in sph_walk_cell and sph_alloc_slot. The attach-time validator sph_validate_header checks the header scalars and region layout against the file size, but does not validate the array contents it then trusts. sph_walk_cell reads entries[buckets[b]] and follows each entry's next link raw, and sph_alloc_slot writes through a file-stored free_head index, none bounded against the entry count (max_entries). A local peer that can write the backing file can leave the header valid while poisoning the bucket chain and free list, so a query reads through an out-of-bounds bucket and next index and an insert writes through an out-of-bounds free-list head, corrupting memory or crashing the process.	N/A	More Details
CVE-2026-16364	Incorrect boundary conditions in the Audio/Video: Playback component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16363	JIT miscompilation in the JavaScript: WebAssembly component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details

CVE-2026-16362	Use-after-free in the WebRTC: Audio/Video component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16360	Memory safety bugs present in Firefox ESR 115.37, Firefox ESR 140.12 and Firefox 152. Some of these bugs showed evidence of memory corruption and we presume that with enough effort some of these could have been exploited to run arbitrary code. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16358	Site isolation issue in the Graphics: WebRender component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16357	Incorrect boundary conditions in the Graphics component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16356	Sandbox escape due to use-after-free in the Disability Access APIs component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16355	JIT miscompilation in the JavaScript Engine: JIT component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16354	Information disclosure in the Graphics: ImageLib component. This vulnerability was fixed in Firefox 153, Firefox ESR 115.38, and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16372	Privilege escalation in the DOM: Content Processes component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-50756	An issue in DayuanJiang next-ai-draw-io 0.4.13 allows a remote attacker to obtain sensitive information via the x-ai-provider component	N/A	More Details
CVE-2026-55081	DHIS2 is a flexible information system for data capture, management, validation, analytics and visualization. The DHIS2 OpenAPI HTML endpoint reflected values from the `scope` query parameter into the generated HTML document without sufficient sanitization. A crafted `scope` value could be rendered as active HTML or JavaScript in the OpenAPI documentation page. An attacker able to get a user to open a crafted OpenAPI HTML URL could execute JavaScript in that user's browser in the DHIS2 origin. Affected versions: DHIS2 2.42 and 2.43 before the 2026-06-09 security patch releases, and the development branch for DHIS2 2.44 before the fix was merged. Patched in 2.42.5.1, 2.43.0.1, the 2.42 and 2.43 line branches, and the 2.44 development branch.	N/A	More Details
CVE-2026-59142	Data::HashMap::Shared versions before 0.14 for Perl allow an out-of-bounds read via an unvalidated arena offset and length in shm_str_copy. The attach-time validator shm_validate_header checks the header scalars and region layout against the file size, but does not validate the array contents it then trusts. shm_str_copy does memcpy(dst, arena + off, len) with off and len read raw from the mmap'd segment and unbounded, on the each, keys, values, pop, shift, take, swap, drain and cursor paths. The get path bounds off and len separately and is not affected. A local peer that can write the backing file can leave the header valid while poisoning a record's offset and length, so iterating or draining the map copies a file-controlled offset and length out of the arena, reading adjacent memory or crashing the process.	N/A	More Details
CVE-2026-16387	Site isolation issue in the Networking component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-55082	DHIS2 is a flexible information system for data capture, management, validation, analytics and visualization. DHIS2 SQL View data endpoints allowed authenticated users with SQL View access to provide crafted filter values that were interpolated into generated SQL. An authenticated user with access to SQL View execution could manipulate SQL generated for SQL View filters and potentially access data outside the intended SQL View result set. This is distinct from CVE-2026-55084, which tracks the related SQL View filter column-name injection. Known affected release lines for this advisory: DHIS2 2.37, 2.38, and 2.39 before the 2026-06-09 EOS security updates. Patched by the 2026-06-09 EOS security updates for 2.37, 2.38, and 2.39. The same value-slot hardening was already present on later supported branches through DHIS2-20174 / PR #22253.	N/A	More Details
CVE-2026-16386	Information disclosure due to uninitialized memory in the Graphics: WebGPU component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-			

2026-16385	Information disclosure due to uninitialized memory in the Graphics: WebGPU component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16384	Information disclosure due to uninitialized memory in the Graphics: WebGPU component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-59695	Improper Validation of Specified Quantity in Input in ZenHive mpp allows an unauthenticated remote client to drain the fee-payer wallet in a single request by naming an arbitrarily high gas price. When the mpp Elixir library is configured as fee payer (fee_payer: true), MPP.Tempo.Transaction.cosign_fee_payer/3 re-signs the client-supplied base fields of the 0x76 AASigned envelope verbatim, including max_fee_per_gas and max_priority_fee_per_gas, without validating that they are within reasonable bounds. A malicious client embeds arbitrarily large values for these fields in the signed envelope. The server co-signs and broadcasts the transaction. The effective_gas_price billed against the fee-payer wallet is derived from the attacker-supplied ceilings, so the server pays those inflated per-gas rates out of its own wallet. A single crafted request can drain the wallet entirely, after which the server can no longer sponsor gas for legitimate payment requests. This issue affects mpp: from 0.2.0 before 0.6.0.	N/A	More Details
CVE-2026-16383	Mitigation bypass in the DOM: Networking component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-59139	Data::ReqRep::Shared versions before 0.05 for Perl allow an out-of-bounds read via an unvalidated arena offset and length in reqrep_recv_locked. The attach-time validator reqrep_validate_header checks the header scalars and region layout against the file size, but does not validate the array contents it then trusts. reqrep_recv_locked does memcpy(copy_buf, req_arena + arena_off, len) with arena_off and len read raw from the mmap'd segment and never bounded against the arena capacity (req_arena_cap). A local peer that can write the backing file can leave the header valid while poisoning a request slot's offset and length, so receiving the request copies a file-controlled offset and length out of the arena, reading adjacent memory or crashing the process.	N/A	More Details
CVE-2026-59140	Data::SortedSet::Shared versions before 0.03 for Perl allow an out-of-bounds read via unvalidated node indices in the rank and min/max query paths. The attach-time validator ss_validate_header bounds only the root index against the node pool (node_capacity). The order-statistics and min/max queries then follow children[], leftmost and rightmost node indices read raw from the mmap'd segment without bounding them against node_capacity. A full structural check (ss_validate_tree) exists but runs only via an explicit validate method, not on attach. A local peer that can write the backing file can leave the header valid while poisoning the tree links, so the next rank, min or max query dereferences an out-of-bounds node index, reading adjacent memory or crashing the process.	N/A	More Details
CVE-2026-59141	Data::RadixTree::Shared versions before 0.02 for Perl allow an out-of-bounds read via unvalidated node and arena indices in rdx_find_locked. The attach-time validator rdx_validate_header checks the header scalars and region layout against the file size, but does not validate the node records it then trusts. rdx_find_locked indexes nodes[cur].children[k] and reads each node's label_off and label_len raw from the mmap'd segment, none bounded against the node count or the arena size. A local peer that can write the backing file can leave the header valid while poisoning the node records, so a lookup dereferences an out-of-bounds node or arena index, reading adjacent memory or crashing the process.	N/A	More Details
CVE-2026-16382	Mitigation bypass in the DOM: Service Workers component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-59254	n8n before 2.28.1 contains an information disclosure vulnerability where external secrets are incorrectly resolved in workflow node expressions outside credentials scope. Authenticated project editors can read plaintext external secret values by referencing them in node expressions without requiring explicit secrets access permissions.	N/A	More Details
CVE-2026-50755	An issue in DayuanJiang next-ai-draw-io 0.4.13 allows a remote attacker to obtain sensitive information via the X-Forwarded-For header value	N/A	More Details
CVE-2026-59236	Authorization Bypass Through User-Controlled Key (CWE-639) in the Excel import handlers (CustomerImport, LeadImport, ProductImport) in Roskus Prospero Flow CRM before 5.14.0 allows a remote, authenticated user of any role or company to create customer, lead, and product records inside another company's tenant via a spreadsheet whose company_id column points to the victim tenant, uploaded to POST /customer/import/excel/save, which maps company_id directly from the file and performs no check that it matches the authenticated user's company.	N/A	More Details
CVE-2026-30632	Directory traversal vulnerability in knowns-dev/knowns 0.11.4 via crafted folder name value to the create_doc tool.	N/A	More Details

CVE-2026-46600	Parsing an invalid SVCB or HTTPS RR can panic when the size of a parameter value overflows the message buffer.	N/A	More Details
CVE-2026-16381	Same-origin policy bypass in the Networking: DNS component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16380	Mitigation bypass in the Networking component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16379	Privilege escalation in the DOM: Content Processes component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16378	Other issue in the DOM: Copy & Paste and Drag & Drop component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16377	Mitigation bypass in the PDF Viewer component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16376	Denial-of-service in the Graphics: WebGPU component. This vulnerability was fixed in Firefox 153.	N/A	More Details
CVE-2026-16375	Site isolation issue in the Networking: HTTP component. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-16374	Information disclosure in the Framework component in DevTools. This vulnerability was fixed in Firefox 153 and Firefox ESR 140.13.	N/A	More Details
CVE-2026-63878	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: check num_entries in GEM_OP GET_MAPPING_INFO kvcalloc(args->num_entries, sizeof(*vm_entries), GFP_KERNEL) at amdgpu_gem.c:1050 uses the user-supplied num_entries directly without any upper bounds check. Since num_entries is a __u32 and sizeof(drm_amdgpu_gem_vm_entry) is 32 bytes, a large num_entries produces an allocation exceeding INT_MAX, triggering WARNING in __kvmalloc_node_noprof(), causing a kernel WARNING, TAINT_WARN, and panic on CONFIG_PANIC_ON_WARN=y systems. Add a size bounds check before we invoke the kvzalloc() to reject oversized num_entries early with -EINVAL. (cherry picked from commit 1fe7bf5457f6efd7be60b17e23163ba54341d73d)	N/A	More Details
CVE-2026-64188	In the Linux kernel, the following vulnerability has been resolved: net: qualcomm: rmnet: fix endpoint use-after-free in rmnet_dellink() rmnet_dellink() removes the endpoint from the hash table with hlist_del_init_rcu() and then immediately frees it with kfree(). However, RCU readers on the receive path (rmnet_rx_handler -> __rmnet_map_ingress_handler) may still hold a reference to the endpoint and dereference ep->egress_dev after the memory has been freed. The endpoint is a kmalloc-32 object, and the stale read at offset 8 corresponds to the egress_dev pointer. BUG: unable to handle page fault for address: ffffffffde942eef Oops: 0002 [#1] SMP NOPTI CPU: 1 UID: 0 PID: 137 Comm: poc_write Not tainted 7.0.0+ #4 PREEMPTLAZY RIP: 0010:rmnet_vnd_rx_fixup (rmnet_vnd.c:27) Call Trace: <TASK> __rmnet_map_ingress_handler (rmnet_handlers.c:48 rmnet_handlers.c:101) rmnet_rx_handler (rmnet_handlers.c:129 rmnet_handlers.c:235) __netif_receive_skb_core.constprop.0 (net/core/dev.c:6096) __netif_receive_skb_one_core (net/core/dev.c:6208) netif_receive_skb (net/core/dev.c:6467) tun_get_user (drivers/net/tun.c:1955) tun_chr_write_iter (drivers/net/tun.c:2003) vfs_write (fs/read_write.c:688) ksys_write (fs/read_write.c:740) </TASK> Add an rcu_head field to struct rmnet_endpoint and replace kfree() with kfree_rcu() so the endpoint memory remains valid through the RCU grace period. Also remove the rmnet_vnd_dellink() call and inline only the nr_rmnet_devs decrement, since rmnet_vnd_dellink() would set ep->egress_dev to NULL during the grace period, creating a data race with lockless readers.	N/A	More Details
CVE-2026-54466	websocket-driver is a WebSocket protocol handler with pluggable I/O. Prior to 0.7.5, the frame format in draft versions of the WebSocket protocol includes a length header that allows an arbitrarily large integer to be encoded as a sequence of bytes with the high bit set. By sending an indefinite sequence of bytes with values 0x80 or above, a client can make the server parse these bytes into an ever-growing integer in lib/websocket/driver/draft75.js; because JavaScript numbers are 64-bit floating point values, this number will eventually lose precision and lead to the subsequent payload being parsed incorrectly. This issue is fixed in version 0.7.5.	N/A	More Details
CVE-	Qinglong is a timed task management platform supporting Python3, JavaScript, Shell, and Typescript. Prior to 2.20.1, the init guard middleware in back/loaders/express.ts checks /api/user/init but not /open/user/init,		More

2026-55445	while rewrite('/open/*', '/api/$1') rewrites the whitelisted /open/* path after JWT authentication and the guard have passed; an unauthenticated attacker can send PUT /open/user/init to reset administrator credentials on an initialized instance. This issue is fixed in 2.20.1.	N/A	Details
CVE-2026-64043	In the Linux kernel, the following vulnerability has been resolved: ovpn: fix race between deleting interface and adding new peer While deleting an existing ovpn interface, there is a very narrow window where adding a new peer via netlink may cause the netdevice to hang and prevent its unregistration. It may happen during ovpn_dellink(), when all existing peers are freed and the device is queued for deregistration, but a CMD_PEER_NEW message comes in adding a new peer that takes again a reference to the netdev. At this point there is no way to release the device because we are under the assumption that all peers were already released. Fix the race condition by releasing all peers in ndo_uninit(), when the netdevice has already been removed from the netdev list. Also ovpn_peer_add() has now an extra check that forces the function to bail out if the device reg_state is not REGISTERED. This way any incoming CMD_PEER_NEW racing with the interface deletion routine will simply stop before adding the peer. Note that the above check happens while holding the netdev_lock to prevent racing netdev state changes. ovpn_dellink() is now empty and can be removed.	N/A	More Details
CVE-2026-64049	In the Linux kernel, the following vulnerability has been resolved: drm/msm/adreno: fix userspace-triggered crash on a2xx-a4xx Before a5xx Adreno driver will not try fetching UBWC params (because those generations didn't support UBWC anyway), however it's still possible to query UBWC-related params from the userspace, triggering possible NULL pointer dereference. Check for UBWC config in adreno_get_param() and return sane defaults if there is none. Patchwork: https://patchwork.freedesktop.org/patch/717778/	N/A	More Details
CVE-2026-63822	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: fix warning when unbinding If there is an error during some initialization related to firmware, the buffers dp->tx_ring[i].tx_status are released. However this is released again when the device is unbinded (ath11k_pci), and we get: WARNING: CPU: 0 PID: 6231 at mm/slub.c:4368 free_large_kmalloc+0x57/0x90 Call Trace: free_large_kmalloc ath11k_dp_free ath11k_core_deinit ath11k_pci_remove ... The issue is always reproducible from a VM because the MSI addressing initialization is failing. In order to fix the issue, just set the buffers to NULL after releasing in order to avoid the double free.	N/A	More Details
CVE-2026-63175	PlaywrightCapture stored capture-specific configuration and runtime data as mutable class-level variables rather than instance-level variables. Consequently, multiple Capture objects running within the same Python process could share state, including HTTP headers, cookies, browser storage, HTTP credentials, proxy configuration, user-agent settings, geolocation information, and captured request data. In a multi-user or concurrent deployment, information supplied during one capture could therefore persist and be reused by a subsequent or parallel capture. This could result in the disclosure of authentication cookies, credentials, browser storage, or captured request data belonging to another user. It could also cause requests to be performed with another capture's authentication context, headers, or proxy configuration, potentially enabling unauthorized access to remote resources or interference with other capture operations. The vulnerability is resolved by initializing all capture-specific settings and request data as instance variables in the Capture constructor, ensuring that state is isolated between capture operations.	N/A	More Details
CVE-2026-64052	In the Linux kernel, the following vulnerability has been resolved: block: bio-integrity: Fix null-ptr-deref in bio_integrity_map_user() pin_user_pages_fast() can partially succeed and return the number of pages that were actually pinned. However, the bio_integrity_map_user() does not handle this partial pinning. This leads to a general protection fault since bvec_from_pages() dereferences an unpinned page address, which is 0. To fix this, add a check to verify that all requested memory is pinned. If partial pinning occurs, unpin the memory and return -EFAULT. Kernel Oops: Oops: general protection fault, probably for non-canonical address 0xdffffc0000000001: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f] CPU: 0 UID: 0 PID: 1061 Comm: nvme-passthroug Not tainted 7.0.0-11783-g90957f9314e8-dirty #16 PREEMPT(lazy) Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.17.0-0-gb52ca86e094d-prebuilt.qemu.org 04/01/2014 RIP: 0010:bio_integrity_map_user.cold+0x1b0/0x9d6	N/A	More Details
CVE-2026-63821	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: usb: fix memory leaks on USB write failures When rtw_usb_write_port() fails to submit a USB Request Block (URB) (e.g., due to device disconnect or ENOMEM), the completion callback is never executed. Currently, the driver ignores the return value of rtw_usb_write_port() in rtw_usb_write_data() and rtw_usb_tx_agg_skb(). Because these functions rely on the completion callback to free the socket buffers (skbs) and the transaction control block (txcb), a submission failure results in: 1. A memory leak of the allocated skb in rtw_usb_write_data(). 2. A memory leak of the txcb structure and all aggregated skbs in rtw_usb_tx_agg_skb(). Fix this by checking the return value of rtw_usb_write_port(). If it fails, explicitly free the skb in rtw_usb_write_data(), and properly purge the tx_ack_queue and free the txcb in rtw_usb_tx_agg_skb(). The issue was discovered in practice during device disconnect/reconnect scenarios and memory pressure conditions. Tested by verifying normal TX operation continues after the fix without regressions.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: net: shaper: reject duplicate leaves in GROUP request net_shaper_nl_group_doit() does not deduplicate NET_SHAPER_A_LEAVES entries. When userspace supplies the same leaf handle twice, the same old-parent pointer lands twice in old_nodes[]. The cleanup loop double frees the parent. Of course the same parent may still be in old_nodes[] twice if we are	N/A	More Details

64054	moving multiple of its leaves. Note that this patch also implicitly fixes the fact that the i >= leaves_count path forgets to set ret.		
CVE-2026-55576	MaaAssistantArknights is a one-click tool for daily Arknights tasks. In the current dev-v2 workflow, .github/workflows/release-preparation.yml inlined attacker-controlled github.event.pull_request.title into a run: shell command during the pull_request opened, reopened, and ready_for_review events, so a non-draft fork PR whose title starts with Release v could execute shell commands on the ubuntu-latest runner during the generate-changelog job. This vulnerability is fixed by commit cafc3946059e6337d2089d4fec8b6885ba17c332.	N/A	More Details
CVE-2026-64059	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix folio->private handling in netfs_perform_write() Under some circumstances, netfs_perform_write() doesn't correctly manipulate folio->private between NULL, NETFS_FOLIO_COPY_TO_CACHE, pointing to a group and pointing to a netfs_folio struct, leading to potential multiple attachments of private data with associated folio ref leaks and also leaks of netfs_folio structs or netfs_group refs. Fix this by consolidating the place at which a folio is marked uptodate in one place and having that look at what's attached to folio->private and decide how to clean it up and then set the new group. Also, the content shouldn't be flushed if group is NULL, even if a group is specified in the netfs_group parameter, as that would be the case for a new folio. A filesystem should always specify netfs_group or never specify netfs_group. The Sashiko auto-review tool noted that it was theoretically possible that the fpos >= ctx->zero_point section might leak if it modified a streaming write folio. This is unlikely, but with a network filesystem, third party changes can happen. It also pointed out that __netfs_set_group() would leak if called multiple times on the same folio from the "whole folio modify section".	N/A	More Details
CVE-2026-54490	websocket-driver is a WebSocket protocol handler with pluggable I/O. Prior to 0.7.5, if this library is used with the permessage-deflate extension, a WebSocket server or client can be made to accept messages that are larger than the configured maximum message size because the limit is checked against the message frames' length headers, which give the size of the compressed data, not the size after decompression in lib/websocket/driver/hybi.js. This can lead to applications accepting larger messages than expected and exceeding their intended resource usage. This issue is fixed in version 0.7.5.	N/A	More Details
CVE-2026-64060	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix leak of request in netfs_write_begin() error handling Fix netfs_write_begin() to not leak our ref on the request in the event that we get an error from netfs_wait_for_read().	N/A	More Details
CVE-2026-64062	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix potential deadlock in write-through mode Fix netfs_advance_writethrough() to always unlock the supplied folio and to mark it dirty if it isn't yet written to the end. Unfortunately, it can't be marked for writeback until the folio is done with as that may cause a deadlock against mmapped reads and writes. Even though it has been marked dirty, premature writeback can't occur as the caller is holding both inode->i_rwsem (which will prevent concurrent truncation, fallocation, DIO and other writes) and ictx->wb_lock (which will cause flushing to wait and writeback to skip or wait). Note that this may be easier to deal with once the queuing of folios is split from the generation of subrequests.	N/A	More Details
CVE-2026-64063	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix streaming write being overwritten In order to avoid reading whilst writing, netfslib will allow "streaming writes" in which dirty data is stored directly into folios without reading them first. Such folios are marked dirty but may not be marked uptodate. If a folio is entirely written by a streaming write, uptodate will be set, otherwise it will have a netfs_folio struct attached to ->private recording the dirty region. In the event that a partially written streaming write page is to be overwritten entirely by a single write(), netfs_perform_write() will try to copy over it, but doesn't discard the netfs_folio if it succeeds; further, it doesn't correctly handle a partial copy that overwrites some of the dirty data. Fix this by the following: (1) If the folio is successfully overwritten, free the netfs_folio struct before marking the page uptodate. (2) If the copy to the folio partially fails, but short of the dirty data, just ignore the copy. (3) If the copy partially fails and overwrites some of the dirty data, accept the copy, update the netfs_folio struct to record the new data. If the folio is now filled, free the netfs_folio and set uptodate, otherwise return a partial write. Found with: fsx -q -N 1000000 -p 10000 -o 128000 -l 600000 \ /xfstest.test/junk --replay-ops=junk.fsxops using the following as junk.fsxops: truncate 0x0 0 0x927c0 write 0x63fb8 0x53c8 0 copy_range 0xb704 0x19b9 0x24429 0x79380 write 0x2402b 0x144a2 0x90660 * write 0x204d5 0x140a0 0x927c0 * copy_range 0x1f72c 0x137d0 0x7a906 0x927c0 * read 0x00000 0x20000 0x9157c read 0x20000 0x20000 0x9157c read 0x40000 0x20000 0x9157c read 0x60000 0x20000 0x9157c read 0x7e1a0 0xcfb9 0x9157c on cifs with the default cache option. It shows folio 0x24 misbehaving if the FMODE_READ check is commented out in netfs_perform_write(): if (//(file->f_mode & FMODE_READ) || netfs_is_cache_enabled(ctx)) { and no fscache. This was initially found with the generic/522 xfstest.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix netfs_invalidate_folio() to clear dirty bit if all changes gone If a streaming write is made, this will leave the relevant modified folio in a not-uptodate, but dirty state with a netfs_folio struct hung off of folio->private indicating the dirty range. Subsequently truncating the file such that the dirty data in the folio is removed, but the first part of the folio theoretically remains will cause the netfs_folio struct to be discarded... but will leave the dirty flag set. If the folio is then read via mmap(), netfs_read_folio() will see that the page is dirty and jump to netfs_read_gaps() to fill in the missing bits. netfs_read_gaps(), however, expects there to be a netfs_folio struct present and can		

CVE-2026-64064	oops because truncate removed it. Fix this by calling folio_cancel_dirty() in netfs_invalidate_folio() in the event that all the dirty data in the folio is erased (as nfs does). Also add some tracepoints to log modifications to a dirty page. This can be reproduced with something like: dd if=/dev/zero of=/xfstest.test/foo bs=1M count=1 umount /xfstest.test mount /xfstest.test xfs_io -c "w 0xbbbf 0xf96c" \ -c "truncate 0xbbbf" \ -c "mmap -r 0xb000 0x11000" \ -c "mr 0xb000 0x11000" \ /xfstest.test/foo with fscaching disabled (otherwise streaming writes are suppressed) and a change to netfs_perform_write() to disallow streaming writes if the fd is open O_RDWR: if (//(file->f_mode & FMODE_READ) || <--- comment this out netfs_is_cache_enabled(ctx)) { It should be reproducible even without this change, but if prevents the above trivial xfs_io command from reproducing it. Note that the initial dd is important: the file must start out sufficiently large that the zero-point logic doesn't just clear the gaps because it knows there's nothing in the file to read yet. Unmounting and mounting is needed to clear the pagecache (there are other ways to do that that may also work). This was initially reproduced with the generic/522 xfstest on some patches that remove the FMODE_READ restriction.	N/A	More Details
CVE-2026-64065	In the Linux kernel, the following vulnerability has been resolved: netfs: fix VM_BUG_ON_FOLIO() issue in netfs_write_begin() call The multiple runs of generic/013 test-case is capable to reproduce a kernel BUG at mm/filemap.c:1504 with probability of 30%. while true; do sudo ./check generic/013 done [9849.452376] page: refcount:3 mapcount:0 mapping:00000000e58ff252 index:0x10781 pfn:0x1c322 [9849.452412] memcg:ffff8881a1915800 [9849.452417] aops:ceph_aops ino:1000058db9e dentry name(?):"f9XXXXXX" [9849.452432] flags: 0x17ffffc0000000(node=0|zone=2|lastcpupid=0x1fffff) [9849.452441] raw: 0017ffffc0000000 0000000000000000 dead000000000122 ffff88816110d248 [9849.452445] raw: 0000000000010781 0000000000000000 00000003ffffffff ffff8881a1915800 [9849.452447] page dumped because: VM_BUG_ON_FOLIO(!folio_test_locked(folio)) [9849.452474] ------------[cut here]------------ [9849.452476] kernel BUG at mm/filemap.c:1504! [9849.478635] Oops: invalid opcode: 0000 [#1] SMP KASAN NOPTI [9849.481772] CPU: 2 UID: 0 PID: 84223 Comm: fsstress Not tainted 7.0.0-rc1+ #18 PREEMPT(full) [9849.482881] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.17.0-9.fc43 06/1 0/2025 [9849.484539] RIP: 0010:folio_unlock+0x85/0xa0 [9849.485076] Code: 89 df 31 f6 e8 1c f3 ff ff 48 8b 5d f8 c9 31 c0 31 d2 31 f6 31 ff c3 cc cc cc cc 48 c7 c6 80 6c d9 a7 48 89 df e8 4b b3 10 00 <0f> 0b 48 89 df e8 21 e6 2c 00 eb 9d 0f 1f 40 00 66 66 2e 0f 1f 84 [9849.493818] RSP: 0018:ffff8881bb8076b0 EFLAGS: 00010246 [9849.495740] RAX: 0000000000000000 RBX: ffffea00070c8980 RCX: 0000000000000000 [9849.498678] RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000000 [9849.500559] RBP: ffff8881bb8076b8 R08: 0000000000000000 R09: 0000000000000000 [9849.501097] R10: 0000000000000000 R11: 0000000000000000 R12: 0000000010782000 [9849.502108] R13: ffff8881935de738 R14: ffff88816110d010 R15: 0000000000001000 [9849.502516] FS: 00007e36cbe94740(0000) GS:ffff88824a899000(0000) knlGS:0000000000000000 [9849.502996] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [9849.503810] CR2: 000000c0002b0000 CR3: 000000011bbf6004 CR4: 0000000000772ef0 [9849.504459] PKRU: 55555554 [9849.504626] Call Trace: [9849.505242] <TASK> [9849.505379] netfs_write_begin+0x7c8/0x10a0 [9849.505877] ? __kasan_check_read+0x11/0x20 [9849.506384] ? __pfx_netfs_write_begin+0x10/0x10 [9849.507178] ceph_write_begin+0x8c/0x1c0 [9849.507934] generic_perform_write+0x391/0x8f0 [9849.508503] ? __pfx_generic_perform_write+0x10/0x10 [9849.509062] ? file_update_time_flags+0x19a/0x4b0 [9849.509581] ? ceph_get_caps+0x63/0xf0 [9849.510259] ? ceph_get_caps+0x63/0xf0 [9849.510530] ceph_write_iter+0xe79/0x1ae0 [9849.511282] ? __pfx_ceph_write_iter+0x10/0x10 [9849.511839] ? lock_acquire+0x1ad/0x310 [9849.512334] ? ksys_write+0xf9/0x230 [9849.512582] ? lock_is_held_type+0xaa/0x140 [9849.513128] vfs_write+0x512/0x1110 [9849.513634] ? __fget_files+0x33/0x350 [9849.513893] ? __pfx_vfs_write+0x10/0x10 [9849.514143] ? mutex_lock_nested+0x1b/0x30 [9849.514394] ksys_write+0xf9/0x230 [9849.514621] ? __pfx_ksys_write+0x10/0x10 [9849.514887] ? do_syscall_64+0x25e/0x1520 [9849.515122] ? __kasan_check_read+0x11/0x20 [9849.515366] ? trace_hardirqs_on_prepare+0x178/0x1c0 [9849.515655] __x64_sys_write+0x72/0xd0 [9849.515885] ? trace_hardirqs_on+0x24/0x1c0 [9849.516130] x64_sys_call+0x22f/0x2390 [9849.516341] do_syscall_64+0x12b/0x1520 [9849.516545] ? do_syscall_64+0x27c/0x1520 [9849.516783] ? do_syscall_64+0x27c/0x1520 [9849.517003] ? lock_release+0x318/0x480 [9849.517220] ? __x64_sys_io_getevents+0x143/0x2d0 [9849.517479] ? percpu_ref_put_many.constprop.0+0x8f/0x210 [9849.517779] ? entry_SYSCALL_64_after_hwframe+0x76/0x7e [9849.518073] ? do_syscall_64+0x25e/0x1520 [9849.518291] ? __kasan_check_read+0x11/0x20 [9849.518519] ? trace_hardirqs_on_prepare+0x178/0x1c0 [9849.518799] ? do_syscall_64+0x27c/0x1520 [9 ---truncated---	N/A	More Details
CVE-2026-64070	In the Linux kernel, the following vulnerability has been resolved: powerpc/hv-gpci: fix preempt count leak in sysfs show paths Four sysfs show() callbacks in hv-gpci take get_cpu_var(hv_gpci_reqb) (which calls preempt_disable()) but only call the matching put_cpu_var() on the error path under the 'out:' label. Every successful read leaks one preempt_disable(): processor_bus_topology_show() processor_config_show() affinity_domain_via_virtual_processor_show() affinity_domain_via_domain_show() (affinity_domain_via_partition_show() was already correct.) On a CONFIG_PREEMPT=y kernel, repeated reads raise preempt_count and eventually return to userspace with preemption still disabled. The next user-mode page fault then hits faulthandler_disabled() == 1, gets forced to SIGSEGV, and the resulting coredump trips 'BUG: scheduling while atomic' in call_usermodehelper_exec -> wait_for_completion_state -> schedule: BUG: scheduling while atomic: <task>/<pid>/0x00000004 ... __schedule_bug+0x6c/0x90 __schedule+0x58c/0x13a0 schedule+0x48/0x1a0 schedule_timeout+0x104/0x170 wait_for_completion_state+0x16c/0x330 call_usermodehelper_exec+0x254/0x2d0 vfs_coredump+0x1050/0x2590 get_signal+0xb9c/0xc80 do_notify_resume+0xf8/0x470 Add an out_success	N/A	More Details

	label that calls put_cpu_var() before returning the byte count, mirroring affinity_domain_via_partition_show().		
CVE-2026-64071	In the Linux kernel, the following vulnerability has been resolved: nvme-pci: fix use-after-free in nvme_free_host_mem() nvme_free_host_mem() frees dev->hmb_sgt via dma_free_noncontiguous() but never clears the pointer afterward. This leads to a use-after-free if nvme_free_host_mem() is called twice in the same error path. This can happen during nvme_probe() when nvme_setup_host_mem() succeeds in allocating the HMB (setting dev->hmb_sgt) but nvme_set_host_mem() fails with an I/O error: nvme_setup_host_mem() nvme_alloc_host_mem_single() -> sets dev->hmb_sgt nvme_set_host_mem() -> fails with -EIO nvme_free_host_mem() -> frees hmb_sgt, but does NOT NULL it return error nvme_probe() error path: nvme_free_host_mem() -> dev->hmb_sgt is stale, use-after-free The second call dereferences the freed sgt, causing a NULL pointer dereference in iommu_dma_free_noncontiguous() when it accesses sgt->sgl->dma_address (the backing memory has been freed and zeroed). This is reproducible on Thunderbolt-attached NVMe devices (e.g., OWC Envoy Express behind a Dell WD22TB4 dock) where the device intermittently returns I/O errors during HMB setup due to PCIe link instability. BUG: kernel NULL pointer dereference, address: 0000000000000010 RIP: 0010:iommu_dma_free_noncontiguous+0x22/0x80 Call Trace: <TASK> dma_free_noncontiguous+0x3b/0x130 nvme_free_host_mem+0x30/0xf0 [nvme] nvme_probe.cold+0xcc/0x275 [nvme] local_pci_probe+0x43/0xa0 pci_device_probe+0xeea/0x290 really_probe+0xf9/0x3b0 __driver_probe_device+0x8b/0x170 driver_probe_device+0x24/0xd0 __driver_attach_async_helper+0x6b/0x110 async_run_entry_fn+0x37/0x170 process_one_work+0x1ac/0x3d0 worker_thread+0x1b8/0x360 kthread+0xf7/0x130 ret_from_fork+0x2d8/0x3a0 ret_from_fork_asm+0x1a/0x30 </TASK> Fix this by setting dev->hmb_sgt to NULL after freeing it, so the second call takes the multi-descriptor path which safely handles the already-cleaned-up state.	N/A	More Details
CVE-2026-64072	In the Linux kernel, the following vulnerability has been resolved: nvme: fix bio leak on mapping failure The local bio is always NULL, so we'd leak the bio if the integrity mapping failed. Just get it directly from the request.	N/A	More Details
CVE-2026-64040	In the Linux kernel, the following vulnerability has been resolved: cachefiles: Fix error return when vfs_mkdir() fails When vfs_mkdir() fails, the error code is not extracted from the returned error pointer. This causes mkdir_error to be reached with ret=0, which leads to returning ERR_PTR(0) (NULL) instead of a proper error pointer. Fix this by extracting the error code from the error pointer when vfs_mkdir() fails.	N/A	More Details
CVE-2026-64038	In the Linux kernel, the following vulnerability has been resolved: hwmon: (lm90) Stop work before releasing hwmon device Sashiko reports: In lm90_probe(), the devm action to cancel the alert_work and report_work (lm90_restore_conf) is registered in lm90_init_client() before devm_hwmon_device_register_with_info() is called. Because devm executes cleanup actions in reverse order during module unbind or probe failure, the hwmon device is unregistered and freed first. If lm90_alert_work() or lm90_report_alarms() runs in the window between the hwmon device being freed and the delayed works being cancelled, lm90_update_alarms() will dereference the freed data->hwmon_dev here. Fix the problem by canceling the workers separately after registering the hwmon device and before registering the interrupt handler. This ensures that the workers are canceled after interrupts are disabled and before the hwmon device is released. Add "shutdown" flag to indicate that device shutdown is in progress to prevent workers from being re-armed.	N/A	More Details
CVE-2026-63826	In the Linux kernel, the following vulnerability has been resolved: fbdev: fix use-after-free in store_modes() store_modes() replaces a framebuffer's modelist with modes from userspace. On success it frees the old modelist with fb_destroy_modelist(). Two fields still point into that freed list. One pointer is fb_display[i].mode, the mode a console is using. fbcon_new_modelist() moves these pointers to the new list. It only does so for consoles still mapped to the framebuffer. An unmapped console is skipped and keeps its stale pointer. Unbinding fbcon, for example, sets con2fb_map[i] to -1 but leaves fb_display[i].mode set. An FBIOPUT_VSCREENINFO ioctl with FB_ACTIVATE_INV_MODE later reaches fbcon_mode_deleted(). That function reads the stale fb_display[i].mode through fb_mode_is_equal(). The read is a use-after-free. The other pointer is fb_info->mode, the current mode. It is set through the mode sysfs attribute. store_modes() does not update fb_info->mode, so it is left pointing into the freed list. show_mode(), the attribute's read handler, dereferences the stale fb_info->mode through mode_string(). The read is a use-after-free. Clear both pointers before freeing the list. Commit a1f305893074 ("fbcon: Set fb_display[i]->mode to NULL when the mode is released") added the helper fbcon_delete_modelist(). It clears every fb_display[i].mode that points into a given list. So far it is called only from the unregister path. Call it from store_modes() too, and set fb_info->mode to NULL.	N/A	More Details
CVE-2026-64028	In the Linux kernel, the following vulnerability has been resolved: tracing: Avoid NULL return from hist_field_name() on truncation hist_field_name() returns "" everywhere except the fully-qualified VAR_REF/EXPR case, where snprintf() truncation returns NULL early and bypasses the bottom NULL->"" guard. Callers don't expect NULL: strcat(expr, hist_field_name(field, 0)) at trace_events_hist.c:1758 and the strcmp() in the sort-key match loop at :4804 both deref it. system and event_name are bounded by MAX_EVENT_NAME_LEN, but the field name on a VAR_REF is kstrdup'd from a histogram variable name parsed out of the trigger string and has no length cap, so a long enough var name in a fully qualified reference can reach the truncation path. Keep the length check but leave field_name as "" on overflow.	N/A	More Details
CVE-2026-46686	Emlog is an open source website building system. In 2.6.13 and earlier, the admin backend user search module's keyword parameter from admin/user.php is processed with addslashes but not HTML-escaped before being rendered into the value attribute in admin/views/user.php, allowing reflected cross-site scripting	N/A	More Details

	in an administrator's backend session. No fixed version is currently identified.		
CVE-2026-46687	Emlog is an open source website building system. In 2.6.13 and earlier, the article publishing interface stores a path-traversal template parameter from api_controller.php without validation, and log_controller.php later checks file_exists and calls include View::getView($template), allowing an authenticated author to include an arbitrary local .php file when an article is viewed. No fixed version is currently identified.	N/A	More Details
CVE-2026-15925	Improper TLS hostname verification in Snowflake Connector for Python versions prior to 4.7.1 and 3.18.1 may have allowed a network-positioned attacker to bypass certificate hostname validation on HTTPS connections made by the connector. An attacker with on-path network access could exploit this by intercepting or redirecting network traffic and presenting a certificate signed by any trusted CA for any domain, causing the connector to accept connections without validating that the certificate matched the requested hostname. Successful exploitation requires an on-path traffic interception capability (e.g. ARP/DNS poisoning, rogue access point, BGP hijacking, or malicious proxy/exit node). This vulnerability may have exposed credentials, query data, and staged file contents to interception and tampering, and may have enabled the attacker to issue arbitrary SQL within the context of the victim's connector session. Impact is limited by the privileges of the affected Snowflake role. The fix is available in Snowflake Connector for Python versions 4.7.1 and 3.18.1. Users must manually upgrade.	N/A	More Details
CVE-2026-64006	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: fix dst corruption in same register operation For lshift and rshift, the shift operations are performed in a loop over 32-bit words. The loop calculates the shifted value and write it to dst, and then immediately reads from src to calculate the carry for the next iteration. Because src and dst could point to the same memory location, the carry is incorrectly calculated using the newly modified dst value instead of the original src value. Adding a temporary local variable to cache the original value before writing to dst and using it for the carry calculation solves the problem. In addition, partial overlap is rejected from control plane for all kind of operations including byteorder. This was tested with the following bytecode: table test_table ip flags 0 use 1 handle 1 ip test_table test_chain use 3 type filter hook input prio 0 policy accept packets 0 bytes 0 flags 1 ip test_table test_chain 2 [immediate reg 1 0x44332211 0x88776655] [bitwise reg 1 = (reg 1 << 0x08000000)] [cmp eq reg 1 0x66443322 0x00887766] [counter pkts 0 bytes 0] ip test_table test_chain 4 3 [immediate reg 1 0x44332211 0x88776655] [bitwise reg 1 = (reg 1 << 0x08000000)] [cmp eq reg 1 0x55443322 0x00887766] [counter pkts 21794 bytes 1917798]	N/A	More Details
CVE-2026-63839	In the Linux kernel, the following vulnerability has been resolved: platform/x86: lenovo-wmi-helpers: Fix memory leak in lwmi_dev_evaluate_int() lwmi_dev_evaluate_int() leaks output.pointer when retval == NULL (found by sashiko.dev [1]). Fix it by moving `ret_obj = output.pointer' outside of the `if (retval)' block so that it is always freed by the __free cleanup callback. No functional change intended.	N/A	More Details
CVE-2026-64012	In the Linux kernel, the following vulnerability has been resolved: net/sched: sch_sfb: Replace direct dequeue call with peek and qdisc_dequeue_peeked When sfb has children (eg qfq qdisc) whose peek() callback is qdisc_peek_dequeued(), we could get a kernel panic. When the parent of such qdiscs (eg illustrated in patch #3 as tbf) wants to retrieve an skb from its child (sfb in this case), it will do the following: 1a. do a peek() - and when sensing there's an skb the child can offer, then - the child in this case(sfb) calls its child's (qfq) peek. qfq does the right thing and will return the gso_skb queue packet. Note: if there wasnt a gso_skb entry then qfq will store it there. 1b. invoke a dequeue() on the child (sfb). And herein lies the problem. - sfb will call the child's dequeue() which will essentially just try to grab something of qfq's queue. [127.594489][T453] KASAN: null-ptr-deref in range [0x0000000000000048-0x000000000000004f] [127.594741][T453] CPU: 2 UID: 0 PID: 453 Comm: ping Not tainted 7.1.0-rc1-00035-gac961974495b-dirty #793 PREEMPT(full) [127.595059][T453] Hardware name: Bochs Bochs, BIOS Bochs 01/01/2011 [127.595254][T453] RIP: 0010:qfq_dequeue+0x35c/0x1650 [sch_qfq] [127.595461][T453] Code: 00 fc ff df 80 3c 02 00 0f 85 17 0e 00 00 4c 8d 73 48 48 89 9d b8 02 00 00 48 b8 00 00 00 00 00 fc ff df 4c 89 f2 48 c1 ea 03 <80> 3c 02 00 0f 85 76 0c 00 00 48 b8 00 00 00 00 00 fc ff df 4c 8b [127.596081][T453] RSP: 0018:ffff88810e5af440 EFLAGS: 00010216 [127.596337][T453] RAX: dffffc0000000000 RBX: 0000000000000000 RCX: dffffc0000000000 [127.596623][T453] RDX: 0000000000000009 RSI: 0000001880000000 RDI: ffff888104fd82b0 [127.596917][T453] RBP: ffff888104fd8000 R08: ffff888104fd8280 R09: 1ffff110211893a3 [127.597165][T453] R10: 1ffff110211893a6 R11: 1ffff110211893a7 R12: 0000001880000000 [127.597404] [T453] R13: ffff888104fd82b8 R14: 0000000000000048 R15: 0000000040000000 [127.597644][T453] FS: 00007fc380cbfc40(0000) GS:ffff88816f2a8000(0000) knlGS:0000000000000000 [127.597956][T453] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [127.598160][T453] CR2: 00005610aa9890a8 CR3: 000000010369e000 CR4: 0000000000750ef0 [127.598390][T453] PKRU: 55555554 [127.598509][T453] Call Trace: [127.598629][T453] <TASK> [127.598718][T453] ? mark_held_locks+0x40/0x70 [127.598890] [T453] ? srso_alias_return_thunk+0x5/0xfbef5 [127.599053][T453] sfb_dequeue+0x88/0x4d0 [127.599174][T453] ? ktime_get+0x137/0x230 [127.599328][T453] ? srso_alias_return_thunk+0x5/0xfbef5 [127.599480][T453] ? qdisc_peek_dequeued+0x7b/0x350 [sch_qfq] [127.599670][T453] ? srso_alias_return_thunk+0x5/0xfbef5 [127.599831][T453] tbf_dequeue+0x6b1/0x1098 [sch_tbf] [127.599988][T453] __qdisc_run+0x169/0x1900 The right thing to do in #1b is to grab the skb off gso_skb queue. This patchset fixes that issue by changing #1b to use qdisc_dequeue_peeked() method instead.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ACPI: button: Fix ACPI GPE handler leak during removal Commit a7e23ec17fee ("ACPI: button: Install notifier for system events as well") changed the		

CVE-2026-64013	ACPI notify handler type for ACPI buttons to ACPI_ALL_NOTIFY, but it forgot to update acpi_button_remove() to reflect that change. This leads to leaking the notify handler past driver removal, which may cause a kernel crash to occur if ACPI notify on the given device is triggered after removing the driver, and causes a subsequent probe of the given device with the same driver to fail. Address this by updating the acpi_remove_notify_handler() call in acpi_button_remove() as appropriate.	N/A	More Details
CVE-2026-64014	In the Linux kernel, the following vulnerability has been resolved: Input: usbtouchscreen - clamp NEXIO data_len/x_len to URB buffer size nexio_read_data() pulls data_len and x_len from a packed __be16 header in the device's interrupt packet and then walks packet->data[0..x_len) and packet->data[x_len..data_len) comparing each byte against a threshold. Both fields are 16-bit on the wire (max 65535). The existing adjustments shave at most 0x100 / 0x80 off, so the loop bound can still reach roughly 0xfeff. The URB transfer buffer for NEXIO is rept_size (1024) bytes from usb_alloc_coherent(), with the first 7 occupied by the packed header — so packet->data[] has 1017 valid bytes. read_data() callbacks are not given urb->actual_length, and nothing else bounds the walk. A device that lies about its length can get a ~64 KiB out-of-bounds read past the coherent DMA allocation. The first index whose byte exceeds NEXIO_THRESHOLD lands in begin_x / begin_y and from there into the reported touch coordinates, so adjacent kernel memory contents leak to userspace as ABS_X / ABS_Y events. Far enough out, the read can also hit an unmapped page and fault. Fix this all by clamping data_len to the buffer's data[] capacity and x_len to data_len.	N/A	More Details
CVE-2026-63838	In the Linux kernel, the following vulnerability has been resolved: ASoC: rsnd: Fix potential out-of-bounds access of component_dais[] component_dais[RSND_MAX_COMPONENT] is initially zero-initialized and later populated in rsnd_dai_of_node(). However, the existing boundary check: if (i >= RSND_MAX_COMPONENT) does not guarantee that the last valid element remains zero. As a result, the loop can rely on component_dais[RSND_MAX_COMPONENT] being zero, which may lead to an out-of-bounds access. Found by Linux Verification Center (linuxtesting.org) with SVACE.	N/A	More Details
CVE-2026-63837	In the Linux kernel, the following vulnerability has been resolved: net: ena: PHC: Check return code before setting timestamp output ena_phc_gettimex64() is setting the output parameter regardless of whether ena_com_phc_get_timestamp() succeeded or failed. When ena_com_phc_get_timestamp() returns an error, the timestamp parameter may contain uninitialized stack memory (e.g., when PHC is disabled or in blocked state) or invalid hardware values. Passing these to userspace via the PTP ioctl is both a security issue (information leak) and a correctness bug. Fix by checking the return code after releasing the lock and only setting the output timestamp on success.	N/A	More Details
CVE-2026-64019	In the Linux kernel, the following vulnerability has been resolved: nvme-pci: fix dma mapping leak on data setup error We're leaking the initial DMA mapping during iteration if we fail to allocate the tracking descriptor for both PRP and SGL. Unmap the iterator directly; we can't use the existing unmap helper because it depends on the tracking descriptor being successfully allocated, so a new one for an in-use iterator is provided. The mappings were also leaking when the driver detects an invalid bio_vec when mapping PRPs, so fix that too.	N/A	More Details
CVE-2026-63836	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tp_meter: avoid divide-by-zero for dec_cwnd The cwnd is always MSS <= cwnd <= 0x20000000. But the calculation in batadv_tp_update_cwnd() assumes unsigned 32 bit arithmetics. ((mss * 8) ** 2) / (cwnd * 8) In case cwnd is actually 0x20000000, it will be shifted by 3 bit to the left end up at 0x100000000 or U32_MAX + 1. It will therefore wrap around and be 0 - resulting in: ((mss * 8) ** 2) / 0 This is of course invalid and cannot be calculated. The calculation should must be simplified to avoid this overflow: (mss ** 2) * 8 / cwnd It will keep the precision enhancement from the scaling (by 8) but avoid the overflow in the divisor. In theory, there could still be an overflow in the dividend. It is at the moment fixed to BATADV_TP_PLEN in batadv_tp_recv_ack() - so it is not an imminent problem. But allowing it to use the whole u32 bit range, would mean that it can still use up to 67 bits. To keep this calculation safe for 32 bit arithmetic, mss must never use more than floor((32 - 3) / 2) bits - or in other words: must never be larger than 16383.	N/A	More Details
CVE-2026-64021	In the Linux kernel, the following vulnerability has been resolved: drm/xe/oa: Fix exec_queue leak on width check in stream open In xe_oa_stream_open_ioctl(), when param.exec_q->width > 1 the function returns -EOPNOTSUPP directly, skipping the existing err_exec_q cleanup path. The exec_queue reference obtained by xe_exec_queue_lookup() is leaked. The exec queue holds a reference on the xe_file, which is only dropped during queue teardown. The leaked lookup ref is not on the file's exec_queue xarray, so file close cannot release it. This keeps both the exec queue and the file private state pinned indefinitely. Jump to err_exec_q instead of returning directly so the reference is released. (cherry picked from commit 339fa0be9e4a5d69fa47e91f4a36574224fb478f)	N/A	More Details
CVE-2026-64022	In the Linux kernel, the following vulnerability has been resolved: gpio: aggregator: remove the software node when deactivating the aggregator The dynamic software node we create for the aggregator platform device when using configfs is leaked when the device is deactivated. Destroy it as the last step in the tear-down path.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: batman-adv: v: prevent OGM aggregation on disabled hardif When an interface gets disabled, the worker is correctly disabled by batadv_hardif_disable_interface() -> ... -> batadv_v_ogm_iface_disable(). In this process, the skb aggr_list is also freed. But batadv_v_ogm_send_meshif() can still queue new skbs (via batadv_v_ogm_queue_on_if()) to	N/A	More Details

63835	the aggr_list. This will only stop after all cores can no longer find the RCU protected list of hard interfaces. These queued skbs will never be freed or consumed by batadv_v_ogm_aggr_work. The batadv_v_ogm_iface_disable() function must block batadv_v_ogm_queue_on_if() to avoid leak of skbs.		
CVE-2026-63834	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tp_meter: restrict number of unacked list entries When the unacked_list is unbound, an attacker could send messages with small lengths and appropriated seqno + gaps to force the receiver to allocate more and more unacked_list entries. And the end either causing an out-of-memory situation or increase the management overhead for the (large) list that significant portions of CPU cycles are wasted in searching through the list. When limiting the list to a specific number, it is important to still correctly add a new entry to the list. But if the list became larger than the limit, the last entry of the list (with the highest seqno) must be dropped to still allow the earlier seqnos to finish and therefore to continue the process. Otherwise, the process might get stuck with too high seqnos which are not handled by batadv_tp_ack_unordered().	N/A	More Details
CVE-2026-46515	Frogman provides headless PBX control through MCP and HTTP API. Prior to 1.6.3, PERM_READ access was sufficient to call fm_list_managers, fm_list_pinsets, fm_show_context, fm_get_mcp_config, fm_backup_status, fm_whos_calling, fm_run_saved_query, and fm_diagnose_trunk, exposing AMI manager secrets, outbound dial PINs, full Asterisk dialplan context, root SSH connection commands, backup artifact paths, CDR history, arbitrary saved GraphQL query execution, and raw AMI endpoint dumps containing SIP fields such as password, md5_cred, and oauth_secret. This issue is fixed in version 1.6.3.	N/A	More Details
CVE-2026-63820	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix missing read bio submission on large folio error f2fs_read_data_large_folio() can keep a read bio across multiple readahead folios. If a later folio hits an error before any of its blocks are added to the bio, folio_in_bio is false and the current error path returns immediately after ending that folio. This can leave the bio accumulated for earlier folios unsubmitted. Those folios then never receive read completion, and readers can wait indefinitely on the locked folios. Route errors through the common out path so any pending bio is submitted before returning. Stop consuming more readahead folios once an error is seen, and only wait on and clear the current folio when it was actually added to the bio.	N/A	More Details
CVE-2026-53446	Wekan is open source kanban built with Meteor. Prior to 9.32, Wekan webhook integration URLs in models/integrations.js are stored from user input and later fetched by server/notifications/outgoing.js without applying the existing validateAttachmentUrl() private-network checks from models/lib/attachmentUrlValidation.js. A board administrator can configure webhook URLs that cause server-side requests to internal or metadata services. This issue is fixed in version 9.32.	N/A	More Details
CVE-2026-64075	In the Linux kernel, the following vulnerability has been resolved: fprobe: Fix unregister_fprobe() to wait for RCU grace period Commit 4346ba1604093 ("fprobe: Rewrite fprobe on function-graph tracer") changed fprobe to register struct fprobe to an rcu-hlist, but it forgot to wait for RCU GP. Thus there can be use-after-free if the fprobe is released right after unregistering. This can be happened on fprobe event and sample module code. To fix this issue, add synchronize_rcu() in unregister_fprobe(). Note that BPF is OK because fprobe is used as a part of bpf_kprobe_multi_link. This unregisters its fprobe in bpf_kprobe_multi_link_release() and it is deallocated via bpf_kprobe_multi_link_dealloc(), which is invoked from bpf_link_defer_dealloc_rcu_gp() RCU callback. For BPF, this also introduced unregister_fprobe_async() which does NOT wait for RCU grace priod.	N/A	More Details
CVE-2026-56679	9Router is an AI router & token saver. Prior to 0.5.4, the PATCH /api/settings endpoint writes the entire request body to persistent settings without a field whitelist, allowing an authenticated user to set security-critical fields such as requireLogin and disable authentication for the whole application, exposing protected routes such as /api/keys and /api/providers to unauthenticated access. This issue is reported as fixed in version 0.5.4.	N/A	More Details
CVE-2026-15449	A time-of-check to time-of-use (TOCTOU) flaw in the illumos data-link pseudo-driver (dld) affects handling of the DLDIOC_GETMACPROP and DLDIOC_SETMACPROP ioctls on /dev/dld. drv_ioc_prop_common() in usr/src/uts/common/io/dld/dld_drv.c copies the dld_ioc_macprop_t ioctl header in once to read its pr_valsize field, sizes and allocates a kernel heap buffer from that value, and then copies the full request in a second time from the same unprivileged user address. A concurrent thread can enlarge pr_valsize between the two copyins, so the second copyin and the subsequent property handling write beyond the end of the undersized allocation and corrupt the kernel heap. An unprivileged local user, including one confined to a non-global zone that owns a datalink, can trigger this to panic the system. The resulting kernel heap corruption may be usable for further compromise.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: f2fs: read COW data with the original inode during atomic write When updating an atomic-write file, f2fs_write_begin() may read the previously written data back from the COW inode: prepare_atomic_write_begin() locates the block in the COW inode and sets use_cow, and the read bio is then built with the COW inode: f2fs_submit_page_read(use_cow ? F2FS_I(inode)->cow_inode : inode, ...); and f2fs_grab_read_bio() decides whether to schedule fs-layer decryption (STEP_DECRYPT) for the bio based on that inode via fscrypt_inode_uses_fs_layer_crypto(). However, the folio being filled belongs to the original inode (folio->mapping->host == inode), and the data stored in the COW block was encrypted (or left as plaintext) using the original inode's context, not the COW inode's -- see f2fs_encrypt_one_page(), which keys off fio->page->mapping->host. fscrypt_decrypt_pagecache_blocks() likewise operates on folio->mapping->host. The COW inode is created as a tmpfile in the parent directory		

CVE-2026-63811	and inherits its encryption policy from there. With test_dummy_encryption the newly created COW inode gets the dummy policy and becomes encrypted, while a pre-existing regular file -- created before the policy applied, e.g. already present in the on-disk image -- stays unencrypted. The read path then sets STEP_DECRYPT based on the encrypted COW inode and calls fscrypt_decrypt_pagecache_blocks() on a folio whose host (the unencrypted original inode) has a NULL ->i_crypt_info, dereferencing it: Oops: general protection fault, probably for non-canonical address ... KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f] RIP: 0010:fscrypt_decrypt_pagecache_blocks+0xa0/0x310 Workqueue: f2fs_post_read_wq f2fs_post_read_work Call Trace: fscrypt_decrypt_bio+0x1eb/0x340 f2fs_post_read_work+0xba/0x140 process_one_work+0x91c/0x1a40 worker_thread+0x677/0xe90 kthread+0x2bc/0x3a0 The COW inode is only needed to locate the on-disk block, and that block address is already resolved into @blkaddr by prepare_atomic_write_begin() via __find_data_block(cow_inode, ...); f2fs_submit_page_read() then reads from that physical @blkaddr directly, so the inode argument only selects the post-read crypto context, not which block is fetched. Reading with @inode therefore returns the same (latest, not-yet-committed) COW data, while making both the fs-layer decryption decision and the inline crypto path use the correct (original inode's) key. With the COW inode no longer used at the read site, the use_cow flag has no remaining consumer; drop it from f2fs_write_begin() and prepare_atomic_write_begin().	N/A	More Details
CVE-2026-63810	In the Linux kernel, the following vulnerability has been resolved: block: Avoid mounting the bdev pseudo-filesystem in userspace The bdev pseudo-filesystem is an internal kernel filesystem with which userspace should not interfere. Unregister it so that userspace cannot even attempt to mount it. This fixes a bug [1] that occurs when attempting to access files, because the system call move_mount() uses pointers declared in the inode_operations structure, which for the bdev pseudo-filesystem are always equal to 0. `inode->i_op = &empty_iops;` [1] BUG: kernel NULL pointer dereference, address: 0000000000000000 #PF: supervisor instruction fetch in kernel mode #PF: error_code(0x0010) - not-present page PGD 23380067 P4D 23380067 PUD 23381067 PMD 0 Oops: 0010 [#1] PREEMPT SMP KASAN NOPTI CPU: 2 PID: 17125 Comm: syz-executor.0 Not tainted 6.1.155-syzkaller-00350-g84221fde2681 #0 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.12.0-1 04/01/2014 RIP: 0010:0x0 Call Trace: <TASK> lookup_open.isra.0+0x700/0x1180 fs/namei.c:3460 open_last_lookups fs/namei.c:3550 [inline] path_openat+0x953/0x2700 fs/namei.c:3780 do_filp_open+0x1c5/0x410 fs/namei.c:3810 do_sys_openat2+0x171/0x4d0 fs/open.c:1318 do_sys_open fs/open.c:1334 [inline] __do_sys_openat fs/open.c:1350 [inline] __se_sys_openat fs/open.c:1345 [inline] __x64_sys_openat+0x13c/0x1f0 fs/open.c:1345 do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_64+0x35/0x80 arch/x86/entry/common.c:81 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 Found by Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2026-44981	CrowdSec offers crowdsourced protection against malicious IPs. From 1.7.0 until 1.7.8, the LAPI router used gin-contrib/gzip with DefaultDecompressHandle globally in pkg/apiserver/controllers/controller.go, causing /v1/watchers and /v1/watchers/login to decompress unauthenticated gzip-compressed JSON request bodies without enforcing a maximum decompressed size and allowing excessive heap allocation that can make LAPI unreachable. This issue is fixed in version 1.7.8.	N/A	More Details
CVE-2026-64119	In the Linux kernel, the following vulnerability has been resolved: l2tp: use list_del_rcu in l2tp_session_unhash An unprivileged local user can pin a host CPU indefinitely in l2tp_session_get_by_ifname() by issuing L2TP_CMD_SESSION_GET on L2TP_ATTR_IFNAME concurrently with L2TP_CMD_SESSION_CREATE and L2TP_CMD_SESSION_DELETE on the same tunnel. All three commands take GENL_UNS_ADMIN_PERM, so CAP_NET_ADMIN in the netns user namespace suffices; on any host that has l2tp_core loaded the trigger is reachable from a standard `unshare -Urn` sandbox. l2tp_session_unhash() removes a session from tunnel->session_list with list_del_init(), but that list is walked by l2tp_session_get_by_ifname() with list_for_each_entry_rcu() under rcu_read_lock_bh(). list_del_init() leaves the deleted entry's next/prev self-pointing; a reader that has loaded the entry and then advances pos->list.next reads &session->list, container_of()s back to the same session, and list_for_each_entry_rcu() never reaches the list head. The CPU stays in strcmp() inside the walker, with BH and preemption disabled, so RCU grace periods on the host stall behind it and the wedged thread cannot be killed (SIGKILL is delivered on syscall return). Use list_del_rcu() to match the existing list_add_rcu() in l2tp_session_register(); the deleted session remains visible to in-flight walkers with consistent next/prev pointers until kfree_rcu() in l2tp_session_free() releases it. tunnel->session_list has exactly one list_del_init() call site; the list_del_init (&session->clist) at l2tp_core.c:533 operates on the per-collision list, which is not walked under RCU. list_empty(&session->list) is not used anywhere in net/l2tp/ after the unhash point, so dropping the post-delete self-init is safe; the fix has no userspace-visible behavior change.	N/A	More Details
CVE-2026-64120	In the Linux kernel, the following vulnerability has been resolved: net: ethtool: fix NULL pointer dereference in phy_reply_size In phy_prepare_data(), several strings such as 'name', 'drvname', 'upstream_sfp_name', and 'downstream_sfp_name' are allocated using kstrdup(). However, these allocations were not checked for failure. If kstrdup() fails for 'name', it returns NULL while the function continues. This leads to a kernel NULL pointer dereference and panic later in phy_reply_size() when it unconditionally calls strlen() on the NULL pointer. While other strings like 'upstream_sfp_name' might be checked before access in certain code paths, failing to handle these allocations consistently can lead to incomplete data reporting or hidden bugs. Fix this by adding proper NULL checks for all kstrdup() calls in phy_prepare_data() and implement a centralized error handling path using goto labels to ensure all previously allocated resources are freed on failure.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net: ifb: report ethtool stats over		

CVE-2026-64121	num_tx_queues ifb_dev_init() allocates dp->tx_private to dev->num_tx_queues entries via kzalloc_objs(*txp, dev->num_tx_queues). Both IFB per-queue RX and TX stats live in those entries: ifb_xmit() updates txp->rx_stats using the skb queue mapping, ifb_ri_tasklet() updates txp->tx_stats, and ifb_stats64() aggregates both over dev->num_tx_queues. The ethtool stats callbacks instead size and walk the per-queue stats with dev->real_num_rx_queues and dev->real_num_tx_queues. With an asymmetric device where the RX queue count exceeds the TX queue count, for example: ip link add name ifb10 numtxqueues 1 numrxqueues 8 type ifb ethtool -S ifb10 ifb_get_ethtool_stats() indexes past the tx_private allocation and copies adjacent slab data through ETHTOOL_GSTATS. Use dev->num_tx_queues consistently for the stats strings, the stats count, and the stats data walks. This reports one RX stats group and one TX stats group for each backing ifb_q_private entry, which is the queue set IFB can actually populate. Reproduced under UML+KASAN at v7.1-rc2: BUG: KASAN: slab-out-of-bounds in ifb_fill_stats_data+0x3c/0xae Read of size 8 at addr 0000000062dbd228 by task ethtool/36 ifb_fill_stats_data+0x3c/0xae ifb_get_ethtool_stats+0xc0/0x129 __dev_ethtool+0x1ca5/0x363c dev_ethtool+0x123/0x1b3 dev_ioctl+0x56c/0x744 sock_do_ioctl+0x15f/0x1b2 sock_ioctl+0x4d5/0x50a sys_ioctl+0xd8b/0xde9 With the patch applied, the same UML+KASAN repro is silent and ethtool -S ifb10 reports only the stats backed by the single allocated tx_private entry.	N/A	More Details
CVE-2026-54728	bunkerweb is an Open-source and next-generation Web Application Firewall (WAF). Prior to BunkerWeb 1.6.12 and BunkerWeb PRO 0.57, authenticated Host header handling in the BunkerWeb UI and API improperly validated and neutralized user-controlled input in a configuration-dependent path, allowing a low-privileged authenticated user to escalate privileges and affect confidentiality, integrity, and availability of the BunkerWeb instance. This issue is fixed in BunkerWeb version 1.6.12 and BunkerWeb PRO version 0.57.	N/A	More Details
CVE-2026-64127	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: ecred_reconfigure: send packed pdu, not stack pointer Commit 1c08108f3014 ("Bluetooth: L2CAP: Avoid -Wflex-array-member-not-at-end warnings") converted the on-stack request PDU in l2cap_ecred_reconfigure() from an explicit packed struct to DEFINE_RAW_FLEX(), but did not adjust the size and source-pointer arguments to l2cap_send_cmd(): - struct { - struct l2cap_ecred_reconf_req req; - __le16 scid; - } pdu; + DEFINE_RAW_FLEX(struct l2cap_ecred_reconf_req, pdu, scid, 1); ... l2cap_send_cmd(conn, chan->ident, L2CAP_ECRED_RECONF_REQ, sizeof(pdu), &pdu); After the conversion, DEFINE_RAW_FLEX() expands to declare an anonymous union pdu_u plus a local pointer "pdu" pointing at it. Therefore: - sizeof(pdu) is now sizeof(struct l2cap_ecred_reconf_req *) = 8 on 64-bit (4 on 32-bit), not the 6 bytes of (mtu, mps, scid[1]). - &pdu is the address of the local pointer's stack storage, not the address of the request payload. l2cap_send_cmd() forwards (data, count) to l2cap_build_cmd(), which calls skb_put_data(skb, data, count). The L2CAP_ECRED_RECONFIGURE_REQ packet body therefore contains 8 bytes copied from the kernel stack starting at &pdu -- the 8 bytes overlap the pdu pointer's value, leaking a kernel stack address to the paired Bluetooth peer. The intended (mtu, mps, scid) fields are not transmitted at all, so the peer rejects the request as malformed and the L2CAP_ECRED_RECONFIGURE feature itself has been broken for the local-side initiator since the introducing commit landed. The sibling site l2cap_ecred_conn_req() in the same commit was converted correctly (sizeof(*pdu) + len, pdu); only this site was missed. Restore the original semantics: pass the full flex-struct size via struct_size(pdu, scid, 1) and the pdu pointer (the struct address) as the source. Validated on a stock 7.0-based host kernel via the real call path: setsockopt(SOL_BLUETOOTH, BT_RCVMTU, ...) on a BT_CONNECTED L2CAP_MODE_EXT_FLOWCTL socket emits an L2CAP_ECRED_RECONFIGURE_REQ whose body is 8 bytes (the on-stack pdu local's value) rather than the expected 6. Three captures from fresh socket / fresh hciemu peer on the same host -- low bytes vary per call, high 0xffff confirms a kernel virtual address (KASLR-randomised stack slot, not a fixed string): RECONF_REQ body (ident=0x02 len=8): 42 fb 54 af 0e ca ff ff RECONF_REQ body (ident=0x02 len=8): 52 3d 2e af 0e ca ff ff RECONF_REQ body (ident=0x02 len=8): b2 fc 5b af 0e ca ff ff After this patch the body is 6 bytes carrying the expected little-endian (mtu, mps, scid).	N/A	More Details
CVE-2026-64128	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: drop ISO_END frames received without prior ISO_START ISO data PDUs carry a packet-boundary flag indicating START, CONT, END or SINGLE. The ISO_CONT branch of iso_recv() guards against a missing ISO_START by checking conn->rx_len before touching conn->rx_skb, but ISO_END does not. If a peer sends an ISO_END as the first packet on a fresh ISO connection, conn->rx_skb is still NULL and conn->rx_len is zero, so skb_put(conn->rx_skb, ...) dereferences NULL and oopses. For BIS, where receivers sync to a broadcaster without pairing, any broadcaster on the air can trigger this. Mirror the ISO_CONT check at the top of ISO_END so a stray end fragment is logged and dropped instead of crashing the host.	N/A	More Details
CVE-2026-64129	In the Linux kernel, the following vulnerability has been resolved: mm/migrate_device: fix spinlock leak in migrate_vma_insert_huge_pmd_page When check_stable_address_space() fails after the PMD spinlock has been acquired via pmd_lock(), the code jumps directly to the abort label, bypassing the spin_unlock() call in unlock_abort. This causes the PMD spinlock to be permanently held, leading to a deadlock. Change the goto target from abort to unlock_abort to ensure the spinlock is always released on this error path.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm/page_alloc: fix initialization of tags of the huge zero folio with init_on_free __GFP_ZEROTAGS semantics are currently a bit weird, but effectively this flag is only ever set alongside __GFP_ZERO and __GFP_SKIP_KASAN. If we run with init_on_free, we will zero out pages during __free_pages_prepare(), to skip zeroing on the allocation path. However, when allocating with __GFP_ZEROTAG set, post_alloc_hook() will consequently not only skip clearing page content, but also skip clearing tag memory. Not clearing tags through __GFP_ZEROTAGS is irrelevant for most pages that will get mapped to user space through set_pte_at() later: set_pte_at() and friends will detect that the tags have		

CVE-2026-64130	not been initialized yet (PG_mte_tagged not set), and initialize them. However, for the huge zero folio, which will be mapped through a PMD marked as special, this initialization will not be performed, ending up exposing whatever tags were still set for the pages. The docs (Documentation/arch/arm64/memory-tagging-extension.rst) state that allocation tags are set to 0 when a page is first mapped to user space. That no longer holds with the huge zero folio when init_on_free is enabled. Fix it by decoupling __GFP_ZEROTAGS from __GFP_ZERO, passing to tag_clear_highpages() whether we want to also clear page content. Invert the meaning of the tag_clear_highpages() return value to have clearer semantics. Reproduced with the huge zero folio by modifying the check_buffer_fill arm64/mte selftest to use a 2 MiB area, after making sure that pages have a non-0 tag set when freeing (note that, during boot, we will not actually initialize tags, but only set KASAN_TAG_KERNEL in the page flags). $ ./check_buffer_fill 1..20 ... not ok 17 Check initial tags with private mapping, sync error mode and mmap memory not ok 18 Check initial tags with private mapping, sync error mode and mmap/mprotect memory ... This code needs more cleanups; we'll tackle that next, like decoupling __GFP_ZEROTAGS from __GFP_SKIP_KASAN. [akpm@linux-foundation.org: s/__GPF_ZERO/__GFP_ZERO/, per David]	N/A	More Details
CVE-2026-64131	In the Linux kernel, the following vulnerability has been resolved: mm/memory: fix spurious warning when unmapping device-private/exclusive pages Device private and exclusive entries are only supported for anonymous folios. This condition is tested in __migrate_device_pages() and make_device_exclusive() using folio_test_anon(). However the unmap path tests this assumption using vma_is_anonymous(). This is wrong because whilst anonymous VMAs can only contain folios where folio_test_anon() is true the opposite relation does not hold. A folio for which folio_test_anon() is true does not imply vma_is_anonymous() is true. Such a condition can occur if for example a folio is part of a private filebacked mapping. In this case vma_is_anonymous() is false as the mapping is filebacked, but folio_test_anon() may be true, thus permitting devices to migrate the folio to device private memory. This can lead to the following spurious warnings during process teardown: [772.737706] ------------[cut here]------------ [772.739201] WARNING: mm/memory.c:1754 at unmap_page_range.cold+0x26/0x18a, CPU#17: hmm-tests/2041 [772.742050] Modules linked in: test_hmm nvidia_uvm(O) nvidia(O) [772.743959] CPU: 17 UID: 0 PID: 2041 Comm: hmm-tests Tainted: G W O 7.0.0+ #387 PREEMPT(full) [772.747104] Tainted: [W]=WARN, [O]=OOT_MODULE [772.748509] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.17.0-0-gb52ca86e094d-prebuilt.qemu.org 04/01/2014 [772.752117] RIP: 0010:unmap_page_range.cold+0x26/0x18a [772.753780] Code: 7e fe ff ff 48 89 4c 24 78 4c 89 44 24 38 e8 f2 ff b1 00 48 8b 4c 24 78 4c 8b 44 24 38 48 8b 44 24 18 48 83 78 48 00 74 04 90 <0f> 0b 90 48 89 ca b8 ff ff 37 00 48 c1 ea 03 48 c1 e0 2a 80 3c 02 [772.759602] RSP: 0018:ffff888112607550 EFLAGS: 00010286 [772.761310] RAX: ffff88811bbf4dc0 RBX: dfffc0000000000 RCX: ffffea03e9bfffd8 [772.763583] RDX: 1ffff1102377e9c1 RSI: 0000000000000008 RDI: ffff88811bbf4e08 [772.765914] RBP: 0000000000000006 R08: ffff8881059f7448 R09: ffffed10224c0e68 [772.768184] R10: ffff888112607347 R11: 0000000000000001 R12: 0000000000000001 [772.770461] R13: ffffea03e9bfffc0 R14: ffff888112607908 R15: ffffea03e9bfffc0 [772.772782] FS: 00007f327caa2780(0000) GS:ffff888427b7d000(0000) knlGS:0000000000000000 [772.775328] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [772.777187] CR2: 00007f327ca89000 CR3: 00000001994d5000 CR4: 00000000000006f0 [772.779135] Call Trace: [772.779792] <TASK> [772.780317] ? dmirror_interval_invalidate+0x1a3/0x290 [test_hmm] [772.781873] ? vm_normal_page_pud+0x2b0/0x2b0 [772.782992] ? __rwlock_init+0x150/0x150 [772.784006] ? lock_release+0x216/0x2b0 [772.785008] ? __mmu_notifier_invalidate_range_start+0x505/0x6e0 [772.786522] ? lock_release+0x216/0x2b0 [772.787498] ? unmap_single_vma+0xb6/0x210 [772.788573] unmap_vmas+0x27d/0x520 [772.789506] ? unmap_single_vma+0x210/0x210 [772.790607] ? mas_update_gap.part.0+0x620/0x620 [772.791834] unmap_region+0x19e/0x350 [772.792769] ? remove_vma+0x130/0x130 [772.793684] ? mas_alloc_nodes+0x1f2/0x300 [772.794730] vms_complete_munmap_vmas+0x8c1/0xe20 [772.795926] ? unmap_region+0x350/0x350 [772.796917] do_vmi_align_munmap+0x36a/0x4e0 [772.798018] ? lock_release+0x216/0x2b0 [772.799024] ? vma_shrink+0x620/0x620 [772.799983] do_vmi_munmap+0x150/0x2c0 [772.800939] __vm_munmap+0x161/0x2c0 [772.801872] ? expand_downwards+0xd60/0xd60 [772.802948] ? clockevents_program_event+0x1ef/0x540 [772.804217] ? lock_release+0x216/0x2b0 [772.805158] __x64_sys_munmap+0x59/0x80 [772.805776] do_syscall_64+0xfc/0x670 [772.806336] ? irqentry_exit+0xda/0x580 [772.806976] entry_SYSCALL_64_after_hwframe+0x4b/0x53 [772.807772] RIP: 0033:0x7f327cbb2717 [772.808323] Code: 73 01 c3 48 8b 0d f9 76 0d 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 b8 0b 00 00 00 0f 05 <48> 3d 01 f0 ff ---truncated---	N/A	More Details
CVE-2026-63804	In the Linux kernel, the following vulnerability has been resolved: gfs2: fix use-after-free in gfs2_qd_dealloc gfs2_qd_dealloc(), called as an RCU callback from gfs2_qd_dispose(), accesses the superblock object sdp through qd->qd_sbd after freeing qd. It does so to decrement sd_quota_count and wake up sd_kill_wait. However, by the time the RCU callback runs, gfs2_put_super() may have already freed sdp via free_sbd(). This can happen when gfs2_quota_cleanup() is called during unmount: it disposes of quota objects via call_rcu() and then waits on sd_kill_wait with a 60-second timeout. If the timeout expires, or if gfs2_gl_hash_clear() triggers additional qd_put() calls that schedule more RCU callbacks after the wait completes, gfs2_put_super() will proceed to free the superblock while RCU callbacks referencing it are still pending. Add an rcu_barrier() before free_sbd() in gfs2_put_super() to ensure all pending RCU callbacks (including gfs2_qd_dealloc) have completed before the superblock is freed.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) widen blackbox-info buffer to I2C_SMBUS_BLOCK_MAX adm1266_nvmem_read_blackbox() declares a 5-byte stack buffer and		

CVE-2026-64135	passes it to i2c_smbus_read_block_data() to retrieve the 4-byte BLACKBOX_INFO response. i2c_smbus_read_block_data() does not honour caller buffer sizes -- it memcpy()s data.block[0] bytes from the SMBus transaction (where data.block[0] is the length byte returned by the slave device, up to I2C_SMBUS_BLOCK_MAX = 32): memcpy(values, &data.block[1], data.block[0]); If the device returns any block length above 5, the call overflows the caller's 5-byte stack buffer before the post-call if (ret != 4) return -EIO; check has a chance to reject the response. Widen the local buffer to I2C_SMBUS_BLOCK_MAX so the helper has room for any well-formed SMBus block response, matching the convention used by the other i2c_smbus_read_block_data() callers in this driver.	N/A	More Details
CVE-2026-55629	Whistle is an HTTP, HTTP2, HTTPS, and WebSocket debugging proxy. Prior to 2.10.3, lib/service/service.js handles GET /cgi-bin/temp/get by reading req.query.filename, joining it to TEMP_FILES_PATH only when it matches the temporary file pattern, and otherwise passing the user-supplied filename directly to getFile, allowing a remote attacker to read arbitrary files such as /etc/passwd. This issue is reported as fixed in version 2.10.3.	N/A	More Details
CVE-2026-63798	In the Linux kernel, the following vulnerability has been resolved: irqchip/imgpdc: Fix resource leak, add missing chained handler cleanup on remove The driver allocates domain generic chips using irq_alloc_domain_generic_chips() during probe and sets up chained handlers using irq_set_chained_handler_and_data(). However, on driver removal, the generic chips are not freed and the chained handlers are not removed. The generic chips remain on the global gc_list and may later be accessed by generic interrupt chip suspend, resume, or shutdown callbacks after the driver has been removed, potentially resulting in a use-after-free and kernel crash. The chained handlers that were installed in probe for peripheral and syswake interrupts are also left dangling, which can lead to spurious interrupts accessing freed memory. Fix these issues by: - Setting IRQ_DOMAIN_FLAG_DESTROY_GC flag in domain->flags, so the core code automatically removes generic chips when irq_domain_remove() is called - Clearing all chained handlers with NULL in pdc_intc_remove()	N/A	More Details
CVE-2026-64110	In the Linux kernel, the following vulnerability has been resolved: igc: fix potential skb leak in igc_fpe_xmit_smd_frame() When igc_fpe_init_tx_descriptor() fails, no one takes care of an allocated skb, leaking it. [1] Use dev_kfree_skb_any() on failure. Tested on an I226 adapter with the following command, while injecting faults in igc_fpe_init_tx_descriptor() to trigger the error path. # ethtool --set-mm $DEV verify-enabled on tx-enabled on pmac-enabled on [1] unreferenced object 0xffff888113c6cdc0 (size 224): ... backtrace (crc be3d3fda): kmem_cache_alloc_node_noprof+0x3b1/0x410 __alloc_skb+0xde/0x830 igc_fpe_xmit_smd_frame.isra.0+0xad/0x1b0 igc_fpe_send_mpacket+0x37/0x90 ethtool_mmsv_verify_timer+0x15e/0x300	N/A	More Details
CVE-2026-15422	The illumos SCTP inbound path performs association lookup for INIT ACK chunks without adequately validating the address parameters carried in the chunk. Since this lookup runs during packet classification (i.e. before SCTP integrity checks or IPsec policy are applied) a remote, unauthenticated attacker can send a crafted SCTP INIT ACK packet with malformed address parameters to cause an out-of-bounds access and kernel heap corruption, which may lead to remote code execution. The flaw has existed since 2010 (illumos-gate commit a5407c02), and affects any illumos distribution prior to illumos-gate commit 53a3efde.	N/A	More Details
CVE-2026-53445	Wekan is open source kanban built with Meteor. Prior to 9.32, the Wekan copyBoard Meteor DDP method in server/publications/boards.js copies a board by caller-supplied board ID without checking this.userId, membership, or admin access. Any authenticated user can copy a private board they are not a member of, including its cards, checklists, custom fields, labels, and rules, while the REST POST /api/boards/:boardId/copy path correctly checks board admin access. This issue is fixed in version 9.32.	N/A	More Details
CVE-2026-64107	In the Linux kernel, the following vulnerability has been resolved: ASoC: codecs: pcm512x: fix null-ptr dereference in pcm512x_overclock_xxx_put() In the pcm512x chipset driver, pcm512x_overclock_xxx_put() is defined as a general mixer kcontrol instead of a DAPM kcontrol, so struct snd_soc_dapm_context must not be accessed via snd_soc_dapm_kcontrol_to_dapm(). This causes a NULL pointer dereference, so it must be modified to use snd_soc_component_to_dapm().	N/A	More Details
CVE-2026-53444	Wekan is open source kanban built with Meteor. Prior to 9.32, Wekan OIDC-related Meteor methods in packages/wekan-oidc/oidc_server.js, server/models/org.js, and server/models/team.js are globally callable without the admin authorization checks used by their non-OIDC counterparts. Authenticated users can call setCreateOrgFromOidc, setOrgAllFieldsFromOidc, setCreateTeamFromOidc, setTeamAllFieldsFromOidc, boardRoutineOnLogin, or groupRoutineOnLogin to create or modify organizations and teams, and groupRoutineOnLogin can grant global admin privileges when PROPAGATE_OIDC_DATA is enabled. This issue is fixed in version 9.32.	N/A	More Details
CVE-2026-52893	Wekan is open source kanban built with Meteor. Prior to 9.32, the Wekan Accounts.onCreateUser hook in server/models/users.js merges OIDC logins into existing accounts when the OIDC email or username matches an existing Wekan user, without verifying ownership or checking email_verified. An attacker using an OIDC provider account with a victim's email or username can cause Wekan to merge the attacker's OIDC credentials into the victim account and then log in as that account. This issue is fixed in version 9.32.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: netfilter: x_tables: allocate hook ops while under mutex arp/ip(6)t_register_table() add the table to the per-netns list via xt_register_table() before		

CVE-2026-64079	allocating the per-netns hook ops copy via kmemdup_array(). This leaves a window where the table is visible in the list with ops=NULL. If the pernet exit happens runs concurrently the pre_exit callback finds the table via xt_find_table() and passes the NULL ops pointer to nf_unregister_net_hooks(), causing a NULL dereference: general protection fault in nf_unregister_net_hooks+0xbc/0x150 RIP: nf_unregister_net_hooks (net/netfilter/core.c:613) Call Trace: ipt_unregister_table_pre_exit iptable_mangle_net_pre_exit ops_pre_exit_list cleanup_net Fix by moving the ops allocation into the xtables core so the table is never in the list without valid ops. Also ensure the table is no longer processing packets before its torn down on error unwind. nf_register_net_hooks might have published at least one hook; call synchronize_rcu() if there was an error. audit log register message gets deferred until all operations have passed, this avoids need to emit another ureg message in case of error unwinding. Based on earlier patch by Tristan Madani.	N/A	More Details
CVE-2026-64083	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) reject short block-read responses in the GPIO accessors adm1266_gpio_get() and adm1266_gpio_get_multiple() both compose the pin-status word as pins_status = read_buf[0] + (read_buf[1] << 8); right after i2c_smbus_read_block_data(), guarding only against an error return. A well-behaved device returns 2 bytes for GPIO_STATUS/PDIO_STATUS, but the helper happily reports a 0- or 1-byte response too. If the device returns 0 bytes, both read_buf slots are uninitialized stack memory; if it returns 1 byte, read_buf[1] is. The composed value then flows through set_bit() into the caller's *bits in adm1266_gpio_get_multiple(), or into the return value of adm1266_gpio_get(), and ends up in userspace via gpiolib (sysfs and the char-dev ioctls). That leaks a few bits of kernel stack per request on any device whose firmware glitch, bus error, or hostile slave produces a short block-read response. Add the missing length check to both call sites and surface a short response as -EIO.	N/A	More Details
CVE-2026-64085	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) bounce blackbox records through a protocol-sized buffer adm1266_pmbus_block_xfer() copies the device-supplied block payload into the caller-provided buffer using the device-supplied length: memcpy(data_r, &msgs[1].buf[1], msgs[1].buf[0]); The helper does not know how large data_r is and trusts the device to return at most one record's worth of bytes. adm1266_nvmem_read_blackbox() violates that contract: it advances read_buff inside data->dev_mem in ADM1266_BLACKBOX_SIZE (64-byte) strides while the helper is willing to write up to ADM1266_PMBUS_BLOCK_MAX (255) bytes. A device that returns more than 64 bytes on the trailing record (read_buff offset 1984 in the 2048-byte dev_mem allocation) overflows dev_mem by up to 191 bytes before the post-call if (ret != ADM1266_BLACKBOX_SIZE) return -EIO; can reject the response. Contain the fix in the caller without changing the helper signature: read each record into a 255-byte local bounce buffer that matches the helper's maximum output, validate the returned length, and only then copy exactly ADM1266_BLACKBOX_SIZE bytes into the dev_mem slot.	N/A	More Details
CVE-2026-64087	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/adm1266) reject implausible blackbox record_count adm1266_nvmem_read_blackbox() loops over a record_count that comes straight from byte 3 of the BLACKBOX_INFO response. The destination buffer is data->dev_mem, sized for the nvmem cell's declared 2048 bytes (ADM1266_BLACKBOX_MAX_RECORDS * ADM1266_BLACKBOX_SIZE = 32 * 64). A device that reports a record_count greater than 32 -- whether due to firmware bugs, bus corruption, or a non-responsive slave returning 0xff -- would walk read_buff past the end of the dev_mem allocation on the trailing iterations. Cap record_count at ADM1266_BLACKBOX_MAX_RECORDS (introduced here) before entering the loop and return -EIO on any larger value, so a malformed BLACKBOX_INFO response cannot drive the loop out of bounds.	N/A	More Details
CVE-2026-49279	WWBN AVideo is an open source video platform. Versions 29.0 and below contain a Stored XSS vulnerability through the autoEvalCodeOnHTML parameter in the MessageSQLite WebSocket Handler. The MessageSQLite.php handler only strips autoEvalCodeOnHTML from $json['msg'], but msgToResourceId() reads from $msg['json'] with higher priority. An attacker can place the XSS payload in the json key instead of msg, bypassing the sanitization entirely. An authenticated attacker can execute arbitrary JavaScript in any connected user's browser session via the WebSocket messaging system, stealing session cookies and authentication tokens, taking over accounts through session hijacking, and chaining with CSRF to perform admin actions on the victim's behalf, in the default SQLite WebSocket backend configuration. This issue has a patch that has yet to be officially released, see https://github.com/WWBN/AVideo/commit/3e0b3ce2bfa766183ff0ae227439394db57b1a23.	N/A	More Details
CVE-2026-64090	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: avoid empty VLAN responses The commit 16116dac2339 ("batman-adv: prevent TT request storms by not sending inconsistent TT TLVLs") added checks to the local (direct) TT response code. But the response can also be done indirectly by another node using the global TT state. To avoid such inconsistency states reported in the original fix, also avoid sending empty VLANs for replies from the global TT state.	N/A	More Details
CVE-2026-64092	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tp_meter: fix tp_vars reference leak in receiver shutdown The receiver shutdown timer handler, batadv_tp_receiver_shutdown(), is responsible for releasing the tp_vars reference it holds. However, the existing logic for coordinating this release with batadv_tp_stop_all() was flawed. timer_shutdown_sync() guarantees the timer will not fire again after it returns, but it returns non-zero only when the timer was pending at the time of the call. If the timer had already expired (and batadv_tp_stop_all() would unsucessfully try to rearm itself), batadv_tp_stop_all() skips its batadv_tp_vars_put(), and batadv_tp_receiver_shutdown() fails to put its own reference as well. Fix this by introducing a new atomic variable receiving that is set to 1 when the receiver is initialized and cleared	N/A	More Details

	atomically with atomic_xchg() by whichever side claims it first. Only the side that observes the transition from 1 to 0 is responsible for releasing the tp_vars timer reference, eliminating the uncertainty.		
CVE-2026-53535	Activepieces is an open source AI workflow automation platform. Prior to 0.82.0, the git-sync feature clones a user-configured Git repository into a temporary directory on the server and then writes flow, table, and connection state into it before pushing back, and two separate weaknesses allowed those writes to escape the intended workspace and land on arbitrary paths on the host filesystem: Git's symbolic-link handling was not disabled on the clone, so an attacker who controlled the remote repository could include symlinks that redirected the writes, and several user-supplied identifiers used to build on-disk paths (the repository slug and the externalId of tables, flows, and connections) were not validated against directory-traversal sequences such as ../. On a self-hosted Enterprise Edition deployment, a user authorized to configure or push to a git-sync repository (holding the WRITE_PROJECT_RELEASE permission) could cause the server to overwrite files anywhere the Activepieces process user can write, which depending on host layout can be leveraged for tampering, denial of service, or remote code execution. This issue is fixed in version 0.82.0.	N/A	More Details
CVE-2026-64094	In the Linux kernel, the following vulnerability has been resolved: batman-adv: bla: avoid NULL-ptr deref for claim via dropped interface Without rtnl_lock held, a hardif might be retrieved as primary interface of a meshif, but then (while operating on this interface) getting decoupled from the mesh interface. In this case, the meshif still exists but the pointer from the primary hardif to the meshif is set to NULL. The mesh_iface must be checked first to be non-NULL before continuing to send an ARP request using meshif.	N/A	More Details
CVE-2026-53536	Activepieces is an open source AI workflow automation platform. Prior to 0.83.0, the /v1/step-files/signed download endpoint verified the supplied JWT against the shared signing secret but did not check the token's audience, and combined with a missing null-check on the decoded fileId, this allowed any caller holding any valid Activepieces JWT (including a freshly created user's own access token) to receive a step-file belonging to another tenant. The file returned was whatever PostgreSQL happened to scan first for type = FLOW_STEP_FILE, varying over time as the database changed, so an authenticated user could obtain step-file attachments belonging to other tenants on the same instance; the attacker could not target a specific victim or file, and the access was read-only with no integrity or availability impact. This issue is fixed in version 0.83.0.	N/A	More Details
CVE-2026-54526	Argo Workflows is an open source container-native workflow engine for orchestrating parallel jobs on Kubernetes. Prior to 3.7.15 and 4.0.6, the allow-list fix for CVE-2026-31892 is incomplete because workflow/util/merge.go ValidateUserOverrides and SanitizeUserWorkflowSpec walk only the top-level fields of WorkflowSpec via reflection, and WorkflowSpec.ArtifactGC is allow-listed wholesale; the struct behind that field, WorkflowLevelArtifactGC, has a PodSpecPatch sub-field whose contents flow unmodified into util.ApplyPodSpecPatch on the artifact-GC pod, the same sink the original fix closed for WorkflowSpec.PodSpecPatch, so a user submitting a Workflow under templateReferencing: Strict or Secure (against a referenced WorkflowTemplate that declares an output artifact and setting spec.artifactGC.strategy: OnWorkflowCompletion) can still inject an arbitrary strategic merge patch into the artifact-GC pod, including hostPath volumes, privileged: true, arbitrary image and command, and hostNetwork: true, defeating the stated purpose of Strict/Secure reference mode. This issue is fixed in versions 3.7.15 and 4.0.6.	N/A	More Details
CVE-2026-64100	In the Linux kernel, the following vulnerability has been resolved: drm/msm: Fix shrinker deadlock With PROVE_LOCKING on an Snapdragon X1 and VM reclaim pressure, we see: == WARNING: possible circular locking dependency detected 7.0.0-debug+ #43 Tainted: G W -- kswapd0/82 is trying to acquire lock: ffff800080ec3870 (reservation_ww_class_acquire){+.+.}-{0:0}, at: msm_gem_shrinker_scan+0x17c/0x400 [msm] but task is already holding lock: ffffc31709b263b8 (fs_reclaim) {+.+.}-{0:0}, at: balance_pgdat+0x88/0x988 which lock already depends on the new lock. the existing dependency chain (in reverse order) is: -> #2 (fs_reclaim){+.+.}-{0:0}: __lock_acquire+0x4d0/0xad0 lock_acquire.part.0+0xc4/0x248 lock_acquire+0x8c/0x248 fs_reclaim_acquire+0xd0/0xf0 dma_resv_lockdep+0x224/0x348 do_one_initcall+0x84/0x5d0 do_initcalls+0x194/0x1d8 kernel_init_freeable+0x128/0x180 kernel_init+0x2c/0x160 ret_from_fork+0x10/0x20 -> #1 (reservation_ww_class_mutex){+.+.}-{4:4}: __lock_acquire+0x4d0/0xad0 lock_acquire.part.0+0xc4/0x248 lock_acquire+0x8c/0x248 dma_resv_lockdep+0x1a8/0x348 do_one_initcall+0x84/0x5d0 do_initcalls+0x194/0x1d8 kernel_init_freeable+0x128/0x180 kernel_init+0x2c/0x160 ret_from_fork+0x10/0x20 -> #0 (reservation_ww_class_acquire){+.+.}-{0:0}: check_prev_add+0x114/0x790 validate_chain+0x594/0x6f0 __lock_acquire+0x4d0/0xad0 lock_acquire.part.0+0xc4/0x248 lock_acquire+0x8c/0x248 drm_gem_lru_scan+0x1ac/0x440 msm_gem_shrinker_scan+0x17c/0x400 [msm] do_shrink_slab+0x150/0x4a0 shrink_slab+0x144/0x460 shrink_one+0x9c/0x1b0 shrink_many+0x27c/0x5c0 shrink_node+0x344/0x550 balance_pgdat+0x2c0/0x988 kswapd+0x11c/0x318 kthread+0x10c/0x128 ret_from_fork+0x10/0x20 other info that might help us debug this: Chain exists of: reservation_ww_class_acquire --> reservation_ww_class_mutex --> fs_reclaim Possible unsafe locking scenario: CPU0 CPU1 ---- ---- lock(fs_reclaim); lock(reservation_ww_class_mutex); lock(fs_reclaim); lock(reservation_ww_class_acquire); *** DEADLOCK *** 1 lock held by kswapd0/82: #0: ffffc31709b263b8 (fs_reclaim){+.+.}-{0:0}, at: balance_pgdat+0x88/0x988 stack backtrace: CPU: 4 UID: 0 PID: 82 Comm: kswapd0 Tainted: G W 7.0.0-debug+ #43 PREEMPT(full) Tainted: [W]=WARN Hardware name: LENOVO 21BX0016US/21BX0016US, BIOS N3HET94W (1.66) 09/15/2025 Call trace: show_stack+0x20/0x40 (C) dump_stack_lvl+0x9c/0xd0 dump_stack+0x18/0x30 print_circular_bug+0x114/0x120	N/A	More Details

	check_noncircular+0x178/0x198 check_prev_add+0x114/0x790 validate_chain+0x594/0x6f0 __lock_acquire+0x4d0/0xad0 lock_acquire.part.0+0xc4/0x248 lock_acquire+0x8c/0x248 drm_gem_lru_scan+0x1ac/0x440 msm_gem_shrinker_scan+0x17c/0x400 [msm] do_shrink_slab+0x150/0x4a0 shrink_slab+0x144/0x460 shrink_one+0x9c/0x1b0 shrink_many+0x27c/0x5c0 shrink_node+0x344/0x550 balance_pgdat+0x2c0/0x988 kswapd+0x11c/0x318 kthread+0x10c/0x128 ret_from_fork+0x10/0x20 kswapd0 holding fs_reclaim calls the MSM shrinker, which calls dma_resv_lock. This in turn acquires fs_reclaim. Fix this deadlock by using dma_resv_trylock() instead, dropping the subsequently unused passed wait-wound lock 'ticket'. Patchwork: https://patchwork.freedesktop.org/patch/723564/ [rob: fixup compile errors, replace lockdep splat with somethin ---truncated---		
CVE-2026-64101	In the Linux kernel, the following vulnerability has been resolved: fwctl: pds: Validate RPC input size before parsing The fwctl core allocates the device-specific RPC input buffer with fwctl_rpc.in_len and passes that buffer to the driver callback. pdsfc_fw_rpc() casts the buffer to struct fwctl_rpc_pds and then calls pdsfc_validate_rpc(), which reads fields from that structure before checking that the input buffer is large enough to contain it. A short in_len can make pds_fwctl read beyond the allocation. Reject pds RPC buffers that are smaller than struct fwctl_rpc_pds before parsing any pds-specific fields.	N/A	More Details
CVE-2026-64103	In the Linux kernel, the following vulnerability has been resolved: scsi: isci: Fix use-after-free in device removal path The ISCI completion tasklet is initialized in isci_host_alloc() (drivers/scsi/isci/init.c:496) and scheduled from both MSI-X and legacy interrupt handlers (drivers/scsi/isci/host.c:223,613). isci_host_deinit() stops the controller and waits for stop completion, but it never kills completion_tasklet before teardown continues. A top-of-function tasklet_kill() is not sufficient here: interrupts are only disabled when isci_host_stop_complete() runs, so until wait_for_stop() returns the IRQ handlers can still requeue the tasklet. The tasklet callback also re-enables interrupts after draining completions, so killing the tasklet before the source is quiesced leaves the same race open. Once wait_for_stop() returns, no further IRQ-driven scheduling can occur. Kill completion_tasklet there so teardown cannot race a queued tasklet running on a dead ihost. On remove or unload, the stale callback can otherwise dereference ihost and touch ihost->smu_registers after the host lifetime ends. A UML + KASAN analogue reproduced the failure class both with no tasklet_kill() and with tasklet_kill() placed before source quiesce, and stayed clean once the kill happened after quiescing the scheduling source. This mirrors commit f6ab594672d4 ("scsi: aic94xx: fix use-after-free in device removal path"), but ISCI needs the kill after wait_for_stop().	N/A	More Details
CVE-2026-64105	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: vgic: Free private_irqs when init fails after allocation Companion to commit 250f25367b58 ("KVM: arm64: Tear down vGIC on failed vCPU creation"), which added the missing kvm_vgic_vcpu_destroy() call to the kvm_share_hyp() failure path in kvm_arch_vcpu_create(). The kvm_vgic_vcpu_init() failure path immediately above it has the same shape and still needs the same cleanup. Call kvm_vgic_vcpu_destroy() when kvm_vgic_vcpu_init() fails so private IRQs allocated before a redistributor iodev registration failure are released before the failed vCPU is freed.	N/A	More Details
CVE-2026-64001	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcm: oss: Fix setup list UAF on proc write error snd_pcm_oss_proc_write() links a newly allocated setup entry into the OSS setup list before duplicating the task name. If the task-name allocation fails, the error path frees the already linked entry and leaves setup_list pointing at freed memory. A later OSS device open can then walk the stale list entry in snd_pcm_oss_look_for_setup() and dereference freed memory. Allocate the task name and initialize the setup entry before publishing the entry on setup_list. Also fetch the initial proc read iterator only after taking setup_mutex, so all setup_list traversal follows the same list lifetime rules.	N/A	More Details
CVE-2026-63999	In the Linux kernel, the following vulnerability has been resolved: ethtool: rss: fix indir_table and hkey leak on get_rxfh failure rss_prepare_get() allocates the indirection table and hash key buffer via rss_get_data_alloc(), then calls ops->get_rxfh() to populate them. If get_rxfh() fails, the function returns an error without freeing the allocation.	N/A	More Details
CVE-2026-63998	In the Linux kernel, the following vulnerability has been resolved: ethtool: module: call ethnl_ops_complete() on module flash errors When validate() fails we are skipping over ethnl_ops_complete() even tho we already called ethnl_ops_begin().	N/A	More Details
CVE-2026-63904	In the Linux kernel, the following vulnerability has been resolved: usb: usbtmc: check URB actual_length for interrupt-IN notifications USBTMC devices can use an optional interrupt endpoint for notification messages. These typically contain two-byte headers indicating the payload format, but the driver does not check if these headers are present before accessing the data buffers. In cases where the URB actual_length is not enough to fit these headers, the driver will either cause an out-of-bounds read, or consume stale leftover data from a previous notification. Fix by checking if actual_data contains enough bytes for the headers, otherwise resubmit URB to the interrupt endpoint.	N/A	More Details
CVE-2026-63907	In the Linux kernel, the following vulnerability has been resolved: uio: uio_pci_generic_sva: fix double free of devm_kzalloc() memory uio_pci_sva allocates struct uio_pci_sva_dev with devm_kzalloc() in probe(), but then calls kfree(udev) both on the probe() error path (label out_free) and again in remove(). Because devm_kzalloc() allocations are devres-managed and are freed automatically when the device is detached (including after a failing probe() and during driver unbind), the explicit kfree() can lead to a double free. If probe() fails after devm_kzalloc(), the error path frees udev and devres cleanup will free it again when the core unwinds the partially bound device. On normal driver removal, remove() frees udev and devres will free	N/A	More Details

	it again when the device is detached. This issue was identified by a static analysis tool I developed and confirmed by manual review. Fix by removing the manual kfree() calls and dropping the now-unused label.		
CVE-2026-63908	In the Linux kernel, the following vulnerability has been resolved: Input: atmel_mxt_ts - fix boundary check in mxt_prepare_cfg_mem When a configuration file provides an object size that is larger than the driver's known mxt_obj_size(object), the driver intends to discard the extra bytes. The loop iterates using for (i = 0; i < size; i++). Inside the loop, the condition to skip processing extra bytes is: if (i > mxt_obj_size(object)) continue; Since i is a 0-based index, the valid indices for the object are 0 through mxt_obj_size(object) - 1. When i == mxt_obj_size(object), the condition evaluates to false, and the code processes the byte instead of discarding it. This causes the code to calculate byte_offset = reg + i - cfg->start_ofs and writes the byte there, overwriting exactly one byte of the adjacent instance or object. Update the boundary check to skip extra bytes correctly by using >=.	N/A	More Details
CVE-2026-54733	The Microsoft 365 and Microsoft Entra ID Plugins for Moodle provide Office 365 and Azure Active Directory integration for Moodle. Prior to 4.5.6, 5.0.5, and 5.1.1, the Microsoft Office 365 Integration plugin local_o365 Teams SSO endpoint sso_login.php base64-decodes a JWT payload and authenticates users from the upn claim without verifying the JWT signature, allowing an unauthenticated attacker to forge a token and obtain a Moodle session as an O365-authenticated user. This issue is fixed in versions 4.5.6, 5.0.5, and 5.1.1.	N/A	More Details
CVE-2026-53597	Prompty is a markdown file format (.prompty) for LLM prompts. From 2.0.0-alpha.1 until 2.0.0-beta.3, the @prompty/core TypeScript loader in runtime/typescript/packages/core/src/core/loader.ts used gray-matter without overriding executable js and javascript frontmatter engines, allowing an attacker-controlled .prompty file with ---js frontmatter to execute arbitrary JavaScript during prompt loading. This issue is fixed in version 2.0.0-beta.3.	N/A	More Details
CVE-2026-12379	An Open Redirect vulnerability (CWE-601) exists in the OAuth/OIDC authentication implementation of the Axivion Dashboard. The login flow did not properly restrict the post-authentication redirect to the application's own origin, so a user who follows a crafted login link can be sent to an untrusted external site after authenticating against the genuine Dashboard. Because the link points at the legitimate Dashboard, this can be abused for phishing, for example credential or second-factor theft via a convincing look-alike page. Exploitation requires the victim to follow the attacker-supplied link and complete the authentication flow.	N/A	More Details
CVE-2026-59863	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.5, Kiota honored a poisoned .kiota/workspace.json workspace configuration without validating per-client or per-plugin outputPath values during kiota client generate and kiota plugin generate, allowing a malicious repository or pull request to use absolute paths, rooted POSIX / paths, UNC \\ or // paths, Windows drive X:\ paths, or .. traversal segments to write generated client files outside the workspace root on a developer or CI host. This issue is fixed in version 1.32.5.	N/A	More Details
CVE-2026-59860	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.3, Kiota is affected by a code-generation injection vulnerability in the C# XML documentation-comment sink (the description, externalDocs label, and externalDocs link fields emitted as /// ... comments). When text from an OpenAPI description is written into single-line XML doc comments without stripping newline and Unicode line-terminator characters, an attacker can break out of the /// comment line and inject additional code into generated C# clients. This issue is fixed in version 1.32.3.	N/A	More Details
CVE-2026-59859	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.4, Kiota's PHP generator embedded OpenAPI description, default fields, property names, and other schema-derived strings into PHP double-quoted literals through SanitizeDoubleQuote() in Writers/StringExtensions.cs without escaping $, allowing attacker-controlled ${...}, $var, or {$obj->prop} interpolation constructs to inject arbitrary PHP code into generated model and request-builder classes. This issue is fixed in version 1.32.4.	N/A	More Details
CVE-2026-59237	Authorization Bypass Through User-Controlled Key (CWE-639) in the Order and OrderItem REST API controllers in Roskus Prospero Flow CRM before 5.5.3 allows a remote, authenticated user to read, modify, and delete orders and order items belonging to any other company (tenant) via a sequential numeric {id} supplied to GET /api/order/{id}, PUT /api/order/{id}, GET /api/order-item/{id}, PUT /api/order-item/{id}, or DELETE /api/order-item/{id}, because the controllers resolve records with Order::find($id) / Item::find($id) without scoping by the authenticated user's company.	N/A	More Details
CVE-2026-14254	A race condition in the account lockout mechanism in Delphix Continous Data allowed the lockout threshold to be bypassed through concurrent authentication requests. Parallel login attempts were processed before the failed-login counter and lockout status were updated, defeating brute-force protections and enabling continued password guessing against a targeted account.	N/A	More Details
CVE-2026-63871	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Fix data-race on iso_pi fields in hci_get_route calls iso_connect_bis(), iso_connect_cis(), iso_listen_bis(), and iso_conn_big_sync() call hci_get_route() using iso_pi(sk)->dst, iso_pi(sk)->src, and iso_pi(sk)->src_type without holding lock_sock(). These fields may be modified concurrently by connect() or setsockopt() on the same socket, resulting in data-races reported by KCSAN. Fix this by snapshotting the required fields under lock_sock() before calling hci_get_route(). BUG: KCSAN: data-race in memcmp+0x45/0xb0 race at unknown origin, with read to 0xffff8880122135cf of 1 bytes by task 333 on cpu 1: memcmp+0x45/0xb0 hci_get_route+0x27e/0x490	N/A	More Details

	iso_connect_cis+0x4c/0xa10 iso_sock_connect+0x60e/0xb30 __sys_connect_file+0xbd/0xe0 __sys_connect+0xe0/0x110 __x64_sys_connect+0x40/0x50 x64_sys_call+0xcad/0x1c60 do_syscall_64+0x133/0x590 entry_SYSCALL_64_after_hwframe+0x77/0x7f		
CVE-2026-14371	The Lenovo XClarity Integrator for Windows Admin Center plugin version 5.1.1 and below running on the WAC Gateway is vulnerable to Powershell Command Injection when establishing remote PowerShell commands.	N/A	More Details
CVE-2026-63868	In the Linux kernel, the following vulnerability has been resolved: net: garp: fix unsigned integer underflow in garp_pdu_parse_attr The receive-side GARP attribute parser computes dlen with reversed operands: dlen = sizeof(*ga) - ga->len; ga->len is the on-wire attribute length and includes the GARP attribute header. For normal attributes with data, ga->len is larger than sizeof(*ga), so the subtraction underflows in unsigned arithmetic. The resulting value is later passed to garp_attr_lookup(), whose length argument is u8. After truncation, the parsed data length usually no longer matches the length stored for locally registered attributes, so received Join/Leave events are ignored. This breaks the GARP receive path for common attributes, such as GVRP VLAN registration attributes. Compute the data length as the attribute length minus the header length.	N/A	More Details
CVE-2026-63928	In the Linux kernel, the following vulnerability has been resolved: USB: serial: omninet: fix memory corruption with small endpoint Make sure that the bulk-out buffers are at least as large as the hardcoded transfer size to avoid user-controlled slab corruption should a malicious device report a smaller endpoint max packet size than expected.	N/A	More Details
CVE-2026-63929	In the Linux kernel, the following vulnerability has been resolved: iio: buffer: Fix DMA fence leak in iio_buffer_enqueue_dmabuf() iio_buffer_enqueue_dmabuf() allocates a struct iio_dma_fence (104 bytes, kmalloc-128) via kmalloc_obj()+dma_fence_init(), which sets the initial kref to 1. It then calls dma_resv_add_fence() which takes a second reference (kref=2), and stores a raw pointer in block->fence. On the success path the function returns without calling dma_fence_put() to release the initial reference, so every buffer enqueue permanently leaks one kmalloc-128 allocation. The iio_buffer_cleanup() work item only releases the temporary reference taken during completion signalling by iio_buffer_signal_dmabuf_done(); the initial reference from dma_fence_init() is never released. With four iio_rwdev instances at 240kHz and 512 samples per buffer, this produces ~1875 kmalloc-128 allocations per second matching the observed slab growth exactly. A test with ftrace confirmed that the dma_fence_destroy event was never triggered. Fix by calling dma_fence_put() after dma_resv_add_fence(), transferring ownership of the fence to the DMA reservation object. The DMA fence then gets properly discarded after being signalled.	N/A	More Details
CVE-2026-63931	In the Linux kernel, the following vulnerability has been resolved: iio: chemical: scd30: fix division by zero in write_raw Add a zero check for val2 before using it as a divisor when setting the sampling frequency. A user writing a zero fractional part to the sampling_frequency sysfs attribute triggers a division by zero in the kernel.	N/A	More Details
CVE-2026-63932	In the Linux kernel, the following vulnerability has been resolved: iio: chemical: mhz19b: reject oversized serial replies mhz19b_receive_buf() appends each serdev chunk into the fixed MHZ19B_CMD_SIZE receive buffer and advances buf_idx by len without checking that the chunk fits in the remaining space. A large callback can therefore overflow st->buf before the command path validates the reply. Reset the reply state before each command and reject oversized serial replies before copying them into the fixed buffer. When an oversized reply is detected, wake the waiter and report -EMSGSIZE instead of overwriting st->buf.	N/A	More Details
CVE-2026-63905	In the Linux kernel, the following vulnerability has been resolved: usbip: vudc: Fix use after free bug in vudc_remove due to race condition This patch follows up Zheng Wang's 2023 report of a use-after-free in vudc_remove(). The original thread stalled on Shuah Khan's request for runtime testing of the unplug/unbind path. This patch supplies that testing and keeps Zheng's original fix shape. In vudc_probe(), v_init_timer() binds udc->tr_timer.timer to v_timer(). usbip_sockfd_store() starts the timer via v_start_timer()/v_kick_timer(). vudc_remove() can then free the containing struct vudc while the timer is still pending or executing. KASAN confirms the race on an unpatched x86_64 QEMU guest with CONFIG_KASAN=y, CONFIG_USBIP_VUDC=y, CONFIG_USB_ZERO=y, and a tight loop that repeatedly writes a socket fd to usbip_sockfd, closes the socket pair, and unbinds/rebinds usbip-vudc.0: BUG: KASAN: slab-use-after-free in __run_timer_base.part.0+0x8ba/0x8e0 Write of size 8 at addr ffff888001b80740 by task trigger_and_unb/239 Allocated by task 239: vudc_probe+0x4d/0xaa0 Freed by task 239: kfree+0x18f/0x520 device_release_driver_internal+0x388/0x540 unbind_store+0xd9/0x100 This lands in the timer core rather than v_timer() itself because the embedded timer_list is being walked after its containing struct vudc has already been freed. The underlying lifetime bug is the same one Zheng reported. With v_stop_timer() called from vudc_remove() and the timer deleted synchronously, the same harness completed 5000 bind/unbind iterations with no KASAN report.	N/A	More Details
CVE-2026-63903	In the Linux kernel, the following vulnerability has been resolved: USB: serial: belkin_sa: validate interrupt status length The Belkin interrupt callback treats interrupt data as a four-byte status report and reads LSR/MSR fields at offsets 2 and 3. The interrupt-in buffer length is derived from endpoint wMaxPacketSize, and short interrupt transfers may complete successfully with a smaller actual_length. Check the completed interrupt packet length before parsing status fields so short interrupt endpoints and short successful packets are ignored instead of causing out-of-bounds or stale status-byte reads. KASAN report as below: BUG: KASAN:	N/A	More Details

	slab-out-of-bounds in belkin_sa_read_int_callback() Read of size 1 Call trace: belkin_sa_read_int_callback() (drivers/usb/serial/belkin_sa.c:202) __usb_hcd_giveback_urb() (drivers/usb/core/hcd.c:1630) dummy_timer() (?:?)		
CVE-2026-63934	In the Linux kernel, the following vulnerability has been resolved: iio: gyro: itg3200: fix i2c read into the wrong stack location itg3200_read_all_channels() takes `__be16 *buf' as a parameter and fills the i2c_msg destination as `(char *)&buf'. Since `buf' is the parameter (a pointer), `&buf' is the address of the local pointer slot on the stack of itg3200_read_all_channels(), not the address of the caller's scan buffer. The (char *) cast hides the type mismatch. i2c_transfer() therefore writes ITG3200_SCAN_ELEMENTS * sizeof(s16) = 8 bytes into the parameter's stack slot, which is discarded when the function returns. The caller's scan buffer in itg3200_trigger_handler() is never written to, so iio_push_to_buffers_with_timestamp() pushes uninitialised stack contents to userspace via /dev/iio:deviceX every scan -- both a functional bug (no actual gyroscope or temperature data is delivered through the triggered buffer) and an information leak. The non-buffered read_raw() path is unaffected: it goes through itg3200_read_reg_s16() which uses `&out' on a local s16 value, where that is correct. Drop the spurious `&' so the i2c read writes into the caller's buffer.	N/A	More Details
CVE-2026-63902	In the Linux kernel, the following vulnerability has been resolved: USB: serial: cypress_m8: validate interrupt packet headers cypress_read_int_callback() parses the interrupt-in buffer according to the selected Cypress packet format. Format 1 has a two-byte status/count header and format 2 has a one-byte combined status/count header. The usb-serial core sizes the interrupt-in buffer from the endpoint descriptor's wMaxPacketSize, and successful interrupt transfers can complete short when URB_SHORT_NOT_OK is not set. Check that the completed packet contains the selected header before reading it. Malformed short reports are ignored and the interrupt URB is resubmitted through the existing retry path, preventing out-of-bounds header-byte reads. KASAN report as below: KASAN slab-out-of-bounds in cypress_read_int_callback+0x240/0x7f0 Read of size 1 Call trace: cypress_read_int_callback() (drivers/usb/serial/cypress_m8.c:1009) __usb_hcd_giveback_urb() dummy_timer() [johan: use constants in header length sanity checks]	N/A	More Details
CVE-2026-63877	In the Linux kernel, the following vulnerability has been resolved: serial: dz: Convert to use a platform device Prevent a crash from happening as the first serial port is initialised: Console: switching to colour frame buffer device 160x64 tgafb: SFB+ detected, rev=0x02 fb0: Digital ZLX-E1 frame buffer device at 0x1e000000 DECstation DZ serial driver version 1.04 CPU 0 Unable to handle kernel paging request at virtual address 000000bc, epc == 8048b3a4, ra == 80470a78 Oops[#1]: CPU: 0 UID: 0 PID: 1 Comm: swapper/0 Not tainted 6.19.0-dirty #35 NONE $ 0 : 00000000 1000ac00 00000004 804707ac $ 4 : 00000000 80e20850 80e20858 81000030 $ 8 : 00000000 8072c81c 00000008 fefefeff $12 : 6c616972 00000006 80c5917f 69726420 $16 : 80e20800 00000000 808f8968 80e20800 $20 : 00000000 807f5a90 808b0094 808d3bc8 $24 : 00000018 80479030 $28 : 80c2e000 80c2fd70 00000069 80470a78 Hi : 00000004 Lo : 00000000 epc : 8048b3a4 __dev_fwnode+0x0/0xc ra : 80470a78 serial_base_ctrl_add+0xa0/0x168 Status: 1000ac04 IEp Cause : 30000008 (ExcCode 02) BadVA : 000000bc PrId : 00000220 (R3000) Modules linked in: Process swapper/0 (pid: 1, threadinfo=(ptrval), task=(ptrval), tls=00000000) Stack : 00400044 00400040 8046f4cc 00000000 808a6148 808a0000 808f8968 8086983c 808e0000 8046fc84 1000ac01 00000028 80e20700 802ba3f8 80e20700 80d34a94 80c1b900 80e20700 80e20700 80e20700 80e20700 80444650 00000000 00000000 00000000 807f5a90 808b0094 80447080 00400040 808e0000 80d34a94 808a6148 80d34a94 00000004 80e20700 00000000 8076974c 80469810 80c2fe3c 1000ac01 ... Call Trace: [<8048b3a4>] __dev_fwnode+0x0/0xc [<80470a78>] serial_base_ctrl_add+0xa0/0x168 [<8046fc84>] serial_core_register_port+0x1c8/0x974 [<808c6af0>] dz_init+0x74/0xc8 [<800470e0>] do_one_initcall+0x44/0x2d4 [<808b111c>] kernel_init_freeable+0x258/0x308 [<8072e434>] kernel_init+0x20/0x114 [<80049cd0>] ret_from_kernel_thread+0x14/0x1c Code: 27bd0018 03e00008 2402ffea <8c8200bc> 03e00008 00000000 27bdffc0 afbe0038 afb30024 ---[end trace 0000000000000000]--- -- where a pointer is dereferenced that has been derived from a null pointer to the port's parent device. Since no device is available with legacy probing and it's not anymore a preferable way to discover devices anyway, switch the driver to using a platform device and use it as the port's parent device. Update resource handling accordingly and only request the actual span of addresses used within the slot, which will have had its resource already requested by generic platform device code. Use platform_driver_probe() not just because the DZ device is fixed with solder on board and not straightforward to remove, but foremost because the associated TTY's major device number is the same as used by the zs driver and the first driver to claim it will prevent the other one from using it. Either one DZ device or some SCC devices will be present in a given system but never both at a time, and therefore we want the major device number to be claimed by the first driver to actually successfully bind to its device and platform_driver_probe() is a way to fulfil that. An unfortunate consequence of the switch to a platform device is we now hand the console over from the bootconsole much later in the bootstrap. The firmware console handler appears good enough though to work so late and in particular with interrupts enabled. Conversely only starting the console port so late lets the reset code fully utilise our delay handlers, so switch from udelay() to fsleep() for transmitter draining so as to avoid busy-waiting for an excessive amount of time.	N/A	More Details
CVE-2026-63882	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: fix NULL pointer bug in svm_range_set_attr The process_info could be NULL if user doesn't call kfd_ioctl_acquire_vm before calling kfd_ioctl_svm. (cherry picked from commit 83a26c812e0529eb040d31a76f73e33e637243d4)	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: serial: zs: Convert to use a platform device Prevent a crash from happening as the first serial port is initialised: Console: switching to mono frame buffer		

CVE-2026-63876	device 160x64 fb0: PMAG-AA frame buffer device at tc0 DECstation Z85C30 serial driver version 0.10 CPU 0 Unable to handle kernel paging request at virtual address 0000002c, epc == 803ab00c, ra == 803aafe0 Oops[#1]: CPU: 0 PID: 1 Comm: swapper Not tainted 6.4.0-rc3-00031-g84a9582fd203-dirty #57 $ 0 : 00000000 10012c00 803aaeb0 00000000 $ 4 : 80e12f60 80e12f50 80e12f58 81000030 $ 8 : 00000000 805ff37c 00000000 33433538 $12 : 65732030 00000006 80c2915d 6c616972 $16 : 80e12f00 807b7630 00000000 00000000 $20 : 00000004 00000348 000001a0 807623b8 $24 : 00000018 00000000 $28 : 80c24000 80c25d60 8078b148 803aafe0 Hi : 00000000 Lo : 00000000 epc : 803ab00c serial_base_ctrl_add+0x78/0xf4 ra : 803aafe0 serial_base_ctrl_add+0x4c/0xf4 Status: 10012c03 KERNEL EXL IE Cause : 00000008 (ExcCode 02) BadVA : 0000002c PrId : 00000440 (R4400SC) Modules linked in: Process swapper (pid: 1, threadinfo=(ptrval), task=(ptrval), tls=00000000) Stack : 80760000 00000cc0 00400044 00400040 803aa02c 80d61ab8 00000000 807b7630 80760000 807623b8 807b7628 803aa644 80386998 00000000 80e17780 80220f68 80e17780 80d61ab8 80c17d80 80e17780 80e17780 8063c798 80e17780 80383fa0 00000010 80e17780 00000000 80386998 807a0000 00000000 00400040 8038f848 807623b8 80d61ab8 00000004 80e17780 00000000 803a68e4 80c25e2c 803bb884 ... Call Trace: [<803ab00c>] serial_base_ctrl_add+0x78/0xf4 [<803aa644>] serial_core_register_port+0x174/0x69c [<8077e9ac>] zs_init+0xc8/0xfc [<800404d4>] do_one_initcall+0x40/0x2ac [<8076cecc>] kernel_init_freeable+0x1e4/0x270 [<80605bec>] kernel_init+0x20/0x108 [<800431e8>] ret_from_kernel_thread+0x14/0x1c Code: 2442aeb0 ae120024 ae0200d0 <8c67002c> 50e00001 8c670000 3c06806e 3c05806e afb30010 ---[end trace 0000000000000000]--- (report at the offending commit) -- where a pointer is dereferenced that has been derived from a null pointer to the port's parent device. Since no device is available with legacy probing and it's not anymore a preferable way to discover devices anyway, switch the driver to using a platform device and use it as the port's parent device. Update resource handling accordingly and only request the actual span of addresses used within the slot, which will have had its resource already requested by generic platform device code. Use platform_driver_probe() not just because SCC devices are fixed with solder on board and not straightforward to remove, but foremost because the associated TTY's major device number is the same as used by the dz driver and the first driver to claim it will prevent the other one from using it. Either one DZ device or some SCC devices will be present in a given system but never both at a time, and therefore we want the major device number to be claimed by the first driver to actually successfully bind to its device and platform_driver_probe() is a way to fulfil that. An unfortunate consequence of the switch to a platform device is we now hand the console over from the bootconsole much later in the bootstrap. The firmware console handler appears good enough though to work so late and in particular with interrupts enabled. Since there is one way only remaining to reach zs_reset() now, remove the port initialisation marker as no longer needed and go through the channel reset unconditionally.	N/A	More Details
CVE-2026-63873	In the Linux kernel, the following vulnerability has been resolved: accel/amdxdna: Fix mm_struct reference leak in aie2_populate_range() aie2_populate_range() jumps back to the again label without calling mmput(mm), leaking a reference to the mm_struct. Add the missing mmput() before jumping to again.	N/A	More Details
CVE-2026-59866	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.5, Kiota emitted x-ms-kiota-info clientClassName and clientNamespaceName values without identifier or path sanitization as both generated client class or namespace names and generated output path components when `kiota generate` ran without -c/--class-name, allowing an attacker-controlled or compromised OpenAPI description to write generated source outside the -o output directory and inject arbitrary text into generated class or namespace declarations. This issue is fixed in version 1.32.5 by GenerationConfiguration.SanitizeClientClassName and SanitizeClientNamespaceName.	N/A	More Details
CVE-2026-63890	In the Linux kernel, the following vulnerability has been resolved: scsi: fcoe: Reject FIP descriptors with zero fip_dlen in CVL walker drivers/scsi/fcoe/fcoe_ctlr.c::fcoe_ctlr_recv_clr_vlink() advanced the descriptor cursor by an attacker-supplied fip_dlen without ever requiring dlen >= sizeof(struct fip_desc) in the default branch. The named descriptor cases (FIP_DT_MAC, FIP_DT_NAME, FIP_DT_VN_ID) checked their per-type minimum lengths, but a FIP_DT_NON_CRITICAL descriptor (fip_dtype >= 128, which the standard requires receivers to silently ignore) skipped that check entirely. An unauthenticated L2 peer on the FCoE control VLAN could hang fcoe_ctlr_recv_work on an fcoe, qedf, or bnx2fc initiator indefinitely by emitting one FIP CVL frame whose single descriptor had fip_dtype == FIP_DT_NON_CRITICAL and fip_dlen == 0: the cursor advanced zero bytes per iteration and the loop condition rlen >= sizeof(*desc) stayed true forever, blocking every subsequent FIP frame on that controller. Tighten the outer dlen guard to also reject dlen < sizeof(struct fip_desc), so a malformed descriptor whose length cannot even cover the descriptor header is rejected before the switch. This is the same lower-bound the named cases already apply and is the minimum scope that closes the loop.	N/A	More Details
CVE-2026-63891	In the Linux kernel, the following vulnerability has been resolved: thunderbolt: property: Cap recursion depth in __tb_property_parse_dir() A DIRECTORY entry's value field is used as the dir_offset for a recursive call into __tb_property_parse_dir() with no depth counter. A crafted peer that chains DIRECTORY entries into a back-reference loop drives the parser until the kernel stack is exhausted and the guard page fires. Any untrusted XDomain peer (cable, dock, in-line inspector, adjacent host) that reaches the PROPERTIES_REQUEST control-plane exchange can trigger this without authentication. Thread a depth counter through tb_property_parse() and __tb_property_parse_dir(), and reject blocks that exceed TB_PROPERTY_MAX_DEPTH = 8. That is comfortably larger than any observed legitimate XDomain layout. Operators who do not need XDomain host-to-host discovery can disable the path entirely with thunderbolt.xdomain=0 on the kernel command line.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: thunderbolt: property: Reject dir_len < 4 to		

CVE-2026-63892	prevent size_t underflow On the non-root path, __tb_property_parse_dir() takes dir_len from entry->length (u16 widened to size_t). Two distinct OOB conditions follow when entry->length < 4: 1. The non-root path begins with kmemdup(&block[dir_offset], sizeof(*dir->uuid), ...) which always reads 4 dwords from dir_offset. tb_property_entry_valid() only enforces dir_offset + entry->length <= block_len, so a crafted entry with dir_offset close to the end of the property block and entry->length in 0..3 passes that gate but lets the UUID copy run off the block (e.g. dir_offset = 497, dir_len = 3 in a 500-dword block reads block[497..501]). 2. After the kmemdup, content_len = dir_len - 4 underflows size_t to ~SIZE_MAX, nentries becomes SIZE_MAX / 4, and the entry walk runs OOB on each iteration until an entry fails validation or the kernel oopses on an unmapped page. Reject dir_len < 4 on the non-root path *before* the UUID kmemdup, which closes both holes. Also move INIT_LIST_HEAD(&dir->properties) up to immediately after the dir allocation so the new error-return path (and the existing uuid-alloc failure path) calling tb_property_free_dir() sees a walkable list rather than the zero-initialized NULL next/prev that list_for_each_entry_safe() would oops on.	N/A	More Details
CVE-2026-59865	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.5, `kiota info` read x-ms-kiota-info.languagesInformation.<language>.dependencyInstallCommand plus dependency name and version values from an OpenAPI description and presented the spec-supplied command as Kiota's recommended install command, allowing an attacker-controlled or compromised description to cause command injection when the suggested command was run manually or through the Kiota VS Code extension's kiota info --json dependency-install flow. This issue is fixed in version 1.32.5.	N/A	More Details
CVE-2026-59864	Kiota is an OpenAPI based HTTP Client code generator. Prior to 1.32.5, `kiota plugin add` and `kiota plugin generate` (with `-t APIPlugin`) emitted attacker-controlled static_template.file values from x-ai-adaptive-card and x-ai-capabilities into generated Microsoft 365 Copilot and Teams plugin manifests without path validation, allowing ../, absolute, rooted, UNC, Windows drive, or URI paths in response_semantics.static_template.file to cause path traversal or out-of-package file inclusion when the generated plugin was deployed. This issue is fixed in version 1.32.5.	N/A	More Details
CVE-2026-63895	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: copy only received bytes on short ep0 read ffs_ep0_read() allocates its control-OUT data buffer with kmalloc() (not kzalloc) at the Length value from the Setup packet, then copies that full len to userspace regardless of how many bytes were actually received: data = kmalloc(len, GFP_KERNEL); ... ret = __ffs_ep0_queue_wait(ffs, data, len); if ((ret > 0) && (copy_to_user(buf, data, len))) ret = -EFAULT; __ffs_ep0_queue_wait() returns req->actual, which on a short control OUT transfer is strictly less than len. The copy_to_user() call still copies len bytes, so on a short OUT the last (len - ret) bytes of the kmalloc() buffer -- uninitialised slab residue -- are delivered to the FunctionFS daemon. Short ep0 OUT completions are specified USB control-transfer behavior and are produced by in-tree UDCs: * dwc2 continues on req->actual < req->length for ep0 DATA OUT (short-not-ok is the only ep0-OUT stall path). * aspeed_udc ends ep0 OUT on rx_len < ep->ep.maxpacket. * renesas_usbf logs "ep0 short packet" and completes the request. * dwc3 stalls on short IN but not on short OUT. A short ep0 OUT is therefore not evidence of a broken UDC; it is a normal condition f_fs has to cope with. The sibling gadgetfs implementation in drivers/usb/gadget/legacy/inode.c already does this correctly via min(len, dev->req->actual) before copy_to_user(). This patch brings f_fs.c to the same safe pattern rather than trimming at a defensive layer. The bug is reached from the FunctionFS device node, which in real deployments is owned by the privileged gadget daemon (adbd, UMS, composite gadget services, etc.); it is not reachable from unprivileged userspace. Linux host stacks normally reject short-wLength control OUTs before they reach the gadget, so reproducing this required a build that bypasses that host-side check. With the bypass in place, a 1-byte payload on a 64-byte Setup produces 63 bytes of non-canary slab residue in the daemon's read buffer. Fix by copying only ret (actually received) bytes to userspace.	N/A	More Details
CVE-2026-63896	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: composite: fix integer underflow in WebUSB GET_URL handling The WebUSB GET_URL handler in composite_setup() narrows landing_page_length to fit the host-supplied wLength using landing_page_length = w_length - WEBUSB_URL_DESCRIPTOR_HEADER_LENGTH + landing_page_offset; If wLength is smaller than WEBUSB_URL_DESCRIPTOR_HEADER_LENGTH the unsigned subtraction wraps, and the subsequent memcpy(url_descriptor->URL, cdev->landing_page + landing_page_offset, landing_page_length - landing_page_offset); ends up copying close to UINT_MAX bytes from cdev->landing_page into cdev->req->buf. KASAN reports a slab-out-of-bounds in composite_setup on the kmalloc-2k gadget_info allocation, and FORTIFY_SOURCE traps the memcpy as a 4294967293-byte field-spanning write into url_descriptor->URL (size 252). A USB host can reach this from a single SETUP packet against any gadget that has webusb/use=1 and a landingPage configured. Handle the small-wLength case before the math: when the host requested fewer bytes than the URL descriptor header, only the header is meaningful and no URL bytes need to be copied. Setting landing_page_length to landing_page_offset makes the existing memcpy a no-op and leaves the descriptor returned to the host unchanged for all larger wLength values.	N/A	More Details
CVE-2026-63897	In the Linux kernel, the following vulnerability has been resolved: USB: serial: mct_u232: fix missing interrupt-in transfer sanity check Add the missing sanity check on the size of interrupt-in transfers to avoid parsing stale or uninitialised slab data (and leaking it to user space).	N/A	More Details
CVE-2026-63898	In the Linux kernel, the following vulnerability has been resolved: USB: serial: mct_u232: fix memory corruption with small endpoint The driver overrides the maximum transfer size for a specific device which only accepts 16 byte packets for its 32 byte bulk-out endpoint. Make sure to never increase the maximum transfer size to prevent slab corruption should a malicious device report a smaller endpoint max packet size	N/A	More Details

	than expected.		
CVE-2026-63899	In the Linux kernel, the following vulnerability has been resolved: USB: serial: mxuport: fix memory corruption with small endpoint Make sure that the bulk-out endpoint max packet size is at least eight bytes to avoid user-controlled slab corruption should a malicious device report a smaller size.	N/A	More Details
CVE-2026-63900	In the Linux kernel, the following vulnerability has been resolved: USB: serial: keyspan: fix missing indat transfer sanity check Add the missing sanity check on the size of usa49wg indat transfers to avoid parsing stale or uninitialised slab data.	N/A	More Details
CVE-2026-63901	In the Linux kernel, the following vulnerability has been resolved: USB: serial: digi_acceleport: fix memory corruption with small endpoints Add the missing bulk-out buffer size sanity checks to avoid out-of-bounds memory accesses or slab corruption should a malicious device report smaller buffers than expected.	N/A	More Details
CVE-2026-63933	In the Linux kernel, the following vulnerability has been resolved: iio: gyro: adis16260: fix division by zero in write_raw Add a validation check for the sampling frequency value before using it as a divisor. A user writing zero to the sampling_frequency sysfs attribute triggers a division by zero in the kernel.	N/A	More Details
CVE-2026-63935	In the Linux kernel, the following vulnerability has been resolved: iio: adc: nxp-sar-adc: fix division by zero in write_raw Add a validation check for the sampling frequency value before using it as a divisor. A user writing zero or a negative value to the sampling_frequency sysfs attribute triggers a division by zero in the kernel. Also prevent unsigned integer underflow when the computed cycle count is smaller than NXP_SAR_ADC_CONV_TIME, which would wrap the u32 inpsamp to a huge value.	N/A	More Details
CVE-2026-63997	In the Linux kernel, the following vulnerability has been resolved: ethtool: module: avoid leaking a netdev ref on module flash errors module_flash_fw_schedule() is missing undo for setting the "in_progress" flag and taking the netdev reference. Delay taking these, the device can't disappear while we are holding rtnl_lock.	N/A	More Details
CVE-2026-63963	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpm: validate VDO count in Discover Identity ACK handlers Properly validate the count passed from a device when calling svdm_consume_identity() or svdm_consume_identity_sop_prime() as the device-controlled value could index off of the static arrays, which could leak data.	N/A	More Details
CVE-2026-63965	In the Linux kernel, the following vulnerability has been resolved: iio: pressure: bmp280: fix stack leak in bmp580 trigger handler bmp580_trigger_handler() declares its scan buffer on the stack without an initializer and then memcpy()s 3 bytes of 24-bit sensor data into each 4-byte __le32 field. The high byte of comp_temp and comp_press is left uninitialized, and the channel storagebits is 32, so two bytes of stack are pushed to userspace per scan. This is a regression from when the buffer lived in the private data, the move to a stack-local struct dropped the implicit zeroing. bme280_trigger_handler() was fixed up to handle this bug, but this driver was not fixed because there was no padding hole, but rather a short-fill issue. Fix this all by just zero-initializing the structure on the stack.	N/A	More Details
CVE-2026-63966	In the Linux kernel, the following vulnerability has been resolved: iio: imu: adis16550: fix stack leak in trigger handler adis16550_trigger_handler() declares the scan data array on the stack without initializing it. The memcpy() at the bottom fills only the first 28 bytes (TEMP + 6 channels of GYRO/ACCEL data), and iio_push_to_buffers_with_timestamp() writes the s64 timestamp at the 8-byte-aligned offset 32. Bytes 28-31 remain uninitialized stack data which leaks to userspace on ever trigger. Fix this all by just zero-initializing the structure on the stack.	N/A	More Details
CVE-2026-63967	In the Linux kernel, the following vulnerability has been resolved: iio: imu: st_lsm6dsx: fix stack leak in tagged FIFO buffer The tagged FIFO path declares iio_buff on the stack with __aligned(8) but no initializer, but there is a hole in the structure, which will then leak to userspace as ST_LSM6DSX_SAMPLE_SIZE bytes (6) will be copied, but the space between that and the timestamp are not initialized. Commit c14edb4d0bdc ("iio:imu:st_lsm6dsx Fix alignment and data leak issues") moved the untagged FIFO path to a kzalloc'd buffer in hw->scan, but for the tagged path it only added the alignment qualifier and not the initializer :(Fix this by just zero-initializing the structure on the stack.	N/A	More Details
CVE-2026-6424	Use-after-free vulnerability in ESET Linux products potentially allowed an attacker to trigger kernel panic on the system	N/A	More Details
CVE-2026-63969	In the Linux kernel, the following vulnerability has been resolved: ipv6: fix possible infinite loop in rt6_fill_node() Sashiko reported this issue [1]. Apply the same fix as commit f8d8ce1b515a ("ipv6: fix possible infinite loop in fib6_info_uses_dev()"). Writers holding tb6_lock can list_del_rcu(&rt->fib6_siblings) without waiting for RCU readers; rt->fib6_siblings.next then still points into the old ring and this softirq-side walker never reaches &rt->fib6_siblings, causing a CPU stall. fib6_del_route() always WRITE_ONCE()s rt->fib6_nsiblings to 0 before list_del_rcu(), so an inside-loop check is a reliable detach signal. [1] https://sashiko.dev/#/patchset/20260526020227.4857-1-jiayuan.chen%40linux.dev	N/A	More Details
CVE-2026-6423	A local privilege escalation vulnerability in ESET Inspect Connector. The vulnerability was caused by improper authentication in an IPC channel.	N/A	More Details

CVE-2026-58078	The Joomla extension Quix Page Builder Pro is vulnerable to an unauthenticated SQL injection.	N/A	More Details
CVE-2026-63973	In the Linux kernel, the following vulnerability has been resolved: net: mana: Add NULL guards in teardown path to prevent panic on attach failure When queue allocation fails partway through, the error cleanup frees and NULLs apc->tx_qp and apc->rxqs. Multiple teardown paths such as mana_remove(), mana_change_mtu() recovery, and internal error handling in mana_alloc_queues() can subsequently call into functions that dereference these pointers without NULL checks: - mana_chn_setxdp() dereferences apc->rxqs[0], causing a NULL pointer dereference panic (CR2: 0000000000000000 at mana_chn_setxdp+0x26). - mana_destroy_vport() iterates apc->rxqs without a NULL check. - mana_fence_rqs() iterates apc->rxqs without a NULL check. - mana_dealloc_queues() iterates apc->tx_qp without a NULL check. Add NULL guards for apc->rxqs in mana_fence_rqs(), mana_destroy_vport(), and before the mana_chn_setxdp() call. Add a NULL guard for apc->tx_qp in mana_dealloc_queues() to skip TX queue draining when TX queues were never allocated or already freed.	N/A	More Details
CVE-2026-63981	In the Linux kernel, the following vulnerability has been resolved: net/sched: act_mirred: Fix blockcast recursion bypass leading to stack overflow tcf_mirred_act() checks sched_mirred_nest against MIRRED_NEST_LIMIT (4) to prevent deep recursion. However, when the action uses blockcast (tcfm_blockid != 0), the function returns at the tcf_blockcast() call BEFORE reaching the counter increment. As a result, the recursion counter never advances and the limit check is entirely bypassed. When two devices share a TC egress block with a mirred blockcast rule, a packet egressing on device A is mirrored to device B via blockcast; device B's egress TC re-enters tcf_mirred_act() via blockcast and mirrors back to A, creating an unbounded recursion loop: tcf_mirred_act -> tcf_blockcast -> tcf_mirred_to_dev -> dev_queue_xmit -> sch_handle_egress -> tcf_classify -> tcf_mirred_act -> (repeat) This recursion continues until the kernel stack overflows. The bug is reachable from an unprivileged user via unshare(CLONE_NEWUSER | CLONE_NEWNET): user namespaces grant CAP_NET_ADMIN in the new network namespace, which is sufficient to create dummy devices, attach clsact qdiscs with shared blocks, and install mirred blockcast filters. BUG: TASK stack guard page was hit at ffffc90000b7fff8 Oops: stack guard page: 0000 [#1] SMP KASAN NOPTI CPU: 2 UID: 1000 PID: 169 Comm: poc Not tainted 7.0.0-rc7-next-20260410 RIP: 0010:xas_find+0x17/0x480 Call Trace: xa_find+0x17b/0x1d0 tcf_mirred_act+0x640/0x1060 tcf_action_exec+0x400/0x530 basic_classify+0x128/0x1d0 tcf_classify+0xd83/0x1150 tc_run+0x328/0x620 __dev_queue_xmit+0x797/0x3100 tcf_mirred_to_dev+0x7b1/0xf70 tcf_mirred_act+0x68a/0x1060 [repeating ~30+ times until stack overflow] Kernel panic - not syncing: Fatal exception in interrupt Fix this by incrementing sched_mirred_nest before calling tcf_blockcast() and decrementing it on return, mirroring the non-blockcast path. This ensures subsequent recursive entries see the updated counter and are correctly limited by MIRRED_NEST_LIMIT.	N/A	More Details
CVE-2026-63982	In the Linux kernel, the following vulnerability has been resolved: net/sched: Fix ethx:ingress -> ethy:egress -> ethx:ingress mirred loop When mirred redirects to ingress (from either ingress or egress) the loop state from sched_mirred_dev array dev is lost because of 1) the packet deferral into the backlog and 2) the fact the sched_mirred_dev array is cleared. In such cases, if there was a loop we won't discover it. Here's a simple test to reproduce: ip a add dev port0 10.10.10.11/24 tc qdisc add dev port0 clsact tc filter add dev port0 egress protocol ip \ prio 10 matchall action mirred ingress redirect dev port1 tc qdisc add dev port1 clsact tc filter add dev port1 ingress protocol ip \ prio 10 matchall action mirred egress redirect dev port0 ping -c 1 -W0.01 10.10.10.10	N/A	More Details
CVE-2026-63983	In the Linux kernel, the following vulnerability has been resolved: net/sched: fix packet loop on netem when duplicate is on When netem duplicates a packet it re-enqueues the copy at the root qdisc. If another netem sits in the tree the copy can be duplicated again, recursing until the stack or memory is exhausted. The original duplication guard temporarily zeroed q->duplicate around the re-enqueue, but that does not cover all cases because it is per-qdisc state shared across all concurrent enqueue paths and is not safe without additional locking. Use the skb tc_depth field introduced in an earlier patch: - increment it on the duplicate before re-enqueue - skip duplication for any skb whose tc_depth is already non-zero. This marks the packet itself rather than mutating qdisc state, therefore it is safe regardless of tree topology or concurrency.	N/A	More Details
CVE-2026-55407	Buffa is a pure-Rust Protocol Buffers implementation with first-class protobuf editions support. Prior to 0.8.0, the decode_unknown_field function in buffa's protobuf decoder allocated heap memory in proportion to untrusted input (unknown fields in the serialized protobuf) without enforcing an allocation budget, affecting any message decoded from untrusted input using code generated with preserve_unknown_fields=true (the default); a small, well-formed payload of nested unknown fields inside a StartGroup could trigger roughly 22x memory amplification (for example a 64 MiB input forcing about 1.4 GB of heap allocation), and length-delimited unknown fields could be sized arbitrarily, so an unauthenticated attacker could crash a process through memory exhaustion because the top-level message size cap did not account for in-decode amplification. This issue is fixed in version 0.8.0.	N/A	More Details
CVE-2026-63986	In the Linux kernel, the following vulnerability has been resolved: ethtool: tsinfo: don't pass ERR_PTR to genlmsg_cancel on prepare failure The goto err label leads to: genlmsg_cancel(skb, ehdr); return ret; If ethnl_tsinfo_prepare_dump() failed, it has not started a genlmsg. There's nothing to cancel, and passing an error pointer to genlmsg_cancel() would cause a crash.	N/A	More Details

CVE-2026-63988	In the Linux kernel, the following vulnerability has been resolved: bridge: Fix sleep in atomic context in sysfs path Since the start of the git history, brport_store() always acquired the bridge lock. Back then this decision made sense: The bridge lock protects the STP state of the bridge and its ports and at that time the function was only used by two STP related attributes (cost and priority). Nowadays, brport_store() processes a lot more attributes and most of them do not need the bridge lock: * Bridge flags: Only require RTNL. Read locklessly by the data path. Annotations can be added in net-next. * FDB port flushing: Only requires the FDB lock. * Multicast attributes: Only require the multicast lock. * Group forward mask: Only requires RTNL. Read locklessly by the data path. Annotations can be added in net-next. * Backup port: Only requires RTNL. Read locklessly by the data path. This is a problem as the bridge calls dev_set_promiscuity() when certain bridge port flags change and this function can sleep since the commit cited below, resulting in a splat such as [1]. Fix this by reducing the scope of the bridge lock and only take it when processing the two STP related attributes that require it. Remove the now stale comment from br_switchdev_set_port_flag(). The SWITCHDEV_F_DEFER flag can be removed in net-next. [1] BUG: sleeping function called from invalid context at net/core/dev_addr_lists.c:1262 in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 372, name: bash preempt_count: 201, expected: 0 RCU nest depth: 0, expected: 0 5 locks held by bash/372: #0: ffff88810c51c3f0 (sb_writers#7){.+.+}-{0:0}, at: ksys_write (fs/read_write.c:740) #1: ffff888115ce9480 (&of->mutex){+.+.}-{4:4}, at: kernfs_fop_write_iter (fs/kernfs/file.c:343) #2: ffff88810b9fd330 (kn->active#37){.+.+}-{0:0}, at: kernfs_fop_write_iter (fs/kernfs/file.c:80 fs/kernfs/file.c:344) #3: ffffffffa59473a0 (rtnl_mutex){+.+.}-{4:4}, at: brport_store (net/bridge/br_sysfs_if.c:326) #4: ffff8881099d2d58 (&br->lock){+...}-{3:3}, at: brport_store (./include/linux/spinlock.h:348 net/bridge/br_sysfs_if.c:345) Preemption disabled at: 0x0 Hardware name: Bochs Bochs, BIOS Bochs 01/01/2011 Call Trace: <TASK> dump_stack_lvl (lib/dump_stack.c:94 lib/dump_stack.c:120) __might_resched.cold (kernel/sched/core.c:9163) netif_rx_mode_run (net/core/dev_addr_lists.c:1262) netif_rx_mode_sync (net/core/dev_addr_lists.c:1428) dev_set_promiscuity (net/core/dev_api.c:289) br_manage_promisc (net/bridge/br_if.c:135 net/bridge/br_if.c:172) br_port_flags_change (net/bridge/br_if.c:242 net/bridge/br_if.c:747) store_learning (net/bridge/br_sysfs_if.c:79 net/bridge/br_sysfs_if.c:235) brport_store (net/bridge/br_sysfs_if.c:346) kernfs_fop_write_iter (fs/kernfs/file.c:352) new_sync_write (fs/read_write.c:595) vfs_write (fs/read_write.c:688) ksys_write (fs/read_write.c:740) do_syscall_64 (arch/x86/entry/syscall_64.c:63 arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121)	N/A	More Details
CVE-2026-63989	In the Linux kernel, the following vulnerability has been resolved: bridge: Fix sleep in atomic context in netlink path Since the introduction of the netlink configuration path for bridge ports in commit 25c71c75ac87 ("bridge: bridge port parameters over netlink"), br_setport() was always called with the bridge lock held around it. Back then this decision made sense: The bridge lock protects the STP state of the bridge and its ports and at that time the function only processed three STP related netlink attributes (cost, priority and state). Nowadays, br_setport() processes a lot more attributes and most of them do not need the bridge lock: * Bridge flags: Only require RTNL. Read locklessly by the data path. Annotations can be added in net-next. * FDB port flushing: Only requires the FDB lock. * Multicast attributes: Only require the multicast lock. * Group forward mask: Only requires RTNL. Read locklessly by the data path. Annotations can be added in net-next. * Backup port and NHID: Only require RTNL. Read locklessly by the data path. This is a problem as the bridge calls dev_set_promiscuity() when certain bridge port flags change and this function can sleep since the commit cited below, resulting in a splat such as [1]. Fix this by reducing the scope of the bridge lock and only take it when processing the three STP related attributes that require it. This is consistent with the multicast attributes where each attribute acquires the multicast lock instead of having one critical section for all relevant attributes. [1] BUG: sleeping function called from invalid context at net/core/dev_addr_lists.c:1262 in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 356, name: bridge preempt_count: 201, expected: 0 RCU nest depth: 0, expected: 0 2 locks held by bridge/356: #0: ffffffff919473a0 (rtnl_mutex){+.+.}-{4:4}, at: rtnetlink_rcv_msg (net/core/rtnetlink.c:80 net/core/rtnetlink.c:7002) #1: ffff888115072d58 (&br->lock){+...}-{3:3}, at: br_setlink (./include/linux/spinlock.h:348 net/bridge/br_netlink.c:1117) Preemption disabled at: 0x0 Hardware name: Bochs Bochs, BIOS Bochs 01/01/2011 Call Trace: <TASK> dump_stack_lvl (lib/dump_stack.c:94 lib/dump_stack.c:120) __might_resched.cold (kernel/sched/core.c:9163) netif_rx_mode_run (net/core/dev_addr_lists.c:1262) netif_rx_mode_sync (net/core/dev_addr_lists.c:1428) dev_set_promiscuity (net/core/dev_api.c:289) br_manage_promisc (net/bridge/br_if.c:135 net/bridge/br_if.c:172) br_port_flags_change (net/bridge/br_if.c:242 net/bridge/br_if.c:747) br_setport (net/bridge/br_netlink.c:1000) br_setlink (net/bridge/br_netlink.c:1118) rtnl_bridge_setlink (net/core/rtnetlink.c:5572) rtnetlink_rcv_msg (net/core/rtnetlink.c:7005) netlink_rcv_skb (net/netlink/af_netlink.c:2550) netlink_unicast (net/netlink/af_netlink.c:1318 net/netlink/af_netlink.c:1344) netlink_sendmsg (net/netlink/af_netlink.c:1894) __sock_sendmsg (net/socket.c:787 (discriminator 4) net/socket.c:802 (discriminator 4)) ____sys_sendmsg (net/socket.c:2698) ___sys_sendmsg (net/socket.c:2752) __sys_sendmsg (net/socket.c:2784) do_syscall_64 (arch/x86/entry/syscall_64.c:63 arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121)	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: bonding: refuse to enslave CAN devices syzbot reported a kernel paging request crash in can_rx_unregister() inside net/can/af_can.c. The crash occurs because a virtual CAN device (vxcan) is being enslaved to a bonding master. During the enslavement process, the bonding driver mutates and modifies the network device states to fit an Ethernet-like aggregation model. However, CAN devices operate on a completely different Layer 2 architecture, relying on the CAN mid-layer private data structure (can_ml_priv) instead of standard Ethernet structures. Since bonding does not initialize or maintain these CAN structures, subsequent operations on the half-enslaved		

CVE-2026-63990	interface (such as closing associated sockets via isotp_release) lead to a null-pointer dereference when accessing the CAN receiver lists. Bonding CAN interfaces is architecturally invalid as CAN lacks MAC addresses, ARP capabilities, and standard Ethernet link-layer mechanisms. While generic loopback devices are blocked globally in net/core/dev.c, virtual CAN devices bypass this check because they do not carry the IFF_LOOPBACK flag, despite acting as local software-loopbacks. Fix this by explicitly blocking network devices of type ARPHRD_CAN from being enslaved at the very beginning of bond_enslave(). This prevents illegal state mutations, eliminates the resulting KASAN crashes, and avoids potential memory leaks from incomplete socket cleanups. As the CAN support has been added a long time after bonding the Fixes-tag points to the introduction of ARPHRD_CAN that would have needed a specific handling in bonding_main.c.	N/A	More Details
CVE-2026-63991	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: 6lowpan: check skb_clone() return value in send_mcast_pkt() The skb_clone() function can return NULL if memory allocation fails. send_mcast_pkt() calls skb_clone() without checking the return value, which can lead to a NULL pointer dereference in send_pkt() when it dereferences skb->data. Add a NULL check after skb_clone() and skip the peer if the clone fails.	N/A	More Details
CVE-2026-63964	In the Linux kernel, the following vulnerability has been resolved: usb: typec: ucsi: ccg: reject firmware images without a ':' record header do_flash() locates the first .cyacd record with p = strnchr(fw->data, fw->size, ':'); while (p < eof) { s = strnchr(p + 1, eof - p - 1, ':'); ... } If the firmware image contains no ':' byte, strnchr() returns NULL. NULL compares less than the valid kernel pointer eof, so the loop body runs and strnchr() is called with p + 1 == (void *)1 and a length of roughly (unsigned long)eof, causing a wonderful crash. The not_signed_fw fallthrough earlier in do_flash() and the chip-state branches in ccg_fw_update_needed() allow an unsigned blob to reach this loop, so a root user who can place a crafted file under /lib/firmware and write the do_flash sysfs attribute can trigger the oops. Bail out with -EINVAL when the initial strnchr() returns NULL.	N/A	More Details
CVE-2026-63962	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpm: bound altmode_desc[] per iteration in svdm_consume_modes() svdm_consume_modes() checks pmdata->altmodes against the array size once before the loop over the count, but forgot to check the bound at every point in the loop. In the well-behaved SVDM discovery flow this is harmless because each of at most SVID_DISCOVERY_MAX SVIDs contributes at most MODE_DISCOVERY_MAX modes, exactly filling altmode_desc[ALTMODE_DISCOVERY_MAX]. But the CMDT_RSP_ACK handler in tcpm_pd_svdm() does not correlate an incoming ACK with any request the port actually sent. Once port->partner is set, an unsolicited Discover Modes ACK is consumed unconditionally. A broken or malicious port partner can therefore drive altmodes to ALTMODE_DISCOVERY_MAX - 1 via the normal flow, and then send one extra Discover Modes ACK with seven VDOs. Because the pre-loop check passes, the loop could then writes up to five entries past altmode_desc[]. For mode_data_prime the next field in struct tcpm_port is the partner_altmode[] pointer array, which then receives partner-chosen SVID/VDO bytes. Move the bound check inside the loop so the array can never be indexed past ALTMODE_DISCOVERY_MAX regardless of how many VDOs the partner supplies or how the function was reached.	N/A	More Details
CVE-2026-63936	In the Linux kernel, the following vulnerability has been resolved: iio: adc: mt6359: fix unchecked return value in mt6358_read_imp In mt6358_read_imp(), the variable val_v is passed to regmap_read() but the return value is not checked. If the read fails, val_v remains uninitialized and its random stack content is subsequently reported as a measurement result. Initialize val_v to zero to ensure a predictable value is reported in case of bus failure and to prevent potential stack data leakage. This also satisfies static analyzers that might otherwise flag the variable as used uninitialized.	N/A	More Details
CVE-2026-63961	In the Linux kernel, the following vulnerability has been resolved: usb: typec: altmodes/displayport: validate count before reading Status Update VDO A broken/malicious device can send the incorrect count for a status update VDO, which will cause the kernel to read uninitialized stack data and send it off elsewhere. Fix this up by correctly verifying the count for the update object.	N/A	More Details
CVE-2026-63943	In the Linux kernel, the following vulnerability has been resolved: Input: xpad - fix out-of-bounds access for Share button xpadone_process_packet() receives len directly from urb->actual_length and uses it to index the share-button byte at data[len - 18] or data[len - 26]. Since both len and data[0] are under the device's control, a broken controller can send a GIP_CMD_INPUT packet with actual_length < 18 (e.g. 5 bytes) and reach this code path, causing accesses beyond the actual array. Fix this by calculating the offset and checking bounds against the packet length.	N/A	More Details
CVE-2026-47751	Claude Code Action is a general-purpose GitHub action that runs Claude Code on GitHub pull requests and issues. Prior to 1.0.74, because the action checked out attacker-controlled pull request head branches, read .mcp.json from the working directory via default setting sources, and unconditionally enabled all project MCP servers via enableAllProjectMcpServers, an attacker who opened a pull request containing a malicious .mcp.json file could achieve arbitrary code execution on the GitHub Actions runner and exfiltrate secrets available to the workflow (such as API keys and tokens) when a privileged user or an automatic trigger invoked the Claude action on the pull request. This issue is fixed in version 1.0.74, which restores .claude/ and .mcp.json from the pull request base branch before the CLI runs.	N/A	More Details
CVE-	stoatchat (delta/Revolt) versions from 20241213-1 before 20250210-1 allow users with only ViewChannel (read) permission on a channel to fetch that channel's webhooks, including their tokens, because the		

2025-71388	webhook fetch endpoint checked for ViewChannel instead of ManageWebhooks. Using a retrieved token, an attacker can send arbitrary messages to the channel, bypassing channel permissions and impersonating a bot or webhook. Fixed in 20250210-1 (0.8.2).	N/A	More Details
CVE-2025-71377	stoatchat (delta) versions before 20250210-1 (0.8.2) contain a logic error in the query messages route. When fetching messages 'nearby' another message, the database query can be given a message limit of zero, which the database interprets as 'no limit'. A remote unauthenticated attacker can craft nearby message fetch requests to download an entire channel's message history in a single expensive request, and can send many such requests in parallel, resulting in denial of service through resource exhaustion.	N/A	More Details
CVE-2026-63948	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: fix chan ref leak in l2cap_chan_timeout() on !conn __set_chan_timer() takes a l2cap_chan reference via l2cap_chan_hold() before scheduling the delayed work. The normal path in l2cap_chan_timeout() drops this reference with l2cap_chan_put() at the end, but the early return when chan->conn is NULL skips the put, leaking the reference. Add the missing l2cap_chan_put() before the early return.	N/A	More Details
CVE-2026-63949	In the Linux kernel, the following vulnerability has been resolved: auxdisplay: line-display: fix OOB read on zero-length message_store() linedisp_display() unconditionally reads msg[count - 1] before checking whether count is zero, so a write of zero bytes to the message sysfs attribute hits msg[-1]: write(fd, "", 0); -> message_store(..., buf, count=0) -> linedisp_display(linedisp, buf, count=0) -> msg[count - 1] == '\n' ; OOB read The kernfs write buffer for that store is a 1-byte allocation (kernfs_fop_write_iter() does kmalloc(len + 1) with len == 0), so msg[-1] is a 1-byte read before the slab object. On a KASAN-enabled kernel this trips an out-of-bounds report and panics; on stock kernels it silently reads adjacent slab data and, if that byte happens to be '\n', the following count-- wraps ssize_t 0 to -1 and is then passed to kmemdup_nul(). linedisp_display() is reached from the message_store() sysfs callback (drivers/auxdisplay/line-display.c message attribute, mode 0644) and from the in-tree initial-message setup with count == -1, so the OOB path is only userspace-triggerable via zero-byte writes; vfs_write() does not short-circuit on count == 0 and kernfs_fop_write_iter() dispatches the store callback regardless. Guard the trailing-newline trim with a count check. The existing if (!count) block then takes the clear-display path unchanged. Affects every auxdisplay driver that registers via linedisp_register() / linedisp_attach(): ht16k33, max6959, img-ascii-lcd, seg-led-gpio.	N/A	More Details
CVE-2026-63862	In the Linux kernel, the following vulnerability has been resolved: PCI: mediatek-gen3: Prevent leaking IRQ domains when IRQ not found In mtk_pcie_setup_irq(), the IRQ domains are allocated before the controller's IRQ is fetched. If the latter fails, the function directly returns an error, without cleaning up the allocated domains. Hence, reverse the order so that the IRQ domains are allocated after the controller's IRQ is found. This was flagged by Sashiko during a review of "[PATCH v6 0/7] PCI: mediatek-gen3: add power control support".	N/A	More Details
CVE-2026-59249	Inconsistent interpretation of HTTP requests (HTTP response smuggling) vulnerability in elixir-mint mint allows a malicious HTTP/1 server to desynchronize a strict intermediary and the Mint client on the same pooled connection, enabling response-queue poisoning against subsequent requests that share the connection. The Mint.HTTP1.decode_body/5 function in lib/mint/http1.ex parses the chunk-size line of a Transfer-Encoding: chunked response with Integer.parse(data, 16). RFC 7230 defines chunk-size = 1*HEXDIG and forbids any sign prefix, but Integer.parse/2 accepts an optional leading + or -. A chunk-size line of +5 is accepted as a five-byte chunk; lines of +0 and -0 are accepted as the terminating zero-length chunk and end the message body early. An RFC-strict intermediary in the response path rejects these forms, so the intermediary and the Mint client disagree on where one response ends and the next begins. On a pooled keep-alive connection, an attacker-influenced origin can inject bytes that the client attributes to the next legitimate response on the same connection, poisoning the response queue and corrupting the responses returned to unrelated in-flight requests. This issue affects mint: from 0.1.0 before 1.9.3.	N/A	More Details
CVE-2026-63861	In the Linux kernel, the following vulnerability has been resolved: spi: mtk-snfi: unregister ECC engine on probe failure and remove() callback mtk_snand_probe() registers the on-host NAND ECC engine, but teardown was missing from both probe unwind and remove-time cleanup. Add a devm cleanup action after successful registration so nand_ecc_unregister_on_host_hw_engine() runs automatically on probe failures and during device removal.	N/A	More Details
CVE-2026-63953	In the Linux kernel, the following vulnerability has been resolved: mm/migrate_device: fix pgtable leak in migrate_vma_insert_huge_pmd_page When migrate_vma_insert_huge_pmd_page() jumps to unlock_abort due to a PMD check failure, the pgtable allocated earlier via pte_alloc_one() is never freed, causing a memory leak. Added free_abort label to release the pgtable in error path.	N/A	More Details
CVE-2026-63859	In the Linux kernel, the following vulnerability has been resolved: net: airoha: Add missing bits in airoha_qdma_cleanup_tx_queue() Similar to airoha_qdma_cleanup_rx_queue(), reset DMA TX descriptors in airoha_qdma_cleanup_tx_queue routine. Moreover, reset TX_DMA_IDX to TX_CPU_IDX to notify the NIC the QDMA TX ring is empty.	N/A	More Details
CVE-	Buffa is a pure-Rust Protocol Buffers implementation with first-class protobuf editions support. Prior to 0.7.0, a soundness bug in the OwnedView<V> type allowed safe Rust code to trigger a use-after-free: the OwnedView::decode constructor transmuted a borrowed slice to &'static [u8], and the Deref implementation exposed the promoted 'static lifetime on borrowed view fields (such as &'static str and &'static [u8]) to		More

2026-55406	callers, so the borrow checker permitted those references to outlive the OwnedView; once the OwnedView was dropped and its backing buffer freed, the references became dangling, enabling memory corruption, information disclosure of freed heap contents, and cross-thread misuse without any unsafe code in the calling application. This issue is fixed in version 0.7.0.	N/A	Details
CVE-2026-63956	In the Linux kernel, the following vulnerability has been resolved: USB: serial: cypress_m8: fix memory corruption with small endpoint Make sure that the interrupt-out endpoint max packet size is at least eight bytes to avoid user-controlled slab corruption or NULL-pointer dereference should a malicious device report a smaller size.	N/A	More Details
CVE-2026-63957	In the Linux kernel, the following vulnerability has been resolved: USB: serial: safe_serial: fix memory corruption with small endpoint Make sure that the bulk-out buffer size is at least eight bytes to avoid user-controlled slab corruption in "safe" mode should a malicious device report a smaller size.	N/A	More Details
CVE-2026-63958	In the Linux kernel, the following vulnerability has been resolved: usb: typec: ucsi: validate connector number in ucsi_connector_change() The connector number in a UCSI CCI notification is a 7-bit field supplied by the PPM. ucsi_connector_change() uses it to index the ucsi->connector[] array without checking it against the number of connectors the PPM reported at init time, so a buggy or malicious PPM (EC firmware, or an I2C-attached UCSI controller on the ccg / stm32g0 / glink transports) can drive schedule_work() on memory past the end of the array. Reject connector numbers that are zero or exceed cap.num_connectors before dereferencing the array.	N/A	More Details
CVE-2026-63959	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpm/tcpci_maxim: validate header NDO against RX_BYTE_CNT A broken/malicious port can transmit a CRC-valid frame whose header advertises up to seven data objects but whose body carries fewer than that. Check for this, and rightfully reject the message, instead of reading from uninitialized stack memory.	N/A	More Details
CVE-2026-63960	In the Linux kernel, the following vulnerability has been resolved: usb: typec: wcove: don't write past struct pd_message in wcove_read_rx_buffer() wcove_read_rx_buffer() copies the PD RX FIFO into the caller's struct pd_message with for (i = 0; i < USBC_RXINFO_RXBYTES(info); i++) regmap_read(wcove->regmap, USBC_RX_DATA + i, msg + i); which has two problems: USBC_RXINFO_RXBYTES() is a 5-bit field (max 31) while struct pd_message is 30 bytes (__le16 header + __le32 payload[PD_MAX_PAYLOAD], packed). The byte count latched in RXINFO is the number of bytes the port partner put on the wire, so a malicious partner that transmits a 31-byte frame can drive the loop one byte past the destination if the WCOVE BMC receiver does not enforce the PD object-count limit in hardware. The existing FIXME flagged this as unverified. Independently, regmap_read() takes an unsigned int * and stores a full unsigned int at the destination. Passing the byte pointer msg + i means each iteration writes four bytes; the high three are zero (val_bits is 8) and are normally overwritten by the next iteration, but the final iteration's high bytes are not. With RXBYTES == 30 the i == 29 iteration already writes three zero bytes past msg, which sits on the IRQ thread's stack in wcove_typec_irq(). Clamp the loop to sizeof(struct pd_message) and read each register into a local before storing only its low byte, so the copy can never exceed the destination regardless of what RXINFO reports.	N/A	More Details
CVE-2026-64139	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix SID memory leak in set_posix_acl_entries_dacl() on overflow Commit 299f962c0b02 ("ksmbd: use check_add_overflow() to prevent u16 DACL size overflow") added check_add_overflow() guards that break out of the ACE-building loops in set_posix_acl_entries_dacl() when the accumulated DACL size would wrap past 65535. However, each iteration allocates a struct smb_sid via kmalloc_obj() at the top of the loop and relies on the kfree(sid) call at the end of the loop body (the 'pass_same_sid' label in the first loop, and the explicit kfree at the tail of the second loop) to release it. The newly introduced 'break' statements bypass those kfree() calls, leaking the sid buffer every time an overflow is detected. A malicious or malformed file with enough POSIX ACL entries to trip the overflow check will leak one or more struct smb_sid allocations on every request that touches the file's DACL, providing a trivial kernel memory exhaustion vector. Free sid before breaking out of the loops to plug the leak.	N/A	More Details
CVE-2026-50124	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase can be exploited by uploading payload.zip through the Excel upload API /datasource/upload, creating an H2 datasource that uses the zip: protocol, and executing an SQL dataset path where CalciteProvider.jdbcFetchResultField calls statement.executeQuery(), causing precompiled Java aliases in test.mv.db to execute arbitrary code. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64143	In the Linux kernel, the following vulnerability has been resolved: platform/x86: uniwill-laptop: Do not enable the charging limit even when forced It seems that on some older models (~2020) the battery charging limit can permanently damage the battery. Prevent users from enabling this feature thru the "force" module parameter to avoid causing permanent hardware damage on such devices.	N/A	More Details
CVE-2026-39385	Frappe LMS is an open source learning management system. In version 2.51.0 and earlier, a user could bypass payment validation for courses by using unrelated batch. This has been patched in 2.52.0 with enrollment now validating that the batch is linked to course.	N/A	More Details
CVE-2026-	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61449. Reason: This candidate is a duplicate of CVE-2026-61449. Notes: All CVE users should reference CVE-2026-61449	N/A	More Details

61841	instead of this candidate.		
CVE-2026-42210	Webmin is a web-based system administration tool for Unix-like servers. Prior to version 2.640, for Webmin accounts that require a second authentication factor (typically TOTP), an attacker with knowledge of the username and password can bypass the 2FA requirement by using Basic authentication. Webmin is a web-based system administration tool for Unix-like servers. As a workaround, apply the patch from commit da18a16c84ae5c0b78cad79609cb0efb174000ec manually.	N/A	More Details
CVE-2025-71398	SurrealDB before 2.2.2 fails to validate HTTP redirects in http functions, allowing authenticated users to bypass deny-net restrictions by redirecting to blocked IP addresses. Attackers can host a public server that redirects to denied network targets, enabling server-side request forgery to access internal endpoints and retrieve sensitive information.	N/A	More Details
CVE-2026-61839	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61451. Reason: This candidate is a duplicate of CVE-2026-61451. Notes: All CVE users should reference CVE-2026-61451 instead of this candidate.	N/A	More Details
CVE-2025-71397	SurrealDB before 2.0.5, 2.1.x before 2.1.5, and 2.2.x before 2.2.2 allows authenticated users with OWNER or EDITOR permissions (at the root, namespace, or database level) to define custom database functions via DEFINE FUNCTION using nested FOR loops. Although a single loop's iteration count is constrained, nesting multiple loops (e.g., each with 1,000,000 iterations) is not, so an attacker can execute a function that consumes all server CPU time. Configured timeouts do not stop the execution, rendering the server unresponsive to other queries and connections until it is manually restarted.	N/A	More Details
CVE-2026-61829	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61452. Reason: This candidate is a duplicate of CVE-2026-61452. Notes: All CVE users should reference CVE-2026-61452 instead of this candidate.	N/A	More Details
CVE-2026-61710	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61453. Reason: This candidate is a duplicate of CVE-2026-61453. Notes: All CVE users should reference CVE-2026-61453 instead of this candidate.	N/A	More Details
CVE-2025-71396	SurrealDB before 2.0.5, 2.1.x before 2.1.5, and 2.2.x before 2.2.2 does not enforce a default execution-time limit on embedded JavaScript scripting functions when the scripting capability is explicitly enabled (via --allow-scripting or --allow-all). An authenticated attacker can submit long-running JavaScript functions to exhaust server resources and cause a denial of service. Scripting is disabled by default.	N/A	More Details
CVE-2026-50743	A CSRF vulnerability exists in the `zone-include.php` script in Revive Adserver 6.0.7. Linking and unlinking banners or campaigns to zones could be triggered via crafted GET or POST requests without any verification of the CSRF token, allowing an attacker to perform these actions on behalf of an authenticated administrator.	N/A	More Details
CVE-2026-61606	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61457. Reason: This candidate is a duplicate of CVE-2026-61457. Notes: All CVE users should reference CVE-2026-61457 instead of this candidate.	N/A	More Details
CVE-2025-71395	SurrealDB versions before 2.2.2 contain a memory exhaustion vulnerability in the string::replace function that fails to restrict resulting string length when using regex patterns. An authenticated attacker can craft a malicious query to exhaust server memory through unbounded string allocations, causing denial of service.	N/A	More Details
CVE-2025-71394	SurrealDB versions before 2.2.2 contain a local file read vulnerability in the DEFINE ANALYZER statement that allows authenticated users to read arbitrary files on the file system. Attackers with root, namespace, or database level privileges can point analyzers to arbitrary file paths and exfiltrate content from two-column tab-separated files.	N/A	More Details
CVE-2025-71393	SurrealDB before 2.2.2 with scripting enabled fails to properly enforce recursion limits when native functions contain embedded JavaScript that issues new queries. Authenticated attackers can bypass the recursion limit by chaining native and JavaScript function calls to trigger infinite recursion and exhaust server memory.	N/A	More Details
CVE-2025-71392	SurrealDB before 2.0.5, 2.1.x before 2.1.5, and 2.2.x before 2.2.2 fails to properly escape table and field names in the command-line export command. An authenticated System User with OWNER or EDITOR roles can create tables or fields with malicious names containing SurrealQL. When a higher-privileged user subsequently imports the exported backup, the injected SurrealQL executes, enabling privilege escalation and root-level takeover of the SurrealDB instance. Applications that let users define custom tables or fields are also exposed to a universal second-order SurrealQL injection even when query parameters are sanitized.	N/A	More Details
CVE-2025-71391	SurrealDB versions before 2.2.2 contain an uncaught exception vulnerability in the net module that allows authenticated users to crash the database. Attackers can send crafted HTTP queries containing null bytes to the /sql endpoint, causing an unhandled exception that crashes the SurrealDB instance and any dependent applications.	N/A	More Details
CVE-2025-	SurrealDB before 2.2.6, 2.3.6, and 2.1.8 (and 3.0.0-alpha.7 and earlier) fails to validate DNS-resolved hostnames against --deny-net network access restrictions in its http::* functions. An authenticated user can invoke http::<fn>(<url>) with a hostname that resolves to a denied IP address, causing the server to issue	N/A	More

71390	the request anyway and return the response. This bypasses network access controls, allowing access to restricted internal endpoints and potentially retrieving or altering sensitive information and credentials, depending on the deployment.		Details
CVE-2026-64187	In the Linux kernel, the following vulnerability has been resolved: xfs: fail recovery on a committed log item with no regions If the first op of a transaction is a bare transaction header (len == sizeof(struct xfs_trans_header)), xlog_recover_add_to_trans() adds an item but no region, leaving it on r_itemq with ri_cnt == 0 and ri_buf == NULL. The header can be split across op records, so later ops may still add regions; the item is only invalid if the transaction commits with none. The runtime commit path never emits such a transaction, so this only happens on a crafted log. It came from an AI-assisted code audit of the recovery parser. xlog_recover_reorder_trans() calls ITEM_TYPE() on the item, which reads *(unsigned short *)item->ri_buf[0].iov_base and faults on the NULL ri_buf. Reject it there, before the commit handlers that also read ri_buf[0]. KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] RIP: 0010:xlog_recover_reorder_trans (fs/xfs/xfs_log_recover.c:1836) xlog_recover_commit_trans (fs/xfs/xfs_log_recover.c:2043) xlog_recover_process_data (fs/xfs/xfs_log_recover.c:2501) xlog_do_recovery_pass (fs/xfs/xfs_log_recover.c:3244) xlog_recover (fs/xfs/xfs_log_recover.c:3493) xfs_log_mount (fs/xfs/xfs_log.c:618) xfs_mountfs (fs/xfs/xfs_mount.c:1034) xfs_fs_fill_super (fs/xfs/xfs_super.c:1938) vfs_get_tree (fs/super.c:1695) path_mount (fs/namespace.c:4161) __x64_sys_mount (fs/namespace.c:4367)	N/A	More Details
CVE-2026-40187	In egroupware version 26.0 and earlier, an authenticated administrator can achieve OS-level Remote Code Execution (RCE) by uploading a malicious eTemplate XML file (`.xet`) to the VFS `/etemplates` mount. The `Widget::expand_name()` method passes template widget attribute values directly into a PHP `eval()` call with only double-quote escaping applied - **backtick characters are not escaped**. In PHP, backticks inside a double-quoted `eval()` string execute shell commands. This allows an admin-level user to escalate from web application access to arbitrary OS command execution on the server.	N/A	More Details
CVE-2026-35591	libvips is a fast image processing library with low memory needs. The `tiffload` operation in libvips versions before and including 8.18.1 could incorrectly determine the number of channels in a JPEG or JPEG2000-encoded tile within a TIFF image, leading to a possible buffer overflow. This has been patched in version 8.18.2.	N/A	More Details
CVE-2026-64189	In the Linux kernel, the following vulnerability has been resolved: netfilter: ipset: fix race between dump and ip_set_list resize The release path of ip_set_dump_do() and ip_set_dump_done() read inst->ip_set_list via ip_set_ref_netlink(), a plain rcu_dereference_raw() of the array pointer. These run from netlink_recvmsg() without the nfnl mutex and without an RCU read-side critical section. A concurrent ip_set_create() can grow the array: it publishes the new array, calls synchronize_net() and then kvfree()s the old one. Since the dump paths read the array outside any RCU reader, synchronize_net() does not wait for them and the old array can be freed while they still index into it, causing a use-after-free. The dumped set itself stays pinned via set->ref_netlink, so only the array load needs protecting. Take rcu_read_lock() around it, matching ip_set_get_byname() and __ip_set_put_byindex(). BUG: KASAN: slab-use-after-free in ip_set_dump_do (net/netfilter/ipset/ip_set_core.c:1697) Read of size 8 at addr ffff88800b5c4018 by task exploit/150 Call Trace: ... kasan_report (mm/kasan/report.c:595) ip_set_dump_do (net/netfilter/ipset/ip_set_core.c:1697) netlink_dump (net/netlink/af_netlink.c:2325) netlink_recvmsg (net/netlink/af_netlink.c:1976) sock_recvmsg (net/socket.c:1159) __sys_recvfrom (net/socket.c:2315) ... Oops: general protection fault, probably for non-canonical address ... KASAN NOPTI KASAN: maybe wild-memory-access in range [0x02d6...d0-0x02d6...d7] RIP: 0010:ip_set_dump_do (net/netfilter/ipset/ip_set_core.c:1698) Kernel panic - not syncing: Fatal exception	N/A	More Details
CVE-2026-35590	libvips is a fast image processing library with low memory needs. The EXIF decoder within libvips versions before and including 8.18.1 was not verifying the range of EXIF tag groups before passing data to libexif, leading to a possible null pointer dereference and crash. This has been patched in version 8.18.2.	N/A	More Details
CVE-2026-14253	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-62177	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-60085. Reason: This candidate is a duplicate of CVE-2026-60085. Notes: All CVE users should reference CVE-2026-60085 instead of this candidate.	N/A	More Details
CVE-2026-45797	HeyForm is an open-source form builder. Prior to version 3.0.0-rc.7, the `/api/upload` endpoint allows unauthenticated file uploads including SVG files. Uploaded SVGs are stored in the static assets directory and served with `Content-Type: image/svg+xml` by Express's serve-static middleware, allowing an attacker to achieve stored cross-site scripting (XSS) on the heyform domain without any authentication. Version 3.0.0-rc.7 contains a patch for the issue.	N/A	More Details
CVE-2026-62174	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61435. Reason: This candidate is a duplicate of CVE-2026-61435. Notes: All CVE users should reference CVE-2026-61435 instead of this candidate.	N/A	More Details
	lettre is a a mailer library for Rust. Starting in version 0.10.1 and prior to version 0.11.22, an inverted-boolean bug in lettre's `boring-tls` integration silently disables TLS hostname verification for callers using the default		

CVE-2026-46428	(strict) configuration. An on-path attacker presenting any chain-valid certificate for any domain can intercept SMTP submission, including PLAIN/LOGIN credentials and message contents, against any lettre user built with the `boring-tls` feature. Other TLS backends (`native-tls`, `rustls`) are unaffected. Version 0.11.22 patches the issue.	N/A	More Details
CVE-2026-15997	Out-of-bounds write vulnerability in Legion of the Bouncy Castle Inc. BC-LTS bcprov-lts8on on ARM allows Overflow Buffers. This vulnerability is associated with program files https://github.Com/bcgit/bc-lts-java/blob/main/native_c/arm/sha/shake.C, https://github.Com/bcgit/bc-lts-java/blob/main/native_c/arm/sha/sha3.C. This issue affects BC-LTS: from 2.73.0 before 2.73.12.1. Issue is only applicable if application involved is accepting memoable SHA3 / SHAKE states from potentially untrusted sources.	N/A	More Details
CVE-2026-44181	Jupyter Enterprise Gateway launches remote Jupyter Notebook kernels across distributed clusters like Apache Spark, Kubernetes, and Docker Swarm. In versions 2.0.0rc2 and above, prior to 3.3.0, the environment variables (KERNEL_XXX) used during the rendering of the Kubernetes manifest are vulnerable to Server Side Template Injection (SSTI). By including Jinja2 template expressions it is possible to execution Python code and OS Commands in the Enterprise Gateway service. The code can use or steal the Kubernetes service account token, which can steal Kubernetes secrets and be used to fully compromise the Kubernetes cluster by scheduling a privileged pod or a pod with a hostPath volume mount. This issue has been fixed in version 3.3.0.	N/A	More Details
CVE-2026-44182	Jupyter Enterprise Gateway launches remote Jupyter Notebook kernels across distributed clusters like Apache Spark, Kubernetes, and Docker Swarm. In versions prior to 3.3.0, the server interpolates untrusted environment variables (e.g., KERNEL_XXX) into Kubernetes manifests without YAML-aware escaping, enabling YAML injection attacks. Attackers can inject new fields, overwrite critical fields (e.g., duplicate securityContext keys, where the last one prevails), and inject document boundaries (--- for new documents, ... for end-of-document) to generate multiple resources, potentially creating arbitrary types, such as privileged pods. The Jinja2 template for the Kubernetes manifest contains several kernel_xxx variables, such as kernel_working_dir that are used when rendering the manifest and are all vectors for YAML injection. This issue has been fixed in version 3.3.0.	N/A	More Details
CVE-2026-62173	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61433. Reason: This candidate is a duplicate of CVE-2026-61433. Notes: All CVE users should reference CVE-2026-61433 instead of this candidate.	N/A	More Details
CVE-2026-16312	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-62172	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61436. Reason: This candidate is a duplicate of CVE-2026-61436. Notes: All CVE users should reference CVE-2026-61436 instead of this candidate.	N/A	More Details
CVE-2026-62169	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61430. Reason: This candidate is a duplicate of CVE-2026-61430. Notes: All CVE users should reference CVE-2026-61430 instead of this candidate.	N/A	More Details
CVE-2026-62168	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61443. Reason: This candidate is a duplicate of CVE-2026-61443. Notes: All CVE users should reference CVE-2026-61443 instead of this candidate.	N/A	More Details
CVE-2026-62165	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61446. Reason: This candidate is a duplicate of CVE-2026-61446. Notes: All CVE users should reference CVE-2026-61446 instead of this candidate.	N/A	More Details
CVE-2026-62164	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-60087. Reason: This candidate is a duplicate of CVE-2026-60087. Notes: All CVE users should reference CVE-2026-60087 instead of this candidate.	N/A	More Details
CVE-2026-33327	libvips is a fast image processing library with low memory needs. The `vipsload` operation in versions before and including 8.18.0 could incorrectly determine image dimensions leading to an integer overflow and a subsequent heap-based buffer overflow. This has been patched in version 8.18.1.	N/A	More Details
CVE-2026-33328	libvips is a fast image processing library with low memory needs. On 32-bit systems in versions before and including 8.18.0, the `gifload` operation could incorrectly determine dimensions leading to an integer overflow. This has been patched in version 8.18.1.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix lock leak on ENOMEM in AMDGPU_GEM_OP_GET_MAPPING_INFO The AMDGPU_GEM_OP_GET_MAPPING_INFO branch of amdgpu_gem_op_ioctl() holds three cleanup-tracked resources before calling kvcalloc(): the drm_gem_object reference from drm_gem_object_lookup(), the drm_exec lock on the looked-up GEM via drm_exec_lock_obj(), and the drm_exec lock on the per-process VM root page directory via amdgpu_vm_lock_pd(). All three are		

CVE-2026-63880	released by the out_exec label that every other error path in this function jumps to. The kvcalloc() failure path returns -ENOMEM directly, skipping out_exec and leaking all three. The leaked per-process VM root PD dma_resv lock is the load-bearing leak: any subsequent operation on the same VM (further GEM ops, command-submission, eviction, TTM shrinker callbacks) blocks on the held lock. DRM_IOCTL_AMDGPU_GEM_OP is DRM_AUTH | DRM_RENDER_ALLOW, so this is an unprivileged-local denial of service against the caller's GPU context, reachable by any process with /dev/dri/renderD* access. Route the failure through out_exec so drm_exec_fini() and drm_gem_object_put() run. Reproduced on stock 7.0.0-10, Ryzen 7 5700U / Radeon Vega (Lucienne): the failing ioctl returns -ENOMEM and a second GET_MAPPING_INFO on the same fd then blocks in drm_exec_lock_obj() on the leaked dma_resv. SIGKILL on the caller does not reap the task; the fd-release path during process exit goes through amdgpu_gem_object_close() -> drm_exec_prepare_obj() on the same lock, leaving the task in D state until the box is rebooted. The patched kernel was not rebuilt and re-tested on this hardware; the fix is mechanical. Tested on a single Lucienne / Vega box only. Ziyi Guo posted an independent INT_MAX-bound check for args->num_entries in the same branch [1]; the two patches are complementary and can land in either order. (cherry picked from commit b69d3256d79de15f54c322986ff4da68f1d65b0a)	N/A	More Details
CVE-2026-64190	In the Linux kernel, the following vulnerability has been resolved: net: team: fix NULL pointer dereference in team_xmit during mode change __team_change_mode() clears team->ops with memset() before restoring safe dummy handlers via team_adjust_ops(). A concurrent team_xmit() running under RCU on another CPU can read team->ops.transmit during this window and call a NULL function pointer, crashing the kernel. The race requires a mode change (CAP_NET_ADMIN) concurrent with transmit on the team device. BUG: kernel NULL pointer dereference, address: 0000000000000000 Oops: 0010 [#1] SMP KASAN NOPTI RIP: 0010:0x0 Call Trace: team_xmit (drivers/net/team/team_core.c:1853) dev_hard_start_xmit (net/core/dev.c:3904) __dev_queue_xmit (net/core/dev.c:4871) packet_sendmsg (net/packet/af_packet.c:3109) __sys_sendto (net/socket.c:2265) The original code assumed that no ports means no traffic, so mode changes could freely memset()/memcpy() the ops. AF_PACKET with forced carrier breaks that assumption. Prevent the race instead of making it safe: replace memset()/memcpy() with per-field updates that never touch transmit or receive. Those two handlers are managed solely by team_adjust_ops(), which already installs dummies when tx_en_port_count == 0 (always true during mode change since no ports are present). WRITE_ONCE/READ_ONCE prevent store/load tearing on the handler pointers. synchronize_net() before exit_op() drains in-flight readers that may still reference old mode state from before port removal switched the handlers to dummies.	N/A	More Details
CVE-2026-62178	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61427. Reason: This candidate is a duplicate of CVE-2026-61427. Notes: All CVE users should reference CVE-2026-61427 instead of this candidate.	N/A	More Details
CVE-2026-60034	The Joomla extension JMedia is vulnerable to a stored XSS vulnerability. Unsanitised SVG uploads served without nosniff, leading to stored/reflected XSS.	N/A	More Details
CVE-2026-61425	The Joomla extension Gridbox is vulnerable an authenticated bypass, potentially leading to full admin access.	N/A	More Details
CVE-2026-61900	The Joomla extension JDownloads is vulnerable to an unauthenticated file upload, leading to full RCE.	N/A	More Details
CVE-2026-62414	The Joomla extension Page Builder CK does not properly apply access control to frontend page list views.	N/A	More Details
CVE-2026-9007	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in HCL Notes from HCL Software allows reflected Cross-Site Scripting (XSS). Successful exploitation allows an attacker to execute arbitrary JavaScript in the context of another user. This issue affects HCL Notes: Release 12.0.2FP5HF8 on Linux 4.18.0-553.52.1.El8_10.X64_64#1.	N/A	More Details
CVE-2026-61828	Nixpkgs is a collection of software packages that can be installed with the Nix package manager. Prior to the 25.11 and 26.05 channel fixes, the NixOS module for MySQL services.mysql initializes the MySQL database in a way that allows local users, such as unprivileged web or CGI processes on the same host, to log in as the root user without a password when the service is used with mysql or percona-server. This issue is fixed in the 25.11 and 26.05.	N/A	More Details
CVE-2026-61605	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-58655. Reason: This candidate is a duplicate of CVE-2026-58655. Notes: All CVE users should reference CVE-2026-58655 instead of this candidate.	N/A	More Details
CVE-2026-15788	BuildKit's cache mount source= selector on Windows Container on Windows (WCOW) workers does not detect NTFS directory junctions placed inside the cache root. A build authored by an untrusted user on a WCOW-configured BuildKit daemon can read arbitrary host files reachable to the BuildKit daemon process.	N/A	More Details

CVE-2026-16337	Improper authorization in the ToolGroupResource and RoleAjax REST/DWR endpoints in dotCMS dotCMS 21.02 through 26.06.22-03 on all platforms allows a low-privileged authenticated backend user to self-assign the administrative layout and self-grant the CMS Administrator role, then achieve remote code execution via a crafted OSGi bundle upload whose BundleActivator executes arbitrary shell commands.	N/A	More Details
CVE-2026-47159	Vaultwarden is a Bitwarden-compatible server written in Rust. Prior to 1.36.0, Vaultwarden's SSO discovery and pre-validation flow returned organization-related SSO metadata including organizationIdentifier values for arbitrary email addresses and allowed a valid pre-validation JWT to be obtained with only the discovered identifier, enabling SSO-enabled organization enumeration and authentication workflow abuse. This issue is fixed in version 1.36.0.	N/A	More Details
CVE-2026-45150	Zen is a firefox-based browser. Prior to 1.19.13b, Zen Browser did not provide a persistent, clearly visible security notification when a webpage entered fullscreen mode, allowing an attacker-controlled page to hide the real browser UI and origin information, imitate a trusted website UI, and combine with long-domain URL eliding to spoof a trusted origin for phishing and credential theft. This issue is fixed in version 1.19.13b.	N/A	More Details
CVE-2026-13380	VSee Clinic 7.1.26 and VSee Clinic API 1.3.0 exposes cleartext SFTP credentials in the HTTP responses of three unauthenticated endpoints. The credentials are present in these responses only when SFTP connections have been configured within the application. No authentication is required to retrieve these credentials. An unauthenticated remote attacker who observes any of these HTTP responses on an instance where SFTP is configured can obtain the credentials and use them to access the associated SFTP server.	N/A	More Details
CVE-2026-13381	VSee Clinic 7.1.26 and API 1.3.0 contain an Insecure Direct Object Reference (IDOR) vulnerability in the /v1.3.0/api/files endpoint. An authenticated attacker can manipulate the 'remark' request parameter to enumerate, retrieve, and delete files belonging to other users on the application server.	N/A	More Details
CVE-2026-44507	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-43617. Reason: This candidate is a duplicate of CVE-2026-43617. Notes: All CVE users should reference CVE-2026-43617 instead of this candidate.	N/A	More Details
CVE-2026-44508	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-43618. Reason: This candidate is a duplicate of CVE-2026-43618. Notes: All CVE users should reference CVE-2026-43618 instead of this candidate.	N/A	More Details
CVE-2026-44509	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-43619. Reason: This candidate is a duplicate of CVE-2026-43619. Notes: All CVE users should reference CVE-2026-43619 instead of this candidate.	N/A	More Details
CVE-2026-44979	@hapi/wreck is an HTTP client utility. Prior to 18.1.1, when @hapi/wreck follows a 3xx redirect to a different hostname, only the Authorization and Cookie headers are stripped, and the standard credential header Proxy-Authorization is forwarded intact to the redirect target, potentially exposing forward-proxy credentials to a host outside the original trust boundary when redirects are enabled through the redirects option or Wreck.defaults({ redirects: ... }). This issue is fixed in version 18.1.1.	N/A	More Details
CVE-2026-62294	Flameshot is powerful yet simple to use screenshot software. Prior to 14.0.0, the Open With feature wrote screenshots to a predictable temporary path and followed symlinks, creating a time-of-check to time-of-use race that allowed a local unprivileged attacker on the same machine to pre-plant a symlink and cause Flameshot to write PNG data through it, overwriting any file the victim user could write. This issue is fixed in version 14.0.0.	N/A	More Details
CVE-2026-61424	The Joomla extension DJ-Classifieds is vulnerable to an unauthenticated file upload, leading to full RCE.	N/A	More Details
CVE-2026-60033	The Joomla extension JMedia is vulnerable to an SSRF vulnerability. Remote-URL download could target internal/reserved addresses.	N/A	More Details
CVE-2026-64191	In the Linux kernel, the following vulnerability has been resolved: i2c: stub: Reject I2C block transfers with invalid length The I2C_SMBUS_I2C_BLOCK_DATA case in stub_xfer() uses data->block[0] as the transfer length. The existing check only clamps it to avoid overrunning the chip->words[256] register array, but does not validate it against I2C_SMBUS_BLOCK_MAX (32), which is the limit of the union i2c_smbus_data.block buffer (34 bytes total). The driver is a development/test tool (CONFIG_I2C_STUB=m, not built by default) that must be loaded with a chip_addr= parameter. A local user with access to /dev/i2c-* can issue an I2C_SMBUS ioctl with I2C_SMBUS_I2C_BLOCK_DATA and data->block[0] > 32, causing stub_xfer() to read or write past the end of the union i2c_smbus_data.block buffer: BUG: KASAN: stack-out-of-bounds in stub_xfer (drivers/i2c/i2c-stub.c:223) Read of size 1 at addr ffff88800abcfd92 by task exploit/81 Call Trace: <TASK> stub_xfer (drivers/i2c/i2c-stub.c:223) __i2c_smbus_xfer (drivers/i2c/i2c-core-smbus.c:593) i2c_smbus_xfer (drivers/i2c/i2c-core-smbus.c:536) i2cdev_ioctl_smbus (drivers/i2c/i2c-dev.c:391) i2cdev_ioctl (drivers/i2c/i2c-dev.c:478) __x64_sys_ioctl (fs/ioctl.c:583) do_syscall_64 (arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:130) </TASK> The bug exists because i2c-stub implements .smbus_xfer directly, bypassing the I2C_SMBUS_BLOCK_MAX validation in	N/A	More Details

	i2c_smbus_xfer_emulated(). The I2C_SMBUS_BLOCK_DATA case in the same function correctly validates against I2C_SMBUS_BLOCK_MAX, but the I2C_SMBUS_I2C_BLOCK_DATA case does not. Fix by rejecting transfers with data->block[0] == 0 or data->block[0] > I2C_SMBUS_BLOCK_MAX with -EINVAL, consistent with both the I2C_SMBUS_BLOCK_DATA case in the same function and the I2C_SMBUS_I2C_BLOCK_DATA validation in i2c_smbus_xfer_emulated().		
CVE-2026-60032	The Joomla extension JMedia is vulnerable to an authenticated arbitrary file upload, leading to RCE. Executable uploads/writes possible (incl. polyglot filenames); chmod didn't strip execute bits.	N/A	More Details
CVE-2026-64192	In the Linux kernel, the following vulnerability has been resolved: bpf: Reject BPF_MAP_TYPE_INODE_STORAGE creation if BPF LSM is uninitialized When CONFIG_BPF_LSM=y is set, BPF inode storage maps (BPF_MAP_TYPE_INODE_STORAGE) are compiled into the kernel. However, if the BPF LSM is not explicitly enabled at boot time (e.g. omitted from the "lsm=" boot parameter), lsm_prepare() is never executed for the BPF LSM. Consequently, the BPF inode security blob offset (bpf_lsm_blob_sizes.lbs_inode) is never initialized and remains at its default compiled size of 8 bytes instead of being updated to a valid offset past the reserved struct rcu_head (typically 16 bytes or more). When a privileged user creates and updates a BPF_MAP_TYPE_INODE_STORAGE map, bpf_inode() evaluates inode->i_security + 8. This erroneously aliases the struct rcu_head.func callback pointer at the beginning of the inode->i_security blob. During subsequent map element cleanup or inode destruction, writing NULL to owner_storage clears the queued RCU callback pointer. When rcu_do_batch() later executes the queued callback, it attempts an instruction fetch at address 0x0, triggering an immediate kernel panic. Fix this by introducing a global bpf_lsm_initialized boolean flag marked with __ro_after_init. Set this flag to true inside bpf_lsm_init() when the LSM framework successfully registers the BPF LSM. Gate map allocation in inode_storage_map_alloc() on this flag, returning -EOPNOTSUPP if the BPF LSM is in turn uninitialized. This fail-fast approach prevents userspace from allocating inode storage maps when the supporting BPF LSM infrastructure is absent, avoiding zombie map states.	N/A	More Details
CVE-2026-64205	In the Linux kernel, the following vulnerability has been resolved: i2c: i801: fix hardware state machine corruption in error path A severe livelock and subsequent Hung Task panic were observed in the i2c-i801 driver during concurrent Fuzzing. The crash is caused by an unconditional hardware register cleanup in the error handling path of i801_access(). When i801_check_pre() fails (e.g., returning -EBUSY because the SMBus controller is actively used by BIOS/ACPI), the kernel does not actually acquire the hardware ownership. However, the code jumps to the 'out' label and executes: iowrite8(SMBHSTSTS_INUSE_STS | STATUS_FLAGS, SMBHSTSTS(priv)); This forcefully clears the INUSE_STS lock and resets the hardware status flags without owning the controller. Doing so interrupts ongoing BIOS/ACPI transactions and totally corrupts the SMBus hardware state machine. Consequently, all subsequent i801_access() calls fail at the pre-check stage, triggering an endless stream of "SMBus is busy, can't use it!" error logs. Over a slow serial console, this printk flood monopolizes the CPU (Console Livelock), starving other processes trying to acquire the mmap_lock down_read semaphore, ultimately triggering the hung task watchdog. Fix this by moving the 'out' label below the hardware register cleanup. If i801_check_pre() fails, we safely bypass the iowrite8() and only release the software locks (pm_runtime and mutex), strictly adhering to the rule of not releasing resources that were never acquired.	N/A	More Details
CVE-2026-64206	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: cancel pending_rx_work before taking conn->lock l2cap_conn_del() takes conn->lock and then calls cancel_work_sync() for pending_rx_work. process_pending_rx() takes the same mutex, so teardown can deadlock against the worker it is flushing. This issue was found by our static analysis tool and then manually reviewed against the current tree. The grounded PoC kept the l2cap_conn_ready() -> queue_work(..., &conn->pending_rx_work) submit path, the l2cap_conn_del() -> cancel_work_sync(&conn->pending_rx_work) teardown path, and the process_pending_rx() -> mutex_lock(&conn->lock) worker edge. Lockdep WARNING: possible circular locking dependency detected process_pending_rx+0x21/0x2a [vuln_msv] l2cap_conn_del.constprop.0+0x3f/0x4e [vuln_msv] *** DEADLOCK *** Cancel pending_rx_work before taking conn->lock, matching the existing lock-before-drain ordering used for the two delayed works in the same teardown path. The pending_rx queue is still purged after the work has been cancelled and conn->lock has been acquired.	N/A	More Details
CVE-2026-64207	In the Linux kernel, the following vulnerability has been resolved: net/sched: dualpi2: fix GSO backlog accounting When DualPI2 splits a GSO skb into N segments, it propagates N additional packets to its parent before returning NET_XMIT_SUCCESS. The parent then accounts for the original skb once more, leaving its qlen one larger than the number of packets actually queued. With QFQ as the parent, after all real packets are dequeued, QFQ still has a non-zero qlen while its in-service aggregate has no active classes. qfq_choose_next_agg() returns NULL and qfq_dequeue() passes the result to qfq_peek_skb(), causing a NULL pointer dereference. Follow the same pattern used by tbf_segment() and taprio: count only successfully queued segments, propagate the difference between the original skb and those segments, and return NET_XMIT_SUCCESS whenever at least one segment was queued.	N/A	More Details
CVE-2024-58367	SurrealDB versions before 2.0.4 fail to properly enforce field permissions during SELECT, UPDATE, and DELETE operations, allowing authorized users to access unauthorized field values through various query techniques. Attackers can exploit SELECT VALUE operations, field aliasing, function arguments, WHERE clause filtering, RETURN BEFORE clauses, and SET clause references to leak protected field contents despite lacking SELECT permissions.	N/A	More Details

CVE-2026-34239	Chamilo version 1.11.40 and earlier are vulnerable to authenticated remote code execution in the main/inc/ajax/lang.ajax.php path. This endpoint is protected only by `api_protect_course_script(true)`, which means any authenticated user enrolled in a course (student, teacher, DRH) can reach it.	N/A	More Details
CVE-2024-58356	SurrealDB before 2.1.4 silently fails to overwrite table definitions when the DEFINE TABLE ... OVERWRITE clause is used on tables defined with TYPE RELATION. Because table definitions include the PERMISSIONS clause, an attempt to tighten a table's permissions via OVERWRITE does not take effect, and the administrator may incorrectly believe the change was applied. As a result, a client authorized to run queries may continue to access data in that table that the updated (but unapplied) permissions were intended to restrict.	N/A	More Details
CVE-2026-46715	Flask-Security-Too allows users to add security features to their Flask applicationa. Version 5.8.0's OAuth reauthentication flow can mark a session as fresh after verifying an OAuth account that belongs to a different user. If an attacker can operate an already-authenticated but stale victim session, they can complete OAuth verification using their own OAuth identity. The victim session is then treated as recently reauthenticated, allowing freshness-protected account actions to proceed. Version 5.8.1 contains a fix for this issue.	N/A	More Details
CVE-2026-47703	AdGuard Home is a network-wide software for blocking ads and tracking. Prior to 0.107.75, AdGuard Home's client-triggered DoQ forwarding path to a udp:// upstream reduced backend UDP DNS state by producing dns_id=0 or txid=0 and exposed a quoted-port ICMP source-port oracle, weakening DNS response matching for forwarded queries. This issue is fixed in version 0.107.75.	N/A	More Details
CVE-2026-8169	ExtremeXOS (EXOS) uses a challenge-response mechanism to authorize access to the privileged debug-mode function. The challenge value is generated using an insufficiently random source, which under certain conditions may allow an attacker to predict the expected response and activate debug-mode without authorization. Depending on device configuration and version, this may enable escalation to root-level access and persistent modification of the device software stack. Exploitation requires either a valid low-privilege account on the device (remote scenario) or physical serial console access (local scenario). This vulnerability is distinct from CVE-2017-14329, which addressed a different issue involving Python script privileges. Extreme would like to thank Hadrien Barral (Université Gustave Eiffel) and Georges-Axel Jaloyan (French Ministry of the Interior) for responsible disclosure of their findings.	N/A	More Details
CVE-2026-8170	The mv, cp, and rm file utilities exposed within the ExtremeXOS (EXOS) shell environment fail to safely canonicalize paths and follow symbolic links outside of the intended privilege boundary. An attacker with low-privilege CLI access can create a symbolic link that references a privileged filesystem location and then invoke the affected utilities to read, modify, or replace security-critical files outside of their authorized scope. Under certain conditions, this may enable escalation to root-level access and persistent modification of the device software stack. Exploitation is possible remotely by an attacker holding a low-privilege account, or locally via the serial console. Extreme would like to thank Hadrien Barral (Université Gustave Eiffel) and Georges-Axel Jaloyan (French Ministry of the Interior) for responsible disclosure of their findings.	N/A	More Details
CVE-2026-60026	The Joomla extension Quix Page Builder Pro is vulnerable to an authenticated PHP code execution. Authenticated builder user (core.create/core.edit) could inject PHP tags in element content, that got executed via view-cache include(). Requires caching on (default).	N/A	More Details
CVE-2026-60027	The Joomla extension Quix Page Builder Pro is vulnerable to a unauthenticated path traversal via form elements. Unauthenticated users frontend users are allowed traversal paths and read arbitrary files. Requires a published page with a Form element.	N/A	More Details
CVE-2026-60028	The Joomla extension Quix Page Builder Pro is vulnerable to an authenticated stored XSS vulnerability. Authenticated builder user could inject scripts, fires for any visitor or admin viewing the page. Unescaped output + unsanitised SVG.	N/A	More Details
CVE-2026-60029	The Joomla extension Quix Page Builder Pro is vulnerable to an authenticated stored XSS vulnerability. Authenticated builder users could break out of id/class fields that render for public users.	N/A	More Details
CVE-2026-60030	The Joomla extension Quix Page Builder Pro is vulnerable to an improper access control. Authenticated users could upload media files regardless of their media management permissions.	N/A	More Details
CVE-2026-60031	The Joomla extension Quix Page Builder Pro is vulnerable to an information disclosure. Raw exceptions reflected in AJAX handler responses.	N/A	More Details
CVE-2026-12228	A stored cross-site scripting (XSS) vulnerability exists in the `POST /api/prompts/share` endpoint of parisneo/lollms (latest version). The endpoint stores attacker-controlled `prompt_content` into `DBDirectMessage.content` without server-side sanitization. When a victim opens the direct message (DM) thread, the message is rendered by the DM UI through `MessageContentRenderer`, which uses `v-html` to insert rendered HTML into the DOM. The frontend sanitizer, which is regex-based, fails to comprehensively sanitize attacker-controlled HTML, allowing malicious payloads to execute in the victim's browser context. This vulnerability enables any authenticated user to send a malicious prompt-share message to another	N/A	More Details

	user's inbox, leading to arbitrary JavaScript execution, authenticated actions as the victim, exposure of same-origin application data, and potential account takeover.		
CVE-2026-62180	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61440. Reason: This candidate is a duplicate of CVE-2026-61440. Notes: All CVE users should reference CVE-2026-61440 instead of this candidate.	N/A	More Details
CVE-2026-64144	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btmtk: fix urb->setup_packet leak in error paths The setup_packet of control urb is not freed if usb_submit_urb fails or the submitted urb is killed. Add free in these two paths.	N/A	More Details
CVE-2026-64167	In the Linux kernel, the following vulnerability has been resolved: kho: skip KHO for crash kernel kho_fill_kimage() unconditionally populates the kimage with KHO metadata for every kexec image type. When the image is a crash kernel, this can be problematic as the crash kernel can run in a small reserved region and the KHO scratch areas can sit outside it. The crash kernel then faults during kho_memory_init() when it tries phys_to_virt() on the KHO FDT address: Unable to handle kernel paging request at virtual address xxxxxxxx ... fdt_offset_ptr+... fdt_check_node_offset_+... fdt_first_property_offset+... fdt_get_property_namelen_+... fdt_getprop+... kho_memory_init+... mm_core_init+... start_kernel+... kho_locate_mem_hole() already skips KHO logic for KEXEC_TYPE_CRASH images, but kho_fill_kimage() was missing the same guard. As kho_fill_kimage() is the single point that populates image->kho.fdt and image->kho.scratch, fixing it here is sufficient for both arm64 and x86 as the FDT and boot_params path are bailing out when these fields are unset.	N/A	More Details
CVE-2026-64169	In the Linux kernel, the following vulnerability has been resolved: spi: ep93xx: fix error pointer deref after DMA setup failure The driver falls back to PIO mode if DMA setup fails during probe. Make sure to the clear the DMA channel pointers on setup failure to avoid dereferencing an error pointer on later probe errors or driver unbind. This issue was flagged by Sashiko when reviewing a devres allocation conversion patch.	N/A	More Details
CVE-2026-64170	In the Linux kernel, the following vulnerability has been resolved: spi: qup: fix error pointer deref after DMA setup failure The driver falls back to PIO mode if DMA setup fails during probe. Make sure to the clear the DMA channel pointers on setup failure to avoid dereferencing an error pointer (or attempting to release a channel a second time) on later probe errors or driver unbind. This issue was flagged by Sashiko when reviewing a devres allocation conversion patch.	N/A	More Details
CVE-2026-64171	In the Linux kernel, the following vulnerability has been resolved: i2c: tegra: fix pm_runtime leak on mutex_lock failure If tegra_i2c_mutex_lock() fails, the function returns without calling pm_runtime_put(), leaking the runtime PM reference acquired by the preceding pm_runtime_get_sync(). This prevents the device from ever entering runtime suspend. Add the missing pm_runtime_put() before returning on lock failure.	N/A	More Details
CVE-2026-62963	Centrifugo is an open-source scalable real-time messaging server. Prior to 6.8.4, Centrifugo unidirectional WebSocket transport with uni_websocket.compression enabled enforced uni_websocket.message_size_limit against compressed wire-frame length in internal/websocket/conn.go advanceFrame, but ReadMessage used io.ReadAll after decompression without an output cap, allowing unauthenticated requests to /connection/uni_websocket to trigger large memory and CPU consumption. This issue is fixed in version 6.8.4.	N/A	More Details
CVE-2026-64173	In the Linux kernel, the following vulnerability has been resolved: tracing: Do not call map->ops->elt_free() if elt_alloc() fails In paths where tracing_map_elt_alloc() failed to allocate objects, the map->ops->elt_alloc() call was never successful. In this case, map->ops->elt_free() should not be called.	N/A	More Details
CVE-2026-64174	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: advance loop vars in cfg80211_merge_profile() cfg80211_merge_profile() reassembles a Multi-BSSID non-transmitted BSS profile that has been split across multiple consecutive MBSSID elements. Its while-loop calls cfg80211_get_profile_continuation(ie, ielen, mbssid_elem, sub_elem) but never advances mbssid_elem or sub_elem inside the body. Each iteration therefore searches for a continuation that follows the same fixed pair; the helper returns the same next_mbssid; and the same next_sub bytes are memcpy()'d into merged_ie at a growing offset until the buffer fills. Advance both mbssid_elem and sub_elem to the just-consumed continuation so the next call to cfg80211_get_profile_continuation() searches for a further continuation beyond it (or returns NULL when none exists). A specially-crafted malicious beacon can take advantage of this bug to cause the kernel to spend an excessive amount of time in cfg80211_merge_profile (up to as much as 2ms per beacon received), which could theoretically be abused in some way.	N/A	More Details
CVE-2026-45534	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase Redshift datasource connections can load attacker-controlled rsjdbc.ini configuration from System.getProperty("java.io.tmpdir"), setting socketFactory=org.springframework.context.support.FileSystemXmlApplicationContext so com.amazon.redshift.Driver#connect, com.amazon.redshift.Driver#getJdbcIniFile, and com.amazon.redshift.util.ObjectFactory#instantiate execute a reflection-based remote code execution chain during a normal JDBC connection through io.dataease.datasource.type.Redshift. This issue is fixed in version 2.10.23.	N/A	More Details
	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase export-center		

CVE-2026-45533	deletion can accept path traversal sequences such as ../ in the bulk delete API endpoint and pass attacker-controlled identifiers to ExportCenterManage.delete, allowing recursive deletion of arbitrary server directories through export task cleanup. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64177	In the Linux kernel, the following vulnerability has been resolved: phonet/pep: disable BH around forwarded sk_receive_skb() The networking receive path is usually run from softirq context, but protocols that take the socket lock may have packets stored in the backlog and processed later from process context. In that case release_sock() -> __release_sock() drops the slock with spin_unlock_bh() and then calls sk->sk_backlog_rcv() with bottom halves enabled. Typical sk_backlog_rcv handlers process the socket whose backlog is being drained, so the BH state at entry is irrelevant for the slocks they touch. pep_do_rcv() is different: when the inbound skb targets an existing PEP pipe, it forwards the skb to a different *child* socket via sk_receive_skb(). That helper takes the child slock with bh_lock_sock_nested(), which is just spin_lock_nested() and assumes BH is already off. The same child slock therefore ends up acquired with BH on (process path) and with BH off (softirq path): process context softirq context ---------- ---------- release_sock(listener) __netif_receive_skb() __release_sock() phonet_rcv() spin_unlock_bh() __sk_receive_skb(listener) [BH now ENABLED] [BH already disabled] sk_backlog_rcv: sk_backlog_rcv: pep_do_rcv() pep_do_rcv() sk_receive_skb(child) sk_receive_skb(child) bh_lock_sock_nested(child) bh_lock_sock_nested(child) => SOFTIRQ-ON-W => IN-SOFTIRQ-W Lockdep flags this as inconsistent lock state, and it can become a real self-deadlock if a softirq on the same CPU tries to receive to the same child socket while its slock is held in the BH-enabled path: WARNING: inconsistent lock state inconsistent {SOFTIRQ-ON-W} -> {IN-SOFTIRQ-W} usage. (slock-AF_PHONET/1){+.?.}-{3:3}, at: __sk_receive_skb+0x1cf/0x900 __sk_receive_skb net/core/sock.c:563 sk_receive_skb include/net/sock.h:2022 [inline] pep_do_rcv net/phonet/pep.c:675 sk_backlog_rcv include/net/sock.h:1190 __release_sock net/core/sock.c:3216 release_sock net/core/sock.c:3815 pep_sock_accept net/phonet/pep.c:879 Wrap the forwarded sk_receive_skb() in local_bh_disable() / local_bh_enable() so the child slock is always acquired with BH off. local_bh_disable() nests safely on the softirq path. Discovered via in-house syzkaller fuzzing; the same root cause also on the linux-6.1.y syzbot dashboard as extid 44f0626dd6284f02663c. Reproduced under KASAN + LOCKDEP + PROVE_LOCKING, reproducer: https://pastebin.com/A3t8xzCR	N/A	More Details
CVE-2026-45419	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase template saves call TemplateManageService#save, StaticResourceServer#saveFilesToServe, and the /de2api/templateManage/save endpoint with attacker-controlled staticResource names and Base64 content, allowing path traversal and arbitrary file writes because only / was used when extracting the file name. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64179	In the Linux kernel, the following vulnerability has been resolved: net: wwan: iosm: fix potential memory leaks in ipc_imem_init() The memory allocated in ipc_protocol_init() is not freed on the error paths that follow in ipc_imem_init(). Fix that by calling the corresponding release function ipc_protocol_deinit() in the error path.	N/A	More Details
CVE-2026-64180	In the Linux kernel, the following vulnerability has been resolved: mm/memory_hotplug: fix memory block reference leak on remove Patch series "mm: Fix memory block leaks and locking", v2. This series fixes two memory block device reference leaks and one locking issue around the per-memory_block hwpoison counter. This patch (of 2): remove_memory_blocks_and_altmaps() looks up each memory block with find_memory_block(), which acquires a reference to the memory block device. That reference is never dropped on this path, resulting in a leaked device reference when removing memory blocks and their altmaps. Drop the reference after retrieving mem->altmap and clearing mem->altmap, before removing the memory block device.	N/A	More Details
CVE-2026-45417	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase datasource connection status checks concatenate configuration.getSchema() into getTablesSql and execute the resulting SQL with executeQuery in io.dataease.datasource.provider.CalciteProvider#checkStatus, allowing SQL injection against DB2, SQL Server, PostgreSQL, and other affected datasources. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64182	In the Linux kernel, the following vulnerability has been resolved: drivers/base/memory: fix memory block reference leak in poison accounting memblk_nr_poison_inc() and memblk_nr_poison_sub() look up a memory block via find_memory_block_by_id(), which acquires a reference to the memory block device. Both helpers use the returned memory block without dropping that reference, leaking the device reference on each successful lookup. Drop the reference after updating nr_hwpoison.	N/A	More Details
CVE-2026-64183	In the Linux kernel, the following vulnerability has been resolved: efi: Allocate runtime workqueue before ACPI init Since commit 5894cf571e14 ("acpi/prmt: Use EFI runtime sandbox to invoke PRM handlers") ACPI PRM calls are delegated to a workqueue which runs in a kernel thread, making it easier to detect and mitigate faulting memory accesses performed by the firmware. Rafael reports that such PRM accesses may occur before efisubsys_init() executes, which is where the workqueue is allocated, leading to NULL pointer dereferences. Since acpi_init() [which triggers the early PRM accesses] executes as a subsys_initcall() as well, and has its own dependencies that may be sensitive to initcall ordering, deferring acpi_init() is not an option. So instead, split off the workqueue allocation into its own postcore initcall, as this is the only missing piece to allow EFI runtime calls to be made. This ensures that EFI runtime call (including PRM calls) are accessible to all code running at subsys_initcall() level.	N/A	More Details

CVE-2026-64184	In the Linux kernel, the following vulnerability has been resolved: mm/damon/sysfs-schemes: call missing mem_cgroup_iter_break() damon_sysfs_memcg_path_to_id() breaks mem_cgroup_iter() loop without calling mem_cgroup_iter_break(). This leaks the cgroup reference. Fix the issue by calling mem_cgroup_iter_break() before the break. The issue was discovered [1] by Sashiko.	N/A	More Details
CVE-2026-64185	In the Linux kernel, the following vulnerability has been resolved: sysfs: don't remove existing directory on update failure When sysfs_update_group() is called for a named group and create_files() fails (e.g. -ENOMEM), internal_create_group() calls kernfs_remove(kn) on the group directory. In the update path, kn was obtained via kernfs_find_and_get() and refers to a directory that already existed before this call. Removing it silently destroys a sysfs group that the caller did not create. Only remove the directory if we created it ourselves. On update failure the directory remains as it is left empty by remove_files() inside create_files(), but can be repopulated by a retry.	N/A	More Details
CVE-2026-64168	In the Linux kernel, the following vulnerability has been resolved: spi: sprd: fix error pointer deref after DMA setup failure The driver falls back to PIO mode if DMA setup fails during probe. Make sure to check the dma.enabled flag before trying to release the DMA channels also on late probe errors to avoid dereferencing an error pointer (or attempting to release a channel a second time). This issue was flagged by Sashiko when reviewing a devres allocation conversion patch.	N/A	More Details
CVE-2026-64166	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_ffa: Check for NULL FF-A ID table while driver registration The bus match callback assumes that every FF-A driver provides an id_table and dereferences it unconditionally. Enforce that contract at registration time so a buggy client driver cannot crash the bus during match.	N/A	More Details
CVE-2026-12484	A vulnerability in keras-team/keras version 3.15.0 allows unsafe deserialization of attacker-controlled PyTorch pickle data through the public `keras.layers.TorchModuleWrapper.from_config` method. This method invokes `torch.load(..., weights_only=False)` without requiring an explicit unsafe opt-in, such as a `safe_mode=False` parameter. When called outside a `SafeModeScope(True)` context, the absence of an ambient safe mode state permits unsafe deserialization by default. This issue can lead to arbitrary code execution if untrusted Keras layer configurations are processed using this method. The vulnerability arises because the method does not enforce safe deserialization practices unless explicitly guarded by Keras safe mode.	N/A	More Details
CVE-2026-64165	In the Linux kernel, the following vulnerability has been resolved: ARM: integrator: Fix early initialization Starting with commit bdb249fce9ad4 ("ARM: integrator: read counter using syscon/regmap"), intcp_init_early calls syscon_regmap_lookup_by_compatible which in turn calls of_syscon_register. This function allocates memory. Since the memory management code has not been initialized at that time, the call always fails. It either returns -ENOMEM or crashes as follows. Unable to handle kernel NULL pointer dereference at virtual address 0000000c when read [0000000c] *pgd=00000000 Internal error: Oops: 5 [#1] ARM Modules linked in: CPU: 0 UID: 0 PID: 0 Comm: swapper Not tainted 6.15.0-rc5-00026-g5fcc9bf84ee5 #1 PREEMPT Hardware name: ARM Integrator/CP (Device Tree) PC is at __kmalloc_cache_noprof+0xec/0x39c LR is at __kmalloc_cache_noprof+0x34/0x39c ... Call trace: __kmalloc_cache_noprof from of_syscon_register+0x7c/0x310 of_syscon_register from device_node_get_regmap+0xa4/0xb0 device_node_get_regmap from intcp_init_early+0xc/0x40 intcp_init_early from start_kernel+0x60/0x688 start_kernel from 0x0 The crash is seen due to a dereferenced pointer which is not supposed to be NULL but is NULL if the memory management subsystem has not been initialized. The crash is not seen with all versions of gcc. Some versions such as gcc 9.x apparently do not dereference the pointer, presumably if tracing is disabled. The problem has been reproduced with gcc 10.x, 11.x, and 13.x. Either case, if the crash is not seen, the call to syscon_regmap_lookup_by_compatible returns -ENOMEM, and sched_clock_register is never called. Fix the problem by moving the early initialization code into the standard machine initialization code.	N/A	More Details
CVE-2026-64145	In the Linux kernel, the following vulnerability has been resolved: wifi: wilc1000: fix dma_buffer leak on bus acquire failure wilc_wlan_firmware_download() allocates dma_buffer with kmalloc() at the top of the function and uses a 'fail:' label to free it via kfree(dma_buffer) on error. All later error paths correctly use 'goto fail' to route through this cleanup. However, the early failure path after the first acquire_bus() call uses a bare 'return ret;', which leaks dma_buffer whenever the bus acquire fails. Replace the early return with goto fail so the existing cleanup path runs. Found via a custom Coccinelle semantic patch hunting for kmalloc'd locals leaked on early-return error paths in driver firmware-download code.	N/A	More Details
CVE-2026-64146	In the Linux kernel, the following vulnerability has been resolved: erofs: fix metabuf leak in inode xattr initialization commit bb88e8da0025 ("erofs: use meta buffers for xattr operations") converted xattr operations to use on-stack erofs_buf instances. erofs_init_inode_xattrs() uses such a metabuf while reading the inline xattr header and shared xattr id array. Some error paths after erofs_read_metabuf() leave through out_unlock without dropping the metabuf, so the folio reference can leak. Consolidate the cleanup at out_unlock. erofs_put_metabuf() is a no-op if no folio has been acquired, and this keeps all paths after taking EROFS_I_BL_XATTR_BIT covered by a single cleanup site.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: pds_core: fix debugfs_lookup dentry leak and error handling debugfs_lookup() returns a dentry with an elevated reference count that must be released with dput(). The current code discards the returned dentry without calling dput(), causing a reference leak on every firmware reset recovery. Additionally, when CONFIG_DEBUG_FS is disabled, debugfs_lookup() returns	N/A	More Details

64147	ERR_PTR(-ENODEV), not NULL. The current check passes for error pointers and would call dput() on an invalid pointer, causing a crash.		
CVE-2026-50030	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase SQL preview exposes DatasetDataApi.previewSql/previewSqlCheck through /de2api/datasetData/previewSql, accepts PreviewSqlDTO.sql, PreviewSqlDTO.datasourceId, and PreviewSqlDTO.isCross, then DatasetDataManage.previewSql stores decoded SQL in datasourceRequest.query and CalciteProvider.fetchResultField executes it with prepareStatement(...).executeQuery(), allowing arbitrary readable datasource tables to be queried and returned in preview responses. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64149	In the Linux kernel, the following vulnerability has been resolved: dma-mapping: move dma_map_resource() sanity check into debug code dma_map_resource() uses pfn_valid() to ensure the range is not RAM. However, pfn_valid() only checks for availability of the memory map for a PFN but it does not ensure that the PFN is actually backed by RAM. On ARM64 with SPARSEMEM (128MB section granularity), MMIO addresses that share a section with RAM will falsely trigger the WARN_ON_ONCE and cause dma_map_resource() to return DMA_MAPPING_ERROR. This causes a WARNING on Raspberry Pi 4 during spi_bcm2835 probe because the SPI FIFO register (0xfe204004) falls in the same sparsemem section as the end of RAM (0xf8000000-0xfbffffff), both in section 31 (0xf8000000-0xffffffff). Move the sanity check from dma_map_resource() into debug_dma_map_phys() and replace the unreliable pfn_valid() with pfn_valid() && !PageReserved(), which correctly identifies actual usable RAM without false positives for MMIO regions that happen to have struct pages. Since dma_map_resource() is dma_map_phys(DMA_ATTR_MMIO), the check applies equally to both APIs. Any non-reserved page represents kernel memory to a sufficient degree that using DMA_ATTR_MMIO on it is almost certainly wrong and risks breaking coherency on non-coherent platforms. ZONE_DEVICE pages used for PCI P2P DMA (MEMORY_DEVICE_PCI_P2PDMA) have PageReserved set, so they will not trigger a false positive. The check no longer blocks the mapping and uses err_printk() to integrate with dma-debug filtering.	N/A	More Details
CVE-2026-49867	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase template static resources let authenticated users submit TemplateManageRequest.staticResource through POST /de2api/templateManage/save or DataVisualizationServer.decompression, after which StaticResourceServer.saveFilesToServe and StaticResourceServer.saveSingleFileToServe write Base64-decoded .svg content to /de2api/static-resource/<name>.svg without validating extension, MIME type, decoded bytes, or SVG scriptability, causing stored same-origin cross-site scripting when a victim loads the resource. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-46684	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase enterprise token handling can let TokenFilter#doFilter() pass X-DE-TOKEN values to TokenUtils.validate(), which checks only token presence and length before userBOByToken(token) uses JWT.decode() without signature verification, allowing forged tokens with chosen uid and oid values to be accepted when licenseValid=true. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-63794	In the Linux kernel, the following vulnerability has been resolved: KVM: SVM: Fix page overflow in sev_dbg_crypt() for ENCRYPT path In sev_dbg_crypt(), the per-iteration transfer length is bounded by the source page offset (PAGE_SIZE - s_off) but not by the destination page offset (PAGE_SIZE - d_off). When d_off > s_off, the encrypt path (__sev_dbg_encrypt_user) performs a read-modify-write using a single-page intermediate buffer (dst_tpage): 1. __sev_dbg_decrypt() expands the size to round_up(len + (d_off & 15), 16) before issuing the PSP command. If len + (d_off & 15) > PAGE_SIZE, the PSP writes beyond the end of the 4096-byte dst_tpage allocation. 2. The subsequent memcpy()/copy_from_user() into page_address(dst_tpage) + (d_off & 15) of 'len' bytes overflows by up to 15 bytes under the same condition. Trigger example: s_off = 0, d_off = 1, debug.len = PAGE_SIZE - the PSP is instructed to write round_up(4097, 16) = 4112 bytes to a 4096-byte buffer. Fix by also bounding len by (PAGE_SIZE - d_off), the same check that sev_send_update_data() already performs for its single-page guest region. == BUG: KASAN: slab-use-after-free in sev_dbg_crypt+0x993/0xd10 [kvm_amd] Write of size 4095 at addr ff110062293bb009 by task sev_dbg_test/228214 CPU: 96 UID: 0 PID: 228214 Comm: sev_dbg_test Tainted: G U W 7.0.0-smp--5ce9b0c48211-dbg #156 PREEMPTLAZY Tainted: [U]=USER, [W]=WARN Hardware name: Google Astoria/astoria, BIOS 0.20250817.1-0 08/25/2025 Call Trace: <TASK> dump_stack_lvl+0x54/0x70 print_report+0xbc/0x260 kasan_report+0xa2/0xd0 kasan_check_range+0x25f/0x2c0 __asan_memcpy+0x40/0x70 sev_dbg_crypt+0x993/0xd10 [kvm_amd] sev_mem_enc_ioctl+0x33c/0x450 [kvm_amd] kvm_vm_ioctl+0x65d/0x6d0 [kvm] __se_sys_ioctl+0xb2/0x100 do_syscall_64+0xe8/0x870 entry_SYSCALL_64_after_hwframe+0x4b/0x53 </TASK> The buggy address belongs to the physical page: page: refcount:1 mapcount:0 mapping:0000000000000000 index:0x7fe72b6a0 pfn:0x62293bb memcg:ff11000112827d82 flags: 0x1400000000000000(node=1|zone=1) raw: 1400000000000000 0000000000000000 dead000000000122 0000000000000000 raw: 00000007fe72b6a0 0000000000000000 00000001ffffffff ff11000112827d82 page dumped because: kasan: bad access detected Memory state around the buggy address: ff110062293bbf00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ff110062293bbf80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 >ff110062293bc000: fa fb fb fb fb fb fb fb fc fc fc fc fc fc fc fc ^ ff110062293bc080: fa fb fb fb fb fb fb fb fc fc fc fc fc fc fc fc ff110062293bc100: fa fb fb fb fb fb fb fb fc fc fc fc fc fc fc fc == Disabling lock debugging due to kernel taint [sean: add sample KASAN splat, Fixes, and stable@]	N/A	More Details

CVE-2026-64154	In the Linux kernel, the following vulnerability has been resolved: drm/msm/adreno: Fix a reference leak in a6xx_gpu_init() In a6xx_gpu_init(), node is obtained via of_parse_phandle(). While there was a manual of_node_put() at the end of the common path, several early error returns would bypass this call, resulting in a reference leak. Fix this by using the __free(device_node) cleanup handler to release the reference when the variable goes out of scope. Patchwork: https://patchwork.freedesktop.org/patch/700661/	N/A	More Details
CVE-2026-64155	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: fix error path leaks in some WMI WOW calls Fix two instances where we used to directly return the result of ath11k_wmi_cmd_send(...). Because we did not check the return value, we also did not free the skb in the error path.	N/A	More Details
CVE-2026-64156	In the Linux kernel, the following vulnerability has been resolved: netfs, afs: Fix write skipping in dir/link writepages Fix netfs_write_single() and afs_single_writepages() to better handle a write that would be skipped due to lock contention and WB_SYNC_NONE by returning 1 from netfs_write_single() if it skipped and making afs_single_writepages() skip also. If a skip occurs, the inode must be re-marked as the VFS may have cleared the mark. This is really only theoretical for directories in netfs_write_single() as the only path to that is through afs_single_writepages() that takes the ->validate_lock around it, thereby serialising it.	N/A	More Details
CVE-2026-64157	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix partial invalidation of streaming-write folio In netfs_invalidate_folio(), if the region of a partial invalidation overlaps the front (but not all) of a dirty write cached in a streaming write page (dirty, but not uptodate, with the dirty region tracked by a netfs_folio struct), the function modifies the dirty region - but incorrectly as it moves the region forward by setting the start to the start, not the end, of the invalidation region. Fix this by setting finfo->dirty_offset to the end of the invalidation region (iend).	N/A	More Details
CVE-2026-45535	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase SQL-type datasets store attacker-controlled SQL variable defaultValue entries such as ${var} and SqlparserUtils.handleVariableDefaultValue() inserts them with String.replace() without escaping or parameterization, causing stored SQL injection whenever a user with dataset read permission accesses the dataset. This issue is fixed in version 2.10.23.	N/A	More Details
CVE-2026-64159	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix zeropoint update where i_size > remote_i_size Fix the update of the zero point[*] by netfs_release_folio() when there is uncommitted data in the pagecache beyond the folio being released but the on-server EOF is in this folio (ie. i_size > remote_i_size). The update needs to limit zero_point to remote_i_size, not i_size as i_size is a local phenomenon reflecting updates made locally to the pagecache, not stuff written to the server. remote_i_size tracks the server's i_size. [*] The zero point is the file position from which we can assume that the server will just return zeros, so we can avoid generating reads. Note that netfs_invalidate_folio() probably doesn't need fixing as zero_point should be updated by setattr after truncation or fallocate. Found with: fsx -q -N 1000000 -p 10000 -o 128000 -l 600000 \ /xfstest.test/junk --replay-ops=junk.fsxops using the following as junk.fsxops: truncate 0x0 0x1bbae 0x82864 write 0x3ef2e 0xf9c8 0x1bbae write 0x67e05 0xcb5a 0x4e8f6 mapread 0x57781 0x85b6 0x7495f copy_range 0x5d3d 0x10329 0x54fac 0x7495f write 0x64710 0x1c2b 0x7495f mapread 0x64000 0x1000 0x7495f on cifs with the default cache option. It shows read-gaps on folio 0x64 failing with a short read (ie. it hits EOF) if the FMODE_READ check is commented out in netfs_perform_write(): if (//(file->f_mode & FMODE_READ) || netfs_is_cache_enabled(ctx)) { and no fscache. This was initially found with the generic/522 xfstest.	N/A	More Details
CVE-2026-64161	In the Linux kernel, the following vulnerability has been resolved: net: ti: icssm-prueth: fix eth_ports_node leak in probe The error path on of_property_read_u32() failure inside icssm_prueth_probe() returns without putting eth_ports_node, which was acquired before the for_each_child_of_node() loop. Drop it before returning.	N/A	More Details
CVE-2026-64163	In the Linux kernel, the following vulnerability has been resolved: test_kprobes: clear kprobes between test runs Running the kprobes sanity tests twice makes all tests fail and eventually crashes the kernel. [root@martin-riscv-1 ~]# echo 1 > /sys/kernel/debug/kunit/kprobes_test/run ... # Totals: pass:5 fail:0 skip:0 total:5 ok 1 kprobes_test [root@martin-riscv-1 ~]# echo 1 > /sys/kernel/debug/kunit/kprobes_test/run ... # test_kprobe: EXPECTATION FAILED at lib/tests/test_kprobes.c:64 Expected 0 == register_kprobe(&kp), but register_kprobe(&kp) == -22 (0xffffffffffffffea) ... Unable to handle kernel paging request ... The testsuite defines several kprobes and kretprobes as static variables that are preserved across test runs. After register_kprobe and unregister_kprobe, a kprobe contains some leftover data that must be cleared before the kprobe can be registered again. The tests are setting symbol_name to define the probe location. Address and flags must be cleared. The existing code clears some of the probes between subsequent tests, but not between two test runs. The leftover data from a previous test run makes the registrations fail in the next run. Move the cleanups for all kprobes into kprobes_test_init, this function is called before each single test (including the first test of a test run).	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: btrfs: tracepoints: fix sleep while in atomic context in btrfs_sync_file() The trace event btrfs_sync_file() is called in an atomic context (all trace events are) and its call to dput(), which is needed due to the call to dget_parent(), can sleep, triggering a kernel splat. This can be reproduced by enabling the trace event and running btrfs/056 from fstests for example. The splat shown in dmesg is the following: [53.919] BUG: sleeping function called from invalid context at fs/dcache.c:970 [53.947] in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 32773, name: xfs_io [53.988]		

CVE-2026-64164	preempt_count: 2, expected: 0 [53.967] RCU nest depth: 0, expected: 0 [53.943] Preemption disabled at: [53.944] [<0000000000000000>] 0x0 [54.078] CPU: 0 UID: 0 PID: 32773 Comm: xfs_io Tainted: G W 7.1.0-rc1-btrfs-next-232+ #1 PREEMPT(full) [54.070] Tainted: [W]=WARN [54.071] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.2-0-gea1b7a073390-prebuilt.qemu.org 04/01/2014 [54.072] Call Trace: [54.074] <TASK> [54.076] dump_stack_lvl+0x56/0x80 [54.079] __might_resched.cold+0xd6/0x10f [54.072] dput.part.0+0x24/0x110 [54.078] trace_event_raw_event_btrfs_sync_file+0x75/0x140 [btrfs] [54.089] btrfs_sync_file+0x1ed/0x530 [btrfs] [54.087] ? __handle_mm_fault+0x8ae/0xed0 [54.089] btrfs_do_write_iter+0x172/0x210 [btrfs] [54.091] vfs_write+0x21f/0x450 [54.094] __x64_sys_pwrite64+0x8d/0xc0 [54.096] ? do_user_addr_fault+0x20c/0x670 [54.099] do_syscall_64+0x60/0xf20 [54.092] ? clear_bhb_loop+0x60/0xb0 [54.094] entry_SYSCALL_64_after_hwframe+0x76/0x7e So stop using dget_parent() and dput() and access the parent dentry directly as dentry->d_parent. This is also what ext4 is doing in its equivalent trace event ext4_sync_file_enter().	N/A	More Details
CVE-2026-64186	In the Linux kernel, the following vulnerability has been resolved: iommu/amd: Remove latent out-of-bounds access in IOMMU debugfs In iommu_mmio_write() and iommu_capability_write(), the variables dbg_mmio_offset and dbg_cap_offset are declared as int. However, they are populated using kstrtou32_from_user(). If a user provides a sufficiently large value, it can become a negative integer. Prior to this patch, the AMD IOMMU debugfs implementation was already protected by different mechanisms. 1. #define OFS_IN_SZ 8 ensures the user string <= 8 bytes, so e.g. 0xffffffff isn't a valid input. if (cnt > OFS_IN_SZ) return -EINVAL; 2. Implicit type promotion in iommu_mmio_write(), dbg_mmio_offset is int and iommu->mmio_phys_end is u64 if (dbg_mmio_offset > iommu->mmio_phys_end - sizeof(u64)) return -EINVAL; 3. The show handlers would currently catch the negative number and refuse to perform the read. Replace kstrtou32_from_user() with kstrtos32_from_user() to parse the input, and check for negative values to explicitly prevent out-of-bounds memory accesses directly in iommu_mmio_write() and iommu_capability_write().	N/A	More Details
CVE-2026-45320	DataEase is an open source data visualization and analysis tool. Prior to 2.10.23, DataEase dashboard SQL variables such as ${deptId} are processed by SqlparserUtils.transFilter(), whose final branch returns raw user input for non-in and non-between operators before SubstitutedSql.replace("${var}", value) splices it into dashboard SQL, allowing authenticated users who can view a dashboard to inject SQL against integrated datasources. This issue is fixed in version 2.10.23	N/A	More Details
CVE-2026-27823	A vulnerability has been identified in EGroupware that may lead to Remote Code Execution (RCE). The issue allows an authenticated attacker to execute arbitrary commands on the server. If user self-registration is enabled, the vulnerability may be exploitable without prior authentication. The vulnerability stems from improper authorization checks combined with a file write primitive and an arbitrary file read vulnerability, which together enable full system compromise. This has been patched in versions 26.2.20260224 and 23.1.20260224.	N/A	More Details
CVE-2026-46421	The SAP Cloud Application Programming Model is a tool for building enterprise-grade cloud applications, and cap-js/cds-dbs is the monorepo for SQL database services for that tool. On April 29, 2026, compromised versions of `@cap-js/sqlite@2.2.2`, `@cap-js/postgres@2.2.2`, and `@cap-js/db-service@2.10.1` were published. The malicious packages harvested credentials and attempted self-propagation. If a compromised version was installed, all credentials accessible on that machine (npm tokens, cloud provider credentials, SSH keys, GitHub PATs) should be considered compromised. User should upgrade to `@cap-js/sqlite` >= 2.4.0, `@cap-js/postgres` >= 2.3.0, `@cap-js/db-service` >= 2.11.0. If a compromised version was ever installed, rotate all affected credentials. No known workarounds are available.	N/A	More Details
CVE-2026-63763	SurrealDB before 2.5.0 and before 3.0.0-beta.3 contains a confused deputy privilege escalation vulnerability. Unprivileged users (e.g., those with the database editor role) can create or modify fields containing futures, functions, or closures. Because these are executed in the context of the invoking/querying user rather than their creator, an attacker can plant malicious logic that executes with a higher-privileged user's permissions when that user reads or writes the affected record. This can lead to full privilege escalation, including creation of a root owner and server takeover.	N/A	More Details
CVE-2026-45334	Kirby is an open-source content management system. In versions prior to 4.9.1 and 5.4.1, the content-locking feature returned lock information without checking the requesting user's access permissions. Kirby's Panel includes a content-locking feature that records which user currently has a model open for editing. This lock prevents conflicting edits by multiple users and displays the locking user's identity in the Panel UI so other users know who to contact. Internally, the locking user's email address and identifier are included in every Panel view payload and in error responses returned when a user attempts to edit a model that is currently locked by another user. This allowed a low-privilege authenticated Panel user, whose role was configured with users.access: false or users.list: false, to learn the email address and identifier of any user who currently had a model open for editing in the Panel, including administrators and other higher-privilege users. Content locks are active for a configurable window (10 minutes by default). The email address can allow admin account enumeration, target phishing, and feed credential-stuffing attacks against the Kirby installation or other sites. The internal user ID can be cross-referenced with other endpoints once the requester has obtained a higher privilege through unrelated means. This issue has been fixed in versions 4.9.1 and 5.4.1.	N/A	More Details
	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-48866. Reason:		

CVE-2026-8055	This candidate is a reservation duplicate of CVE-2026-48866. Notes: All CVE users should reference CVE-2026-48866 instead of this candidate. All references and descriptions in this candidate have been removed to prevent accidental usage.	N/A	More Details
CVE-2026-45368	Kirby is an open-source content management system. In versions prior to 4.9.1 and 5.4.1, the underlying URL methods for the KirbyTags and image blocks components did not filter out malicious URL values that resolve to script execution. The vulnerability affects four first-party Kirby renderers that produce `<a href="...">` output from editor-supplied field values: the (`link: ...)` KirbyTag, the `link`: parameter of the `(image: ...)` KirbyTag when it does not resolve to a known file or `self`, the `link` field of the built-in image block, and the HTML importer for the `blocks` field (which accepted the same malicious input as the image block `link` field). While simple `avascript:` URLs were already deactivated by treating them as a relative path and prepending a single slash to the URL, the use of URLs of the format `javascript://x%0A...` bypasses this protection. The `vbscript:`, `data:`, `livescript:`, `mocha:` and `jar:` schemes are affected by the same underlying gap. This issue has been fixed in versions 4.9.1 and 5.4.1.	N/A	More Details
CVE-2026-57309	A Blind SQL injection vulnerability has been identified in Windu CMS. A remote unauthenticated attacker is able to inject SQL syntax into URL path in HTTP header resulting in Blind SQL Injection. Because vendor contact attempts were unsuccessful, the vulnerability has only been confirmed in version 4.1 but may also affect other versions.	N/A	More Details
CVE-2026-57310	Windu CMS uses hashing algorithm based on MD5 and SHA1 with static salt to store user passwords. This allows an attacker who obtain password hash to decode user credentials. Because vendor contact attempts were unsuccessful, the vulnerability has only been confirmed in version 4.1 but may also affect other versions.	N/A	More Details
CVE-2026-57311	Windu CMS does not validate types of uploaded files. An authenticated attacker can upload arbitrary files, including PHP. This can lead to Remote Code Execution. Because vendor contact attempts were unsuccessful, the vulnerability has only been confirmed in version 4.1 but may also affect other versions.	N/A	More Details
CVE-2026-46410	FileBrowser Quantum is a free, self-hosted, web-based file manager. Versions prior to 1.3.2-stable and 1.4.1-beta may leak some sensitive info, such as source and path. Versions 1.3.2-stable and 1.4.1-beta fix the issue. No known workarounds are available.	N/A	More Details
CVE-2026-46516	Frogman provides headless FreePBX control. Prior to version 1.6.6, Frogman's chat-console markdown formatter (`assets/js/chat.js`'s `formatMarkdown`) inserted regex capture groups as raw HTML in four template patterns: inline code, bold, markdown links, and download links. Tool responses that reflect user-controlled fields — extension names, ring-group descriptions, IVR names, queue descriptions, etc. — could carry an HTML/JavaScript payload that executes when another admin views the response through Frogman chat. The payload runs in the viewer's session, with the viewer's permissions. FreePBX's own admin GUI escapes these same fields via `freepbx_htmlspecialchars()` throughout its view templates. The chat formatter was the leaky side. Version 1.6.6 patches the issue.	N/A	More Details
CVE-2026-51386	Rejected reason: DO NOT USE THIS CVE RECORD. ConsultIDs: CVE-2026-46409. Reason: This record is a reservation duplicate of CVE-2026-46409. Notes: All CVE users should reference CVE-2026-46409 instead of this record. All references and descriptions in this record have been removed to prevent accidental usage.	N/A	More Details
CVE-2026-50562	FastGPT is a knowledge-based AI application platform. At commit 22ebfacbb43311e9b73294040ae0eb87390c6bba and earlier, artifacts built from untrusted pull request code in .github/workflows/preview-docs-build.yml and .github/workflows/preview-fastgpt-build.yml can be downloaded by privileged workflow_run jobs in .github/workflows/preview-docs-push.yml and .github/workflows/preview-fastgpt-push.yml, allowing attacker-controlled Docker images from the document/tree or FastGPT build context to be pushed to GHCR and, for documentation previews, deployed with secrets.KUBE_CONFIG_CN.	N/A	More Details
CVE-2026-59238	Stored Cross-site Scripting (CWE-79) in the client-side report rendering functions (renderPreview, renderEditor, renderAuditData in js/app.js) in maalfer Pentestify before 1.1.0 allows a remote, authenticated attacker to execute arbitrary JavaScript in the browser of any user who views an affected report via a payload stored in a finding's images array or a report's client_logo array, which is interpolated into an <img> src attribute without escaping.	N/A	More Details
CVE-2026-11740	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-62287	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61873. Reason: This candidate is a duplicate of CVE-2026-61873. Notes: All CVE users should reference CVE-2026-61873 instead of this candidate.	N/A	More Details
CVE-2026-62248	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-61438. Reason: This candidate is a duplicate of CVE-2026-61438. Notes: All CVE users should reference CVE-2026-61438 instead of this candidate.	N/A	More Details

CVE-2026-26197	HDF5 is a high-performance library and a file format specification that implements the HDF5 data model. If a file is corrupted such that an array datatype's size, the number of elements, and the element size are not in agreement it can trigger an out of bounds read. The array datatype stores the full size of the datatype (`dt->shared->size`) separately from the number of elements (`dt->shared->u.array.nelem`) and the element size (`dt->shared->parent->shared->size`). If any one of these are corrupted so that they don't align with the others (element size * nelem = full size), it can lead to an out of bounds read. Depending on what is corrupted, it can alter the type of out of bounds read triggered. The vulnerability is present only in files that have been maliciously altered, as its generally not possible to independently alter the full size of the datatype, the element count and the element size. As such, this is only present if a malicious actor is altering files, and won't appear in regular usage.	N/A	More Details
CVE-2026-26199	HDF5 is a high-performance library and a file format specification that implements the HDF5 data model. If `H5Iget_name` is invoked on a group id with `0` for the size parameter, it will underflow when trying to place a null terminator in the buffer. This can occur if `H5Iget_name` is invoked in a way where `size` can be forced to zero, and there is important data before the `name` buffer.	N/A	More Details
CVE-2026-63762	SurrealDB before v2.6.1 (and before v3.0.0-beta.3) contains a denial of service vulnerability in its embedded JavaScript scripting engine, which is enabled via the --allow-scripting capability (disabled by default). Any user able to execute arbitrary queries — including unauthenticated guests when --allow-guests is enabled — can use built-in string functions to construct a large string and pass it to the JavaScript runtime for compilation, triggering a null pointer dereference in the underlying QuickJS-NG engine. This causes the server process to terminate immediately without graceful shutdown, requiring a manual restart. The issue was fixed by updating the rquickjs dependency from v0.9.0 to v0.11.0.	N/A	More Details
CVE-2026-49987	Repomix is a tool that packs repositories into AI-friendly files. Prior to 1.14.1, src/core/git/gitCommand.ts execGitShallowClone passes the --remote-branch value directly to git fetch and git checkout without validation or --end-of-options, allowing --upload-pack or other Git option injection that bypasses validateGitUrl() dangerous parameter checks and can execute commands through local or SSH-style transports. This issue is fixed in version 1.14.1.	N/A	More Details
CVE-2026-53403	In the Linux kernel, the following vulnerability has been resolved: fbdev: Fix fb_new_modelist to prevent null-ptr-deref in fb_videomode_to_var info->var, a framebuffer's current mode, is expected to have a matching entry in info->modelist. var_to_display() relies on this and treats a failed fb_match_mode() as "This should not happen". fb_set_var() keeps it true by adding the mode to the list on every change, and do_register_framebuffer() does the same at registration. store_modes() replaces the modelist from userspace. fb_new_modelist() validates the new modes but does not check that info->var still has a match. It relies on fbcon_new_modelist() to re-point consoles, but that only handles consoles mapped to the framebuffer. With fbcon unbound there are none, so info->var is left describing a mode that is no longer in the list. A later console takeover runs var_to_display(), where fb_match_mode() returns NULL and leaves fb_display[i].mode NULL. fbcon_switch() passes it to display_to_var(), and fb_videomode_to_var() dereferences the NULL mode. Keep the current mode in the list in fb_new_modelist(), the same way fb_set_var() does.	N/A	More Details
CVE-2026-49988	Repomix is a tool that packs repositories into AI-friendly files. Prior to 1.14.1, the Repomix MCP server attach_packed_output and read_repomix_output flow can register and read arbitrary local .json, .txt, .md, or .xml files without the file_system_read_file runSecretLint() safety check or Repomix packed-output validation, allowing MCP callers to bypass the local file-read secret-scanning boundary. This issue is fixed in version 1.14.1.	N/A	More Details
CVE-2026-9770	Kasa EC71 v4 and EC70 v4 firmware contains a static cryptographic private key stored in a read-only filesystem that is shared across devices. An attacker with access to the firmware image can extract the embedded key. Successful exploitation may allow an unauthenticated attacker on the same network to use this key in the web management service, compromising the confidentiality of encrypted communications. This may enable passive decryption of traffic or active man-in-the-middle (MITM) attacks	N/A	More Details
CVE-2026-53393	In the Linux kernel, the following vulnerability has been resolved: nfsd: reset write verifier on deferred writeback errors nfsd_vfs_write() and nfsd_commit() both call filemap_check_wb_err() to detect deferred writeback errors, but neither rotates the server's write verifier (nn->writeverf) when this check fails. Every other durable-storage-failure path in these functions calls commit_reset_write_verifier() before returning an error. The missing rotation means clients holding UNSTABLE write data under the current verifier will COMMIT, receive the unchanged verifier back, and conclude their data is durable — silently dropping data that failed writeback. This violates the UNSTABLE+COMMIT durability contract (RFC 1813 §3.3.7, RFC 8881 §18.32). Add commit_reset_write_verifier() calls at both filemap_check_wb_err() error sites, matching the pattern used by adjacent error paths in the same functions. The helper already filters -EAGAIN and -ESTALE internally, so the calls are unconditionally safe.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: vc_screen: fix null-ptr-deref in vcs_notifier() during concurrent vcs_write A KASAN null-ptr-deref was observed in vcs_notifier(): BUG: KASAN: null-ptr-deref in vcs_notifier+0x98/0x130 Read of size 2 at addr qmp_cmd_name: qmp_capabilities, arguments: {} The issue is a race condition in vcs_write(). When the console_lock is temporarily dropped (to copy data from userspace), the vc_data pointer obtained from vcs_vc() may become stale. After re-acquiring the lock,	N/A	More

53385	vcs_vc() is called again to re-validate the pointer. If the vc has been deallocated in the meantime, vcs_vc() returns NULL, and the while loop breaks (with written > 0). However, after the loop, vcs_scr_updated(vc) is still called with the now-NULL vc pointer, leading to a null pointer dereference in the notifier chain (vcs_notifier dereferences param->vc). Fix this by adding a NULL check for vc before calling vcs_scr_updated().		Details
CVE-2026-53382	In the Linux kernel, the following vulnerability has been resolved: media: vidtv: fix NULL pointer dereference in vidtv_mux_push_si syzbot reported a general protection fault in vidtv_psi_ts_psi_write_into [1]. vidtv_mux_get_pid_ctx() can return NULL, but vidtv_mux_push_si() does not check for this before dereferencing the returned pointer to access the continuity counter. This leads to a general protection fault when accessing a near-NULL address. The root cause is that vidtv_mux_pid_ctx_init() does not check the return value of vidtv_mux_create_pid_ctx_once() for PMT section PIDs. If the allocation fails, the PID context is never created, but init returns success. The subsequent vidtv_mux_push_si() call then gets NULL from vidtv_mux_get_pid_ctx() and crashes. Fix both the root cause (add error check in vidtv_mux_pid_ctx_init for PMT PIDs) and add defensive NULL checks in vidtv_mux_push_si for all vidtv_mux_get_pid_ctx() calls. [1] Oops: general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] Workqueue: events vidtv_mux_tick RIP: 0010:vidtv_psi_ts_psi_write_into+0x54a/0xbc0 drivers/media/test-drivers/vidtv/vidtv_psi.c:197 Call Trace: <TASK> vidtv_psi_table_header_write_into drivers/media/test-drivers/vidtv/vidtv_psi.c:799 [inline] vidtv_psi_pmt_write_into+0x3b2/0xa70 drivers/media/test-drivers/vidtv/vidtv_psi.c:1231 vidtv_mux_push_si+0x932/0xe80 drivers/media/test-drivers/vidtv/vidtv_mux.c:196 vidtv_mux_tick+0xe9b/0x1480 drivers/media/test-drivers/vidtv/vidtv_mux.c:408	N/A	More Details
CVE-2026-53379	In the Linux kernel, the following vulnerability has been resolved: media: i2c: ov8856: free control handler on error in ov8856_init_controls() The control handler wasn't freed if adding controls failed, add an error exit label and convert the existing error return to use it.	N/A	More Details
CVE-2026-53378	In the Linux kernel, the following vulnerability has been resolved: drm/colorop: Fix blob property reference tracking in state lifecycle The colorop state blob property handling had memory leaks during state duplication, destruction, and reset operations. The implementation failed to follow the established pattern from drm_crtc's handling of DEGAMMA/GAMMA blob properties. Issues fixed: - drm_colorop_atomic_destroy_state() was freeing state memory without releasing the blob reference, causing a leak - drm_colorop_reset() was directly freeing old state with kfree() instead of properly destroying it, leaking blob references - drm_colorop_cleanup() had duplicate blob cleanup code Changes: - Add __drm_atomic_helper_colorop_destroy_state() helper to properly release blob references before freeing state memory - Update drm_colorop_atomic_destroy_state() to call the helper - Fix drm_colorop_reset() to use drm_colorop_atomic_destroy_state() for proper cleanup of old state - Simplify drm_colorop_cleanup() to use the common destruction path This matches the well-tested pattern used by drm_crtc since 2016 and ensures proper reference counting throughout the state lifecycle. Co-developed by Claude Sonnet 4.5.	N/A	More Details
CVE-2026-53377	In the Linux kernel, the following vulnerability has been resolved: drm/msm: always recover the gpu Previously, in case there was no more work to do, recover worker wouldn't trigger recovery and would instead rely on the gpu going to sleep and then resuming when more work is submitted. Recover_worker will first increment the fence of the hung ring so, if there's only one job submitted to a ring and that causes an hang, it will early out. There's no guarantee that the gpu will suspend and resume before more work is submitted and if the gpu is in a hung state it will stay in that state and probably trigger a timeout again. Just stop checking and always recover the gpu. Patchwork: https://patchwork.freedesktop.org/patch/704066/	N/A	More Details
CVE-2026-53376	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Add upper bound check for num_of_nodes drm/amdkfd: Add upper bound check for num_of_nodes in kfd_ioctl_get_process_apertures_new. (cherry picked from commit 98ff46a5ea090c14d2cdb4f5b993b05d74f3949f)	N/A	More Details
CVE-2026-44174	Kirby is an open-source content management system. Prior to 4.9.1 and 5.4.1, Kirby did not validate the model attributes that were used in its collection queries, allowing attackers to include arbitrary model methods in their queries. This includes methods with sensitive data such as password() (disclosing the password hash) or root() (disclosing the absolute filesystem path on the server) as well as methods that perform impactful actions such as loginPasswordless() (causing a privilege escalation to another user) or delete() (deleting all queried models in one go if the authenticated user has appropriate permissions). This issue has been fixed in versions 4.9.1 and 5.4.1.	N/A	More Details
CVE-2026-44175	Kirby is an open-source content management system. In versions prior to 4.9.1 and 5.4.1, Kirby did not securely sanitize the contents of the list field on save, leaving it vulnerable to cross-site scripting (XSS). Kirby's list field stores its formatted content as HTML, and unlike other field types, its HTML special characters cannot be escaped without losing the formatting. Sanitization was only enforced client-side in the Panel, while the server did not sanitize the content on save. As a result, an attacker could bypass the Panel and send malicious HTML directly to Kirby's API, storing unsanitized markup in the content file. That markup would then be rendered on the site frontend and executed in the browsers of site visitors and logged-in users browsing the site, resulting in persistent XSS. This issue has been fixed in versions 4.9.1 and 5.4.1.	N/A	More Details

CVE-2026-44176	Kirby is an open-source content management system. Versions prior to 4.9.1 and 5.4.1 do not check the `pages.access` permission during page draft rendering. Permissions are defined for each user role in the user blueprint (site/blueprints/users/...). It is also possible to customize the permissions for each target model in the model blueprints (such as in site/blueprints/pages/...) using the options feature. The permissions and options together control the authorization of user actions. Kirby provides the pages.access and pages.list permissions (among others). The list permission controls whether affected models appear in lists throughout the Panel and REST API. The access permission has the same effect but also disables direct access to the affected models. This vulnerability affects the path resolver for the main CMS router. The resolver takes an input path from the requested URL and determines which model (page or file) should be rendered. When a path is requested that points to a page draft, the resolver checks that the request either contains a valid preview token or is authenticated by a valid user. In affected releases, Kirby allowed page drafts to be rendered if any valid user was authenticated, even if that user did not have access to the specific page model. Authenticated attackers with knowledge of the full path to an existing page draft could then access the rendered frontend page. This could lead to the disclosure of sensitive information, e.g. ahead of the launch of a new product or post. This issue has been fixed in versions 4.9.1 and 5.4.1.	N/A	More Details
CVE-2026-53372	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Block PASID attachment to nested domain with dirty tracking Kernel lacks dirty tracking support on nested domain attached to PASID, fails the attachment early if nesting parent domain is dirty tracking configured, otherwise dirty pages would be lost.	N/A	More Details
CVE-2026-53371	In the Linux kernel, the following vulnerability has been resolved: RDMA/ionic: bound node_desc sysfs read with %.64s node_desc[64] in struct ib_device is not guaranteed to be NUL- terminated. The core IB sysfs handler uses "%.64s" for exactly this reason (drivers/infiniband/core/sysfs.c:1307), since node_desc_store() performs a raw memcpy of up to IB_DEVICE_NODE_DESC_MAX bytes with no NUL termination: memcpy(desc.node_desc, buf, min_t(int, count, IB_DEVICE_NODE_DESC_MAX)); If exactly 64 bytes are written via the node_desc sysfs file, the array contains no NUL byte. The ionic hca_type_show() handler uses unbounded "%s" and will read past the end of node_desc into adjacent fields of struct ib_device until it encounters a NUL. ionic supports IB_DEVICE_MODIFY_NODE_DESC, so this is triggerable by userspace. Match the core handler and bound the format specifier.	N/A	More Details
CVE-2026-53370	In the Linux kernel, the following vulnerability has been resolved: perf/x86/intel: Improve validation and configuration of ACR masks Currently there are several issues on the user space ACR mask validation and configuration. - The validation for user space ACR mask (attr.config2) is incomplete, e.g., the ACR mask could include the index which belongs to another ACR events group, but it's not validated. - An early return on an invalid ACR mask caused all subsequent ACR groups to be skipped. - The stale hardware ACR mask (hw.config1) is not cleared before setting new hardware ACR mask. The following changes address all of the above issues. - Figure out the event index group of an ACR group. Any bits in the user-space mask not present in the index group are now dropped. - Instead of an early return on invalid bits, drop only the invalid portions and continue iterating through all ACR events to ensure full configuration. - Explicitly clear the stale hardware ACR mask for each event prior to writing the new configuration. Besides, a non-leader event member of ACR group could be disabled in theory. This could cause bit-shifting errors in the acr_mask of remaining group members. But since ACR sampling requires all events to be active, this should not be a big concern in real use case. Add a "FIXME" comment to notice this risk.	N/A	More Details
CVE-2026-44177	Kirby is an open-source content management system. In versions 5.3.0 and above but prior to 5.4.1, Kirby did not correctly validate the provided user ID, resulting in a path traversal vulnerability. Version 5.3.0 introduced a performance improvement to the Users collection that loaded user objects lazily when first needed. Users were queried by their ID, which was then used to locate the corresponding account directory under site/accounts. This affected the authentication API (accessible to unauthenticated requests), the users API (accessible only to authenticated users), and any other place that uses $users->find() to look up an individual user by a request-provided email or ID. As a result, an attacker could trigger arbitrary PHP file inclusion of files named index.php (for example, the main PHP files of plugins), the impact of which depends on the logic those files contain. It also allowed probing for the existence of arbitrary directories on the server, letting attackers fingerprint the server and site setup, including installed plugins and the content structure. This issue has been fixed in version 5.4.1.	N/A	More Details
CVE-2026-53367	In the Linux kernel, the following vulnerability has been resolved: selinux: fix avdcache auditing The per-task avdcache was incorrectly saving and reusing the audited vector computed by avc_audit_required() rather than recomputing based on the currently requested permissions and distinguishing the denied versus allowed cases. As a result, some permission checks were not being audited, e.g. directory write checks after a previously cached directory search check. [PM: line wrap tweaks]	N/A	More Details
CVE-2026-54443	Dashy is a self-hostable personal dashboard. From 1.9.4 until 3.2.0, the Dashy RSS Widget in src/components/Widgets/RssFeed.vue does not sanitize RSS item link values before rendering feed item titles and Read More links as anchor href attributes, allowing an attacker-controlled feed to provide a javascript: URI that executes when clicked in the Dashy origin. This issue is fixed in version 3.2.0.	N/A	More Details
CVE-2026-	An information disclosure vulnerability was identified in TP-Link Kasa EC70 v4 and EC71 v4 in the local discovery mechanism, which exposes sensitive geolocation information without requiring authentication. This issue allows an attacker on the same local network to retrieve geolocation-related data through crafted	N/A	More Details

13230	responses. The vulnerability impacts confidentiality only, with no evidence of integrity of availability impact.		