SingCERT
Singapore Cyber Emergency Response Team

Security Bulletin 25 November 2026

Generated on 25 November 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-2026-71499	cf-manager is an all-in-one dashboard for managing multiple Cloudflare accounts and services. Prior to 1.1.2, worker/src/routes/workers.ts uses bare fetch(url) calls to retrieve user-supplied URLs and deploys the returned content as a Worker script using stored Cloudflare credentials. The multipart and JSON branches of POST /api/:accountId/workers and POST /api/batch-deploy accept URLs without protocol, destination, redirect, content-type, or response-size validation. An attacker with deployment API access, or unauthenticated access when API_SECRET is not configured, can make cf-manager access internal resources or fetch attacker-controlled JavaScript and execute it in the victim Cloudflare Workers account. This issue is fixed in version 1.1.2.	10.0	More Details
CVE-2026-95787	Vulnerabilities in the PAPI protocol of AOS-10 could allow an unauthenticated remote attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to gain administrative access, modify system configurations, and access or manipulate sensitive data.	10.0	More Details
CVE-2026-77312	Flyto2 Core is an execution kernel for automation and AI-agent workflows. Prior to 2.26.8, the earlier CVE-2026-67429 remediation applied validate_path_with_env_config() to only a hand-picked set of file-writing modules, leaving data.csv.write in src/core/modules/atomic/data/csv_write.py, data.json_to_csv in src/core/modules/atomic/data/json_to_csv.py, file.copy in src/core/modules/atomic/file/copy.py, and file.move in src/core/modules/atomic/file/move.py outside the FLYTO_SANDBOX_DIR confinement boundary. The advisory also identifies unguarded modules across the broader data.*, file.*, and report-writer families. data.csv.write relies on a .. substring denylist that does not reject absolute paths, while the other named modules accept caller-controlled output or destination paths without the canonical guard. A caller able to execute modules or recipes can therefore write attacker-controlled content outside the configured sandbox, potentially overwriting application code, recipes, startup files, scheduled-task files, or authorization files and gaining code execution on the host. This issue is fixed in version 2.26.8.	10.0	More Details
CVE-2026-130778	Incus before 7.6.0 contains an authentication bypass vulnerability that allows remote attackers to authenticate as any trusted certificate by appending its public certificate to their presented TLS client chain. Attackers can append the server's or a cluster member's certificate to gain cluster-level access, bypass authorization, modify trust store entries, and influence scriptlet-based authorization decisions.	10.0	More Details
CVE-2026-133323	A vulnerability in the Audio, Web, and Video Conferencing (AWV) component of Mitel MiCollab through 10.3 (10.3.0.18) could allow an unauthenticated attacker to conduct a blind command injection attack due to insufficient parameter sanitization. A successful exploit could allow an attacker to execute arbitrary commands within the context of the system. This is MTLVULN-1864.	10.0	More Details

CVE-2026-131253	In Splunk SOAR versions below 8.8.0, an authenticated user that holds a role with the Apps edit, Containers edit, and Playbooks execute permissions could use the App Wizard to execute arbitrary code as the user account that runs Splunk SOAR. This could compromise the confidentiality, integrity, and availability of the Splunk SOAR instance. The vulnerability exists because the App Wizard does not validate user-supplied identifiers before incorporating them into generated connector code.	9.9	More Details
CVE-2026-130780	Incus before 7.6.0 contains a path traversal vulnerability that allows malicious migration peers to write outside intended directories by supplying unvalidated dependent volume and snapshot names in the migration offer header. Attackers controlling a migration source can create root-owned volumes with attacker content at arbitrary host paths, compromising hosts using dir storage pools, bypassing restricted.storage-pools.access, or crashing the daemon.	9.9	More Details
CVE-2026-129987	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, the database password validation used by the MySQL, PostgreSQL, MariaDB, MongoDB, and Redis change-password procedures permits backticks, and the procedures interpolate attacker-controlled passwords into Bash command strings. An authenticated user who can update a database password can trigger command substitution before the database command runs, resulting in arbitrary operating-system command execution with the Dokploy process privileges, which are typically root. This issue is fixed in version 0.31.1.	9.9	More Details
CVE-2026-130771	Incus before 7.6.0 fails to reject special files, hard links, or symlinks in backup or image metadata, allowing restricted project users to read or overwrite host files. Attackers can import crafted backups containing device nodes, a hard-linked backup.yaml, or symlinks to corrupt host storage, overwrite host files, or leak other projects' data.	9.9	More Details
CVE-2026-132139	A weakness has been identified in DrayTek Vigor2960 up to 1.5.1.6. Affected is the function setSWMProfile of the file /cgi-bin/mainfunction.cgi. Executing a manipulation of the argument filename can lead to stack-based buffer overflow. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. Upgrading to version 1.5.1.7 is able to address this issue. It is recommended to upgrade the affected component. The vendor recommends: "1.5.1.6 isn't the latest version of 2960, the latest version fixes many security issues. And we have another one in our release plan which should solve the issues reported."	9.9	More Details
CVE-2026-130773	Incus before 7.6.0 contains an authorization bypass that allows restricted project users to create resources in the unrestricted default project by targeting projects lacking features.* flags. Attackers can create bridge networks with raw.dnsmasq dhcp-script injection to run scripts as root, or custom volumes with block.create_options mkfs arguments to read host files.	9.9	More Details
CVE-2026-132143	A vulnerability has been found in DrayTek Vigor2960 up to 1.5.1.6. This vulnerability affects the function doWebBackup of the file /cgi-bin/mainfunction.cgi of the component Backup Handler. Such manipulation of the argument path/value leads to stack-based buffer overflow. The attack can be launched remotely. The exploit has been disclosed to the public and may be used. Upgrading to version 1.5.1.7 is able to resolve this issue. The affected component should be upgraded. The vendor recommends: "1.5.1.6 isn't the latest version of 2960, the latest version fixes many security issues. And we have another one in our release plan which should solve the issues reported."	9.9	More Details
CVE-2026-129316	Dokploy is a free, self-hostable Platform as a Service (PaaS). From 0.29.13 until 0.31.0, the Docker terminal, container-log, and monitoring WebSocket handlers accept a caller-controlled serviceId as authorization but do not consistently bind the requested containerId or appName to that service. A member with read access to any service can consequently target unrelated containers, obtain a root shell in the Dokploy panel container, read other tenants' logs and statistics, and use the mounted Docker socket to compromise the host. This issue is fixed in version 0.31.0.	9.9	More Details
CVE-2026-129986	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, mount.deleteFileMount passes a path containing the user-controlled filePath to the local removeFileOrDirectory helper, which interpolates it unquoted into rm -rf, while the remote branch is quoted. An authenticated member with file-mount create and delete permissions can execute arbitrary commands in the host context immediately when deleting a local file mount, without requiring a running container. This issue is fixed in version 0.31.1.	9.9	More Details
CVE-2026-129989	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, the Nixpacks publishDirectory build setting accepts shell metacharacters and is interpolated into a docker cp command executed through child_process.exec. An authenticated user with access to an application can execute arbitrary commands on the Dokploy host when triggering a Nixpacks deployment, leading to host compromise through the privileged Docker environment. This issue is fixed in version 0.31.1.	9.9	More Details
CVE-2026-130784	Incus before 7.6.0 contains an authorization bypass vulnerability that allows restricted clients to create and access default project resources by exploiting entitlement checks against the requested project. Attackers can create bridge networks with raw.dnsmasq directives to execute code as root on the host, or manipulate inherited resources via the ?project= parameter.	9.9	More Details
	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, the backup prefix field is only checked for a nonempty value, and normalizeS3Path preserves shell metacharacters before the		

CVE-2026-129991	prefix is interpolated into rclone upload and retention-cleanup command strings. An authenticated organization member with backup permission for a service can execute arbitrary commands as root on the Dokploy host or configured remote server when a PostgreSQL, MySQL, MariaDB, MongoDB, libSQL, Compose, or web-server backup runs manually, on schedule, or performs retention cleanup. This issue is fixed in version 0.31.1.	9.9	More Details
CVE-2026-132141	A vulnerability was detected in DrayTek Vigor2960 up to 1.5.1.6. Affected by this issue is the function json_load_lock of the file /cgi-bin/mainfunction.cgi of the component Jsonstatus. The manipulation of the argument usefile results in stack-based buffer overflow. It is possible to launch the attack remotely. The exploit is now public and may be used. Upgrading to version 1.5.1.7 can resolve this issue. It is advisable to upgrade the affected component. The vendor recommends: "1.5.1.6 isn't the latest version of 2960, the latest version fixes many security issues. And we have another one in our release plan which should solve the issues reported."	9.9	More Details
CVE-2026-130786	Incus before 7.6.0 contains a configuration injection vulnerability that allows users who can launch instances from controlled images to inject LXC directives by embedding line breaks in OCI config.json values or oci.entrypoint and oci.cwd keys. Attackers can add liblxc hooks or crafted tmpfs lxc.mount.entry fields to run commands as root on the host or bind-mount host paths, even from restricted projects.	9.9	More Details
CVE-2026-132803	A malicious actor with access to the network and low privileges could exploit a series of authenticated SQL Injection vulnerabilities found in UniFi Access Application to escalate privileges on the host device.	9.9	More Details
CVE-2026-132804	A malicious actor with access to the network and low privileges could exploit a Server-Side Request Forgery (SSRF) vulnerability found in UniFi Access Application to escalate privileges within UniFi OS.	9.9	More Details
CVE-2026-132808	A malicious actor with access to the network and low privileges could exploit an Incorrect Authorization vulnerability found in UniFi Access Application to execute arbitrary commands on certain adopted UniFi Access devices.	9.9	More Details
CVE-2026-129048	Improper Control of Generation of Code ('Code Injection') vulnerability in Jacob N. Breetvelt WP Photo Album Plus wp-photo-album-plus allows Remote Code Inclusion.This issue affects WP Photo Album Plus: from n/a through 9.3.06.003.	9.9	More Details
CVE-2026-130759	Fleet from 4.71.0 before 4.88.2 contains an authentication bypass vulnerability in the SAML SSO callback that looks up sessions using unvalidated client-supplied identifiers. Users able to run live queries or enroll secret holders can plant attacker-controlled identity provider metadata in the session store, then present a self-signed SAML response to log in as a global admin.	9.9	More Details
CVE-2026-132234	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via argument injection in Mercurial VCS root settings	9.9	More Details
CVE-2026-132233	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via build attribute injection	9.9	More Details
CVE-2026-132232	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via argument injection in Perforce settings	9.9	More Details
CVE-2026-132231	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via best-fit character mapping in Perforce settings	9.9	More Details
CVE-2026-132802	A malicious actor with access to the network and low privileges could exploit an Improper Link Resolution Before File Access vulnerability found in UniFi Access Application to execute arbitrary code with root privileges on the host device.	9.9	More Details
CVE-2026-130756	pnpm exe before 10.34.8 embeds an outdated Node.js runtime that still trusts compromised eMudhra intermediate CA certificates during registry TLS connections. Network attackers holding a recovered intermediate key can intercept standalone pnpm registry fetches and serve forged package metadata or tarballs.	9.8	More Details
CVE-2026-130754	All versions of LaraDashboard contain an authentication bypass in SocialAuthService::findOrCreateUser(), which links a social login to an existing account by email without checking that the OAuth provider verified that email. Unauthenticated attackers using an OAuth/OIDC provider that asserts a victim's email can log in as any user, including Superadmin, and keep persistent access through a linked social_accounts entry.	9.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix buffer overflow in SMB2_tcon() In SMB2_tcon(), unc_path is allocated with a fixed buffer size of MAX_SHARENAME_LENGTH * 2 bytes (672 bytes). However, the user-controlled UNC path (tree) length is passed to cifs_strtoUTF16(), which assumes the destination buffer is large enough and writes strlen(tree) + 1 UTF-16 code units. Commit 62a1a439e0fd ("cifs: clean up handling of unc= option") removed the length check on UNC mount options, allowing mount paths to exceed MAX_SHARENAME_LENGTH while retaining the fixed-size buffer allocation in SMB2_tcon(). Consequently, an oversized UNC string causes	9.8	More

99385	cifs_strtoUTF16() to write past the end of the heap object, resulting in a slab-out-of-bounds write. Furthermore, if a tree exceeds MAX_TREE_SIZE characters, strscpy() truncates it in cifs_tcon::tree_name, causing subsequent reconnects or failovers to fail. Fix this by first verifying that the length of tree does not exceed MAX_TREE_SIZE, returning -EINVAL if it does. Then, dynamically allocate unc_path based on the length of tree so that it safely accommodates all UTF-16 code units for paths within MAX_TREE_SIZE. [pc: remove unused MAX_SHARENAME_LENGTH]		Details
CVE-2026-129892	IO::Socket::SSL::PublicSuffix versions through 2.099 for Perl allow code injection via a fetched list line that closes the built-in data heredoc when update_self_from_url rewrites the module's source. The rewrite copies the module's source down to the built-in list's heredoc and appends each fetched line below it. Nothing checks that a line looks like a domain. A line that is not a comment loses only what follows a `//` and its trailing whitespace, so a line reading END_BUILTIN_DATA is appended as it stands and ends the data block in the rewritten file. Every fetched line after it is Perl. Arbitrary Perl runs in every process that loads IO::Socket::SSL, which loads this module unconditionally, until the file is reinstalled. The attacker must control the fetched list, through the cleartext default URL or a caller-supplied one, and the updater must run with write access to the installed module.	9.8	More Details
CVE-2026-131852	OwnTracks Recorder before 1.0.5 contains a format string vulnerability in json_response() and related output functions that allows unauthenticated attackers to execute code by storing a malicious tid value. Attackers can publish locations via POST /pub or MQTT with printf specifiers in tid, then read them back through GET /api/0/last to leak addresses and write memory.	9.8	More Details
CVE-2026-128974	temporary redacted	9.8	More Details
CVE-2026-73969	AnsibleForms is a self-hosted web application that turns automation into self-service forms. Prior to 6.2.1, AnsibleForms uses a public hard-coded fallback JWT signing secret from server/config/auth.config.js when ACCESS_TOKEN_SECRET is unset. An unauthenticated attacker can sign a self-contained token with privileged claims because server/src/auth/auth_jwt.js verifies the signature without reloading the user from the database, and server/src/lib/middleware.js trusts the token's roles and options for authorization. A forged token can therefore be accepted as an administrator session, resulting in authentication bypass and full privileged API access. This issue is fixed in version 6.2.1.	9.8	More Details
CVE-2026-126755	ArangoDB before 3.12.11 contains an authentication bypass vulnerability caused by a parser differential between the raw path authentication gate and the URL-decoded dispatch logic. Attackers can send requests with a %5f-encoded leading underscore to endpoints like /_db/_system/%5fapi/simple/first-example to reach /_api/* actions as superuser without authentication. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-09-16.	9.8	More Details
CVE-2026-128138	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.10.3 until 3.10.4, the NLTK 3.10.3 JVM hardening validates Java launcher tokens in the options channel but appends the caller-controlled cmd sequence without requiring its first token to be a main class. Supplying -jar or an @argfile through nltk.internals.java or StanfordSegmenter.java_class can therefore execute attacker-selected Java bytecode without the unsafe-JAR opt-in. This issue is fixed in version 3.10.4.	9.8	More Details
CVE-2026-131914	Deserialization of untrusted data in double-submit cookie vs. form field validation in Idunno.AntiCsrf versions through 0.9.1 Beta on ASP.NET allows remote unauthenticated attackers to execute code via object injection	9.8	More Details
CVE-2026-14492	The WP Duplicate – WordPress Migration Plugin plugin for WordPress is vulnerable to Remote Code Execution in all versions up to, and including, 1.1.11 via the save_single_batch_file function. This is due to a PHP type-juggling authentication bypass via loose != comparison with no is_string() guard, combined with unsanitized base64-decoded path construction in save_single_batch_file() lacking any directory traversal containment. This makes it possible for authenticated attackers, with custom-level access and above, to execute code on the server. Exploitation requires that the site owner has completed the plugin's initial setup so that the site_type and prod_key_random_id options are non-empty, though any non-empty secret value is bypassable via the type-juggling flaw.	9.8	More Details
CVE-2026-132891	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via external pull requests	9.8	More Details
CVE-2026-116439	Dell PowerStore SDNAS contains a path traversal vulnerability that could allow an unauthenticated attacker to read, modify, or delete arbitrary files on the NAS node.	9.8	More Details
CVE-2026-116440	Dell PowerStore SDNAS contains a buffer overflow vulnerability that could allow an unauthenticated remote attacker to cause a denial of service or potential code execution on the NAS node.	9.8	More Details
CVE-2026-123890	VMware vCenter contains an authentication bypass vulnerability in the VMware Directory Service. Affected versions: 9.1.x (fixed in 9.1.1.0200); 9.0.x (fixed in 9.0.2.0200); 8.0 (fixed in 8.0 U3m); 5.x (fixed in 5.2.5). Telco Cloud Platform 3.0, 4.x, 5.0.x, 5.1.x, 5.2 and Telco Cloud Infrastructure 3.0 are also	9.8	More Details

	affected (see KB463116).		
CVE-2026-99356	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Do not credit control-record octets to the RPC stream svc_tcp_sock_recv_cmsg() receives up to two octets into a local buffer, and returns that count for any record type other than TLS_RECORD_TYPE_ALERT. Nothing reached the caller's buffer, but svc_tcp_read_marker() adds the count to sk_tcplen and svc_tcp_read_msg()'s caller adds it to sk_datalen. The RPC stream advances over octets it never received. The fragment marker is assembled from stale sk_marker octets. The message body comes from pages nothing wrote. A conforming client reaches this. RFC 8446 Section 4.6.3 lets either peer send KeyUpdate once it has sent its Finished, and svcsock has no rekey path. kTLS leaves the partially consumed record on ctx->rx_list, so the body drains two octets per svc_tcp_recvfrom() call. Each pair is credited the same way. Return -EAGAIN for a record that is not an alert. That is what svc_tcp_sock_process_cmsg()'s default arm returned before the receive moved into a local buffer.	9.8	More Details
CVE-2026-99362	In the Linux kernel, the following vulnerability has been resolved: nfsd: preflight SEQUENCE replies before accepting a slot nfsd4_sequence() narrows the reply buffer to the session's cached reply limit before it accepts the slot seqid. A client may negotiate ca_maxresponsesize_cached down to NFSD_MIN_HDR_SEQ_SZ, and nfsd4_alloc_slot() then gives every slot a zero-length sl_data[]. A COMPOUND tag can fill that narrowed buffer until it holds the SEQUENCE opcode but not the status word that follows. nfsd4_encode_operation() returns without running nfsd4_encode_sequence(), so cstate.data_offset stays zero. It leaves op->status at nfs_ok as well, so the COMPOUND is treated as having succeeded. nfsd4_store_cache_entry() declines to cache a lone SEQUENCE that returned an error. That test reads the status the operation reported, so it passes here. The copy starts at offset zero and takes the whole reply, RPC and COMPOUND headers included, into the zero-length sl_data[]. The COMPOUND tag is copied along with it, so the client picks most of the bytes written past the end of the slot: BUG: KASAN: slab-out-of-bounds in read_bytes_from_xdr_buf+0x1bc/0x390 Write of size 80 at addr ffff888003a549cd by task kunit_try_catch/24 __asan_memcpy+0x38/0x60 read_bytes_from_xdr_buf+0x1bc/0x390 nfsd4_sequence_done+0x5b0/0x810 nfs4svc_encode_compoundres+0x1bf/0x240 Check that the fixed-size SEQUENCE result, plus room for a following operation's error status, fits the negotiated limit before narrowing the buffer and consuming the slot seqid. The slot and its reply cache are left unchanged, as RFC 8881 Section 2.10.6.1.2 requires of an error returned from SEQUENCE.	9.8	More Details
CVE-2026-99366	In the Linux kernel, the following vulnerability has been resolved: nfsd: don't modify a session slot when replaying its cached reply nfsd4_sequence() claims a session slot by setting NFSD4_SLOT_INUSE under nn->client_lock. A reply served from the slot's reply cache does not claim it, and that distinction lived only in cstate->status, which nfsd4_sequence() set to nfserr_replay_cache. Commit cc028a10a48c ("NFSD: Hoist status code encoding into XDR encoder functions") moved nfsd4_proc_compound()'s cstate->status assignment below the out: label, and the replay path's goto out was the one path that relied on skipping it. The test in nfsd4_sequence_done() therefore no longer identifies a replay, and every replay now stores its reply and clears NFSD4_SLOT_INUSE as though it owned the slot. NFSD4_SLOT_INUSE is the interlock: check_slot_seqid() rejects every sequence id for a slot that is in use, before replay_matches_cache() runs. A replay never sets it, so two replays can be in flight on the same slot at once. One can read slot->sl_cred under nn->client_lock while the other's completion frees and rebuilds it from the XDR encoder without that lock. Two completions can also collide with each other and release the same group_info twice. free_svc_cred() leaves cr_uid, cr_gid and cr_flavor intact, so the reader passes every earlier test in same_creds() and dereferences a NULL cr_group_info. From a 6.12.91 production server: BUG: kernel NULL pointer dereference, address: 0000000000000004 CPU: 63 UID: 0 PID: 39015 Comm: nfsd RIP: 0010:same_creds+0x38/0xa0 [nfsd] RDX: 0000000000000000 Call Trace: nfsd4_sequence+0x6a8/0x910 [nfsd] nfsd4_proc_compound+0x345/0x670 [nfsd] nfsd_dispatch+0x100/0x220 [nfsd] svc_process_common+0x311/0x700 [sunrpc] svc_process+0x131/0x1c0 [sunrpc] svc_recv+0x7ef/0x9c0 [sunrpc] nfsd+0xa3/0x100 [nfsd] Kernel panic - not syncing: Fatal exception Since v6.14 a live session's slot table can shrink, and the unowned store becomes a use-after-free write. nfsd4_sequence() defers the shrink while a slot is in use, but it tests NFSD4_SLOT_INUSE, which a replay does not set, so free_session_slots() can kfree() the slot the replay still holds in cstate->slot. Record the claim in cstate->slot_owned when nfsd4_sequence() accepts a request and test that in nfsd4_sequence_done(), so only the request that claimed the slot updates its cached reply and clears NFSD4_SLOT_INUSE. svc_generic_init_request() zeroes the compound response before each request, so the flag starts clear.	9.8	More Details
CVE-2026-63189	Dialoqbase is an application for creating chatbots. Prior to 1.11.5, the unauthenticated GET /api/v1/view/* route in server/src/routes/api/v1/view/root.ts resolves its attacker-controlled wildcard path without normalization or containment, allowing traversal outside the intended uploads directory and arbitrary server file reads. Reading environment files can disclose DB_SECRET_KEY, the application’s JWT signing key, which permits forged administrator tokens and takeover of bots, conversations, integrations, and stored third-party credentials. This issue is fixed in version 1.11.5.	9.8	More Details
CVE-2026-84236	IBM Guardium Data Protection 12.2 is vulnerable to insecure deserialization in the SharePoint Tap MulticastTransceiver component. An unauthenticated remote attacker could exploit this vulnerability to execute arbitrary code on the affected system.	9.8	More Details

CVE-2026-126752	Dagu 1.23.0 before 2.0.0 defaults auth.mode to none, allowing unauthenticated attackers to execute OS commands through the /api/v2/dag-runs endpoint. Attackers can send a POST request with an inline spec defining steps containing arbitrary shell commands to achieve remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-03-11.	9.8	More Details
CVE-2026-116448	Dell PowerStore SDNAS contains an OS Command Injection vulnerability that could allow an attacker with internal network access to execute arbitrary commands as root on the NAS node.	9.8	More Details
CVE-2026-110811	The Trust Payments Gateway for WooCommerce (JavaScript Library) plugin for WordPress is vulnerable to Authentication Bypass in all versions up to, and including, 1.3.8. The vulnerability exists because the `new_customer_account()` function stores the raw return value of `wp_create_user()` — including a `WP_Error` object returned on username collision — into the order's `tp_userid` post meta before the `is_wp_error()` guard is evaluated; the `add_customer_details_to_thankyou_page()` callback subsequently reads that meta and casts it with `(int)`, causing PHP to coerce any object to integer `1`, after which `wp_set_current_user(1)` and `wp_set_auth_cookie(1)` are called unconditionally. This makes it possible for unauthenticated attackers to authenticate as WordPress user ID 1 — the bootstrap site administrator or, on multisite, a network super-admin — by supplying the `new_username` parameter matching any existing login during checkout and visiting the resulting thank-you page. Exploitation requires the Trust Payments gateway to be active and the attacker to complete a paid order on their own account; the ID 1 administrator username is additionally enumerable via standard WordPress user enumeration techniques.	9.8	More Details
CVE-2026-132079	Insufficient Session Expiration vulnerability in the Apache Iggy Go SDK. The TCP client method CreatePersonalAccessToken wrote the 32-bit expiry into the upper half of the 64-bit expiry field that the server reads. The server read a nonzero expiry as 2^32 times the requested value. A token created through the Go TCP client therefore stayed valid about four billion times longer than the caller asked for. This issue affects Apache Iggy Go SDK: from 0.5.0 before 0.10.0. Users are recommended to upgrade to version 0.10.0, which fixes the issue. Tokens that an earlier version created keep their long expiry on the server. Delete those tokens and create new ones.	9.8	More Details
CVE-2026-130735	Grav from 2.0.0 before 2.2.18 contains a Twig injection vulnerability that allows media-only accounts and anonymous uploaders to execute Twig by placing delimiters in image titles, alt text, class, id, or file names. Attackers can combine injected Twig with the autoescape tag to serve script to every page visitor, including super admins, or expose unpublished page content.	9.8	More Details
CVE-2026-130733	Grav from 2.0.0 before 2.2.18 contains an authorization bypass vulnerability that allows publishers to overwrite pages their per-page permissions deny by abusing FolderStorage key delimiter collisions and PageStorage folder reconstruction. Attackers can rename a folder like locked*** or save from a page with an unusual prefix width through a front-end Flex pages form, replacing protected content and permissions to access the protected subtree.	9.8	More Details
CVE-2026-87977	IBM Guardium Data Protection 12.2 is vulnerable to SQL injection in the edge-controller component. An unauthenticated remote attacker could exploit this vulnerability to access or modify sensitive database information.	9.8	More Details
CVE-2026-114516	Etherpad is a real-time collaborative editor. From 1.8.14 until 3.3.3, the pre-authentication /static/* plugin asset handler in src/node/utils/Minify.ts performs a backslash-to-forward-slash conversion after src/node/utils/sanitizePathname.ts has completed its POSIX path checks. On POSIX systems the sanitizer correctly treats backslashes as ordinary filename bytes, but the later conversion turns encoded backslash-separated parent-directory segments into traversal components without revalidation or containment. The expressPreSession registration and the always-present ep_etherpad-lite core pseudo-plugin make this path reachable without credentials, allowing a remote attacker to read any file available to the Etherpad process. Disclosed configuration or environment secrets can expose administrator credentials, after which the administrator plugin installer can provide in-process code execution and full server compromise. This issue is fixed in version 3.3.3.	9.8	More Details
CVE-2026-133325	Key exchange without entity authentication vulnerability in Apache Tomcat allowed an unauthenticated node to join a cluster when the cluster was secured with TLS 1.3 and TLS-PSK. This issue affects Apache Tomcat: 11.0.28-11.0.29, 10.1.62, 9.0.124. Users are recommended to switch to TLS 1.2 and TLS-PSK or upgrade to version 11.0.30, 10.1.63 or 9.0.125 (when released), which fix the issue.	9.8	More Details
CVE-2026-130675	Craft CMS from 5.0.0 before 5.11.4 contains a remote code execution vulnerability in ImageTransforms::normalizeTransform() that allows attackers to instantiate arbitrary classes by supplying class:null with a __class key. Attackers can reach the method through Craft Commerce's payments/pay redirect template calling customer.photo.getUrl() with a crafted transform to instantiate classes like yii\rbac\PhpManager and execute PHP.	9.8	More Details
CVE-2026-130674	Craft CMS from 5.0.0 before 5.11.4 contains an incomplete blocklist fix that allows attackers who can author entry type title formats to write arbitrary files using ZipArchive or XMLWriter classes. Attackers can abuse the non-sandboxed renderObjectTemplate() renderer to write attacker-controlled content into the web docroot, achieving remote code execution.	9.8	More Details

CVE-2026-121385	The User Registration Aide WordPress plugin through 1.5.3.8 does not verify the current password or a reset token before changing a user's password, allowing unauthenticated attackers who know a user's username and email address to change the password of any account, including administrators, and take it over.	9.8	More Details
CVE-2026-99590	In the Linux kernel, the following vulnerability has been resolved: ceph: clamp the inline data length in ceph_fill_inline_data() The MDS decides how much inline data to attach to a reply, and the client parses inline_len without any upper bound: ceph_decode_need() only verifies that the message actually carries that many bytes. ceph_fill_inline_data() then copies the data into a single folio with no length check, so a malicious or buggy MDS returning more inline data than the folio can hold writes past its end. Clamp the length to folio_size() so both callers, handle_cap_grant() and fill_inode(), are covered.	9.8	More Details
CVE-2026-110486	In Eclipse Mosquitto versions 2.1 through 2.1.2, a listener configured with `protocol websockets` mishandles empty WebSocket data frames. `net__read_ws()` in `lib/net_ws.c` returns the number of framing bytes consumed (the 4 masking-key bytes) for a data frame with a zero-length payload, without writing any bytes to the caller’s buffer. The MQTT packet reader in `packet__read_single()` (`lib/packet_mosq.c`) trusts that return value as a payload byte count. A remote client can therefore advance the read position over an uninitialised heap allocation, causing the broker to process and forward uninitialised heap memory (information disclosure, including heap and libc pointers), and can underflow the unsigned `to_process` counter so that subsequent attacker-controlled frame data is written past the end of the allocation (heap out-of-bounds write). The out-of-bounds write is reachable from the CONNECT packet before any authentication check, so an unauthenticated client can corrupt the heap and crash the broker even when `allow_anonymous false` is set. Combining the disclosure and the write allows remote code execution as the broker’s user by any client able to establish a session.	9.8	More Details
CVE-2026-99634	In the Linux kernel, the following vulnerability has been resolved: mptcp: do not use Fast Open on MP_JOIN subflows tcp_fastopen_create_child() hands the SYN packet itself to subflow_syn_recv_sock(). For an MP_JOIN request this takes the fatal fallback: the cloned child is destroyed and handed back with drop_req flagged, but tcp_fastopen_create_child() does not check the flag and queues it, so accept() can expose the freed child. MP_JOIN cannot use Fast Open: data on a subflow requires the completed HMAC exchange (RFC 8684, sec. 3.2). Refuse it (only the TFO path passes a SYN skb here) and let tcp_conn_request() fall back to the regular MP_JOIN handshake; also strip the Fast Open cookie from MP_JOIN SYN/ACKs. [Moving code to avoid clearing TS and consolidating checks.]	9.8	More Details
CVE-2026-126708	In Eclipse Mosquitto versions 2.0.0 to 2.0.22 and 2.1.0 to 2.1.2, the return value of password hash functions for sha512-pbkdf2 and legacy sha512 were not checked. Starting with OpenSSL 4.0, sha512-pbkdf2 hashing will fail when using the FIPS provider for passwords shorter than 8 characters. This has the effect that newly created passwords would match against any other password with fewer than 8 characters.	9.8	More Details
CVE-2026-110513	Insufficient Session Expiration vulnerability in the Apache Iggy server HTTP transport. A self-issued JSON Web Token (JWT) identified the user only by numeric user ID, and a new user takes the lowest free ID. After an administrator deleted a user, that user's unexpired token, or a token refreshed from it while the ID was free, authenticated as the next account created with that ID and had that account's permissions. Only deployments with HTTP enabled are affected. This issue affects Apache Iggy: through 0.9.0. Users are recommended to upgrade to version 0.10.0, which fixes the issue.	9.8	More Details
CVE-2026-99676	In the Linux kernel, the following vulnerability has been resolved: net: restore refcount ordering for lockless socket access Before commit fb041bb7c0a9 ("locking/refcount: Consolidate implementations of refcount_t"), default builds implemented refcount_inc_not_zero() using atomic_add_unless(), which provided full memory ordering on success. The commit removed that ordering on weakly ordered architectures, allowing accesses after a successful reference acquisition to be reordered before it. Use refcount_inc_not_zero_acquire() in the IPv4 and IPv6 established lookups and in inet_twsk_purge(), where the socket is validated again after taking a reference. bpf_sk_assign() can receive a socket pointer obtained through BTF without a reference. Use acquire ordering when it takes its own reference for a refcounted socket. Use refcount_set_release() when setting refcounts in the SYN cookie, TCP Fast Open, request migration, and TIME_WAIT paths. Keep the existing TIME_WAIT smp_wmb(), which separately orders the refcount store before ehash publication.	9.8	More Details
CVE-2026-128981	** UNSUPPORTED WHEN ASSIGNED ** Improper Authentication in Apache SeaTunnel Web 1.0.2 allows a user whose account has been disabled to authenticate with valid credentials and obtain a new login token. The login query does not check the account status, so disabling an account does not prevent new logins or access to protected APIs using the newly issued token. This issue requires valid credentials for the disabled account; it is not a credential-free authentication bypass. Apache SeaTunnel Web is no longer maintained. The affected 1.0.2 release does not contain the account-status check added to the main branch in pull request #264. There is no maintained fixed release. Users should discontinue use of this unsupported component. NOTE: This vulnerability only affects products that are no longer supported by the maintainer.	9.8	More Details
	The Probono Form Basic WordPress plugin before 1.3.3 does not validate the real type of files uploaded through its public form, trusting the client-declared MIME type instead, which could allow		

CVE-2026-128851	unauthenticated attackers to upload PHP files to a public directory and execute them before they are deleted, on servers that do not honour the Probono Form Basic WordPress plugin before 1.3.3's Apache-only access restriction.	9.8	More Details
CVE-2026-127021	The Flamix: Bitrix24 and WooCommerce Products Sync WordPress plugin before 1.7.0 does not restrict or validate an unauthenticated file-exchange request, allowing anyone to write a file with an attacker-chosen name, extension and path to a web-served location and thereby execute arbitrary PHP.	9.8	More Details
CVE-2026-116446	Dell PowerStore SDNAS contains an out-of-bounds memory vulnerability in the NAS NFS service that an unauthenticated remote attacker could exploit to cause a denial of service.	9.8	More Details
CVE-2026-110009	The Unzer Payments plugin for WordPress is vulnerable to Authentication Bypass leading to Account Takeover in all versions up to, and including, 2.2.1. The vulnerability exists because the Google Pay Express Checkout endpoint (`/?wc-api=google_pay_express`), handled by `startCheckoutDataWithNormalizedCustomerData`, accepts an attacker-controlled `paymentData.email` value and passes it to `createOrGetGuestCustomer()`, which resolves the submitted email to any existing WordPress user via `get_user_by('email', ...)` and then calls `wp_set_current_user()` and `wp_set_auth_cookie()` for the resolved user ID without verifying ownership of the email address, a password, a nonce, an authentication token, or any proof of payment from the payment provider. This makes it possible for unauthenticated attackers to log in as any existing WordPress user — including administrators — by submitting that user's known email address, receiving valid `wordpress_logged_in_*` session cookies for the targeted account and achieving full site takeover. The endpoint is entirely unprotected with no capability check, nonce validation, or payment-provider signature verification applied, and administrator email addresses are frequently discoverable via the WordPress REST API user enumeration endpoint (`/wp-json/wp/v2/users`), author archive pages, Gravatar lookups, or site contact footers.	9.8	More Details
CVE-2012-20007	ThinkPHP 2.0 through 2.2 and 3.0 Lite mode contains an eval injection vulnerability in the Dispatcher PATH_INFO parsing, which uses preg_replace with the /e modifier and a double-quoted replacement. Unauthenticated attackers can send crafted requests to index.php with a malicious 's' parameter containing PHP expressions to execute arbitrary code on the server. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2023-54422	iDocView before 13.10.1_20231115 contains an unrestricted file upload vulnerability in the /html/2word endpoint that allows unauthenticated attackers to write arbitrary files by supplying a malicious url parameter. Attackers can send a crafted GET request with the url parameter to download linked resources with an unsanitized path and unblocked .jsp extension, achieving remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-11-22.	9.8	More Details
CVE-2026-124990	The Masteriyo LMS – LMS Course Builder, Quizzes & Certificates plugin for WordPress is vulnerable to PHP Object Injection in versions up to, and including, 3.5.1 via deserialization of untrusted input in the course importer. The set_post_metas() method calls maybe_unserialize() on attacker-controlled postmeta string values parsed from the uploaded JSON import file. This makes it possible for unauthenticated attackers to inject arbitrary PHP Objects. Exploitation by unauthenticated attackers requires that the site has enable_instructor_registration=true AND enable_guest_checkout=true AND at least one purchasable course exists. Install.php sets enable_instructor_registration=false on fresh installs, so this is a non-default configuration. On a default installation, where that is not configured, the vulnerability requires attackers with author-level access and above.	9.8	More Details
CVE-2026-123637	When BIG-IP APM is provisioned and a per-request policy that includes the Client Information agent and Endpoint Security agents is applied to a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE). Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system and may allow the attacker to perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	9.8	More Details
CVE-2026-121805	Description: When an iRule that processes client traffic using the PEM::subscriber command is configured on a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE). Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system and may perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	9.8	More Details
CVE-2026-120722	Description: When an iRule that processes client traffic using the DNS::scrape command is configured on a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE). Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system. In addition, the attacker may be able to perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical	9.8	More Details

	Support (EoTS) are not evaluated.		
CVE-2026-120292	Description: When an HTTP profile and an iRule using the MR::message route command are configured on a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE) Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system and may allow the attacker to perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	9.8	More Details
CVE-2026-120285	Description When an HTTP profile and an iRule containing the PSC::ip_address command are configured on a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE). Impact Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system and may allow the attacker to perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	9.8	More Details
CVE-2026-116953	Description: When an iRule that processes client traffic using the GENERICMESSAGE::route command is configured on a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate and may lead to remote code execution (RCE). Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system and may allow the attacker to perform RCE. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	9.8	More Details
CVE-2026-116449	Dell PowerStore SDNAS contains an OS Command Injection vulnerability that could allow an attacker with internal network access to execute arbitrary commands as root on the NAS node.	9.8	More Details
CVE-2026-116443	Dell PowerStore SDNAS contains a stack-based buffer overflow in the NAS TFTP service that could allow an unauthenticated remote attacker to cause a denial of service or potential code execution.	9.8	More Details
CVE-2023-54419	nginxWebUI 3.0.0 before 3.5.8 contains an authentication bypass vulnerability that allows unauthenticated attackers to reach the /AdminPage/conf/runCmd endpoint by using a case-variant path that evades the case-sensitive /adminPage/ login filter. Attackers can pass OS commands through the cmd parameter, which are executed via /bin/sh with insufficient filtering of & and $() characters, achieving remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2021-48025	OpenSNS 3.2.1 contains an unauthenticated code injection vulnerability in the shareBox method of ShareController that allows attackers to invoke arbitrary model methods via the query parameter. Attackers can send crafted GET requests to /index.php?s=weibo/Share/shareBox abusing parse_str and ShareModel::getInfo to reach Schedule::runSchedule and execute code through assert. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2021-48023	Weaver E-cology 9.0 contains an unrestricted file upload vulnerability in /page/exportImport/uploadOperation.jsp that allows unauthenticated attackers to upload arbitrary JSP files. Attackers can send a crafted multipart POST request with a JSP webshell written to /page/exportImport/fileTransfer/ to achieve remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-20.	9.8	More Details
CVE-2018-25443	Baidu UEditor 1.4.0 through 1.4.3.3 contains an unrestricted file upload vulnerability in the catchimage action of /ueditor/net/controller.ashx that only validates the Content-Type header. Unauthenticated attackers can supply a source[] URL like x.jpg?.aspx to retain a malicious extension and write an executable .aspx file for remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-17.	9.8	More Details
CVE-2019-25782	ThinkPHP 5.0.0 before 5.0.24 contains a code injection vulnerability that allows unauthenticated attackers to execute commands by setting _method=__construct in POST data to overwrite the Request filter property. Attackers can send POST requests to endpoints reached via ?s=captcha or ?s=index/index/index with filter[]=system and server[REQUEST_METHOD] values to execute arbitrary system commands. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2026-99513	In the Linux kernel, the following vulnerability has been resolved: lockd: Fix use-after-free in nlmsvc_retry_blocked nlmsvc_retry_blocked() examines the block at the head of nlm_blocked under nlm_blocked_lock, then releases the lock before calling nlmsvc_grant_blocked() or retry_deferred_block(). The nlm_blocked list reference is all that keeps the block alive across that window. nlmsvc_grant_blocked() does take one of its own, but not until after the lock has been dropped. Unmounting the nfsd filesystem while a lock request is still blocked reaches nlmsvc_traverse_blocks(), which drops the list reference and frees the block along with the nlm_rqst hanging off it. BUG: KASAN:	9.8	More Details

	slab-use-after-free in nlm_async_call+0xd6/0x230 Read of size 8 at addr ffff88811b04c808 by task lockd/8377 nlm_async_call+0xd6/0x230 nlmsvc_retry_blocked+0x61c/0x800 lockd+0x144/0x1c0 Freed by task 8392: nlmsvc_release_block+0x231/0x290 nlmsvc_traverse_blocks+0x139/0x1b0 nlm_traverse_files+0x1aa/0xa00 nlmsvc_free_host_resources+0x12/0x60 nlm_shutdown_hosts_net+0x127/0x280 lockd_down+0xd5/0x1c0 Take a reference before releasing nlm_blocked_lock and drop it once the retry has run.		
CVE-2022-51027	MeterSphere versions 1.13.0 before 1.14.7 and 1.15.0 before 1.16.4 contain a pre-authentication code injection vulnerability due to the /plugin/** endpoints being exposed with anonymous access. Attackers can upload a malicious jar via POST /plugin/add and invoke it through POST /plugin/customMethod using the 'entry' and 'request' parameters, which are passed to Class.forName and method invocation for remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2026-130682	SiYuan before 3.8.10 persists a table block's caption attribute into IAL properties without HTML sanitization, causing stored cross-site scripting upon import or sync. Attackers can inject an entity-encoded caption with an onerror handler that executes in the Electron client's nodeIntegration context to achieve remote code execution.	9.6	More Details
CVE-2026-130783	Incus before 7.6.0 exposes certificate and cluster join token secrets through the default project event stream and skips authorization when cancelling operations without associated permissions. Restricted clients can subscribe to operation events to capture certificate tokens and enroll unrestricted certificates, gaining full server access, or cancel token operations by UUID.	9.6	More Details
CVE-2026-130712	NocoDB before 2026.08.1 contains an authorization bypass vulnerability that allows authenticated Creators to access another tenant's external database by reusing a victim source id when adding a source. Because NcConnectionMgrv2 caches connections by bare source id, attackers can reflect the victim's cached database schema and read, update, and insert records through their own base models.	9.6	More Details
CVE-2026-77369	Daptin is a GraphQL/JSON-API headless CMS. From 0.11.4 until 0.12.30, the generic GET /yjs/:documentName WebSocket route in InitializeYjsResources requires a session but does not authorize the database resource selected by a typename.referenceId.columnName room name. The YJS initial-content provider then reads the selected file.* column without CanRead or CanUpdate checks, and bidirectional synchronization allows a low-privileged authenticated user to read or modify restricted row content. Unauthenticated requests are rejected, and standalone non-canonical room names remain shared-by-name rooms. This issue is fixed in version 0.12.30.	9.6	More Details
CVE-2026-133189	A command injection vulnerability was found in OpenStack Cinder's SSH-based storage-driver support. Incomplete validation of shell metacharacters allows an authenticated user who owns a volume on an affected backend to supply crafted connection information and inject additional commands into the storage array's management interface. Commands execute with the privileges of Cinder's configured backend account. This can disclose storage configuration information and, where that account has administrative permissions, could allow changes that affect the integrity or availability of other tenants' storage.	9.6	More Details
CVE-2026-55506	Wavelog is web-based amateur radio logging software. Prior to 2.5, the filename parameter accepted by POST /qsl/addQsoToQsl in application/controllers/Qsl.php addQsoToQsl() is stored without rejecting directory separators or relative path components, and POST /qsl/delete reaches application/models/Qsl_model.php deleteQsl(), which passes the planted qsl_images filename to unlink(). A standard authenticated user can associate the malicious filename with the user's own QSO, satisfy the ownership check, escape the per-user QSL image directory, and delete files writable by the PHP worker. Deleting configuration or operational files can cause denial of service, while deleting installer gate files can reopen the unauthenticated installer and enable application identity takeover through attacker-controlled database configuration. This issue is fixed in version 2.5.	9.6	More Details
CVE-2026-116441	Dell PowerStore SDNAS contains an out-of-bounds write vulnerability in the storage-processor interconnect driver that could allow a malicious peer to execute code with system privileges.	9.6	More Details
CVE-2026-130779	Incus before 7.6.0 contains a path traversal vulnerability that allows users permitted to import custom volumes to store unvalidated snapshot names from the backup index configuration section. Attackers can craft a backup with traversal sequences in snapshot names so later deletion removes arbitrary host directories as root, while renaming or mounting operates on arbitrary paths.	9.6	More Details
CVE-2026-131287	In Splunk SOAR versions below 8.8.0, a user that holds a role with the Playbooks view, edit, execute, and delete permissions could cause existing cross-repository playbook references to run attacker-controlled playbook code. The vulnerability is possible because Splunk SOAR can bind existing cross-repository playbook references to a different source control repository when a deleted repository name is reused.	9.6	More Details
CVE-2026-130680	SiYuan from 3.3.0 before 3.8.10 contains a stored cross-site scripting vulnerability that allows attackers to inject HTML via unescaped checkbox column names in Gallery and Kanban AttributeView renderers. Attackers can share a crafted .sy.zip document whose rendering in the Electron desktop client with nodeIntegration enabled executes Node.js and host OS commands.	9.6	More Details

CVE-2026-77367	Daptin is a GraphQL/JSON-API headless CMS. Prior to 0.12.30, the generic GET /api/{entity} list endpoint accepts attacker-controlled sort values in PaginatedFindAllWithoutFilters in server/resource/resource_findallpaginated.go without validating the requested column against the table schema. The code uses those values for SELECT and ORDER BY identifier construction, so an embedded identifier quote can terminate the quoted identifier and make the remainder parse as SQL, allowing an unauthenticated attacker to extract unauthorized table data through blind or error-based behavior or repeatedly trigger query failures. This issue is fixed in version 0.12.30.	9.4	More Details
CVE-2026-99679	In the Linux kernel, the following vulnerability has been resolved: xfs: refuse a pNFS LAYOUTCOMMIT of a range without blocks A client can commit a hole in a LAYOUTCOMMIT, because nfsd does not check the extents against its layouts. xfs_iomap_write_unwritten() then allocates written blocks there. A large hole overruns the block reservation and shuts XFS down. A small one exposes stale data. Refuse such a range with -EIO, using XFS_BMAPI_CONVERT_ONLY.	9.4	More Details
CVE-2026-76536	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could use an empty clustering shared secret to read and modify clustering configuration through the cluster configuration Representational State Transfer (REST) API endpoint on a standalone instance. This could expose clustering configuration, persist attacker-controlled clustering settings, and disrupt clustering operations. The vulnerability is possible because Splunk Enterprise accepts the empty clustering shared secret as a shared-authentication candidate on a standalone instance when clustering is turned off.	9.4	More Details
CVE-2026-99680	In the Linux kernel, the following vulnerability has been resolved: xfs: mark the inode preallocated for pNFS write layouts past EOF xfs_fs_map_blocks() marks the inode XFS_DIFLAG_PREALLOC only when it allocates blocks for a write layout. A write layout past EOF can also hand out allocated blocks, such as speculative preallocation. blockgc then frees them under the layout, and another file may get them. Mark the inode whenever a write mapping ends past EOF.	9.4	More Details
CVE-2026-99370	In the Linux kernel, the following vulnerability has been resolved: NFSD: copy SETXATTR data from all XDR buffer segments nfsd4_vbuf_from_vector() manually copies the head followed by whole pages starting at offset zero. This assumes that all remaining data starts at the beginning of the page array and that no data is stored in the tail. However, nfsd4_decode_setxattr() passes an xdr_buf returned by xdr_stream_subsegment(). Such a sub-buffer can have a nonzero page_base and can span its page and tail segments. Ignoring that layout can copy the wrong bytes into the value passed to vfs_setxattr(). Use read_bytes_from_xdr_buf(), which observes the complete xdr_buf layout.	9.4	More Details
CVE-2026-55508	Wavelog is web-based amateur radio logging software. Prior to 2.5, the callsign parameter accepted by POST /oqrs/search_log in application/controllers/Oqrs.php search_log() reaches application/models/Oqrs_model.php search_log(), where it is concatenated directly into a raw SQL WHERE clause without parameter binding. An unauthenticated remote attacker can use UNION-based SQL injection to read arbitrary application database records, including account data, password hashes, API keys, station credentials, reset tokens, QSO data, and station metadata. Disclosure of password reset tokens can be combined with the forgot_password workflow to take over accounts, including administrator accounts. This issue is fixed in version 2.5.	9.4	More Details
CVE-2026-130156	Bazarr is a companion application to Sonarr and Radarr. Prior to 1.6.3-beta.11, the /api/subtitles/subtitles-info handler uses SubtitleNameContents.get() in subtitles_contents.py and reads the caller-supplied subtitlePath after checking only its subtitle extension; because the resolved path is not confined to files indexed by Bazarr, traversal can read files outside the media library. The /api/subtitles/subtitles handler in subtitles.py passes caller-supplied path and reference to sync and translation without per-path authorization, allowing an out-of-library reference file to be read and output to overwrite a file outside the media library. The authenticate decorator gates both endpoints when API authentication is enabled, while auth = None allows access without credentials. Both paths can affect only files that the Bazarr process can read or write. This issue is fixed in version 1.6.3-beta.11.	9.4	More Details
CVE-2026-76807	Todoist MCP Server provides tools that allow AI agents to access and modify a Todoist account. Prior to 10.3.3, the optional todoist-mcp-http server accepts browser-origin requests to POST /mcp without validating the Host and Origin headers or authenticating the browser caller. A malicious website can use DNS rebinding to reach the loopback listener and invoke MCP tools through StreamableHTTPServerTransport, which uses the server-held Todoist API key. This exposes Todoist account data and permits operations such as task creation, updates, completion, and deletion. The stdio transport is not browser-reachable and is not affected by this attack path. This issue is fixed in version 10.3.3.	9.3	More Details
CVE-2026-22309	The visitor tracking of JTL-Shop automatically registers unknown crawler user agents without filtering the value, and the stored value is rendered without output encoding on several pages of the administration backend (search engine statistics, crawler dashboard widget, crawler edit form, sitemap export download list and campaign details). An unauthenticated remote attacker can store script code with a single request. The code is executed in the browser session of an administrator who opens one of these pages, allowing actions to be performed with the administrator's privileges.	9.3	More Details
	The preview mode of the OPC (OnPage Composer) page builder of JTL-Shop is activated by a request parameter alone and renders page data supplied in the request without checking for an authenticated		

CVE-2026-22308	backend session with the corresponding permission. Several portlet properties are output without output encoding. An unauthenticated remote attacker can thereby cause script code to be executed in the browser of a visitor or administrator who opens a crafted link; in the case of an administrator, actions can be performed with the administrator's privileges. Third-party plugins whose portlets evaluate property values as templates can additionally be made to execute template code through the same entry point. This is a separate flaw in such plugins and does not affect the portlets of JTL-Shop itself.	9.3	More Details
CVE-2026-99544	In the Linux kernel, the following vulnerability has been resolved: Drivers: hv: vmbus: Bounds check the shared ring buffer indices In a CoCo VM the host is untrusted. Since the VMBus ring buffer read and write indices live in shared memory, the guest has to treat both as potentially malicious. hv_ringbuffer_write() copies into the ring at write_index with no bounds checking, so the host can make the guest write packet data at any offset up to 4 GiB past the start of the ring buffer. read_index matters too: the available space derived from both indices is the only bound on how much is copied, and an out-of-range index can make it far larger than the ring. The fix is to validate both indices before use and to use the validated snapshot instead of re-accessing. For invalid indices, hv_ringbuffer_write() logs and returns -EIO. One re-read of write_index is left unchecked in this patch: hv_get_ring_bufferindices() fetches it again for the prev_indices trailer. The value is only returned to the host and is not used as an offset or length for indexing. Since this is not a security issue, it is addressed later in the series. Open-coding the bytes available arithmetic keeps the fix free of prerequisites; a later patch refactors it into a helper function.	9.3	More Details
CVE-2026-76535	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could run arbitrary Search Processing Language (SPL) searches through RapidDiag task collection as an authenticated user. With this access, the user could read or modify data and configuration available to that authenticated user. The vulnerability is possible because Splunk Enterprise accepts an empty shared secret when it authenticates RapidDiag task-collection requests. For more information see Using RapidDiag (https://help.splunk.com/en/splunk-enterprise/administer/troubleshoot/9.3/contact-splunk-support/using-rapiddiag) in the Splunk documentation.	9.1	More Details
CVE-2026-130650	WWBN AVideo through c3edcc274c389816d434acadac07ee78eaf330c1 contains a cross-site scripting vulnerability because the default node socket server relays autoEvalCodeOnHTML without the PHP-only filter. Unauthenticated attackers can obtain an anonymous socket token and send a spoofed MSG_BATCH message to trigger client eval() in administrators' browsers for account takeover.	9.1	More Details
CVE-2026-118028	Execution After Redirect (EAR) vulnerability in Apache Tomcat in RedirectValve affecting flags [F], [G] and [R]. This issue affects Apache Tomcat: from 11.0.0-M18 through 11.0.27, from 10.1.20 through 10.1.61, from 9.0.87 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.100. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62, 9.0.124, which fixes the issue. Note: Tomcat 11.0.28 was not released	9.1	More Details
CVE-2026-116561	Improper cleanup on thrown exception vulnerability in Apache Tomcat meant that an HTTP 200 response was sent rather than an HTTP 500 response after a AuthException was thrown during Jakarta Authentication secureResponse(). This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.0.M1 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.0.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	9.1	More Details
CVE-2026-111508	Dell PowerStore contains an OS Command Injection vulnerability. Exploitation could lead to arbitrary command execution with root privileges on the appliance.	9.1	More Details
CVE-2026-130706	NocoDB before 2026.09.0 contains an sql injection vulnerability that allows authenticated Creator or Owner users to inject raw SQL by persisting a crafted formula parsed_tree reused without revalidation. Attackers can bypass the SQLite metadata path guard with an alias to plant the tree, escalating to super administrator or executing OS commands on external PostgreSQL hosts.	9.1	More Details
CVE-2026-129705	In Eclipse Mosquitto versions 2.1.0 to 2.1.2, a broker with a listener that is configured to use the PROXY protocol v2 is vulnerable to an untrusted proxy providing details that cause the broker to accept fake SSL information. This can be mitigated by following the documented requirements for PROXY listeners: It is very important that only trusted proxies are allowed to connect to any listener using the PROXY protocol, and that clients are not allowed to connect to the listener at all.	9.1	More Details
CVE-2026-82339	IBM Guardium Data Protection 12.2 is vulnerable to missing authentication in the SharePoint S-TAP RelayService, allowing an unauthenticated network attacker to inject forged audit events and access the active monitoring policy.	9.1	More Details
CVE-2026-	Zyrabit Platform provides sovereign AI infrastructure for enterprise retrieval-augmented generation. Prior to 2.4.4, import_to_vault(source_path, destination_name) accepts both caller-controlled paths, but the content-pattern guard does not confine destination_name. The vulnerable tool is implemented in zyrabit-slm/api-rag/app/domain/services/mcp_service.py, where Path(DOCS_DIR) / destination_name is passed to shutil.copy2, allowing traversal or an absolute destination to write outside the vault. An unrestricted source_path can be copied into the vault and retrieved through secure_query, disclosing	9.1	More

130154	files readable by the service; files that do not match the executable-pattern filter pass that filter. When ENABLE_LEGACY_EXTENSIONS=true, zyrabit-slm/api-rag/app/main.py registers POST /mcp/rpc without get_current_user, allowing unauthenticated network tool calls; this optional exposure is disabled by default and is not required for local invocation. If an attacker can stage chosen source content and write to an execution-triggered destination, the advisory describes conditional process execution. This issue is fixed in version 2.4.4.		Details
CVE-2026-95517	Reliance on obfuscation or encryption of Security-Relevant inputs without integrity checking, Improper validation of integrity check value, Cleartext transmission of sensitive information vulnerability in the EncryptInterceptor of Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.13 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.38 through 8.5.100, from 7.0.100 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124 and switch to TLS-PSK to encrypt cluster communication, which fixes the issue. Note: Tomcat 11.0.28 was not released	9.1	More Details
CVE-2026-111509	Dell PowerStore SDNAS contains a cryptographically weak PRNG vulnerability that could allow a network attacker to decrypt or tamper with NAS SSH/SFTP sessions.	9.1	More Details
CVE-2026-131202	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_admin" Splunk Enterprise Security role could execute arbitrary code on the search head as the user account running Splunk Enterprise by combining lookup file upload and Configuration Checker functionality. The vulnerability is possible because Splunk Enterprise Security does not restrict lookup upload file types or keep Configuration Checker module paths within the intended directory. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) and Create and manage lookups in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/managing-security-content/create-and-manage-lookups-in-splunk-enterprise-security) in the Splunk documentation.	9.1	More Details
CVE-2026-104099	Fabio is an HTTP(S) and TCP router for deploying applications managed by consul. Prior to 1.7.3, route/table.go selects targets using raw req.URL.Path, while proxy/http_proxy.go applies per-route ACL and basic-auth checks to the selected target and forwards the raw path to the origin. Dot-segment, encoded-dot, or repeated-slash variants can fall through to a less restrictive route that reaches the same backend; the backend then normalizes the path and serves the protected resource without the configured access controls. A less restrictive fallback route must reach the same backend; without one, the request returns 404. This issue is fixed in version 1.7.3.	9.1	More Details
CVE-2026-95645	Zohocorp Manageengine DDI Central build below 6500 are affected by a authenticated command injection vulnerability in the Windows DDI Agent's clear_blocked_domains functionality, allowing specially crafted blocked-domain values to trigger arbitrary command execution on Windows systems.	9.1	More Details
CVE-2026-132142	A flaw has been found in DrayTek Vigor2960 up to 1.5.1.6. This affects the function dataflowstatus of the file /cgi-bin/mainfunction.cgi of the component CGI Process. This manipulation of the argument option causes stack-based buffer overflow. The attack can be initiated remotely. The exploit has been published and may be used. Upgrading to version 1.5.1.7 is able to mitigate this issue. You should upgrade the affected component. The vendor recommends: "1.5.1.6 isn't the latest version of 2960, the latest version fixes many security issues. And we have another one in our release plan which should solve the issues reported."	9.1	More Details
CVE-2026-132622	A vulnerability was found in GearUP Hyperev up to 20261020. This impacts the function get_wlan_sitesurvey of the file ld/utils.lua of the component WLAN Site Survey. Performing a manipulation of the argument line results in os command injection. The attack is possible to be carried out remotely. The vendor was contacted early about this disclosure but did not respond in any way.	9.1	More Details
CVE-2026-95643	Zohocorp Manageengine DDI Central build below 6500 are affected by a authenticated PowerShell command injection vulnerability in the Telnet connectivity feature, allowing specially crafted IPv6 scope identifiers to execute arbitrary commands on managed Windows nodes.	9.1	More Details
CVE-2026-65094	NVIDIA VIRTIO-Net for Bluefield contains a vulnerability in Virtio-Net where a VM user may cause a Write-What-Where condition by crafted message. A successful exploit of this vulnerability may lead to Code Execution in Virtio-Net scope.	9.0	More Details
CVE-2026-133303	A vulnerability in the Audio, Web, and Video Conferencing (AWV) component of Mitel MiCollab through 10.3 (10.3.0.18) could allow an unauthenticated attacker to leverage pre-staged malicious content to conduct an XML injection attack due to insufficient parameter validation. A successful exploit could allow an attacker to execute arbitrary commands within the context of the system. This is MTLVULN-1862.	9.0	More Details
CVE-2026-130678	SiYuan before 3.8.10 contains a stored cross-site scripting vulnerability that allows attackers to execute script by storing malicious Attribute View column names rendered unescaped in col.ts. Attackers can inject HTML via addAttributeViewKey, .sy.zip import, or sync, and when a victim adds or duplicates the column, insertAdjacentHTML runs JavaScript with Node.js privileges for desktop code execution.	9.0	More Details

CVE-2026-130679	SiYuan before 3.8.10 contains a stored cross-site scripting vulnerability in genAVValueHTML() that allows attackers who can store Attribute View asset values to inject script through an unescaped image src attribute. Attackers can store an image URL containing a double quote and onerror handler that executes OS commands via Node.js in the desktop app when victims open the bound note or row attribute panel.	9.0	More Details
CVE-2026-130681	SiYuan before 3.8.10 contains a stored cross-site scripting vulnerability, an incomplete fix for CVE-2026-44670, because renderAVAttribute() inserts an Attribute View name unescaped into the database row attribute panel header. Attackers who can persist the name, such as via the setAttrViewName action, can inject an onerror handler that runs OS commands through Node.js when a desktop user opens the panel.	9.0	More Details
CVE-2026-123058	The web installer of JTL-Shop executes its state-changing installation tasks without checking whether the shop is already installed, and was not removed or locked after the installation. An unauthenticated remote attacker who can reach the installer directory can supply their own database credentials, create administrator and synchronisation accounts and overwrite the shop configuration. This can lead to a full takeover of the shop, including code execution on the server.	9.0	More Details
CVE-2026-130158	RepoManager is a self-hosted deb/rpm package repository. Prior to 5.13.3, a registered host or a party with its host/API credential can submit package metadata. Inventory reporting accepts package name and version, while available-update reporting accepts package name, version, and repository; the repository field exists only on that path. RepoManager stores these values, and stored values are rendered on the host's package detail page without preventing stored HTML or JavaScript from executing. An administrator must later view the page for script execution. The script runs in the administrator's authenticated browser session and can expose administrative data or forge privileged requests; the advisory describes possible escalation through admin-only functionality. Fixed in 5.13.3.	9.0	More Details
CVE-2026-132807	A malicious actor with access to the network and under certain conditions could exploit an Improper Certificate Validation vulnerability found in UniFi Access Application to obtain sensitive information from and make unauthorized changes to an adopted UniFi Access device.	9.0	More Details
CVE-2026-61532	Chamilo LMS is an open-source learning management system. Prior to 1.11.40 and 2.0.3, src/Chamilo/CourseBundle/Component/CourseCopy/CourseRestorer.php and src/CourseBundle/Component/CourseCopy/CourseRestorer.php use paths supplied by an imported legacy course backup without consistently confining them to the extraction directory. A course teacher can import a crafted backup containing parent-directory path segments, causing file operations to write attacker-controlled content outside the intended destination, including PHP-executable or automatically loaded locations. This can result in code execution as the PHP process and compromise course data, user data, credentials, and the host. This issue is fixed in versions 1.11.40 and 2.0.3.	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2025-64441	The ownCloud Desktop Client synchronizes files between an ownCloud server and a desktop. Prior to 6.0.2, an authenticated attacker can share a specially crafted file with a victim. If the file is synced by the client and revealed in the macOS Finder, it can cause operating-system command injection on the victim’s computer. This issue is fixed in version 6.0.2.	8.9	More Details
CVE-2026-111506	Dell PowerStore contains a Missing Authentication vulnerability that a network-adjacent attacker could exploit to bypass authentication and authorization, leading to information disclosure and denial of service.	8.8	More Details
CVE-2026-133018	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via a UI patches callback in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-133027	In JetBrains TeamCity before 2026.2.1, 2026.1.5 privilege escalation was possible via username handling in token authentication	8.8	More Details
CVE-2026-133080	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via Maven repository settings	8.8	More Details
CVE-2026-133086	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via a URL stream handler in the Kotlin DSL sandbox	8.8	More Details
CVE-			

2026-133087	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via a time zone callback in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-133092	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via Git submodule URLs	8.8	More Details
CVE-2026-133172	Cozy Cloud cozy-stack 1.6.62 before 1.6.69 contains an authorization bypass vulnerability in the shared-drive move handler that allows authenticated users to access other instances' files. Attackers owning a shared drive can name another instance's domain and any directory ID as the move source to copy and delete the victim's file tree.	8.8	More Details
CVE-2026-99322	In the Linux kernel, the following vulnerability has been resolved: ksmbd: rebind completed CREATE replays without draining the old channel A completed DurableHandleReqV2 CREATE can be replayed through another channel while its original channel is exiting. The existing replay path returns the open but leaves its connection and oplock references pointing at the old channel. Its open count also stays on that channel, and a later CLOSE through the surviving channel decrements the wrong counter. After checking the owner, lease key and persistent flag, verify that an active open is still published in the requesting session's file table. Do this even when no oplock was granted, and before processing a consumed replay. Do not read a borrowed opinfo session pointer for this check. Keep a healthy original binding. If the original channel is exiting or releasing, move the open and its own oplock references to the requesting channel. Serialize this with CLOSE using the session file-table lock, and with detach, reconnect and notification snapshots using m_lock. The earlier lock index and lease-table changes make their lifetimes independent of this connection change. No old-channel r_count wait is needed, including when an unrelated request remains active there. Account final close against fp->conn in one place, balance durable detach and reconnect, and transfer the count with a replay binding change. A final global lookup put may run after CLOSE has removed the volatile id; unlink the open from the inode list on that final close as well. A live completed CREATE replay may follow a break that lowered the durable open's batch oplock. Apply the batch eligibility check only to detached reconnects, while retaining the owner, lease key and persistent flag checks for active replays.	8.8	More Details
CVE-2026-99346	In the Linux kernel, the following vulnerability has been resolved: KVM: SVM: Trigger new ASID allocation in both VMCBs on pCPU switch When the pCPU where the VMCB was most recently used is switched, reset the ASID generation in both VMCBs, triggering new ASID allocation for the immediate VMRUN as well as the next VMRUN on the other VMCB. The ASID is shared between vmcb01 and vmcb02, and gets flushed on every nested transition. However, since pCPU tracking is done per VMCB, it is possible for one VMCB to allocate a new ASID when migrated to a new pCPU, and then the other VMCB reuses that ASID on the old pCPU. This can result in the same ASID being used by multiple vCPUs on the old pCPU. Example scenario: - vCPU runs on pCPU A, vmcb01 is active, asid=1. - vCPU migrates to pCPU B, vmcb01 pCPU changes, new asid=2. - Another vCPU runs on pCPU A and allocates asid=2 as well. - vCPU migrates back to pCPU A, and then switches to vmcb02 before it runs again with vmcb01. - No pCPU switch is detected for vmcb02, so VMRUN is done with asid=2. - Two vCPUs end up using asid=2 on the same pCPU. Keep the VMCB dirtying to the active VMCB only. Clean bits are tracked by a pCPU for each VMCB, so do not unnecessarily dirty a VMCB if its pCPU does not change. Additionally, initialize the tracked pCPU for vmcb02 to -1 on nested enablement as hardening, so that a new ASID allocation is always triggered when nested is disabled and re-enabled. No performance regression was noticed when overcommitting L1 vCPUs in L0 (to force rescheduling), pinning L1 <-> L2 vCPUs, and running CPUID in a tight loop bouncing between 2 vCPUs in L2. An alternative (and perhaps more proper) fix would be tracking the ASID per-VMCB instead (e.g. [1]). However, that's a more involved change, and it would result in having different ASIDs for L1 and L2 without actually properly maintaining them. It would probably work because all TLB flushes target the current VMCB, and the other VMCB is always flushed on nested transitions, but the code ends up in an arguably more fragile state. Punt a proper clean fix to an incoming (and overdue) overhaul of SVM's ASID usage [2]. [1] https://lore.kernel.org/lkml/20250205182402.2147495-2-yosry.ahmed@linux.dev [2] https://lore.kernel.org/kvm/20260728003557.1136583-1-yosry@kernel.org	8.8	More Details
CVE-2026-99388	In the Linux kernel, the following vulnerability has been resolved: smb: smbdirect: fix buffer length underflow in RDMA transfers When a remote buffer descriptor is larger than the remaining local buffer, smbdirect_connection_rdma_xmit() clamps its length and sets buf_len to zero. The common accounting code then subtracts the clamped length from buf_len, causing it to underflow. If another descriptor follows, the loop processes it with a wrapped remaining length and includes memory beyond the local buffer in the RDMA transfer. A client can trigger an out-of-bounds kernel write with an SMB WRITE request, or an out-of-bounds kernel read with an SMB READ request, by supplying multiple RDMA buffer descriptors with an oversized descriptor before the last one. Remove the premature assignment to buf_len. Subtracting the clamped length already consumes the remaining buffer, so the next iteration stops and only descriptors within the local buffer are used to build work requests.	8.8	More Details
CVE-2026-130741	Grav from 2.0.0 before 2.2.18 contains a Twig sandbox bypass vulnerability that allows publishers to read any server file by abusing allowlisted array_key_value() and page accessor setters such as filePath(). Attackers can repoint page objects or overwrite the assets collections to render private keys and the API signing secret, then forge a super-admin API token.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2/dlm: hold the lookup reference across		

CVE-2026-99432	convert AST enqueue dlm_convert_lock_handler() drops the reference it took when it found the lock on res->granted at its leave label, one statement before it calls dlm_queue_ast(). dlm_free_dead_locks() takes res->spinlock and drops the list and base references, so a peer death landing in that gap leaves the handler's own put as the one that frees the lock, and __dlm_queue_ast() then reads lock->lockres from freed memory. The handler is registered next to dlm_create_lock_handler(), so a convert message buffered on a dying peer's socket arrives on the same o2net receive work that o2net_disconnect_node() does not wait for. The two can interleave as follows, with the reference counts as they are before this fix: Socket data-ready Node-down thread Receive worker ---------- ---------- ---------- o2net_data_ready() select sc under callback read lock o2net_disconnect_node() queue shutdown begin flush; new color queue sc_rx_work in the new color release callback lock shutdown socket flush returns read buffered convert_lock dlm_convert_lock_handler() find lock on res->granted dlm_lock_get() refs = 3 lock res->spinlock __dlmconvert_master() call_ast = 1 unlock res->spinlock DLM node-down cleanup lock res->spinlock dlm_free_dead_locks() unlink lock put list ref refs = 2 put base ref refs = 1 unlock res->spinlock dlm_lock_put() refs = 0 free lock dlm_queue_ast() __dlm_queue_ast() read lock->lockres (use-after-free) __dlmconvert_master() only moves the lock between lists, so the count entering the gap is base, list and lookup. Drop the lookup reference after the AST is queued instead of before it and the race ends with the one __dlm_queue_ast() took, which dlm_flush_asts() releases later with ast_pending cleared and ast_list unlinked. No extra get is needed, and the normal case ends with the same count as before. dlmconvert_master() has the same shape but is only reached from dlmlock() for a lock this node owns, which dlm_free_dead_locks() never matches.	8.8	More Details
CVE-2026-99433	In the Linux kernel, the following vulnerability has been resolved: ocfs2/dlm: hold a lock reference across grant AST enqueue A newly granted remote lock must remain alive until its grant AST is queued. dlmlock_master() publishes the lock on res->granted with its base and queue references, then drops res->spinlock before calling dlm_queue_ast(). The resource's reserved AST count does not retain the lock itself. This can fail when a peer dies with an immediately grantable CREATE buffered on its socket. O2NET's node-down callback flushes receive work before DLM cleanup, but the flush only waits for work queued on entry. A data-ready callback that already selected the socket container can queue new receive work after the flush advances its work color. Socket shutdown does not discard bytes already queued for TCP receive. The late receive work and DLM node-down cleanup can interleave as follows (the lock reference operations shown are before this fix): Socket data-ready Node-down thread Receive worker ---------- ---------- ---------- o2net_data_ready() select sc under callback read lock o2net_disconnect_node() queue shutdown begin flush; new color queue sc_rx_work in the new color release callback lock shutdown socket flush returns read buffered CREATE dlmlock_master() lock res->spinlock add lock to granted unlock res->spinlock DLM node-down cleanup lock res->spinlock dlm_free_dead_locks() unlink lock put queue/base refs free lock unlock res->spinlock dlm_queue_ast() __dlm_queue_ast() read lock->lockres (use-after-free) DLM cleanup holds res->spinlock while dlm_free_dead_locks() drops the queue and base references. The CREATE handler no longer holds that spinlock when __dlm_queue_ast() reads the lock, before taking its AST list reference. Take a temporary lock reference under res->spinlock when a grant AST will be queued, and drop it after dlm_queue_ast(). This keeps the lock alive across cleanup until the AST path has taken its own reference, without reversing the existing AST/resource lock order. KASAN report as below: BUG: KASAN: slab-use-after-free in __dlm_queue_ast+0x8a4/0xc00 Read of size 8 at addr ffff88810ead5940 by task kworker/u16:3/52 CPU: 3 UID: 0 PID: 52 Comm: kworker/u16:3 Not tainted 7.3.0-rc4-next-20260921-pmb-bt-functional-v1+ #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: o2net o2net_rx_until_empty Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 ? __dlm_queue_ast+0x8a4/0xc00 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x20e/0x420 ? __dlm_queue_ast+0x8a4/0xc00 kasan_report+0xe0/0x110 ? __dlm_queue_ast+0x8a4/0xc00 __dlm_queue_ast+0x8a4/0xc00 ? srso_alias_return_thunk+0x5/0xfbef5 ? srso_alias_return_thunk+0x5/0xfbef5 ? do_raw_spin_lock+0x130/0x270 ? __pfx___dlm_queue_ast+0x10/0x10 ? __pfx_do_raw_spin_lock+0x10/0x10 ? _raw_spin_unlock_irqrestore+0x45/0x60 ? srs ---truncated---	8.8	More Details
CVE-2026-99441	In the Linux kernel, the following vulnerability has been resolved: nvmet-auth: serialize access to DH-HMAC-CHAP state to fix UAF A REPLY AUTH_SEND command can derive a dhchap_skey key while a concurrent AUTH_RECEIVE command observing this REPLY state executes the failure path and frees that key. The REPLY command then continues and writes into freed memory, causing a use-after-free. KASAN reports it as: BUG: KASAN: slab-use-after-free in __sha256_final Write of size 4 at addr ffff88800359c540 ... nvmet_auth_sq_free (drivers/nvme/target/auth.c:237) nvmet_execute_auth_receive (drivers/nvme/target/fabrics-cmd-auth.c:621) Add a mutex to each submission queue to serialize authentication commands and timeout handling. Unlock the mutex before completing requests and ignore outdated timeout callbacks. Limit timeout delays to prevent immediate expiry on 32-bit kernels.	8.8	More Details
CVE-2026-99454	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix protected RX frame validation in decrypt path The RX software decrypt path mishandles crafted protected frames from a malicious AP in two ways: 1) decryptor() never checks that a protected frame is long enough to hold the 802.11 header plus the per-cipher trailer. All of rtw_wep_decrypt(), rtw_tkip_decrypt() and rtw_aes_decrypt() compute length = hdr.len - hdrlen - iv_len; and act on the result, and recvframe_chkmic() similarly computes datalen = hdr.len - hdrlen - iv_len - icv_len - 8; for TKIP (the 8-byte Michael MIC is not accounted for in icv_len). A frame shorter than hdrlen + iv_len + icv_len (+ 8 for TKIP) underflows these unsigned subtractions, turning into a ~4 GiB iteration count in aes_decipher()/arc4_crypt()/rtw_seccalctkipmic() or a ~4 GiB crc32_le() length, causing out-of-bounds reads/writes of up to ~4 GiB starting from the rx skb. Reject such frames in decryptor() before touching the IV; this covers all three ciphers and the TKIP MIC check, and	8.8	More Details

	also keeps the unconditional iv[3] access within the frame. 2) validate_80211w_mgmt() keeps using the frame - two memcpys through the stale rx_data pointer and a pkt_len update - before checking whether decryptor() returned NULL. On decrypt failure (e.g. a unicast protected deauth/disassoc/action frame with a bad MIC, which this path always software-decrypts) the skb has already been freed, so this is a use-after-free read/write, and the caller's error path frees the frame a second time, drifting free_recvframe_cnt. Bail out immediately when decryptor() fails, freeing the temporary mgmt_DATA buffer first. Discovered by Atuin - Automated Vulnerability Discovery Engine.		
CVE-2026-99313	In the Linux kernel, the following vulnerability has been resolved: HID: logitech-hidpp: fix race condition when accessing stale stack pointer The driver uses hidpp->send_receive_buf to point to a stack-allocated buffer in the synchronous command path (__do_hidpp_send_message_sync). However, this pointer is not cleared when the function returns. If an event is processed (e.g. by a different thread) while the send_mutex is held by a new command, but before that command has updated send_receive_buf, the handler (hidpp_raw_hidpp_event) will observe that the mutex is locked and dereference the stale pointer. This results in an out-of-bounds access on a different thread's kernel stack (or a NULL pointer dereference on the very first command). Fix this by: 1. Clearing hidpp->send_receive_buf to NULL before releasing the mutex in the synchronous command path. 2. Moving the assignment of the local 'question' and 'answer' pointers inside the mutex_is_locked() block in the handler, and adding a NULL check before dereferencing.	8.8	More Details
CVE-2026-99312	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Unconditionally reset MMU context when emulating INIT Force an MMU reset when emulating an INIT even if paging was already disabled, as the MMU page role includes guest_mode (L2), i.e. an INIT while L2 is running with CR0.PG=0 will result in KVM running L1 with a stale MMU context. Alternatively, the reset could be scoped strictly to paging being enabled or L2 being active, but SMM is also encoded in the page role. Even though it _should_ be impossible for the vCPU to be INIT'd while SMM is active, there's zero value in not handling SMM. At that point, the only scenario where KVM _doesn't_ need to reset the MMU context is on an INIT with the vCPU running in Real Mode or Unpaged Protect Mode, and while INIT in those modes is somewhat common, it's far from a hot path and re-initializing the MMU is not expensive, especially when the role isn't changing (KVM elides most of the re-configuration in that case). Note, like the Fixes commit, this bug only affects nested SVM as INIT is an unconditional intercept on VMX, i.e. it's impossible to emulate INIT while L2 is active. Temporarily keep the guest TLB flush when CR0.PG is being cleared even though it is fully redundant with the flush that is emitted on all INIT events (KVM itself only does RESET when a vCPU is being created, i.e. when CR0.PG must be '0').	8.8	More Details
CVE-2026-113497	A low-privileged remote attacker can change the password of arbitrary users using the REST API, including administrative accounts. By changing the administrator password, an attacker can obtain full administrative access to the device and modify all configuration settings.	8.8	More Details
CVE-2026-125820	Deserialization of Untrusted Data vulnerability in Rymera Web Co Product Feed PRO for WooCommerce woo-product-feed-pro allows Object Injection.This issue affects Product Feed PRO for WooCommerce: from n/a through 13.5.8.	8.8	More Details
CVE-2025-2235	Improper Verification of Cryptographic Signature vulnerability in OpenText™ Connected MX allows Privilege Escalation. The vulnerability could allow a local, low-privileged user of a macOS endpoint to gain root privileges through the Connected MX Agent update process. This issue affects Connected MX macOS Agent: before 25.2.	8.8	More Details
CVE-2026-75795	The Edwiser Bridge WordPress plugin before 4.3.5 does not correctly scope one of its REST API authentication checks, disabling WordPress's REST nonce validation for every route on the site, which allows attackers to make a logged-in administrator perform arbitrary privileged actions, such as creating a new administrator account, via CSRF.	8.8	More Details
CVE-2026-5765	The Essential Real Estate plugin for WordPress is vulnerable to PHP Object Injection in all versions up to, and including, 5.3.3 via deserialization of untrusted input in the 'saved_search_check_result' function. This makes it possible for authenticated attackers, with Subscriber-level access and above, to inject a PHP Object. No 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.	8.8	More Details
CVE-2026-99540	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix parent dentry lifetime in directory enumeration ksmbd_populate_dot_dotdot_entries() builds the ".." entry from the open directory's d_parent and passes that pointer to ksmbd_vfs_fill_dentry_attrs() without taking a reference. That pointer is only valid while the directory stays under this parent. If a concurrent rename moves the directory elsewhere, the old parent loses its last reference; a following rmdir then turns it into a negative dentry with d_inode cleared, and after the RCU grace period the dentry itself is freed. The worker still holds the raw pointer and dereferences d_inode through vfs_getattr() and vfs_getxattr(). A remote SMB client, authenticated or guest, triggers this by listing a directory while that directory is moved under a new parent and the old parent is deleted. It can be done by the client itself with write access or by other concurrent writers, triggering null-ptr-deref and slab-use-after-free on a KASAN build. KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] Workqueue: ksmbd-io handle_ksmbd_work RIP: 0010:xattr_permission+0x9d/0x160	8.8	More Details

	Call Trace: vfs_getxattr+0x9c/0x1b0 ksmbd_vfs_getxattr+0x82/0x100 ksmbd_vfs_get_dos_attrib_xattr+0x7f/0x130 ksmbd_vfs_fill_dentry_attrs+0x435/0x650 ksmbd_populate_dot_dotdot_entries+0x25e/0x630 smb2_query_dir+0xcf9/0x1a20 handle_ksmbd_work+0x5f1/0x1c80 process_one_work+0x90b/0x19d0 Fix by using the dentry reference counting APIs: take the parent with dget_parent(), take a matching dget() on ".", and dput() the dentry once ksmbd_vfs_fill_dentry_attrs() has returned or the pattern does not match.		
CVE-2026-123132	The Otter Blocks WordPress plugin before 3.2.10 does not use URL-aware escaping when it resolves its dynamic data tags inside HTML attributes, which could allow unauthenticated users to post comments containing links that execute arbitrary JavaScript when clicked, for example by a logged-in administrator.	8.8	More Details
CVE-2026-99556	In the Linux kernel, the following vulnerability has been resolved: ksmbd: serialize directory emptiness checks with readdir SMB2 QUERY_DIRECTORY holds fp->readdir_lock while it calls iterate_dir() on fp->filp. An SMB2 SET_INFO request with FILE_DISPOSITION_INFORMATION can resolve the same FileId and call ksmbd_vfs_empty_dir() on the same filp without that lock. iterate_dir()'s inode read lock permits concurrent readers and does not serialize per-file iterator state. ext4 stores its dx readdir cursor and fname RB tree in file->private_data. Concurrent reset or refill can therefore free a fname node while the other worker still traverses it or retains its RB-tree link, resulting in use-after-free accesses. An authenticated SMB client with list and delete rights on a directory in a writable KSMBD share backed by ext4 can trigger this when the directory takes the dx readdir path. The client overlaps an SMB2 QUERY_DIRECTORY request with an SMB2 SET_INFO request carrying FILE_DISPOSITION_INFORMATION on one live FileId. In one x86_64 SMP KASAN multi-shot guest using ksmbd-for-next at 171df27077fb, the unmodified kernel reported a cross-worker use-after-free before a panic. The first report's access and free stacks show the two request paths directly (trimmed): BUG: KASAN: slab-use-after-free in ext4_htree_store_dirent+0x488/0x550 Read of size 8 by task kworker/2:0/29 Workqueue: ksmbd-io handle_ksmbd_work Call Trace: ext4_htree_store_dirent+0x488/0x550 htree_dirblock_to_tree+0x535/0x990 ext4_htree_fill_tree+0x319/0xba0 ext4_readdir+0x175a/0x2f40 iterate_dir+0x1b6/0x630 smb2_query_dir+0x1003/0x1cc0 Freed by task 279: free_rb_tree_fname+0x67/0xc0 ext4_readdir+0x13a9/0x2f40 iterate_dir+0x1b6/0x630 ksmbd_vfs_empty_dir+0xa2/0x100 smb2_set_info+0x2ffa/0x45d0 Take fp->readdir_lock across the entire emptiness check. This places both iterators over that handle's filp in the same lock domain while preserving the disposition status handling. The exact code change completed a warning-free x86_64 KASAN bzImage build. Tracing three request pairs showed that the two iterate_dir() calls no longer overlapped and the expected statuses were preserved. A further 1,600 reproducer iterations completed with expected statuses and no KASAN report, Oops, or panic. The reproducer and complete guest configuration are available privately to maintainers on request.	8.8	More Details
CVE-2026-133307	ClassiCube before 1.4.0 contains a heap-based out-of-bounds write vulnerability in Png_Decode in src/Bitmap.c that allows malicious servers to corrupt client memory by sending a PNG with a duplicate IHDR chunk. Attackers can supply crafted PNGs via texture pack URLs or player skins so the unreset bufferIdx underflows remaining space, letting IDAT data overflow the buffer to crash clients or execute code.	8.8	More Details
CVE-2026-130160	RepoManager is a self-hosted deb/rpm package repository. Prior to 5.13.3, an authenticated user can supply a user-controlled, array-valued GPG key identifier to the key deletion action. Validate::string() filters HTML characters but not shell metacharacters. The user-controlled identifier reaches the delete action without shell-safe encoding. controllers/Gpg.php contains the vulnerable deletion path. Process::execute() passes the command string through proc_open() and /bin/sh -c, allowing the attacker to execute arbitrary OS commands as the web/PHP process. This can expose secrets or poison packages. Authentication is required, but the advisory makes an administrator-only restriction conditional on the intended access policy and does not establish an administrator role as a prerequisite. Fixed in 5.13.3.	8.8	More Details
CVE-2026-131875	fqterm before 0.9.11 contains a heap out-of-bounds write vulnerability in FQTermDecode that allows malicious servers to corrupt memory by sending CSI sequences with unbounded parameter indexes. Attackers controlling a telnet or SSH server, or intercepting cleartext telnet, can send over 30 semicolon-separated CSI parameters to write chosen 32-bit values past the decoder object on the heap.	8.8	More Details
CVE-2026-131199	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the mc_investigation_read capability could cause the Risk Event Timeline Representational State Transfer (REST) API to run arbitrary Search Processing Language (SPL) commands with elevated permissions. With this access, the user could expose stored credentials and modify indexes and lookup data available to those permissions. The vulnerability is possible because the affected REST API does not validate request values before including them in search macro arguments. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) and Risk analysis (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/risk-analysis) in the Splunk documentation.	8.8	More Details
CVE-2026-133019	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via a context parameters callback in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-133016	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via script definitions in the Kotlin DSL settings repository	8.8	More Details

CVE-2026-131246	In Splunk SOAR versions below 8.8.0, an authenticated user with edit permission for the container label on a pending approval could read an external-prompt capability through an Approval details Representational State Transfer (REST) API response and use it to view the prompt content, change the pending approval decision, and modify its action parameters. The vulnerability is possible because an Approval details response exposes the external-prompt capability without checking approval ownership. For more information see REST External Prompt (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/8.5.0/external-prompt-endpoints/rest-external-prompt) and Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.5.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	8.8	More Details
CVE-2026-130692	NocoDB before 2026.10.1 contains a template injection vulnerability that allows authenticated Creator-role users to execute server-side JavaScript by supplying array-form property paths to Handlebars helpers in webhook templates. Attackers can bypass the string-only prototype guard to reach the function constructor, running operating-system commands with server privileges or exfiltrating process-internal values via outbound requests.	8.8	More Details
CVE-2026-130739	Grav from 2.0.0 before 2.2.18 contains a privilege escalation vulnerability that allows user managers and Flex editors to overwrite, create, or delete accounts by supplying client-controlled storage keys. Attackers can abuse the create form username, cross-type form flash replay, or __META storage key rename to gain super-admin accounts or delete account files.	8.8	More Details
CVE-2026-130743	Grav from 2.0.0 before 2.2.18 contains an insecure blueprint trust vulnerability that allows publisher accounts to read secrets and execute PHP by planting or selecting attacker-controlled blueprint files. Attackers can abuse the import@ directive, uploaded YAML media named as the page template, or form.flash_folder to leak password hashes and keys or invoke static PHP methods for code execution.	8.8	More Details
CVE-2026-130744	Grav from 2.0.0 before 2.2.18 contains a remote code execution vulnerability that allows page editors to escape the Twig sandbox by sending requests with a Host: pages header. Attackers can write configuration YAML through page folders and media .meta.yaml sidecars to register an unsandboxed Twig namespace and run shell commands as the web server user.	8.8	More Details
CVE-2026-130736	Grav from 2.0.0 before 2.2.18 contains a remote code execution vulnerability that allows publishers to bypass the media action allowlist using scoped method names in Markdown image URLs. Attackers can chain actions like self::addMetaFile into static calls such as Process::fromShellCommandline to read files and run shell commands as the web server user.	8.8	More Details
CVE-2026-82952	Latvijas Valsts radio un televīzijas centrs (LVRTC) eParakstītājs 3.0 version 1.10.0-resolute for Linux allows an unauthenticated remote attacker to execute arbitrary commands with the privileges of the current user by convincing the user to open and validate a signed PDF with a crafted filename and select Open with Adobe Reader. The application embeds the filename in a sh -c command without neutralizing shell command substitution.	8.8	More Details
CVE-2026-87789	Latvijas Valsts radio un televīzijas centrs (LVRTC) eParakstītājs 3.0 version 1.10.0 for macOS allows an unauthenticated remote attacker to execute arbitrary commands with the privileges of the current user by convincing the user to open a PDF with a crafted filename and select Send. The application embeds the attachment path in AppleScript source without neutralizing AppleScript syntax, allowing injected do shell script expressions to execute.	8.8	More Details
CVE-2026-87790	Latvijas Valsts radio un televīzijas centrs (LVRTC) eParakstītājs 3.0 version 1.10.0-resolute allows an unauthenticated remote attacker to cause HTTP GET requests from a user's workstation by convincing the user to open a crafted signed eDOC file. Automatic certificate validation retrieves a signer-controlled Authority Information Access (AIA) caIssuers URL and follows redirects without restricting the destination, enabling blind SSRF to loopback and private-network services. Additionally, unrestricted loading of local file URIs (such as file:///dev/zero) causes unbounded memory buffering resulting in Denial of Service (OutOfMemoryError). In automated server verification pipelines with UI:N, impact scales up to CVSS 8.6 High.	8.8	More Details
CVE-2026-107515	Dell PowerStore contains an Improper Privilege Management vulnerability. An authenticated attacker could exploit this vulnerability, potentially leading to privilege escalation to administrator.	8.8	More Details
CVE-2026-116447	Dell PowerStore SDNAS contains a local privilege-escalation vulnerability that could allow a low-privileged attacker to execute code as root on the NAS node.	8.8	More Details
CVE-2012-20008	openSIS 5.1 before 9.0 contains a local file inclusion vulnerability that allows authenticated users to include arbitrary files through the modname parameter of ajax.php. Attackers can prefix a payload with misc/ and use ../ directory traversal to read sensitive files like /etc/passwd or execute local scripts in the web server context. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	8.8	More Details
CVE-2026-	NocoDB before 2026.09.0 contains a server-side template injection vulnerability that allows authenticated Creator or Owner users to execute JavaScript by abusing Handlebars webhook helpers. Attackers can supply property paths like constructor.constructor to pluck, sortBy, filter, or withSort helpers to reach the Function	8.8	More Details

130709	constructor and run operating-system commands.		
CVE-2026-130704	NocoDB before 2026.10.0 contains an improper access control vulnerability that allows Creator-role users to attach NocoDB's internal database by supplying is_meta and is_local flags. Attackers can create sources or bases with these flags to expose user, token and integration tables, then read and modify records including user roles.	8.8	More Details
CVE-2026-130702	NocoDB before 2026.10.0 contains an account takeover vulnerability that allows users with the base creator role to hijack other accounts by resending base invitations. Attackers can trigger baseUserInviteResend to mint a new invite_token, read it from the base member list, and reset the victim's password through signup, taking over owner or super admin accounts.	8.8	More Details
CVE-2026-130690	NocoDB before 2026.10.1 contains an incorrect authorization vulnerability that allows organization Creators to reach routes restricted to other roles because the aclFn middleware never enforced allowedRoles. Attackers can generate a super admin password reset link via org-users endpoints to take over the instance, or read and modify SSO client secrets.	8.8	More Details
CVE-2026-133015	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via boot classpath libraries in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-130689	NocoDB before 2026.10.0 contains an sql injection vulnerability that allows authenticated users with the Creator role to inject SQL by supplying a crafted formula parsed_tree during table creation. Attackers can set a malicious binary-expression operator that is interpolated into a knex.raw fragment, reading user and token tables and escalating to super administrator on SQLite installs.	8.8	More Details
CVE-2026-125274	An authenticated iControl SOAP user with at least Guest access can inject an unauthorized management request. Impact: This vulnerability may allow an authenticated attacker with network access and at least Guest role privileges to the iControl SOAP utility to execute arbitrary system commands, create or delete files, or disable services. There is no data plane exposure; this is a control plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.8	More Details
CVE-2026-132236	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via path traversal in the pipeline YAML file path	8.8	More Details
CVE-2026-132237	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via path traversal in the external artifact cache	8.8	More Details
CVE-2026-130753	heym before 0.0.142 contains a remote code execution vulnerability in the Playwright step-mode node where _resolve_expressions_in_code substitutes workflow values directly into generated Python source. Authenticated non-admin users can craft step field values that break out of string literals to execute arbitrary Python as root inside the backend container.	8.8	More Details
CVE-2026-130774	Incus before 7.6.0 contains a privilege escalation vulnerability: rebuilding a never-started unprivileged container makes it boot without a user namespace. Root processes inside the container run as host uid 0, keeping capabilities like sys_module, with no device cgroup restriction, letting them load kernel modules and compromise the host.	8.8	More Details
CVE-2026-128668	An Unrestricted File Upload vulnerability in the admin panel of PHPGurukul News Portal Project V4.1 allows authenticated attackers to upload files with arbitrary content by bypassing weak filename extension validation in add-post.php and change-image.php. Although files are renamed with an image extension upon storage, the lack of MIME type and content validation enables content spoofing, storage of malicious markup, and creates a vector for Remote Code Execution (RCE) with server misconfigurations or secondary Local File Inclusion.	8.8	More Details
CVE-2026-63129	mssql-mcp-core is the shared library that powers MSSQL MCP servers. Prior to core and server 0.6.0 and reader and writer 0.5.0, multiple database tools interpolate caller-controlled identifiers and predicates into T-SQL batches passed to Request.query() while parameterizing only data values. update_data, delete_data, insert_data, execute_transaction, create_table, create_index, and drop_table accept unsafe tableName, column, index, schema, or free-form whereClause input, while reader-tier list_databases interpolates stateFilter without runtime allowlist enforcement. Row-count and preview queries execute before confirmUpdate or confirmDelete, allowing malicious arguments to bypass confirmation and maxRows controls and execute stacked DDL or DML, read or destroy data, and potentially run operating-system commands when the SQL login is privileged and xp_cmdshell is enabled. This issue is fixed in core and server version 0.6.0 and reader and writer version 0.5.0.	8.8	More Details
CVE-2026-	NotrinosERP is an open-source web-based enterprise resource planning system written in PHP and MySQL. From 1.0.0 until 1.1.0, the employee document handler in hrm/manage/employees.php allows an authenticated user with SA_EMPLOYEE permission to submit doc_file through tab_documents() and stores the raw client filename, including an executable extension, under company/0/documents/employees/ without	8.8	More Details

76913	extension, MIME, or content validation. Uploading a PHP file places executable content beneath the web root, where requesting the stored file results in server-side code execution. This issue is fixed in version 1.1.0.		
CVE-2026-132806	A malicious actor with access to the network and low privileges could exploit an authenticated SQL Injection vulnerability found in UniFi Access Application to obtain sensitive information and make unauthorized changes within the UniFi Access Application.	8.8	More Details
CVE-2026-132906	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via native code loading in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-132909	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via Mercurial repository contents	8.8	More Details
CVE-2026-132917	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via a file write bypass in the Kotlin DSL sandbox	8.8	More Details
CVE-2026-131236	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user who can influence indexed event data could store attacker-controlled values in event fields. If another authenticated user who reviews the event selects the View Adaptive Response Invocations or View Adaptive Response Results workflow action from Event Actions, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the action. The unauthenticated user could then access or modify sensitive data and disrupt services available to the user who selects the action. The vulnerability is possible because the affected Common Information Model workflow actions insert the orig_sid and orig_rid event-field values directly into searches without safely validating or quoting the values. For more information see About workflow actions in Splunk Web (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.2/workflow-actions/about-workflow-actions-in-splunk-web) in the Splunk documentation.	8.8	More Details
CVE-2026-130742	Grav 2.0.0 before 2.2.18 resolves Flex sort and filter field names by first calling them as public methods, so a field named delete invokes FlexObject::delete() without any permission check. Authenticated users with only list access can send GET /api/v1/flex-objects/user-accounts?sort=delete requests to delete every account, then use the reopened API first-run setup to create a super admin.	8.8	More Details
CVE-2026-131252	In Splunk SOAR versions below 8.8.0, a user that holds a role with the Assets edit or delete permission and access to one or more tenants could change or delete assets assigned to other tenants through the Representational State Transfer (REST) API. With this access, the user could redirect connector traffic and expose the credentials and operational data associated with an asset, or disrupt integrations and data ingestion for other tenants. The vulnerability is possible because Splunk SOAR verifies the Assets edit or delete permission but does not verify that the user can access every tenant assigned to the target asset before Splunk SOAR changes or deletes the asset. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises), Configure multiple tenants on your Splunk SOAR (On-premises) instance (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/configure-product-settings-for-your-splunk-soar-on-premises-instance/configure-multiple-tenants-on-your-splunk-soar-on-premises-instance), Add and configure apps and assets to provide actions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.6.0/manage-your-splunk-soar-on-premises-apps-and-assets/add-and-configure-apps-and-assets-to-provide-actions-in-splunk-soar-on-premises), and Using the REST API reference for Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/8.5.0/using-the-splunk-soar-rest-api/using-the-rest-api-reference-for-splunk-soar-on-premises) in the Splunk documentation.	8.8	More Details
CVE-2026-131851	ILIAS before 9.25, 10.13, and 11.6 contains an SQL injection vulnerability in ilPRGAssignmentFilter::toConditions() that allows authenticated users to inject SQL through the StudyProgramme members list Name and Status filters. Attackers holding only the OrgUnit permission to view members can break out of LIKE literals and execute stacked statements, such as overwriting the root password hash to gain administrative control.	8.8	More Details
CVE-2026-133476	In Dolibarr before 24.0.3, dolForgeSQLCriteriaCallback() in htdocs/core/lib/functions.lib.php allows account takeover because it strips only a single table-name prefix from a Universal Search Filter (sqlfilters) operand before checking it against the forbidden-field denylist (pass, pass_crypted, api_key). MariaDB and MySQL accept a two-level qualification such as schema.alias.column, so an operand of the form dolibarr.u.pass_crypted survives the check and is emitted into the WHERE clause. A remote authenticated user holding only the holiday read and readall permissions, which confer no user-management or administrative access, can use the resulting boolean oracle at GET /api/index.php/holidays to recover any user's bcrypt password hash and api_key one character at a time, including an administrator's. Because api_key is accepted directly by the REST API as a bearer credential, this yields full administrative control of the application. This bypasses the fix for CVE-2026-89012, which added only a strtolower() call and did not	8.8	More Details

	address prefix depth. The same denylist also omits pass_temp, which stores a password in cleartext, as well as national_registration_number, salary, birth, address, personal_mobile and note_private.		
CVE-2026-131295	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to create and run playbooks could start child playbooks on containers outside the authorized label, or repository scope for that user. This could expose container data and affect automation outside that scope. The vulnerability is possible because Splunk SOAR does not bind the identity used to authorize the child playbook run to the parent playbook execution context.	8.8	More Details
CVE-2026-131271	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to view and edit containers for one label could use a slash command through the Representational State Transfer (REST) API to access sensitive product data and credentials outside the authorized container scope. The vulnerability is possible because Splunk SOAR does not enforce authorization for related objects that slash commands access from an authorized artifact.	8.8	More Details
CVE-2026-131270	In Splunk SOAR versions below 8.8.0, an authenticated user with repository edit access could cause Splunk SOAR to write files outside the configured repository, which lets the user modify files that the user account running Splunk SOAR can write to. The vulnerability is possible because Splunk SOAR does not require source control repository paths to remain inside the configured repository.	8.8	More Details
CVE-2026-131269	In Splunk SOAR versions below 8.8.0, a user that holds a role with the Apps edit permission could cause Splunk SOAR to execute arbitrary code on the host when the user validates an installed app. The vulnerability is possible because the app validation workflow does not isolate untrusted code from the trusted environment.	8.8	More Details
CVE-2026-76538	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could cause Splunk Enterprise to run arbitrary Search Processing Language (SPL) searches with elevated privileges through data model acceleration. The vulnerability is possible because Splunk Enterprise does not correctly neutralize caller-supplied data model acceleration values before it places them into generated summarization searches.	8.8	More Details
CVE-2026-131261	In Splunk SOAR versions below 8.8.0, a user that holds a role with the Apps edit permission could execute arbitrary code as the user account that runs Splunk SOAR by installing a malicious app. The vulnerability is possible because Splunk SOAR does not restrict app dependencies to validated binary packages before installing them.	8.8	More Details
CVE-2026-131263	In Splunk SOAR versions below 8.8.0, a user that holds a role with only the apps:edit permission could install a connector app that causes Splunk SOAR to write attacker-controlled files outside the app directory. This could let the user modify files used by Splunk SOAR and execute arbitrary code as the user account that runs Splunk SOAR. The vulnerability is possible because Splunk SOAR constructs the dependency installation command from app-supplied metadata without preventing that metadata from changing its arguments. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) and Add and configure apps and assets to provide actions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.6.0/manage-your-splunk-soar-on-premises-apps-and-assets/add-and-configure-apps-and-assets-to-provide-actions-in-splunk-soar-on-premises) in the Splunk documentation.	8.8	More Details
CVE-2026-76540	In Splunk Enterprise versions below 9.4.17, 10.0.12, 10.2.9, and 10.4.5, a user who does not hold the "admin" or "power" Splunk roles can configure a metric alert through the Representational State Transfer (REST) API that runs a search with elevated privileges. The vulnerability lets the user access or modify data and disrupt the Splunk Enterprise instance. The vulnerability occurs because Splunk Enterprise does not enforce the required role-based authorization when it processes metric alert configuration through the REST API. Splunk Enterprise schedules the resulting search with elevated authorization and runs attacker-controlled search processing with that authorization. For more information see About configuring role-based user access (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/manage-splunk-platform-users-and-roles/about-configuring-role-based-user-access) in the Splunk documentation.	8.8	More Details
CVE-2026-132224	Optimizely Web Experimentation before 2026-11-17 mishandles the account signup flow, leading to presence of attacker-controlled code in the projectJS (aka Project JavaScript) field of certain cdn.optimizely.com assets.	8.7	More Details
CVE-2026-107518	Dell PowerStore contains a Server-Side Request Forgery (SSRF) vulnerability. An authenticated user with limited privileges could potentially exploit this vulnerability to make the appliance connect to attacker-chosen internal systems.	8.7	More Details
CVE-2026-	AnsibleForms is a self-hosted web application that turns automation into self-service forms. Prior to 6.2.1, AnsibleForms stores attacker-influenceable job output that server/src/lib/common.js processes through formatOutput without HTML escaping and client/src/components/AppAnsibleOutput.vue renders through v-html. An authenticated user can submit job input that is echoed into persistent output, causing script execution when another user views the job. Additional designer-controlled HTML and expression values reach	8.7	More

73967	v-html in client/src/components/AppForm.vue and client/src/components/BsInputForForm.vue without sanitization. Because Content Security Policy is disabled and the JWT is stored in localStorage, script execution can steal a victim's token and take over the account, including an administrator account. This issue is fixed in version 6.2.1.		Details
CVE-2026-130648	WWBN AVideo from 6.5 contains a stored cross-site scripting vulnerability that allows users with upload permission to inject script by submitting a javascript: scheme videoLink for embed videos. The unvalidated URL is rendered into an iframe src in videoEmbeded.php and PlayerSkins.php, executing for viewers including administrators and enabling creation of new administrator accounts.	8.7	More Details
CVE-2026-130235	In the Okta .NET SDK Abstractions library, outbound HTTP requests are not restricted to the configured Okta organization domain when processing certain URL formats. If a URL parameter is modified to specify an external host, the library may transmit the configured SSWS API token to that unintended destination.	8.7	More Details
CVE-2026-130716	NocoDB before 2026.08.0 contains a stored cross-site scripting vulnerability that allows authenticated Creator or Editor users to execute script by injecting markup into fields rendered through URL() formulas. The canvas grid renderer assigns unsanitized formula output to innerHTML, letting attackers steal session tokens from localStorage and take over accounts of users viewing the table.	8.7	More Details
CVE-2026-125061	The administration page "Offene Bestellungen" (open orders) of the JTL PayPal Checkout plugin for JTL-Shop outputs the customer name of pending orders without output encoding. A registered customer (free self-registration) can submit an HTML-entity-encoded name, which is decoded to markup when the billing address of an order is written. If the PayPal payment is left unfinished, the order is listed on the page and the name is rendered as markup in the browser of an administrator who opens the page. The script code runs with the privileges of that administrator and can, for example, create a new administrator account. Guest orders are not affected, because the name is encoded again.	8.7	More Details
CVE-2026-131258	In Splunk SOAR versions below 8.8.0, a user with edit permission for a container or workbook task could store attacker-controlled Hypertext Markup Language (HTML) in a note. When another user with view permission for the affected container and its notes opens the affected note in Mission Control, the note could run unauthorized JavaScript code in the browser of that user, expose data available through the authenticated session, and perform actions available through the authenticated session. The vulnerability is possible because Splunk SOAR does not neutralize active content in all notes before Mission Control renders them.	8.7	More Details
CVE-2026-131260	In Splunk SOAR versions below 8.8.0, a user with the containers:edit permission could import a container that includes a malicious Hypertext Markup Language (HTML) note. When another user views the imported note, Splunk SOAR could run unauthorized JavaScript in the browser of that user, exposing data and allowing actions available to that user. The vulnerability is possible because Splunk SOAR does not sanitize HTML note content during container import before it renders the content in Mission Control.	8.7	More Details
CVE-2026-130730	NocoDB before 2026.10.1 contains an authentication bypass vulnerability that allows OAuth token holders to act as full sessions because OAuth access tokens are signed with the instance JWT secret. Attackers can present a narrowly scoped OAuth token to the session strategy, short-lived-token exchange, or NocoSocket handshake to gain the granting user's full privileges.	8.7	More Details
CVE-2026-133528	BeamNG.drive before 0.40.6 contains an out-of-bounds write vulnerability because JIT-compiled LuaJIT traces omit FFI array bounds checks, letting untrusted Lua escape the sandbox. Attackers supplying a malicious mod can JIT-compile an FFI store loop, write out of range using a tostring() address leak, re-enable ffi.load, and run commands as the game user.	8.6	More Details
CVE-2026-130151	Okta Privileged Access (OPA) Gateway includes a version of the FreeRDP library affected by memory management issues in its smartcard redirection handling. When an RDP session is established meeting specific session criteria, command execution within the context of the local Gateway service account may be allowed.	8.5	More Details
CVE-2026-130236	In the Auth0 Authorization Extension, request authentication handling can process claims from query parameters independently of header validation. Under specific preconditions, this may allow an authenticated tenant user to issue requests to the extension's administrative endpoints. Also, input handling during rule generation may permit modification of the tenant login rule published by the extension.	8.5	More Details
CVE-2026-130724	NocoDB before 2026.08.0 contains a server-side request forgery vulnerability that allows authenticated base Editors to make backend HTTP requests by supplying arbitrary URLs to the integration remote fetch operation. Attackers can control the method, headers, and body to reach internal APIs, loopback services, and cloud metadata endpoints, reading responses and echoed decrypted integration secrets.	8.5	More Details
CVE-2026-130179	The Okta .NET SDK Abstractions library systematically writes per-request authentication credentials to a shared HTTP client instance, rather than being isolated to individual outbound request objects. When using scoped child clients, a subsequent scoped request can overwrite the shared client configuration. Because the original parent client credential has not been restored, subsequent outbound requests may be transmitted using the newly written credential, which can result in subsequent actions being executed under the context of an incorrect security principal in the Okta Management API and improper attribution in audit logs.	8.5	More Details

CVE-2026-130785	Incus from 7.3.0 before 7.6.0 contains an authentication bypass vulnerability in the /dev/incus guest API that allows a process in one container to impersonate another by creating a cgroup named with a forged lxc.monitor. component. Attackers can read the target container's user.* and cloud-init.* configuration and devices, subscribe to its events, mark it ready, and export images across projects.	8.5	More Details
CVE-2026-122676	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in PublishPress PublishPress Authors publishpress-authors allows Blind SQL Injection.This issue affects PublishPress Authors: from n/a through 4.18.0.	8.5	More Details
CVE-2026-77315	Flyto2 Core is an execution kernel for automation and AI-agent workflows. Prior to 2.26.8, http.paginate in src/core/modules/atomic/http/paginate.py validates only the initial base_url, and _paginate_link_header then follows a server-supplied Link: rel="next" URL without calling the SSRF guard again. A low-privileged caller can make an allowed first page point the next request at an internal or metadata service, causing a readable second-order SSRF whose response body is aggregated into all_items and returned. This issue is fixed in version 2.26.8.	8.5	More Details
CVE-2026-130770	Incus before 7.6.0 creates the NBD servers used for virtual machine migration and disk export on abstract unix sockets, which have no permission checks and whose names are visible in /proc/net/unix. A local unprivileged user can connect while an operation runs to read or modify the VM disk, gaining code execution in the guest on a migration's receiving side.	8.4	More Details
CVE-2026-132913	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the P4 Code Review build tab was possible	8.4	More Details
CVE-2026-99685	In the Linux kernel, the following vulnerability has been resolved: iommu/virtio: Reset device before deleting virtqueues on probe failure viommu_probe() marks the device DRIVER_OK before populating the event virtqueue and registering the IOMMU device in sysfs. viommu_fill_evtq() hands the device a set of device-writable buffers through virtqueue_add_inbuf(), so from that point on the device may write into them and into the rings. If either step fails, the error path deletes the virtqueues without resetting the device first. The event buffers are allocated with devm_kmalloc_array() and are released as probe unwinds, so the device can go on writing to memory that has been freed. Reset the device before deleting the virtqueues, the way viommu_remove() already does.	8.4	More Details
CVE-2026-130701	NocoDB before 2026.10.1 contains an SQL injection vulnerability because select-field option titles are embedded into MySQL or MariaDB enum/set column definitions with only single-quote doubling, leaving backslashes unescaped. A base creator or owner can end an option title with a backslash to escape the literal and append ALTER TABLE clauses, adding, dropping or renaming columns and altering indexes in the connected database.	8.3	More Details
CVE-2026-132805	A malicious actor with access to the network could exploit a Server-Side Request Forgery (SSRF) vulnerability found in UniFi Access Application to obtain information and make unauthorized changes on the host device.	8.3	More Details
CVE-2026-129990	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, every network.* procedure is only authentication-protected and does not require the Docker permission, allowing a standard member to list, inspect, create, import, recreate, resynchronize, and delete Docker networks. A member can disrupt service connectivity or alter network configuration, and create/import operations can target an unchecked serverId. This issue is fixed in version 0.31.1.	8.3	More Details
CVE-2018-25442	Baidu UEditor 1.4.0 before 1.4.3.1 contains a server-side request forgery vulnerability that allows unauthenticated attackers to make arbitrary requests via the source[] parameter in /ueditor/jsp/controller.jsp?action=catchimage. Attackers can supply crafted URLs that the ImageHunter component fetches through url.openConnection() to reach internal services and bypass network restrictions. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-07-03.	8.3	More Details
CVE-2026-99589	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: clear ROAM flag when re-adding a pending local client When a ROAMING_ADV is received for an announced local client, batadv_tt_local_mark_removed() sets BATADV_TT_CLIENT_ROAM, queues a DEL|ROAM change and marks the entry as BATADV_TT_CLIENT_PENDING. If the client sends a frame again before the next TTVN commit, batadv_tt_local_add_existing() only clears the PENDING flag and returns early. The ADD cancels the queued DEL, but the entry keeps the ROAM flag: * batadv_is_my_client() reports the client as not served by this node * batadv_tt_local_client_is_roaming() reroutes unicast frames for it to the node which sent the ROAMING_ADV The ROAM flag is only removed with the next frame of the client. Since BATADV_TT_CLIENT_ROAM is set together with BATADV_TT_CLIENT_PENDING, it has to be cleared also at the same time with BATADV_TT_CLIENT_PENDING when the client roamed back.	8.3	More Details
CVE-2026-131293	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to edit and run playbooks could read or overwrite persisted data for playbook runs outside the authorized label, or ownership scope for that user. The affected data can include playbook inputs, prompt and block results, outputs, and secrets that playbooks save in run data. Overwriting the data could corrupt results that paused, scheduled, or parent playbook runs later use. The vulnerability is possible because the playbook automation APIs do not bind data	8.3	More Details

	operations to the active playbook run.		
CVE-2026-128664	An insecure password reset vulnerability in the Admin Password Recovery Module of PHPGurukul News Portal Project v4.1 allows unauthenticated remote attackers to change an administrator or sub-administrator account password by providing the corresponding username and email address. The application does not require verification of account ownership through an email token, OTP, or another authentication mechanism before updating the password, potentially resulting in account takeover.	8.3	More Details
CVE-2026-131218	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the mc_investigation_read capability could inject arbitrary Search Processing Language (SPL) into Mission Control dashboard search requests. Mission Control could then run the injected SPL with elevated permissions. With this access, the user could cause Mission Control to read stored credentials and write data to indexes and lookups. The vulnerability is possible because Mission Control does not validate reserved search filters before adding them to dashboard searches. For more information see Overview of Mission Control in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.1/mission-control/overview-of-mission-control-in-splunk-enterprise-security) and Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.2/installation/capability-reference-for-splunk-enterprise-security) in the Splunk documentation.	8.3	More Details
CVE-2026-131215	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the ueba_finding_exclusion_edit capability could store attacker-controlled values in a User and Entity Behavior Analytics (UEBA) finding exclusion rule. Splunk Enterprise Security could later run attacker-controlled Search Processing Language (SPL) commands from the rule with elevated permissions. With this access, the user could cause Splunk Enterprise Security to read stored credentials and modify lookups and knowledge objects. The vulnerability is possible because the UEBA finding exclusion workflow does not escape user-controlled rule values before using them in generated SPL. For more information see Create a finding exclusion rule using UEBA configuration page (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/user-and-entity-behavior-analytics/finding-exclusions-in-splunk-enterprise-security/create-a-finding-exclusion-rule-using-ueba-configuration-page) and Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) in the Splunk documentation.	8.3	More Details
CVE-2026-131209	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the ueba_ers_execution_post capability could cause the User and Entity Behavior Analytics (UEBA) Content App to run arbitrary Search Processing Language (SPL) commands with elevated permissions through an Entity Risk Score calculation. With this access, the user could expose stored credentials and other data outside the permissions assigned to that role. The vulnerability is possible because the UEBA Content App does not validate Entity Risk Score calculation values before including them in SPL searches. For more information see Installing the UEBA Content App for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.5/user-and-entity-behavior-analytics/installing-the-ueba-content-app-for-splunk-enterprise-security), Entity risk scoring in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.4/risk-based-alerting/entity-risk-scoring-in-splunk-enterprise-security), and Roles and knowledge objects in UEBA for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.8/user-and-entity-behavior-analytics/ueba-content-app-for-on-premises/roles-and-knowledge-objects-in-ueba-for-splunk-enterprise-security) in the Splunk documentation.	8.3	More Details
CVE-2026-99541	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix OOB read and cross-share confusion in ksmbd_validate_name_reconnect() ksmbd_validate_name_reconnect() indexes ab_pathname[share->path_sz + 1] without checking that the path is longer than share->path_sz + 1 (d_path() fills the buffer from the end). A durable fp is not bound to its original share, so a client can preserve a file from a short-path share and reconnect it on a long-path share, making strcmp() read past the PATH_MAX buffer - with a long enough share path this oopses the kernel, and on a mapped page it also gives a byte-equality oracle on adjacent heap (the reconnect succeeds exactly when the out-of-bounds string equals the requested name). A same-share variant with the durable fp at the share root reads 1 byte past the buffer (strlen(ab_pathname) == share->path_sz). Check before comparing: the path must be longer than share->path_sz, it must start with share->path, and the next character must be '/'. The share root itself is accepted only for an empty name.	8.3	More Details
CVE-2026-131203	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_analyst" Splunk Enterprise Security role could modify a search macro associated with risk exclusions. Scheduled Risk data model acceleration searches could later run attacker-controlled Search Processing Language (SPL) commands from that macro with elevated permissions. With this access, the user who modified the search macro could expose stored credentials and modify indexes or lookup data available to those searches. The vulnerability is possible because the Splunk Enterprise Security app metadata grants the "ess_analyst" role write access to the risk exclusion search macro. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security), Finding exclusions in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.3/user-and-entity-behavior-analytics/finding-exclusions-in-splunk-enterprise-security), and Use search macros in searches (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.4/search-macros/use-search-macros-in-searches) in the Splunk documentation.	8.3	More Details

CVE-2026-103035	TeamPass is a self-hosted collaborative password manager. Prior to 3.2.1.2, the update_folder action in app/sources/folders.queries.php allowed an authenticated, non-read-only user to re-parent a folder without checking whether the target folder or submitted parentId belonged to the caller's user-accessible_folders set. Selecting a personal folder as the new parent passed the personal-folder guard and set personal_folder on the target, transferring the credential subtree into the attacker's accessible space after the next login and hiding it from the organization, including administrators. The change exposed credential labels, logins, URLs, email addresses, and descriptions, while sharekey-protected passwords remained protected; hiding the subtree also denied other users access. This issue is fixed in version 3.2.1.2.	8.3	More Details
CVE-2026-133254	In JetBrains IntelliJ IDEA before 2026.3, 2026.2.4 code execution was possible in untrusted projects via the Code Provenance plugin	8.3	More Details
CVE-2022-51034	Tangshan Pingsheng Reservoir Safety Supervision Platform builds through 2022-11-30 contain an unauthenticated SQL injection vulnerability in the simcardId parameter of /WebServices/SIMMaintainService.asmx/GetAllRechargeRecordsBySIMCardId. Attackers can send crafted POST requests with time-based MSSQL payloads to extract database information including user, version, and database name. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-07-30.	8.2	More Details
CVE-2026-127064	@fastify/reply-from is a Fastify plugin that forwards requests to an upstream HTTP server. In versions before 12.8.4, its buildURL function rejects parent-directory path segments by an exact string match, so a segment that carries a semicolon path parameter, such as a dot-dot segment followed by a semicolon and an optional parameter, is not equal to the plain dot-dot form and passes the traversal check, and the WHATWG URL parser keeps it as an ordinary segment so the base-prefix containment check also passes. Backends that strip semicolon path parameters before resolving dot-segments, such as Apache Tomcat, Jetty, WildFly, and Spring Boot, turn the segment back into a parent-directory reference and resolve it, so an unauthenticated client can escape the intended upstream prefix and reach internal paths such as admin or internal API endpoints, including with state-changing methods. Deployments that proxy to Node.js, nginx, or Express upstreams are not affected. The issue is fixed in @fastify/reply-from 12.8.4, and users should upgrade to 12.8.4 or later.	8.2	More Details
CVE-2026-130683	SiYuan before 3.8.10 uses an incomplete deny-list in IsSensitivePath that omits shell-initialization files such as .zshenv and .bash_profile, as well as ~/.local. Authenticated administrators can call /api/file/copyFile with an absolute destination outside the workspace to write payloads there, which execute as the user on the next shell or login.	8.2	More Details
CVE-2021-48021	Chamilo LMS 1.11.10 before 1.11.16 contains an unauthenticated SQL injection vulnerability in the searchOptionsFromTags function through the options parameter of extra_field.ajax.php. Attackers can send POST requests with a malicious JSON options array to inject SQL into the IN clause and extract database information. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	8.2	More Details
CVE-2019-25784	FineCMS versions 5.0.4 before 5.0.11 contain an unauthenticated SQL injection vulnerability in the data2 API action handled by list_tag in libraries/Template.php. Attackers can send GET requests to index.php?c=api&m=data2 with a valid auth hash and the param action=sql sql='...' to execute arbitrary SELECT queries and extract database information. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	8.2	More Details
CVE-2026-125060	The checkout of JTL-Shop stores the delivery address ID submitted by the customer in the session and adopts it into the order without verifying that the address belongs to the ordering customer. A remote attacker with access to the checkout, including a guest, can thereby attach the delivery address of another customer to their own order. Name and address of the other customer, and depending on the shop configuration further contact data, then appear in the order confirmation and customer account. As the IDs are sequential, addresses can be enumerated. Goods may also be shipped to the wrong address.	8.2	More Details
CVE-2026-130886	@fastify/reply-from is a Fastify plugin that forwards requests to an upstream HTTP server. In versions 12.8.1 through 12.8.3, its buildURL function checked the attacker-supplied source string for parent-directory segments, but when the source begins with the same special scheme as the base, such as http or https, the WHATWG URL parser drops the scheme and resolves the remainder as a relative path. A payload that glues the scheme to a percent-encoded parent-directory sequence therefore passes the check, because the inspected segment reads as the scheme plus the dots rather than a bare parent-directory token, while the proxy forwards the encoded traversal unchanged. An upstream that percent-decodes the request path before resolving dot-segments then serves a path outside the configured prefix, so an unauthenticated client scoped to an upstream prefix, or relying on Fastify route guards in front of the proxy, can escape that prefix and reach sibling or parent upstream endpoints. The per-segment check added in 12.8.1 is what the scheme prefix bypasses, and releases up to 12.8.0 already rejected the decoded form. The issue is fixed in @fastify/reply-from 12.8.4, and users should upgrade to 12.8.4 or later.	8.2	More Details
CVE-2026-107950	An XML External Entity (XXE) vulnerability affects a Storefront REST API in HCL Commerce versions 8, 9.0.1 and 9.1.x. A remote attacker, using a specially crafted XML, could potentially view internal files or consume memory resources.	8.2	More Details

CVE-2026-59988	django-docutils provides reStructuredText support for Django. Prior to 0.30.0, public rendering paths including publish_doctree, publish_html_from_source, publish_html_from_doctree, publish_parts_from_doctree, the {% rst %} template tag, the rst template filter, DocutilsTemplates, and DocutilsView pass lower-trust reStructuredText to Docutils without disabling file_insertion_enabled, raw_enabled, ambient docutils.conf processing, or unsafe URI schemes and without consistently sanitizing doctrees. An attacker who can influence rendered reStructuredText can use include or csv-table directives to read local files or cause server-side URL requests, or use raw HTML and unsafe URI output to inject stored script-capable content. Applications that render only trusted static reStructuredText are not affected unless attackers can influence it. This issue is fixed in version 0.30.0.	8.2	More Details
CVE-2026-133117	WebToffee Product Feed before 2.4.4 for WordPress contains a use of hard-coded cryptographic key vulnerability that allows unauthenticated attackers to bypass the URL-cron signature check by computing a valid hash using the publicly known constant salt embedded in the plugin source. Attackers can send forged GET requests with attacker-controlled parameters to the unauthenticated init-hooked endpoint to repeatedly trigger product-feed generation, causing CPU/database resource exhaustion and overwriting of the generated feed file.	8.2	More Details
CVE-2026-30502	A Cross-Site Scripting (XSS) vulnerability in OpenKM v6.3.12 allows attackers to execute arbitrary Javascript in the context of the user's browser via a crafted input injected into the Content parameter.	8.2	More Details
CVE-2022-51028	CmsEasy 7.7.5_20210919 contains an unauthenticated SQL injection vulnerability in the execsql_action function at /?case=crossall&act=execsql that executes attacker-supplied queries decrypted with the hardcoded key 'cmseasy_sql'. Attackers can send a crafted 'sql' parameter to run arbitrary SQL and dump the admin username and password hash. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-22.	8.2	More Details
CVE-2026-130008	sqlite-mcp is a Model Context Protocol server for SQLite. Prior to 0.1.2, the default read-only mode can be bypassed through execute_query, which accepts a leading SELECT while the SQLite driver executes trailing statements, and through describe_table, which interpolates the caller-controlled table_name into PRAGMA table_info. Either path can issue ATTACH DATABASE, which opens an attached file read-write despite mode=ro and can make the protected database writable, read other SQLite database files, or create a SQLite-format file at a path writable by the server process. Existing non-empty files that are not SQLite databases are refused and left unchanged. The server has no built-in authentication and listens on all interfaces at :8080 by default, so any client that can reach the port can invoke these tools unless an authenticating proxy is in place. This issue is fixed in version 0.1.2.	8.2	More Details
CVE-2024-58402	Azon Dominator releases through 2024-06-03 contain an SQL injection vulnerability that allows unauthenticated attackers to manipulate database queries through the cid parameter in /fetch_products.php. Attackers can send crafted POST requests with boolean-based or time-based blind payloads to extract sensitive database information. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-07-30.	8.2	More Details
CVE-2026-61528	Chamilo LMS is an open-source learning management system. From 2.0.0 until 2.0.3, the BigBlueButton plugin is vulnerable to XML External Entity (XXE) injection when an authenticated administrator configures a malicious BigBlueButton server URL. Processing attacker-controlled XML from that server can read arbitrary local files and perform Server-Side Request Forgery (SSRF). Sensitive configuration, credentials, source code, and internal network resources may be exposed. This issue is fixed in version 2.0.3.	8.2	More Details
CVE-2026-99586	In the Linux kernel, the following vulnerability has been resolved: iommu/arm-smmu-v3: Handle ARM erratum for CONT under invalidation with SVA The erratum (MMU-700: #3777127, S3: #3673557) deals with under invalidation of a CONT PTE grouping in the SMMU. The recommended work around is to use a range invalidation that spans the entire CONT. The only user of CONT in the kernel right now is through SVA sharing a CPU page table that contains a CONT created by the mm. Previously it was thought that this errata was dealt with because the driver always uses range invalidation. However, there is a subtle detail in the errata that the invalidation range must fully enclose the entire CONT for it to work. It seems that two sequential range invalidations, with a split point falling inside a CONT grouping, will not prevent the errata. The SMMU's range invalidation generation algorithm does not produce a single invalidation for a single SVA invalidation request, nor does the mm carefully align the SVA invalidation ranges to accommodate the splitting of the invalidation into several ranges. Thus, when processing a SVA invalidation, the range invalidation splitting routine can generate a range invalidation that is split in the middle of the CONT and risk under invalidation from this errata. This condition could be triggered by a malicious userspace manipulating the TLB gathers via mmap/mprotect/munmap. Update the errata list to the include the S3 variation, detect the IOMMUs that have it, and then have SVA invalidations use a simplified version of the over invalidation algorithm from the tlbi rework series. This ensures that a single range invalidation is issued for a single MMU notifier callback and now the range invalidation is guarenteed to cover any posible CONT. Future work to add CONT to iommu_domain page tables should either use this one-invalidate/one-range invalidation algorithm or disable CONT support in the iommu_domain.	8.2	More Details
CVE-2026-	node-soap before 1.14.1 contains a prototype pollution vulnerability in WSDL.xmlToObject that allows unauthenticated attackers to pollute Object.prototype via an XML element with id="__proto__". Attackers can send a single crafted SOAP request to a soap.listen() server, or return a malicious response to a client,	8.2	More Details

133243	causing all subsequent SOAP parsing to fail until process restart.		
CVE-2026-59961	mini-snmpd is a minimal SNMP agent implementation. Prior to 1.7, the ASN.1 BER parser in protocol.c does not verify that an attacker-controlled field length is no greater than the remaining packet size before decode_str(), decode_oid(), and decode_ptr() evaluate size - len + 1. When len exceeds size, unsigned size_t arithmetic wraps and bypasses the bounds check, allowing the decoder to read beyond the 2048-byte packet buffer. Parsing occurs before the SNMP community string is checked, so a remote unauthenticated attacker can send one crafted request to UDP port 161 to crash the process and potentially disclose adjacent heap memory. This issue is fixed in version 1.7.	8.2	More Details
CVE-2026-120718	Description: BIG-IP iControl REST tokens may not expire as expected in all scenarios. Impact: An attacker who has obtained a token may be able to extend its expiry lifetime unless the token owner deletes it. There is no data plane exposure; this is a control plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details
CVE-2026-132907	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 project viewers could redirect GitHub App connection credentials of parent projects	8.1	More Details
CVE-2026-132229	A flaw was found in hypershift-addon. Due to an overly permissive role-based access control (RBAC) configuration that grants permissions to all authenticated users rather than restricting them to the addon service, a low-privileged authenticated user can view, modify, or delete managed cluster resources. This improper authorization allows unauthorized users to alter or disrupt clusters across the entire managed fleet.	8.1	More Details
CVE-2026-132244	In JetBrains TeamCity before 2026.1.3, 2025.11.7 stored XSS on the Disk Usage and Build Time pages was possible	8.1	More Details
CVE-2026-123633	When NGINX Plus or NGINX Open Source is configured to use HTTP/2 for client connections and proxy client requests to a backend, a specially crafted request can result in use-after-free in the worker process, if this process is running low on worker_connections. Additionally, attackers can execute code on systems with Address Space Layout Randomization (ASLR) disabled or when the attacker can bypass ASLR. Impact: This vulnerability may allow remote attackers to cause a denial-of-service (DoS) on the NGINX system or may allow a code execution. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details
CVE-2026-130791	Finsys Dockhand before 1.0.55 contains a path traversal vulnerability that allows users with stacks:edit permission to move arbitrary server directories by supplying a crafted oldDir parameter. Attackers can send POST requests to /api/stacks/[name]/relocate with an unvalidated source path to relocate directories into stack paths, disrupting services and exposing their contents.	8.1	More Details
CVE-2026-130700	NocoDB before 2026.10.1 contains a consent phishing vulnerability that allows unauthenticated attackers to register OAuth clients with arbitrary HTTPS redirect URIs through the unauthenticated dynamic client registration endpoint. Attackers can impersonate trusted applications on a consent screen lacking unverified-client warnings or redirect host display, obtaining tokens to read and write victim tables and records.	8.1	More Details
CVE-2026-123634	This issue occurs when NGINX Plus or NGINX Open Source is configured as a proxy, the proxied server name includes variables, and a Server Side Includes (SSI) or mirror directive sends a subrequest to a location that resolves the same server name. Under these conditions, a specially crafted HTTP request may cause a use-after-free condition in a worker process. An attacker may also be able to execute code if Address Space Layout Randomization (ASLR) is disabled or the attacker can bypass ASLR. Impact: An attacker could leverage this vulnerability to crash the NGINX worker process, corrupt its memory, or perform remote code execution (RCE) Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details
CVE-2026-76007	IBM Engineering AI Hub 1.0.0 through 1.4.0 could allow a remote attacker to obtain sensitive information and perform unauthorized actions due to improper validation of a redirect URL.	8.1	More Details
CVE-2026-113425	When NGINX Plus or NGINX Open Source is configured as a mail proxy, undisclosed traffic can cause a use-after-free in an nginx worker process. Additionally, attackers can execute code on systems with Address Space Layout Randomization (ASLR) disabled or when the attacker can bypass ASLR. Impact: A remote unauthenticated attacker may crash an nginx worker process or corrupt its memory, disrupting service. Remote code execution is a potential impact. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details
CVE-2026-121810	Description: This issue affects NGINX Plus and NGINX Open Source configurations that use HTTP/3. A remote, unauthenticated attacker can cause a use-after-free condition in the NGINX worker process, leading to a restart. Exploitation also depends on conditions beyond the attacker’s control. Additionally, attackers can execute code on systems with Address Space Layout Randomization (ASLR) disabled or when the attacker can bypass ASLR. Impact: This vulnerability may allow remote attackers to cause a denial-of-service (DoS) on the NGINX system or to possibly trigger a code execution. There is no control plane exposure; this is a data	8.1	More Details

	plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.		
CVE-2026-127178	A stored cross site scripting attack when viewing email aliases in postfixadmin 4.0 before 4.0.7, and postfixadmin 3.3 before 3.3.17 could be caused by signed in admins able to inject aliases.	8.1	More Details
CVE-2026-132904	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 missing authorization allowed project viewers to remove VCS roots	8.1	More Details
CVE-2026-132903	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 missing authorization allowed any authenticated user to start and stop cloud agent instances	8.1	More Details
CVE-2026-132892	In JetBrains TeamCity before 2026.2 privilege escalation was possible via the authorization flow	8.1	More Details
CVE-2026-129446	The WPO365 | SEAMLESS WORDPRESS + MICROSOFT INTEGRATION (WPO365 | LOGIN) WordPress plugin before 45.4 does not ensure the signature validated on a SAML response covers the element the user identity is read from, allowing unauthenticated attackers to forge a response via signature wrapping and log in as any user, including administrators, when they can supply a single artifact signed by the configured identity provider.	8.1	More Details
CVE-2026-108687	The SAML SSO Login WordPress plugin through 1.5.2 does not verify that the SAML assertion whose signature it validates is the same assertion it reads the user identity from, allowing unauthenticated attackers who hold a validly signed assertion for any identity to authenticate as any existing user, including an administrator, through XML signature wrapping.	8.1	More Details
CVE-2026-116442	Dell PowerStore SDNAS contains an authentication-bypass-by-replay vulnerability in the NAS NDMP service that could allow a network attacker to gain backup and restore control of the appliance.	8.1	More Details
CVE-2026-130751	OpenWrt LuCI luci-app-travelmate on master before 3.0.0-1, on openwrt-25.12 before 2.4.9-6, and on openwrt-24.10 and openwrt-23.05 (no fix available) grants file exec permissions for state-mutating commands in the read scope of its rpcd ACL, letting authenticated read-only users run those commands as root. Attackers can invoke ubus file exec on /etc/init.d/travelmate setup, start, stop, restart, or /sbin/ifup to modify and commit network and firewall configuration or disrupt the travelmate daemon.	8.1	More Details
CVE-2026-61798	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. Prior to 0.0.23.Final, BoringSSLAsymmetricCipherKeyPair.toString in codec-ohttp-hpke-classes-boringssl/src/main/java/io/netty/incubator/codec/hpke/boringssl/BoringSSLAsymmetricCipherKeyPair.java includes its private parameter, BoringSSLAsymmetricKeyParameter.toString in codec-ohttp-hpke-classes-boringssl/src/main/java/io/netty/incubator/codec/hpke/boringssl/BoringSSLAsymmetricKeyParameter.java renders the complete private-key byte array when isPrivate is true, and BoringSSL.EVP_HPKE_KEY_init_or_throw in codec-ohttp-hpke-classes-boringssl/src/main/java/io/netty/incubator/codec/hpke/boringssl/BoringSSL.java includes rejected privateKeyBytes in an exception message. Applications that log affected key-pair objects or initialization exceptions can persist complete HPKE private key material in logs or telemetry. Anyone who can read those records can recover the key and may compromise OHTTP message confidentiality or authenticity expectations while the key remains usable. This issue is fixed in version 0.0.23.Final.	8.1	More Details
CVE-2026-82342	IBM Guardium Data Protection 12.2 is vulnerable to a stack-based buffer overflow when processing a specially crafted PROXY protocol v1 header. A remote attacker can exploit this vulnerability to cause a denial of service or potentially execute arbitrary code.	8.1	More Details
CVE-2026-76544	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could cause scheduled metric alert searches to run arbitrary Search Processing Language (SPL) commands and access credentials stored in cleartext. This requires at least one non-internal metric index on the instance. The vulnerability is possible because Splunk Enterprise does not fully validate metric alert filter input before using it in a scheduled search.	8.1	More Details
CVE-2026-132121	Supergateway 2.4.0 through 4.2.2 contains a cross-site WebSocket hijacking vulnerability in stdio-to-WebSocket mode that allows malicious web pages to connect by bypassing missing Origin header validation. Attackers can open the /message WebSocket from a victim's browser, even against localhost-bound gateways, to initialize an MCP session and invoke wrapped stdio MCP server tools with victim privileges.	8.1	More Details
CVE-2026-131305	In Splunk AI Assistant versions below 2.0.0, a user who does not hold the "admin" or "power" Splunk roles could submit client-controlled feature settings through the Representational State Transfer (REST) API. This could modify limited response metadata outside the expected server-side configuration. The vulnerability is possible because the REST API processes client-supplied feature settings without enforcing the expected	8.1	More Details

	server-side configuration boundary.		
CVE-2026-99460	In the Linux kernel, the following vulnerability has been resolved: mptcp: hold MP_JOIN msk ref when cloning reqsk TCP request migration clones pending request sockets with inet_reqsk_clone(). For MPTCP MP_JOIN requests this byte-copies subflow_req->msk, but the clone does not own a reference. The original and cloned requests can consequently drop the same msk reference, leaving one request with a dangling pointer. This manifests as a KASAN slab-use-after-free in subflow_req_destructor(): BUG: KASAN: slab-use-after-free in subflow_req_destructor (net/mptcp/subflow.c:45) Write of size 4 at addr ff1100010e008d40 by task swapper/0/0 CPU: 0 UID: 0 PID: 0 Comm: swapper/0 Not tainted 7.2.0-rc5-00353-gc27e36054537 #10 PREEMPT(full) Call Trace: subflow_req_destructor (net/mptcp/subflow.c:45) subflow_v4_req_destructor (net/mptcp/subflow.c:694) __reqsk_free (net/ipv4/inet_connection_sock.c:906) reqsk_timer_handler (include/net/request_sock.h:137 net/ipv4/inet_connection_sock.c:1147) Allocated by task 10014: sk_prot_alloc (net/core/sock.c:2246) sk_clone (net/core/sock.c:2488) mptcp_sk_clone_init (include/net/sock.h:1848 net/mptcp/protocol.c:3564) subflow_syn_recv_sock (net/mptcp/subflow.c:883) tcp_check_req (net/ipv4/tcp_minisocks.c:934) Freed by task 0: slab_free_after_rcu_debug (include/linux/kasan.h:235 mm/slub.c:2677 mm/slub.c:6439) rcu_core (kernel/rcu/tree.c:2645 kernel/rcu/tree.c:2897) Last potentially related work creation: kmem_cache_free (mm/slub.c:2638 mm/slub.c:6377 mm/slub.c:6504) __sk_destruct (net/core/sock.c:2289 net/core/sock.c:2391) sk_destruct (net/core/sock.c:2419) __sk_free (net/core/sock.c:2430) sk_free (net/core/sock.c:2441) mptcp_close (include/net/sock.h:2020 net/mptcp/protocol.c:3399) inet_release (net/ipv4/af_inet.c:442) The buggy address belongs to the cache MPTCP of size 2968 The buggy address is 128 bytes inside a freed 2968-byte region A non-NULL subflow_req->msk means that the request owns one reference. The third ACK can concurrently transfer the original request reference to the child and release it, so taking an unconditional hold on the copied pointer is unsafe. MPTCP sockets use SLAB_TYPESAFE_BY_RCU and all current inet_reqsk_clone() callers run in an RCU read-side critical section. Document and lockdep-check that calling-context requirement. Read the pointer from the original request and acquire a reference only if it is still live. Use refcount_inc_not_zero_acquire() so the subsequent ownership validation cannot be reordered before the try-get, then re-read the original request to verify that it still owns the same msk. If either check fails, clear the clone pointer; otherwise its normal destructor balances the new reference. Mark the ownership-transfer store with WRITE_ONCE() to match the lockless reads. The following patch completes the clone fixup for MP_CAPABLE token ownership. Both patches carry the same Fixes tag and are required for stable backports.	8.1	More Details
CVE-2026-116764	The The GDPR Framework By Data443 plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.5.2. 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 delete all postmeta rows for every Easy Digital Downloads payment order on the site, destroying order metadata belonging to all customers. Exploitation requires that Easy Digital Downloads 2.x is active and the GDPR Framework's EDD compatibility option is enabled with a delete or delete_and_notify action configured.	8.1	More Details
CVE-2026-85018	IBM Guardium Data Protection 12.2 is vulnerable to improper TLS certificate validation in the GDP Terraform provider, allowing a network attacker with a man-in-the-middle position to intercept sensitive credentials transmitted to the Guardium appliance.	8.1	More Details
CVE-2026-107519	Dell PowerStore contains an integer underflow vulnerability. An attacker could exploit this vulnerability, potentially leading to denial of service of a cryptographic service.	8.1	More Details
CVE-2026-131297	In Splunk SOAR versions below 8.8.0, an operator of a configured identity provider could authenticate as an existing user associated with another configured identity provider. If that account has the "Administrator" Splunk SOAR role, the operator could control the Splunk SOAR deployment. Splunk SOAR does not bind external user accounts to the identity provider that provisions them and authenticates accounts by username without verifying provider ownership. The vulnerability requires more than one configured external identity provider with distinct trust domains and a target username that an operator of a configured identity provider can authenticate through a different provider. For more information see Configure single sign-on authentication (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/manage-your-splunk-soar-on-premises-users-and-accounts/configure-single-sign-on-authentication-for-splunk-soar-on-premises) and Manage roles and permissions (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	8.1	More Details
CVE-2026-123419	The immonex Kickstart plugin for WordPress is vulnerable to Local File Inclusion in all versions up to, and including, 1.19.0 via the 'inx-r-cidata (JSON blob containing 'template' and 'template_folders')' parameter parameter. 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. The filename-level sanitization that strips dot-slash sequences applies only to the template filename component, not to the attacker-supplied template_folders directory paths, rendering it ineffective against directory traversal via that parameter.	8.1	More Details

CVE-2026-130721	NocoDB from 2026.06.2 before 2026.08.0 contains an authorization bypass vulnerability that allows authenticated base Editors to write to data-read-only sources by abusing the internal batch endpoint. Attackers can name a writable decoy table on the batch envelope and target a read-only source table in dataUpdate, dataInsert or dataDelete sub-operations to modify or destroy records.	8.1	More Details
CVE-2025-15716	An incorrect implementation of authentication algorithm vulnerability in the SSO API in Synology DiskStation Manager (DSM) before 7.2.2-72806-5 and 7.3.1-86003-1 (7.2.1-69057 is not affected) allows remote attackers to read or write arbitrary files and conduct denial-of-service attacks with prior knowledge of the distinguished name (DN).	8.1	More Details
CVE-2026-130747	Grav before 1.7.53.7 and 2.x before 2.2.18 contains an authentication bypass vulnerability that allows anonymous visitors to create super-admin accounts because UserCollection::count() uses an unescaped glob() pattern. On installs whose path contains brackets or a backslash, the account count returns zero, so attackers can send requests to POST /api/v1/auth/setup or the classic admin registration form to create or overwrite administrator accounts.	8.1	More Details
CVE-2026-111507	Dell PowerStore contains an authorization bypass vulnerability. A remote attacker could exploit this vulnerability to act under another user's session, potentially leading to denial of service.	8.1	More Details
CVE-2026-131282	In Splunk SOAR versions below 8.8.0, a user who holds the "Observer" Splunk SOAR role could access a password-reset credential in a diagnostic log bundle and reset the password for an active local account. The vulnerability is possible because Splunk SOAR does not redact password-reset email bodies before it logs communications between product processes.	8.1	More Details
CVE-2026-133310	Wing FTP Server before 8.2.8 contains a file mask bypass vulnerability allowing attackers to upload restricted files via an embedded NUL byte in Upload Link filenames. Attackers can plant a Lua session file with a chosen UID and send a matching UID cookie to execute operating system commands as the service account.	8.1	More Details
CVE-2026-105929	A path traversal vulnerability (CWE-22) was found in ansible-core's ansible-galaxy role installation. An explicit role "name" taken from a role's meta/requirements.yml is used to construct the on-disk installation path without validating that the resulting destination stays within the selected roles directory. A name containing "../" escapes that directory, and an absolute name replaces it entirely. Because ansible-galaxy processes transitive dependencies from an installed role's meta/requirements.yml, a malicious role publisher can cause attacker-supplied files to be written outside the operator's chosen roles directory when the operator installs the role, and, when --force is used, can cause an existing role directory outside that path to be removed and replaced. The impact is bounded by the invoking user's filesystem permissions and requires the operator to install the attacker-influenced content.	8.1	More Details
CVE-2026-99380	In the Linux kernel, the following vulnerability has been resolved: smb: client: validate NT transaction response layout and offsets validate_ntransact() assumes that the response uses the fixed 18-word NT transaction layout. A different WordCount makes get_bcc() locate the byte count at a different offset from pSMBr->ByteCount, allowing the subsequent bounds checks to use inconsistent response boundaries. The function also turns the 32-bit ParameterOffset and DataOffset fields into pointers before checking their bounds. On a 32-bit client, a large offset can wrap the pointer below end_of_smb and pass the one-sided pointer comparison. The caller then dereferences the parameter pointer or copies from the data pointer. Require the expected WordCount before reading the byte count. Check each offset and count against the received SMB length as integers before forming either pointer. Use subtraction after validating the offset so that the range checks cannot overflow.	8.1	More Details
CVE-2026-99355	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: fix oversized GSS proxy token copy gss_read_proxy_verf() allocates one page per token chunk, but it copies the entire linear head of an RPC request into in_token->pages[0]. A UDP datagram can leave more than PAGE_SIZE bytes in the linear head, so a large RPCSEC_GSS INIT token overwrites following pages. Copy the linear head in page-sized chunks before switching to the request page array. The linear head itself can span multiple pages.	8.1	More Details
CVE-2016-20105	Lucee Server versions 5.0.0.0 before 5.3.4.46 contain an insecure default configuration that allows unauthenticated attackers to set the initial admin password via the first-run form on /lucee/admin/web.cfm and /lucee/admin/server.cfm. Attackers can access the 'New Password' set-password form before any password is configured to seize administrative control of the server. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-24.	8.1	More Details
CVE-2026-84273	IBM Guardium Data Protection 12.2 is vulnerable to path traversal in the datasource functionality. An authenticated attacker could exploit this vulnerability to delete arbitrary files on the appliance.	8.1	More Details
CVE-2026-133089	In JetBrains TeamCity before 2026.2.2, 2026.1.6 the TFS helper process exposed an unauthenticated command server	8.1	More Details
	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the ueba_finding_exclusion_edit capability could create a finding exclusion rule that references an attacker-		

CVE-2026-131210	controlled lookup definition if the role also grants write access to globally shared lookup definitions, as the "ess_analyst" Splunk Enterprise Security role does. The User and Entity Behavior Analytics (UEBA) Content App could later include the lookup name in a search and run arbitrary Search Processing Language (SPL) commands with elevated permissions. With this access, the user who created the finding exclusion rule could expose stored credentials and other data outside the permissions assigned to the "ess_analyst" role. The vulnerability is possible because UEBA finding exclusion processing does not safely quote lookup names before including them in SPL searches. For more information see Finding exclusions in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.3/user-and-entity-behavior-analytics/finding-exclusions-in-splunk-enterprise-security) and Roles and knowledge objects in UEBA for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.8/user-and-entity-behavior-analytics/ueba-content-app-for-on-premises/roles-and-knowledge-objects-in-ueba-for-splunk-enterprise-security) in the Splunk documentation.	8.1	More Details
CVE-2026-131200	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_analyst" Splunk Enterprise Security role could create an attacker-controlled saved search. Threat intelligence lookup generation could later run the saved search with elevated permissions, which could let the user who created the saved search run arbitrary Search Processing Language (SPL) commands with those permissions. The vulnerability is possible because Splunk Enterprise Security does not restrict threat intelligence lookup generation to the intended saved searches. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) and Turn on threat matching searches in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.0/threat-intelligence/turn-on-threat-matching-searches-in-splunk-enterprise-security) in the Splunk documentation.	8.1	More Details
CVE-2026-130644	A path traversal flaw was found in ansible-core. When installing a collection with ansible-galaxy collection install, the install destination directory is derived from the namespace and name fields declared inside the collection artifact's own metadata, which are not validated as well-formed collection identifiers or checked for path containment. A malicious collection artifact -- delivered as a local file, a remote URL, a source directory, a git repository, or as a dependency -- can use a crafted name value to redirect the install destination outside the configured collections path, causing ansible-galaxy to delete an existing directory writable by the invoking user and then write attacker-controlled files to that location.	8.1	More Details
CVE-2026-133095	In JetBrains TeamCity before 2026.2.2, 2026.1.6 a CSRF attack was possible on the Hub integration settings	8.1	More Details
CVE-2026-126088	In Autonomy Logic OpenPLC Runtime versions 4.2.4 and prior, a race condition allows multiple concurrent requests to pass the existence check before any user is committed, resulting in several unauthenticated admin accounts being created.	8.1	More Details
CVE-2026-123891	SDDC Manager contains an SQL-injection vulnerability due to improper sanitization of user-supplied input. Affected versions: 9.1.x.x (fixed in 9.1.1.0200); 9.0.x.x (fixed in 9.0.2.0100); 5.x (fixed in 5.2.5).	8.1	More Details
CVE-2026-104129	pREST (PostgreSQL REST) is an API that delivers an application on top of a Postgres database. Prior to 2.4.0, the vulnerable parameter is _groupby, handled by GroupByClause in adapters/postgres/postgres.go. adapters/postgres/postgres.go uses the weaker isSafeSQLExpression check for this function-expression branch instead of the stricter single-field validator used for _select. That check validates only the outer function name against a 13-function allowlist and permits nested subqueries, trailing UNION syntax, and inner pg_ functions. controllers/crud.go concatenates the GROUP BY expression while parameterizing only WHERE values, allowing attacker-controlled SQL structure into the clause. Default configuration permits unauthenticated access; a restricted installation may require a low-privilege table reader. UNION-based extraction requires compatible column counts and types and was not verified end to end, but a delay-producing function can cause denial of service without that alignment. This issue is fixed in 2.4.0.	8.1	More Details
CVE-2026-103026	Paymenter is a free and open-source webshop solution for management of hosting services. Prior to 1.5.7, the HasConfirmation trait in app/Livewire/Traits/HasConfirmation.php, included by app/Livewire/Component.php across 36 client-facing Livewire components, exposed executeConfirmedAction as a public action that invoked any component method accepted by method_exists. The action could call inherited InteractsWithProperties::fill(), whose data_set() writes arbitrary public properties, including state marked with #[Locked] and dotted paths into hydrated Eloquent models. Depending on the methods and properties exposed by an individual component, an unauthenticated caller could use this path to bypass intended authorization and modify server-side state intended to be immutable or protected from client-side manipulation. This issue is fixed in version 1.5.7.	8.1	More Details
CVE-2026-123007	This issue affects NGINX Plus or NGINX Open Source configurations that use HTTP/2. When NGINX processes undisclosed traffic, the flow control window can exceed the allowed threshold, potentially resulting in a heap buffer overflow typically followed by an NGINX worker segmentation fault. Additionally, attackers can execute code on systems with Address Space Layout Randomization (ASLR) disabled or when the attacker can bypass ASLR. Impact: An attacker could leverage this vulnerability to crash the NGINX worker process, corrupt its memory, or potentially execute code remotely. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	8.1	More Details

CVE-2026-117659	The InfoScan Electron desktop client improperly handles TLS certificate errors by accepting certificates without validating their trustworthiness. The client also configures Electron with nodeIntegration enabled, contextIsolation disabled, and webSecurity disabled. These settings weaken the client's security boundaries. An attacker who is able to intercept or manipulate network communications may be able to present an untrusted certificate without the application rejecting the connection.	8.1	More Details
CVE-2026-131304	In Splunk AI Assistant versions below 2.0.0 a user who does not hold the "admin" or "power" Splunk roles could access data associated with another tenant and cause actions to be attributed to that tenant through the Representational State Transfer (REST) API. The vulnerability is possible because Splunk AI Assistant accepted a client-controlled tenant identifier when selecting the authorization context for a request.	8.1	More Details
CVE-2026-22310	When the setting for saving requests to non-existent pages (404) is enabled, JTL-Shop stores the raw requested URL in its redirect table, and the administration backend renders it without output encoding in the form used to edit a redirect. An unauthenticated remote attacker can store script code with a single request, which is executed in the browser session of an administrator who opens the affected redirect, allowing actions to be performed with the administrator's privileges. The setting is disabled by default in new installations but may still be enabled in older ones.	8.0	More Details
CVE-2026-130766	Obot from 0.26.0 before 0.27.0 fails to validate the resource identifier in remote MCP server protected resource metadata, allowing malicious MCP servers to steal access tokens. Attackers can name a different service and authorization server with chosen scopes, causing Obot to obtain the victim's token for that service and forward it to the attacker.	8.0	More Details
CVE-2026-117520	The Dynific Addons for Elementor (formerly Anywhere Elementor) WordPress plugin before 1.3.1 does not perform a capability check on an AJAX action that stores a configuration value, allowing any authenticated user, including Subscribers, to store a value that is later output without proper escaping in a quoted attribute on an administrator settings page, resulting in Stored Cross-Site Scripting that executes in the administrator's session.	8.0	More Details
CVE-2026-133083	In JetBrains TeamCity before 2026.2.2, 2026.1.6 stored XSS in Server Health reports was possible	8.0	More Details
CVE-2026-133091	In JetBrains TeamCity before 2026.2.2, 2026.1.6 stored XSS on the Users page was possible	8.0	More Details
CVE-2026-73971	Music Assistant is an open-source media library manager that integrates streaming services with connected speakers. Prior to 2.8.0, parse_arguments() passes the caller-controlled return_type parameter for config/core/get_value, config/providers/get_value, and config/players/get_value to parse_value() and Python eval() when resolving type parameters in music_assistant/helpers/api.py. Any authenticated user, regardless of role, can reach these commands through the WebSocket /ws endpoint or the REST JSON-RPC /api endpoint, causing arbitrary Python code to run with the Music Assistant server process privileges before command-handler validation, even if the request later fails. A default installation requires no special configuration, and compromise can expose secrets and configuration, alter library data, and control the application container. This issue is fixed in version 2.8.0.	8.0	More Details
CVE-2026-133017	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via Git VCS root settings	8.0	More Details
CVE-2026-132914	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Issue Log tab was possible	8.0	More Details
CVE-2026-82171	Dell PowerStore contains an XML External Entity (XXE) Injection vulnerability. A high-privileged attacker could exploit this vulnerability, potentially leading to disclosure of sensitive files and root-level compromise of the appliance.	8.0	More Details
CVE-2026-130234	The Okta Swift Management SDK does not percent-encode forward slash (/) characters during path parameter interpolation. If an application supplies input containing relative path sequences (e.g., ../) to an SDK method, standard URL normalization may direct the outbound request to an alternative Okta Management API endpoint.	8.0	More Details
CVE-2026-62340	X-Road is open-source data exchange layer software. Prior to .8.2, 7.7.3, and 7.6.5, the Security Server writes each downloaded federated global-configuration part under the X-Road instance identifier declared by the part itself without checking that the identifier matches the source ecosystem that signed and delivered it. A malicious or compromised federated ecosystem can provide a validly signed SHARED-PARAMETERS part naming a victim ecosystem, causing the part to be written as the victim's authoritative shared-params.xml. The poisoned configuration can persist across refreshes, add attacker-controlled trust material, and enable impersonation of victim members and Security Servers across every ecosystem federated with the malicious source. This issue is fixed in versions 7.8.2, 7.7.3, and 7.6.5.	8.0	More Details

CVE-2026-109555	The WP STAGING and WP STAGING Pro plugins for WordPress are vulnerable to PHP Object Injection in all versions up to, and including, 4.17.0 and 6.17.0, respectively, via deserialization of untrusted input during database restore and staging reset. This is due to the database importer treating attacker-controlled text beginning with {WPSTG_BINARY} as trusted internal binary data and decoding its hexadecimal suffix into serialized bytes. This makes it possible for authenticated attackers, with Subscriber-level access and above, to plant a malicious serialized object in their own profile. When an Administrator later performs an ordinary database-including restore or staging reset containing that value, and the restored profile metadata is loaded, either during the restore or in a later request, a built-in POP chain can be used to execute arbitrary code on the server. Additional built-in gadget chains allow arbitrary local file read in both WP STAGING and WP STAGING Pro and arbitrary file deletion in WP STAGING Pro. The attacker does not require Administrator credentials or a copy of a backup, but exploitation does require the later privileged restore or reset.	8.0	More Details
CVE-2026-99674	In the Linux kernel, the following vulnerability has been resolved: can: dev: can_dropped_invalid_skb: drop CAN XL frames on non-CAN XL devices Sending a PF_PACKET bypasses the CAN framework logic and can directly reach a CAN driver's xmit() function. The PF_PACKET framework only checks that skb->len does not exceed the net_device MTU. For a CAN device that is not CAN XL capable, anything above CANFD_MTU (72 bytes) is therefore dropped before it reaches the driver. However, CAN XL frames are variable length. can_is_canxl_skb() accepts lengths in the range CANXL_HDR_SIZE + CANXL_MIN_DLEN up to CANXL_MTU, i.e. 13 to 2060 bytes. As a result, an ETH_P_CANXL skb with a length between 13 and 72 bytes can pass both the MTU and the can_dropped_invalid_skb() checks. A driver that does not support CAN XL will interpret canxl_frame->flags as a length because of the overlap with can_frame->len. And because CANXL_XLF is set, the resulting length is between 128 and 255. For drivers that do not check can_frame->len before copying can_frame->data, as most drivers do, this results in a buffer overflow of up to 247 bytes. Drop ETH_P_CANXL skbs if the device does not have the CAN_CAP_XL capability. Keep can_is_canxl_skb() for the validation of CAN XL skbs.	7.8	More Details
CVE-2026-99672	In the Linux kernel, the following vulnerability has been resolved: can: protect the receive lists with a per-device lock The receive lists of a CAN device are located in the device's can_ml_priv and travel with the device when it is moved to another network namespace. But they are protected by the per-netns rcvlists_lock. So can_rx_unregister() needs dev_net(dev) to pick the right lock and silently returns when the passed namespace does not match dev_net(dev) anymore. isotp_release() and bcm_release() leave their notifier list before they take any lock that serializes against the netdev notifier, which runs under RTNL. When the bound device is moved to another namespace in this window, the notifier skips the socket and the release path has to remove the filters itself. isotp_release() passes the socket's namespace, which does not match dev_net(dev) anymore. bcm_rx_unreg() passes dev_net(dev), which can_rx_unregister() reads again for its check, so a namespace change between both reads makes the check fail as well. In both cases can_rx_unregister() silently returns. As can_rx_register() stores rcv->sk without a reference, the remaining receivers point to the freed socket (isotp) or the freed op (bcm): BUG: KASAN: use-after-free in isotp_rcv+0x1570/0x24d0 Read of size 1 at addr ffff888118130552 by task isotp_ns_uaf/578 can_rcv_filter+0x4af/0x8c0 can_receive+0x28d/0x3c0 can_rcv+0x2a9/0x310 __netif_receive_skb_one_core+0x21a/0x260 process_backlog+0x210/0x760 Besides this, the lists of a device are protected by the rcvlists_lock of the namespace that is passed in. Callers which pass dev_net(dev) read it at different times, so the same lists may be modified under the locks of different namespaces while the device is moved. Taking RTNL in these release paths is no option: raw_bind(), raw_setsockopt() and raw_release() already hold RTNL while adding and removing up to CAN_RAW_FILTER_MAX filters per socket, which unprivileged users can trigger repeatedly and which makes RTNL a global bottleneck. Instead fix the root cause and add a spinlock to struct can_dev_rcv_lists, which protects the receive lists of a CAN device or the 'any' CAN device lists of a namespace. It is initialized in can_set_ml_priv(), which is used by all CAN network devices. With this, can_rx_unregister() finds the subscription independently of the current namespace of the device. The protocols keep the device valid with their tracked references. The per-netns rcvlists_lock only protects the receive list statistics anymore, and CAN_RAW, CAN_BCM and CAN_ISOTP pass the namespace that has been used for can_rx_register(). So isotp_release() and bcm_release() now remove their filters correctly without RTNL, and raw_release(), raw_setsockopt(), raw_bind() and isotp_bind() do not need RTNL anymore and add and remove their filters under lock_sock(). To serialize bind() with NETDEV_UNREGISTER, bind() publishes the device inside the RCU section of its lookup, and raw_notify() and isotp_notify() recheck the bound device under lock_sock(). As RTNL does not keep NETDEV_DOWN out of bind() anymore, bind() checks IFF_UP after publishing the bound device, with a barrier pair against raw_notify() and isotp_notify(), so a down transition during bind() is still reported with ENETDOWN.	7.8	More Details
CVE-2026-99668	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: fix zapping migration entries for device-private THPs A PMD migration entry can refer to a device-private folio. zap_huge_pmd_folio() checks the folio type instead of the entry type, so it treats such a migration entry as a device-private entry. Installing the migration entry already removed the reverse mapping and dropped the mapping reference. Zapping it must not repeat those operations. Concurrent MADV_DONTNEED during migration back to system memory can trigger this. Check the PMD entry type using pmd_is_device_private_entry(). Fix generated by an LLM and sent by Chris Mason in private. [david: simplify description]	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: fix truncate racing with PMD zapping zap_huge_pmd_folio() removes the rmap before calling folio_mark_dirty(). If this was the last		

CVE-2026-99669	mapping, truncation can skip unmapping the folio and proceed without waiting for the PMD lock, racing with dirtying. ifs_set_range_dirty() can oops if MADV_DONTNEED races with ftruncate() on XFS, when we get the folio set dirty after it was found to not be mapped anymore. Move folio_mark_dirty() and folio_mark_accessed() before rmap removal, matching the PTE zap path. This keeps the folio mapped while updating its state, so truncation must wait for the page table lock. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify comment and description]	7.8	More Details
CVE-2026-132150	A flaw was found in the file-xmc plugin in GIMP. When processing a specially crafted XMC image file, the plugin does not properly validate the comment subtype before using it as an array index. This missing validation leads to out-of-bounds memory access, invalid free and use-after-free issues. This vulnerability can result in an application crash, leading to a denial of service, or potentially allow for arbitrary code execution.	7.8	More Details
CVE-2026-132038	Rapid7 Metasploit Framework versions prior to 6.5.10 suffer from an Improper Access Control issue in the msfdb database initialisation component whereby a PostgreSQL cluster initialised via msfdb init is configured with trust authentication and a world-accessible Unix domain socket in /tmp, allowing any local user on the same host to authenticate as the passwordless postgres superuser and execute arbitrary operating system commands as the msfdb owner. This issue is fixed in version 6.5.10.	7.8	More Details
CVE-2026-99707	In the Linux kernel, the following vulnerability has been resolved: mm/khugepaged: flush deferred unmaps before dropping a failed folio collapse_file() can fail its reference-count or dirty-folio check after unmapping with TTU_BATCH_FLUSH. Both paths put back the isolated folio, then unlock it and drop the lookup reference before reaching the common try_to_unmap_flush(). The page-cache reference does not keep the folio stable once the lock is released. Another collapse can replace and free it. With its PTEs already gone, that collapse cannot flush the first task's per-task TLB batch, and retract_page_tables() skips short or unaligned VMAs. A CPU can therefore retain a user translation to the freed folio. This has been reproduced with unprivileged MADV_COLLAPSE on a memfd. Flush at out_unlock while the lookup reference and folio lock are still held. The common flush continues to cover the accumulated pagelist on both success and rollback, preserving batching on successful collapses.	7.8	More Details
CVE-2026-99724	In the Linux kernel, the following vulnerability has been resolved: ntfs: don't allow compression on volumes with clusters larger than 4K On a volume made with "mkntfs -c 8192", a user can set FILE_ATTR_COMPRESSED on an empty file through system.ntfs_attrib. Writing to the file and reading it back after dropping the page cache then crashes: BUG: KASAN: slab-out-of-bounds in ntfs_read_compressed_block+0xa2c/0x2040 Read of size 8192 at addr ffff88804355a000 by task dd/106 The read path only handles clusters up to 4K. It copies each cluster out of a single page, and a compression block of 16 8K clusters does not fit in the 64K ntfs_compression_buffer. Loading an inode already rejects compressed data on such volumes. Reject the flag change the same way.	7.8	More Details
CVE-2026-99697	In the Linux kernel, the following vulnerability has been resolved: mm: memcontrol: fix NULL pointer deref and UAF from high_work_func() on dying memcg When memory.high is exceeded in non-task context, try_charge_memcg() schedules memcg->high_work without taking a reference on the memcg, and high_work_func() invokes reclaim_high() without pinning the memcg. With obj_cgroup reparenting, introduced in commit bf4f059954dc ("mm: memcg/slab: obj_cgroup API"), objects charged to a child memcg are reparented to its parent in mem_cgroup_css_offline() instead of holding a reference on the child until uncharged. Subsequent uncharges of those objects go to the parent, which can leave the child's memory use permanently above its high limit, so any later charge on the offlined memcg in non-task context schedules high_work. This was later widened from kernel memory to LRU folios when commit f1cf8d2f36dc ("mm: memcontrol: eliminate the problem of dying memory cgroup for LRU folios") converted LRU folios to use obj_cgroup reparenting. cancel_work_sync(&memcg->high_work) only runs at the end of teardown in mem_cgroup_css_free(), so high_work_func() can start or still be running reclaim_high() after mem_cgroup_css_released() or shrink_many() has called lru_gen_release_memcg() and unhashed lruvec->lrugen.list. If reclaim tries to rotate the memcg, lru_gen_rotate_memcg() dereferences a NULL pprev pointer. This issue is specific to MGLRU and was introduced with the memcg LRU in commit e4dde56cd208 ("mm: multi-gen LRU: per-node lru_gen_folio lists"). With either traditional LRU or MGLRU, css_free_rwork_fn() calls css_rstat_exit() which frees per-CPU rstat data before calling mem_cgroup_css_free(). If reclaim_high() is still running, rstat updates via mod_lruvec_state() -> __css_rstat_updated() cause a use-after-free on the per-CPU rstat data. This was introduced when commit 5da3bfa029d6 ("cgroup: use separate rstat trees for each subsystem") moved css_rstat_exit() ahead of ss->css_free(). Fix both problems by acquiring a reference with css_tryget_online() in high_work_func() before calling reclaim_high(), and by resetting memory.high and swap.high to PAGE_COUNTER_MAX in mem_cgroup_css_offline() so that charges via references on an offlined memcg do not trigger high_work on the dead child and skip checking live ancestors in try_charge_memcg(). Acquiring the reference in high_work_func() rather than before scheduling high_work is safe because mem_cgroup_css_free() calls irq_work_sync() and cancel_work_sync() which cancels pending work and waits for an already running high_work_func() to complete before mem_cgroup_free() frees the memcg. css_tryget_online() uses percpu_ref_tryget_live() and is safe to call after percpu_ref_exit() before the cgroup is freed unlike css_tryget(). Oops: general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] CPU: 1 UID: 0 PID: 27 Comm: kworker/1:0 Tainted: G B W 7.3.0-rc5-g2bba28e98578 #284 PREEMPT(full) Tainted: [B]=BAD_PAGE, [W]=WARN <-- Taint from lru_gen_rotate_memcg() VM_WARN_ON_ONCE Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.17.0-debian-1.17.0-1 04/01/2014 Workqueue: events high_work_func RIP: 0010:lru_gen_rotate_memcg+0x1a3/0xe90 Call Trace:	7.8	More Details

	<TASK> ? read_ctrl_pos+0x22f/0x3f0 shrink_lruvec+0x16cb/0x1b80 ? ret_from_fork_asm+0x1a/0x30 ? lapic_next_event+0x15/0x20 ? clockevents_program_event+0x2bb/0x6f0 ? lapic_next_event+0x15/0x20 ? __pfx_shrink_lruvec+0x10/0x10 ? __pfx_clockevents_program_event+0x10/0x10 ? _raw_spin_lock+0x84/0xe0 ? __pfx__raw_spin_lock+0x10/0x10 ? kasan_save_stack+0x30/0x50 ? psi_group_change+0x381/0x740 ? __pfx_sched_clock_cpu+0x10/0x10 ? save_fpregs_to_fpstate+0xc4/0x2c0 ? finish_task_switch.isra.0+0x1de/0xac0 ? __schedule+0xf62/0x42c0 ? __pfx___schedule+0x10/0x10 ? __pfx___irq_work_queue_lo ---truncated---		
CVE-2026-99690	In the Linux kernel, the following vulnerability has been resolved: ocfs2: cache the mount context for dentry lock callbacks ocfs2_drop_dentry_lock() drops the inode reference before releasing the DENTRY lock. This lets inode eviction release the local inode and open locks before a remote unlink can proceed past the DENTRY lock. A DLM callback can still arrive before the DENTRY lock has been dropped. ocfs2_get_dentry_osb() obtains the mount context through dl_inode->i_sb, which can access a freed inode when the downconvert thread drops the last inode reference in its post-unlock callback: Downconvert thread DLM callback thread ---------- ---------- ocfs2_process_blocked_lock() ocfs2_dentry_post_unlock() ocfs2_dentry_lock_put() ocfs2_drop_dentry_lock() iput(dl->dl_inode) [inode reclaimed after RCU] ocfs2_blocking_ast() ocfs2_get_dentry_osb() dl->dl_inode->i_sb ocfs2_simple_drop_lockres() Store the mount context in the DENTRY lock resource's l_priv field and use the default mount-context lookup. Recover the containing dentry lock with container_of(), as already done for refcount locks, and remove ocfs2_get_dentry_osb(). ocfs2_dismount_volume() shuts down the DLM before freeing the mount context, so callbacks can use it without accessing the inode. Preserve the existing iput-before-unlock order to avoid delaying local inode and open lock release until after the DENTRY lock is released.	7.8	More Details
CVE-2026-109908	A heap buffer overflow in reading encoding files in X.Org libfontenc before 1.11.0 could be used by local attackers able to inject enc files to the X server could crash or escalate privileges.	7.8	More Details
CVE-2026-99688	In the Linux kernel, the following vulnerability has been resolved: net/sched: sch_qfq: prevent aggregate self-replacement qfq_change_class() snapshots the current aggregate settings while holding the qdisc tree lock, but drops the lock before selecting the destination aggregate. During that gap, qfq_enqueue() can move the class to an aggregate matching the requested settings. When qfq_change_class() resumes, qfq_find_agg() then returns cl->agg. If it is a singleton, qfq_deact_rm_from_agg() frees the aggregate before qfq_add_to_agg() immediately accesses the same pointer, causing a use-after-free. While holding the tree lock, skip the replacement when the destination is already the current aggregate. Any estimator replacement has already completed, so the class change can finish normally.	7.8	More Details
CVE-2026-99662	In the Linux kernel, the following vulnerability has been resolved: mm/gup: fix out-of-bounds write for nr_pages == 0 in GUP-fast With nr_pages == 0, GUP-fast still enters its page table walk and can write past the caller's page array. The old end <= start check caught empty ranges, but was replaced with an overflow check. Users such as OrangeFS can pass a zero page count derived from a userspace-supplied length. So let's return zero for nr_pages == 0 before starting the walk, just like GUP-slow. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify description]	7.8	More Details
CVE-2026-99686	In the Linux kernel, the following vulnerability has been resolved: vhost: fix used ring index wrap in vhost_add_used_n_in_order vhost_add_used_n_in_order() advances start by nheads[i] but wraps it with a single subtraction. If start + nheads[i] >= 2 * vq->num, start can remain outside the ring. With used-ring logging enabled, start > vq->num makes (vq->num - start) underflow. log_write() then sets bits beyond the dirty-log bitmap, potentially corrupting adjacent userspace memory. This is reachable on the vhost-net TX path with VIRTIO_F_IN_ORDER: done_idx can reach VHOST_NET_BATCH (64), triggering the issue with a 32-entry ring. Wrap start with a mask to keep it within the ring. VHOST_SET_VRING_NUM enforces a nonzero power-of-two ring size, allowing start to be wrapped using & (vq->num - 1).	7.8	More Details
CVE-2026-130893	Tanium addressed an unauthorized code execution vulnerability.	7.8	More Details
CVE-2026-99657	In the Linux kernel, the following vulnerability has been resolved: mm/mmu_notifier: set ->mm only after mmgrab() Radeon uses a non-NULL notifier->mm to decide whether to call mmu_interval_notifier_remove() when freeing a buffer object. __mmu_interval_notifier_insert() sets ->mm before validating the range, so a wrapping userptr range can leave ->mm set after insertion fails. Cleanup then tries to remove an unregistered notifier for which no mm reference was acquired. Set ->mm only after mmgrab(), once insertion can no longer fail. While we could fix radeon's code instead, setting ->mm on success only makes the code more robust. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify description]	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: mm: fix unmap_mapping_range() beyond 4 GiB on 32-bit Commit 9eab0421fa94 ("mm: fix unmap_mapping_range high bits shift bug") made holebegin and holelen unsigned before they are shifted, by casting them to pgoff_t. That is an unsigned long: on a 32-bit kernel the cast cuts off the upper half of the two loff_t values, so a hole which begins at 4 GiB or beyond is taken for one in the first 4 GiB of the file. truncate_pagecache() passes the new size of the file and asks for private COWed pages to be unmapped as well. If a file of 6 GiB is truncated to 5 GiB, everything from 1 GiB on is unmapped, and what a process has written to a private mapping of the range between 1 GiB and 5 GiB	7.8	More

99654	is lost: it reads the contents of the file again. Where the range has an end, the wrong range is unmapped instead of the right one. The page cache copes with that, as it unmaps every folio that it removes. The callers which only have unmap_mapping_range() to get rid of user mappings do not, for example drm_vma_node_unmap(), whose fake offsets reach beyond 4 GiB on 32-bit, and revoke_iomem() with physical addresses of more than 32 bits. Cast to u64 instead. Nothing changes for 64-bit kernels. Fix generated by an LLM and sent by Josef Bacik in private.		Details
CVE-2026-133553	A flaw was found in crun. When copying files onto a tmpfs mounted with tmpcopyup, crun opens the destination a second time by name. A container that shares a writable volume with the container being created can change that path between the two lookups, so crun copies image files onto host files, including files under /proc/sys.	7.8	More Details
CVE-2026-99650	In the Linux kernel, the following vulnerability has been resolved: mm/ksm: fix use of a freed anon_vma at swap-in A folio read from a swap slot previously shared by KSM can remain in the swap cache after the process that first mapped it exits. Its anon_vma can then be freed, so ksm_might_need_to_copy() must not dereference it to compare roots. Doing so can result in some false-positives, resulting in the reverse mapping missing the new folio mapping, preventing migration and reclaim. SLAB_TYPESAFE_BY_RCU allows individual anon_vma objects to be reused immediately after freeing. A stale pointer can then refer to another anon_vma whose root matches, even though the target VMA is not linked to that anon_vma. Move the anonymous rmap compatibility checks into folio_anon_rmap_compatible(). Keep the page-index check and compare the folio's anon_vma pointer against the VMA's own anon_vma and root first. If neither matches, bite the bullet and walk the VMA's anon_vma_chain. The VMA keeps these anon_vmas alive, and finding a match guarantees that reverse mapping can find the new mapping. Lorenzo Stoakes says in a private discussion: "This isn't quite equivalent to the existing check. Every anon_vma in an anon_vma_chain shares the same root, but if you split a VMA then fork both of them like this: VMA A / \ <split> VMA B VMA C \ / root <- anon_vma 1 ----^ ^---- <fork> / parent \ anon_vma 2 anon_vma 3 | | VMA D VMA E VMA's D and E both get their own new anon_vma's such that: VMA D: anon_vma 2 -> anon_vma 1 -> root VMA E: anon_vma 3 -> anon_vma 1 -> root And if you found a folio referencing anon_vma 2, the existing check: anon_vma->root == vma->anon-vma->root Would be true for VMA's B, C, D & E. However after your change it'd be true for VMA D only, as it's the only one with anon_vma 2 in its anon_vma_chain. Whether or not that could actually happen in practice... I don't know. [...] after your change it'd be true for VMA D only, as it's the only one with anon_vma 2 in its anon_vma_chain. [...] I would [...] say that you are making the test stricter but more correct. " If this affects performance when KSM was never even involved, we could just track whether KSM ever merged anything and just skip these checks. There isn't really a lot we can do otherwise: for mapped folios we might be able to use folio_get_anon_vma(), but for unmapped folios there really is no easy way to see if the anon_vma is still valid. Document the (confusing) lifetime requirements of folio_anon_vma().	7.8	More Details
CVE-2026-99611	In the Linux kernel, the following vulnerability has been resolved: ALSA: line6: Reject capture buffers smaller than two packets line6_capture_copy() splits the copy into the ALSA ring buffer once at the buffer end and wraps pos_done once, which is only correct when one packet holds at most buffer_size frames. The packet size comes from the device's endpoint, and the per-device hw structs allow 64-byte periods with periods_min 1, so a packet can be larger than the whole buffer. audio_in_callback() logs an oversized packet and then copies it anyway. The packet size is known at probe time, so constrain SNDRV_PCM_HW_PARAM_BUFFER_BYTES in snd_line6_capture_open() to at least twice the larger of the two endpoints' packet sizes. Two rather than one, because pos_done advances by the packet and wraps once, so a buffer of exactly one packet leaves the position unchanged after every packet. hw_params then rejects a buffer that cannot hold two packets, and the single wrap is valid for every configuration that gets through. Keep a guard in line6_capture_copy() that drops an oversized packet with a ratelimited error, for the case that should no longer be reachable. Reproduced on 7.3-rc7 with a USB/IP device presenting the GuitarPort interface layout and 1024-byte packets into a 128-byte buffer: ALSA logged "invalid position: pos = 224, buffer size = 32" and the driver logged "driver bug: len" on every following packet. With this patch a 1024-byte buffer is refused at hw_params, and a 2048-byte buffer streams for 30 s with no report and with the position moving inside the buffer.	7.8	More Details
CVE-2026-99422	In the Linux kernel, the following vulnerability has been resolved: ovl: use the pinned upper dir in ovl_fix_origin() Renames made directly on the upper filesystem during a lookup crash the kernel. A dentry that is passed to vfs_setxattr() must be pinned. ovl_fix_origin() passes upper->d_parent to ovl_set_impure() without a reference, and the lookup holds no lock of the upper directory. A rename on the upper filesystem changes d_parent, and only the child pins the new parent, so a second rename and an rmdir free it. With delays in ovl_fix_origin() and ovl_set_impure(), renaming the upper directory twice and removing the intermediate parent during the lookup of a merge directory without an origin xattr gave a KASAN use-after-free in vfs_setxattr(). Fix this by passing the upper dentry that the overlay parent holds. [amir: add helper ovl_set_impure_dir()]	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: ovl: set OVL_WHITEOUTS before publishing the upper dir Listing a directory while it is copied up causes a use-after-free of the readdir cache. A merge directory must never be iterated as a real upper directory: that mode frees the cache in OVL_I(inode)->cache without looking at its refcount, while files opened in merged mode still use it. ovl_copy_up_workdir() publishes the upper dentry before it sets OVL_WHITEOUTS, and ovl_dir_open() samples both without a lock, so an open between the two stores finds the upper directory and takes it for a real one. With a delay between	7.8	More

99424	the two stores, an unprivileged process that listed a directory during its copy-up saw only "." and "..", and its next getdents freed the readdir cache of another descriptor. Fix this by setting OVL_WHITEOUTS before ovl_inode_update(), so that ovl_dir_reset() agrees with ovl_dir_open(), and by testing OVL_TYPE_MERGE() of the path type in ovl_dir_open(), so that an open that observes the upper dentry never iterates a merge dir as a real dir, regardless of the order in which it observes the flag. [amir: test OVL_TYPE_MERGE() instead of smp_rmb() in ovl_dir_open()]		Details
CVE-2026-99425	In the Linux kernel, the following vulnerability has been resolved: ovl: hold a reference on the parent in readdir Listing a directory on an overlay uses a freed dentry and crashes the kernel when a lookup moves the directory at the same time. Code that can sleep may only use d_parent without a reference while something stops a move of the dentry. ovl_iterate_real(), ovl_need_adjust_d_ino() and ovl_cache_update() use the d_parent of the listed directory and hold only the lock of that directory. When a lookup finds the directory under another parent after a change to a layer, d_splice_alias() moves it without that lock and drops its reference on the old parent. With a delay in ovl_iterate_real(), getdents64 on an overlay directory raced a rename in the lower layer and a lookup that moved its dentry, and after an rmdir of the old parent KASAN reported a use-after-free in security_inode_getattr(). Fix this by using dget_parent() in the three functions. [amir: fix comments in ovl_need_adjust_d_ino()]	7.8	More Details
CVE-2026-99429	In the Linux kernel, the following vulnerability has been resolved: ocfs2: pin refcount trees before dropping the lookup lock A refcount tree returned to ocfs2_lock_refcount_tree() must stay alive until that caller releases its operation reference. However, ocfs2_get_refcount_tree() returns the cached pointer after dropping osb_lock, and the caller only takes its kref afterwards. The registry reference does not protect this gap from concurrent removal. This can happen when another node frees a refcount root and reuses its block with a different generation while this mount retains the old cached tree. Two local operations on different inodes sharing the new root, such as data CoW and an external xattr update, can then look up the same stale tree. Their inode locks are independent, and both the REFCOUNT DLM lock and rf_sem are acquired after the first operation reference. The following ordering is possible, with no other operation references: Data CoW Xattr update -------- ------------ Look up the old tree under osb_lock, then drop osb_lock Look up the same tree and take a kref Acquire the tree locks and read the new on-disk generation Erase the old tree under osb_lock Unlock and drop the operation kref Drop the registry kref, freeing tree Take a kref through the freed tree pointer The final kref_get() accesses freed rf_getcnt storage, and subsequent accesses in ocfs2_lock_refcount_tree() also use the freed allocation. Take the operation reference in ocfs2_get_refcount_tree() before releasing osb_lock, including when an insertion loses to another lookup. Removal must acquire the same lock, so it cannot drop the registry reference before the lookup has pinned the tree. Keep the existing unlock and error path puts. The two callers that already hold the refcount tree lock keep their borrowed lookup behavior and do not acquire an extra reference. KASAN report as below: BUG: KASAN: slab-use-after-free in ocfs2_lock_refcount_tree+0x94a/0x9b0 Read of size 4 at addr ffff888117212820 by task python3/497 [Kernel and hardware identification omitted.] Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 ? ocfs2_lock_refcount_tree+0x94a/0x9b0 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x20e/0x420 ? ocfs2_lock_refcount_tree+0x94a/0x9b0 kasan_report+0xe0/0x110 ? ocfs2_lock_refcount_tree+0x94a/0x9b0 ocfs2_lock_refcount_tree+0x94a/0x9b0 ? __pfx_ocfs2_lock_refcount_tree+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? ocfs2_refcount_cow+0x3b6/0xd90 ? srso_alias_return_thunk+0x5/0xfbef5 ? __kmalloc_cache_noprof+0x28d/0x610 ? ocfs2_refcount_cow+0x3b6/0xd90 ocfs2_refcount_cow+0x410/0xd90 ? __pfx_ocfs2_refcount_cow+0x10/0x10 ? __pfx___might_resched+0x10/0x10 ? down_write+0x11a/0x1f0 ? __pfx_down_write+0x10/0x10 ? ocfs2_inode_unlock+0x1a9/0x260 ? __pfx_ocfs2_inode_unlock+0x10/0x10 ? ocfs2_file_write_iter+0x1413/0x1d30 ocfs2_file_write_iter+0x157b/0x1d30 ? __pfx_ocfs2_file_write_iter+0x10/0x10 ? lock_acquire+0x190/0x300 ? __x64_sys_pwrite64+0x1b7/0x220 ? __pfx___might_resched+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? security_file_permission+0x26/0x80 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_is_held_type+0x8f/0x100 vfs_write+0x639/0x1010 ? __pfx_ocfs2_file_write_iter+0x10/0x10 ? lock_vma_under_rcu+0x154/0x4a0 ? __pfx_vfs_write+0x10/0x10 ? lock_release+0x121/0x290 ? srso_alias_return_thunk+0x5/0xfbef5 __x64_sys_pwrite64+0x1b7/0x220 ? __pfx___x64_sys_pwrite64+0x10/0x10 ? irqentry_exit+0xbe/0x820 ---truncated---	7.8	More Details
CVE-2026-99448	In the Linux kernel, the following vulnerability has been resolved: ufs: use u64 for directory size in ufs_last_byte ufs1_read_inode() and ufs2_read_inode() copy the untrusted 64-bit on-disk size into inode->i_size. ufs_last_byte() then truncates that value to an unsigned int before calculating the valid extent of a directory folio. A directory size of exactly 4 GiB is truncated to zero. In ufs_find_entry(), subtracting the requested record length from that zero forms an endpoint almost 4 GiB beyond the mapped folio. A crafted UFS2 image produces: BUG: KASAN: use-after-free in ufs_find_entry+0x583/0x6e0 [ufs] Read of size 1 ... ufs_find_entry ufs_inode_by_name ufs_lookup __lookup_slow path_lookupat The KASAN classification reflects that the adjacent physical page was free; the source-level operation is an out-of-folio read. A controlled UFS1 test with CONFIG_UFS_FS_WRITE enabled also made ufs_find_entry() return a directory entry from an adjacent anonymous page. unlink() then passed that pointer to ufs_delete_entry(), which cleared the adjacent page's 3 2-bit d_ino. This write result reproduced three out of three times. UFS normally requires a privileged mount path. A relevant boundary is a privileged automounter or image-processing service handling an attacker-supplied filesystem. The write primitive additionally requires UFS1 to be mounted read-write with CONFIG_UFS_FS_WRITE enabled; the read is reachable with read-only UFS2. Keep the intermediate size arithmetic at 64 bits so non-final pages are bounded at PAGE_SIZE. With this change, the same controlled	7.8	More Details

	tests reach the real on-disk directory entry and produce no adjacent-page write or KASAN report. A reproducer and full validation logs are available on request. The vulnerable helper is present in v7.2-rc3, v7.1, v6.18.38, v6.12.95, v6.6.111, and upstream master at 1229e2e57a5c. Runtime reproduction and fix validation were performed in an x86-64 QEMU guest running v7.2-rc3 with generic KASAN.		
CVE-2026-99457	In the Linux kernel, the following vulnerability has been resolved: net: sched: prevent scheduling deactivated qdiscs TAPRIO can schedule a deactivated qdisc after some_qdisc_is_busy() has reported it idle. Commit 47e8dbb6e763 ("net/sched: do not reset queues in graft operations") made graft deactivation skip queue reset, leaving TAPRIO's timer armed. A late callback schedules NET_TX softirq work, whose dequeue path can race the unlocked child-queue reset in taprio_reset(). The same skb can be queued for deferred freeing by RTNL and dequeued for transmission. In the reproducer, the TAP reader frees it before __rtnl_unlock() reads its next pointer. KASAN reports a slab-use-after-free in __rtnl_unlock+0xa3/0xd0, with a read of size 8 from the freed skb. Reject deactivated qdiscs in __netif_schedule(). Callers' RCU protection lets the existing grace period join earlier scheduling calls before the busy check. Restart nonempty transmit queues on reactivation so suppressed scheduling does not strand packets.	7.8	More Details
CVE-2026-99462	In the Linux kernel, the following vulnerability has been resolved: mm/truncate: look up the end-edge straddler by index In truncate_inode_partial_folio(), after the first split at the start edge, folio_split() unlocks and drops the refcount of the after-split sub-folios. The sub-folio straddling the end of the truncation range is therefore unlocked and only transiently ref'd in the page cache while the code still derives it from a page pointer inside the original folio. Between the first split finishing and page_folio() resolving split_at2, that tail page can be reclaimed, freed and reallocated as a new large folio in the same mapping at a different file offset. folio2 then points at a folio that does not cover the end boundary, yet folio2->mapping == folio->mapping still holds, so the stale pointer passes the mapping check and folio_split_or_unmap() splits a folio at a wrong position (or, with a transient refcount, a use-after-free window opens between try_get and the split). __folio_split()'s own folio != page_folio(split_at) check cannot catch this either since split_at2 has been reallocated as part of the new folio, so page_folio(split_at2) resolves back to folio2. Look the straddler up by its page index instead. __filemap_get_folio() returns the folio currently covering the boundary, ref'd and locked, with the mapping validated under the lock, so the split target is always the real folio at the end edge.	7.8	More Details
CVE-2026-99463	In the Linux kernel, the following vulnerability has been resolved: mm/rmap: don't discard lazyfree THPs at PMD level __discard_anon_folio_pmd_locked() clears a lazyfree THP's PMD before it determines whether the folio can be discarded, and restores it if the folio is dirty or has extra references. It is done under only the PMD lock and the anon_vma read lock inside try_to_unmap_one() and no mmap lock is held. The temporarily none PMD causes UAF. A parallel munmap() or mremap() is supposed to serialize with it on the PMD lock, but zap_pmd_range() (called by munmap()) and move_page_tables() (called by mremap()) do not take the lock when they see the temporarily none PMD and skip the PMD. There is no more page table walk after the none PMD is restored: - munmap(): free_pgtables() waits in unlink_anon_vmas() for the anon_vma lock, then free_pmd_range() drops the restored PMD as bad without removing the rmap, so the folio keeps a stale mapcount. - whole-VMA MREMAP_DONTUNMAP: dontunmap_complete() unlinks the source VMA from its anon_vma while the restored PMD still maps the folio. - plain mremap(): the THP is not moved but unmapped, so the THP's data is lost. In the first two cases, the folio stays mapped, but no VMA that maps it is linked to its anon_vma. When the anon_vma is freed, it is still reachable when one rmap walks the folio. Racing lazyfree reclaim against MREMAP_DONTUNMAP as an unprivileged user reproduces a KASAN use-after-free in folio_lock_anon_vma_read() as reported in the Closes link. One solution is to use an invalidated PMD instead of a cleared one until the discard is committed, but it needs arch fixes first. 1. GUP-fast still follows an invalidated PMD on riscv, LoongArch and s390. 2. sparc64 miscounts thp_pte_count when an invalidated PMD is set again or cleared. To keep the fix simple for backport, stop discarding lazyfree THPs at PMD level instead. try_to_unmap_one() splits the PMD and discards the folio at PTE level. unmap_huge_pmd_locked() and __discard_anon_folio_pmd_locked() are left in place, unused, to minimize this fix for backport and the invalidated PMD approach is planned as a follow-up series to use them again.	7.8	More Details
CVE-2026-99471	In the Linux kernel, the following vulnerability has been resolved: KVM: PPC: Book3S HV: Add SRCU protection for virtual-mode HPT hcalls The virtual-mode HPT hcall handlers (H_ENTER, H_REMOVE, H_BULK_REMOVE, H_CLEAR_REF, H_CLEAR_MOD) call into book3s_hv_rm_mmu.c, which accesses memslots via kvm_memslots_raw(). kvm_memslots_raw() uses rcu_dereference_raw_check() to bypass SRCU lockdep annotation checking. This is safe in real mode because the entire guest entry/exit is wrapped in srcu_read_lock/unlock inside kvmppc_run_core() and kvmhv_run_single_vcpu(). However, since commit 6165d5dd99db ("KVM: PPC: Book3S HV: add virtual mode handlers for HPT hcalls and page faults"), these same handlers are also executed in virtual mode via kvmppc_pseries_do_hcall(), which runs after SRCU has already been released on guest exit. A concurrent KVM_SET_USER_MEMORY_REGION deletion or move can therefore call synchronize_srcu_expedited() — which does not wait for this thread — and then kfree(slot) and vfree(slot->arch.rmap) while the hcall handler still holds a raw pointer to the memslot. lock_rmap() then writes to freed vmalloc memory, leading to use-after-free and memory corruption. Because kvm_memslots_raw() suppresses lockdep checks, this race is entirely silent. Fix this by factoring out the 7 HPT hcall handlers (H_REMOVE, H_ENTER, H_READ, H_CLEAR_MOD, H_CLEAR_REF, H_PROTECT, H_BULK_REMOVE) into a helper function, kvmppc_pseries_do_hpt_hcall(), and wrapping its call site in srcu_read_lock(&kvm->srcu) / srcu_read_unlock(&kvm->srcu, idx). Targeting only the HPT hcalls avoids wrapping non-HPT hcalls that sleep (such as H_CONFER, H_REGISTER_VPA, H_PAGE_INIT) or handlers that already acquire SRCU internally (such as H_RTAS).	7.8	More Details

CVE-2026-99472	In the Linux kernel, the following vulnerability has been resolved: KVM: pfncache: use a dedicated invalidation sequence for cache refresh The gfn_to_pfn_cache refresh path guards against mmu notifier invalidations which complete while it has dropped gpc->lock for the HVA->PFN lookup: hva_to_pfn_retry() samples kvm->mmu_invalidate_seq and retries if it changed, or if mn_active_invalidate_count is still elevated. That is insufficient for HVA-based caches. mmu_invalidate_seq is only advanced by kvm_mmu_invalidate_end() when the invalidated range overlaps a memslot, and an HVA-based cache (e.g. the Xen shared_info page mapped with KVM_XEN_ATTR_TYPE_SHARED_INFO_HVA) need not be backed by any memslot at all. An invalidation of the cached HVA which starts and ends entirely within the lookup window is thus invisible to the retry check: mn_active_invalidate_count is back to zero and the sequence never moved. The refresh then publishes a mapping of a page which has already been freed, and the next reader dereferences it: BUG: KASAN: use-after-free in kvm_xen_shared_info_init+0x3c6/0x440 Read of size 4 at addr ffff8880599c2900 by task syz.2.383/7257 Since gfn_to_pfn_cache_invalidate_start() deliberately skips caches which are not currently valid (including one whose refresh is in progress, as the refresh clears the valid flag before dropping the lock), the retry check is the only line of defence, and it must fire for *any* invalidation, not just those hitting a memslot. Add a dedicated kvm->gpc_invalidate_seq, incremented by every kvm_mmu_notifier_invalidate_range_end() under mn_invalidate_lock before mn_active_invalidate_count is decremented, and check it in hva_to_pfn_retry() instead of mmu_invalidate_seq. Incrementing in range_end() in the same critical section as the in-progress count also closes the variant where the cache is activated with the contested HVA only after invalidate_range_start() has run. The same bug is also reachable through the per-vCPU vcpu_info cache (KVM_XEN_VCPU_ATTR_TYPE_VCPU_INFO_HVA), where the stale mapping is then dereferenced by kvm_setup_guest_pvclock() on the next KVM_RUN: BUG: KASAN: use-after-free in kvm_setup_guest_pvclock+0x5bf/0x660 This intentionally makes refresh retry on *unrelated* mmu notifier events; restoring precision (and reworking the GPC locking more generally) is left for a subsequent series. Reproducers: https://david.woodhou.se/xen_shinfo_race.c https://david.woodhou.se/vcpu_info_race.c	7.8	More Details
CVE-2026-99477	In the Linux kernel, the following vulnerability has been resolved: fs/proc/task_mmu: pin mm_users in do_pagemap_scan() The PAGEMAP_SCAN ioctl takes mmap_read_lock() on an mm obtained from /proc/<pid>/pagemap without holding a reference to mm_users. proc_mem_open() deliberately grabs only mm_count, so a consumer that touches the address space has to call mmget_not_zero() first. Every other consumer in this file does -- m_start(), do_procmap_query(), show_smaps_rollup() and pagemap_read() -- but do_pagemap_scan() does not. Without that reference the target can reach exit_mmap() concurrently, and exit_mmap() calls unmap_vmas() while holding mmap_lock only for reading, which does not exclude the scanner. For ordinary VMAs this is harmless, as free_pgtables() and remove_vma() run afterwards under mmap_write_lock. A shared hugetlb VMA is different: __hugetlb_zap_end() frees the VMA's hugetlb_vma_lock inside the read-locked phase. walk_hugetlb_range() can therefore load vma->vm_private_data, have the lock freed underneath it, and then down_read() a freed rw_semaphore. The load happens before the freeing side clears the pointer, so this is a use-after-free and not a NULL dereference. Any unprivileged user can reach it with no huge pages configured, as hugetlb_reserve_pages() allocates the lock before the VM_NORESERVE early return. Pin mm_users across the scan, the way the other consumers do. mmput() from the exiting task can then no longer reach exit_mmap() while a walk is in progress. [akpm@linux-foundation.org: remove unneeded & undesirable test for !mm]	7.8	More Details
CVE-2026-99478	In the Linux kernel, the following vulnerability has been resolved: filelock: keep superblock active through delegation wait A caller may drop its path after try_break_deleg() returns a delegated inode but before break_deleg_wait() finishes. The extra inode reference held in struct delegated_inode does not keep the superblock active. After a lazy unmount, closing the delegation holder can release the last mount reference and shut down the superblock while another task waits for the delegation break. The waiter's final iput() can then enter the filesystem's eviction code after its superblock-private state was freed. This has been observed as a KASAN use-after-free in shmem_evict_inode(). When try_break_deleg() saves a delegated inode, take an active reference to its superblock while the caller still pins the mount. Keep that reference through break_deleg_wait() and the final iput(), then release it with deactivate_super(). This protects every caller of the helpers, including the mknod, mkdir, symlink and rmdir paths that drop their path before waiting.	7.8	More Details
CVE-2026-99480	In the Linux kernel, the following vulnerability has been resolved: drm/panthor: Fix TOCTOU race between GROUP_REGISTERED check and erase panthor_group_destroy() checks the GROUP_REGISTERED mark and then calls xa_erase() without holding the xarray lock across both. If the handle is destroyed by another thread in between, and a concurrent GROUP_CREATE reuses the freed ID, the first destroy erases and releases the new, still-initializing group, leading to a use-after-free in panthor_group_create(). Do the mark check and erase atomically under xa_lock().	7.8	More Details
CVE-2026-99484	In the Linux kernel, the following vulnerability has been resolved: ALSA: emux: Fix UAF of soundfont zones at note-on snd_emux_note_on() looks up the matching zones via snd_soundfont_search_zone(), which takes sflist->lock only during the lookup and returns the raw zone pointers after releasing it. The caller then stores them in vp->zone without any protection. Meanwhile, a concurrent SNDRV_EMUX_IOCTL_REMOVE_LAST_SAMPLES ioctl may unlink and free those zones in snd_soundfont_remove_unlocked(); this leads to a use-after-free in snd_emux_note_on(), and later at accessing the cached vp->zone, e.g. in snd_emux_send_effect(). Fix it by holding sflist->lock in snd_emux_note_on() over the whole zone lookup and voice setup, and let snd_soundfont_search_zone() require the caller to hold the lock, instead. snd_soundfont_lock_preset() takes sflist->lock for setting	7.8	More Details

	presets_locked, so the removal path now syncs with the note-on finish.		
CVE-2026-99506	In the Linux kernel, the following vulnerability has been resolved: mm/hugetlb: fix use-after-free of a new anon folio For a fault in a private mapping with nothing in the page cache, hugetlb_no_page() allocates a folio, locks it and maps it. The folio has one reference, which belongs to the page table entry from then on. After dropping the page table lock the function still sets the migratable flag of the folio and unlocks it. Whoever zaps the entry in between frees the folio. The fault holds the fault mutex and the lock of the vma. MADV_DONTNEED and the first unmap of a hole punch, hugetlb_vmdelete_list(), do not take the fault mutex, and a private vma only has a lock if it owns a reservation map: not with MAP_NORESERVE, and not in the child after fork(). Then the flag is set in, and folio_unlock() is called on, a folio that is free or has a new owner. A surplus page goes back to the page allocator while it is locked: BUG: Bad page state in process init pfn:108e00 flags: 0x200000000000059(locked|uptodate|dirty|head|node=0|zone=2) page dumped because: PAGE_FLAGS_CHECK_AT_FREE flag(s) set A page of the pool can be given to the next user, of any file, who locks it and loses the lock to the late folio_unlock(): kernel BUG at mm/filemap.c:1532! hugetlb_no_page+0x111b/0x23e0 hugetlb_fault+0x3a9/0x14a0 An unprivileged process is enough: threads that write to a MAP_PRIVATE | MAP_NORESERVE hugetlb mapping, threads that call MADV_DONTNEED on it, and one that takes the freed pages for another file. Until commit 3a47c54f09c4 ("hugetlbfs: revert use i_mmap_rwsem for more pmd sharing synchronization") the fault held i_mmap_rwsem for reading, which everybody who zaps takes for writing. The lock that commit bf4916922c60 ("hugetlbfs: extend hugetlb_vma_lock to private VMAs") gave to private vmas closed the window again, but only for those with a reservation map. Be done with the folio before the page table lock is dropped, as hugetlb_wp() is with its copy. The function already unlocks page cache folios under that lock. Original fix generated by an LLM and sent in private by Josef. [david@kernel.org: simplify+move comment]	7.8	More Details
CVE-2026-99507	In the Linux kernel, the following vulnerability has been resolved: mm/hugetlb: fix use-after-free in hugetlb_mfill_atomic_pte() Let's fix another instance where we use a folio after dropping the PTL, where someone else could e.g., MADV_DONTNEED it once we drop the PTL, making us trigger a use-after-free on the folio.	7.8	More Details
CVE-2026-99523	In the Linux kernel, the following vulnerability has been resolved: RDMA/uverbs: Fix DMABUF use-after-free during ufile cleanup The DMABUF release can race with ufile cleanup. When the cleanup holds hw_destroy_rwsem for write, uverbs_uobject_release() fails its trylock and leaves the uobject on the ufile list, but the dma_buf is freed as soon as the release returns. The cleanup then destroys the uobject and uses the freed dma_buf: BUG: KASAN: slab-use-after-free in uverbs_dmabuf_fd_destroy_uobj+0x46/0x200 Read of size 8 at addr ff11000109cf14b0 by task repro/435 Call Trace: uverbs_dmabuf_fd_destroy_uobj+0x46/0x200 destroy_hw_fd_uobject+0x3f/0x50 uverbs_destroy_uobject+0xc5/0x380 __uverbs_cleanup_ufile+0x1a8/0x260 uverbs_destroy_ufile_hw+0x54/0x160 ib_uverbs_close+0x34/0x110 Allocated by task 435: dma_buf_export+0x1b5/0x530 ib_uverbs_handler_UVERBS_METHOD_DMABUF_ALLOC+0x223/0x3f0 Freed by task 436: kfree+0x301/0x530 dentry_kill+0x14c/0x560 __fput+0x23a/0x4a0 rdma_user_mmap_entry_remove() dereferences the same pointer. Protect the stored dma_buf pointer with a mutex and clear it before release returns. Use a shared revoke helper so uobject destruction and mmap entry removal both skip a cleared pointer. Hold a temporary uobject reference across uverbs_uobject_release() and the pointer update. Use a separate mutex because uobject destruction can free the mmap entry before release finishes. Map and unmap use attachment->dmabuf directly, so they do not need to take the new mutex while holding the reservation lock.	7.8	More Details
CVE-2026-99528	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: fix netns use-after-free in write_gssp() While fuzzing with a customized syzkaller, I hit a refcount warning in xprt_alloc(), reached from a write to /proc/net/rpc/use-gss-proxy: refcount_t: addition on 0; use-after-free. WARNING: lib/refcount.c:25 at refcount_warn_saturate+0xf8/0x120 CPU: 0 PID: 10404 Comm: syz.0.17 Not tainted 7.3.0-rc1-00096-gcfebfd3db73d Call Trace: <TASK> xprt_alloc+0x83f/0x9d0 xs_setup_xprt+0xaf/0x3c0 xs_setup_local+0x47/0x7f0 xprt_create_transport+0x16c/0x730 rpc_create+0x38e/0x7f0 gssp_rpc_create+0xe2/0x180 set_gssp_clnt+0xba/0x1b0 write_gssp+0x200/0x310 proc_reg_write+0x240/0x330 vfs_write+0x2aa/0x1050 ksys_write+0x12a/0x250 do_syscall_64+0x117/0x750 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> create_use_gss_proxy_proc_entry() stores the struct net pointer as the proc entry's private data without taking a reference on it, and an open file descriptor does not pin the namespace either, as procfs only pins the proc_dir_entry. write_gssp() reads that pointer back with pde_data() and hands it to set_gssp_clnt(), which reaches xprt_init(): xprt->xprt_net = get_net_track(net, &xprt->ns_tracker, GFP_KERNEL); get_net_track() -> get_net() -> refcount_inc() is unconditional, so rpc_create() assumes its caller holds a reference on args.net. write_gssp() does not. Dropping the last ns.count reference does not make the file go away. __put_net() only queues cleanup_net() on a workqueue, and the remove_proc_entry() that fences off further writes runs from the rpcsec_gss pernet .exit method, that is, from inside cleanup_net(). A write landing between those two points increments a refcount that is already zero. cleanup_net() does not re-read ns.count, so the namespace is freed anyway and the rpc_xprt -- along with the AF_LOCAL socket opened for it -- is left pointing at freed memory. The resurrected count is visible to the rest of that teardown too: with CONFIG_IPV6_MROUTE the same cleanup_net() worker then trips !mr_can_free_table() in ip6mr_free_table(), which tests check_net() on the namespace it is freeing. It reduces to opening /proc/net/rpc/use-gss-proxy inside a new network namespace, calling setns() back to the original one to drop the last reference, and then writing "1" to the still-open descriptor. Take the reference in write_gssp() itself and refuse the write when the namespace is already gone. This is what procfs does for every other /proc/net file: get_proc_net() is maybe_get_net(PDE_NET(PDE(inode))) and seq_open_net() returns	7.8	More Details

	-ENXIO when it fails. PDE_NET() cannot be reused here because the parent directory /proc/net/rpc carries no namespace pointer. Adding a .pre_exit to rpcsec_gss_net_ops so the entry is removed earlier would only narrow the window, since pre_exit also runs from cleanup_net(). Live namespaces are unaffected: a write still returns the error from set_gssp_clnt(), a write of "2" still returns -EINVAL, and reads are untouched.		
CVE-2026-133309	@scalar/typebox before 0.1.4 contains a code injection vulnerability in TypeCompiler that allows attackers to execute JavaScript by supplying crafted schema strings. Attackers can inject single-quote breakouts into string format, Kind names, property keys or Literal values to run arbitrary JavaScript in the host Node.js process during Compile() and Check().	7.8	More Details
CVE-2026-99400	In the Linux kernel, the following vulnerability has been resolved: xfs: reject out-of-range attribute value lengths in xfs_attr_copy_value xfs_attr_copy_value() takes the value length as a signed int and, for a remote xattr, is handed args->rmtvaluelen. That field is filled from the on-disk __be32 xfs_attr_leaf_name_remote.valuelen in xfs_attr3_leaf_getvalue(), so a crafted length such as 0x80000000 is stored into the signed rmtvaluelen as a negative number. The "buffer too small" guard in xfs_attr_copy_value() is a signed comparison: if (args->valuelen < valuelen) return -ERANGE; A negative valuelen therefore compares as smaller than the caller's buffer size, skips the -ERANGE path, and is then used as a copy length, leading to an out-of-bounds copy of a full remote block into a small getxattr(2) buffer on a mounted crafted image. Reject a value length that is negative or larger than the maximum xattr size before it is used, so a bogus on-disk length can no longer slip through the value copier.	7.8	More Details
CVE-2026-99395	In the Linux kernel, the following vulnerability has been resolved: xfs: bound inode fork length against the fork size during log recovery xlog_recover_inode_commit_pass2() copies an inode log item's data/attr fork region into the inode buffer using the on-log region length without bounding it against the fork capacity, e.g.: len = item->ri_buf[2].iov_len; memcpy(XFS_DFORK_DPTR(dip), src, len); The only guard is an ASSERT, which is a no-op on production kernels (CONFIG_XFS_DEBUG off), and xfs_dinode_verify() runs only after the copy. A crafted image with a dirty log can therefore drive a heap out-of-bounds write at mount time. The XFS_ILOG_DBROOT sibling already passes XFS_DFORK_DSIZE as a bound; the DDATA/DEXT and ADATA/AEXT memcpy paths did not. Bound each logged fork region against the destination fork size before copying it, and reject the log item with -EFSCORRUPTED when it does not fit. Because the recovered inode is only verified after the fork data has been copied in, the checks are done up front, before any memcpy into the on-disk inode.	7.8	More Details
CVE-2026-99345	In the Linux kernel, the following vulnerability has been resolved: KVM: nSVM: Reject KVM_SET_NESTED_STATE if L1 has EFER.LMA=1 && EFER.LME=0 Reject KVM_SET_NESTED_STATE if the incoming L1 host state has what is effectively an impossible EFER combination of LMA=1 but LME=0, i.e. if the state says long mode is active but not enabled. Unlike VMX, SVM doesn't have an explicit consistent check for the illegal combination; presumably hardware simply ignores EFER.LMA if EFER.LME=0. Unfortunately, KVM doesn't ignore EFER.LMA in this case and consumes the illegal state when constructing the shadow MMU for L2. E.g. if userspace also clears CR4.PAE, then kvm_calc_cpu_role() will compute a role with 4 or 5 levels of paging, but shadow_mmu_init_context() will wire up the MMU to use the paging32 template, which maxes out its levels at 2. Note, the "real badness" is effectively the same as what happened with the nVMX bug fixed by commit 112e66017bff ("KVM: nVMX: add missing consistency checks for CR0 and CR4"). Unfortunately, the sanity check added by commit 72e2fb24a0b0 ("KVM: x86/mmu: Bug the VM if a vCPU ends up in long mode without PAE enabled") doesn't work for this case, since L2 state is active at the time of the page fault, but it's L1 that has the bad state.	7.8	More Details
CVE-2026-132751	yt through 4.4.4 contains a code injection vulnerability that allows attackers to execute Python code by supplying a crafted con_args attribute in a ytdata HDF5 file. SavedDataset._parse_parameter_file() passes con_args to eval() unrestricted, so loading a malicious .h5 file with yt.load() runs attacker code as the user.	7.8	More Details
CVE-2026-99609	In the Linux kernel, the following vulnerability has been resolved: ALSA: usx2y: Require a buffer of two URBs usx2y_urb_capt_retire() wraps the ring buffer once per packet and usx2y_urb_play_prepare() once per URB of nr_of_packs() packets of up to 50 frames. Both assume that the data fits in the buffer, and the hw constraints only require about two 1 ms periods, so a packet larger than the buffer on capture, or a URB larger than the buffer on playback, violates that assumption. Constrain SNDRV_PCM_HW_PARAM_BUFFER_SIZE in snd_usx2y_pcm_open() to at least two URBs, 2 * nr_of_packs() * 50 frames, for both directions. The bound is in frames and derived from the driver's own packet limit, so it does not read the endpoint, whose altsetting is only selected in prepare. Two URBs rather than one, because the position advances by a URB on playback and only moves when a URB is less than the whole buffer. Keep a guard in usx2y_urb_capt_retire() that drops a packet larger than the buffer, for the case the constraint makes unreachable. With nrpacks above about 163 (S16_LE) or 109 (S24_3LE) the bound exceeds the 64 KiB buffer the substream preallocates (twice that with the preallocation raised to its 128 KiB maximum), so hw_params fails for those module settings. That is intended: the driver cannot be safe with a URB larger than half the largest buffer. Reproduced on 7.3-rc7 with a USB/IP device presenting the US-122 interface layout and 1024-byte packets into a 512-byte capture buffer: BUG: KASAN: out-of-bounds in i_usx2y_urb_complete+0x1285/0x1ac0 Read of size 18446744073709551104 at addr ffff88800e448800 by task vhci_rx/83 Tested on 7.3-rc7 with the same device, S16_LE at 44100 Hz, nrpacks 4, so the bound is 400 frames. A 128-frame capture buffer and a 200-frame playback buffer are refused at hw_params with EINVAL. A 400-frame buffer streams for 30 s in each direction with the position advancing inside the buffer and no KASAN report; the capture guard did not fire. The same 400-frame capture buffer with S24_3LE, which selects altsetting 2 in prepare, streams for 30 s with the same result.	7.8	More Details

CVE-2026-99603	In the Linux kernel, the following vulnerability has been resolved: mptcp: pm: kernel: use RCU to free endpoint listeners mptcp_pm_nl_create_listen_socket() starts the TCP listener directly instead of going through mptcp_listen(). It consequently misses setting SOCK_RCU_FREE on the parent MPTCP socket. tcp_child_process() can invoke the listener's subflow_data_ready() callback without holding the MPTCP parent lock or a reference to the parent. An authenticated MP_JOIN racing with endpoint teardown can therefore access the parent after its last reference has been dropped. The ADD_ADDR error path can also close an already published listener. Set SOCK_RCU_FREE before __inet_listen_sk() publishes the TCP listener, matching mptcp_listen(). This protects the parent on every subsequent close path, including a failed endpoint addition.	7.8	More Details
CVE-2026-99601	In the Linux kernel, the following vulnerability has been resolved: udf: check the unallocated space table when it is loaded udf_table_new_block() hands out the first block of the extent closest to its goal and trusts that the extent covers at least one whole block inside the partition. Nothing checks that. A zero-length extent makes it return the block the extent points at, which need not be free, and the subtraction of one block then underflows the length into the type bits, turning the entry into a bogus continuation. An extent running past the end of the partition makes it return blocks behind the partition. Such tables are not only found on crafted images: until the previous commit, udf_table_new_block() itself left zero-length extents behind in tables written by mkudffs, and then converted them into extents running past the partition. Check every extent once when the table is loaded, and refuse read-write access if one is empty, is not a whole number of blocks, or does not lie inside the partition, in the same way as for other allocation information the kernel cannot use. Check first that the table is an Unallocated Space Entry at all: the walk starts at the offset of the allocation descriptors in one, and for a file entry that offset lies before the inode's in-memory copy of the descriptors. Treat a chain of allocation extent descriptors that leads back to itself the same way: the walk would otherwise return the same extents forever, and the mount would never finish. Let a fatal signal interrupt the walk and fail the mount. Refusing read-write access to an empty extent is deliberate. It shows that an old kernel took the last block but left the entry behind. If consumed again, the entry returns an in-use block and turns into a huge bogus extent. Refusing the table before that happens avoids this second allocation; the filesystem can still be mounted read-only. The check cannot detect every overlap that may already exist in a damaged table. This patch should be backported together with the previous one: that patch stops creating new empty extents, and this one rejects existing ones before they are consumed again.	7.8	More Details
CVE-2026-99592	In the Linux kernel, the following vulnerability has been resolved: x86/mm: Fix global ASID handed out to two processes On AMD CPUs with INVLPGB, two live processes can get the same MAX_ASID_AVAILABLE global ASID value. The two then hit each other's TLB entries, so each can read or corrupt the other's memory. The global ASID space is managed lazily. When a global ASID is freed, its bit is set in the global_asid_freed bitmap. Once the free ASID search reaches the end of the global_asid_used bitmap, a global flush is done to really reclaim the lazily freed global ASIDs. However, the current code can assign the same global ASID to two different processes, if one process is assigned MAX_ASID_AVAILABLE - 1, another process MAX_ASID_AVAILABLE, and both stick around (keeping their global ASID numbers in use) until the search wraps around again. At that point a third process could also get assigned MAX_ASID_AVAILABLE. the highest ASID stays in use, so it is not handed out again last_global_asid stays below it, so no reset the clear bits run out find_next_zero_bit() returns MAX_ASID_AVAILABLE global_asid_available > 0, so the failure is not caught MAX_ASID_AVAILABLE is handed out as an ASID A search only looks at the bits below MAX_ASID_AVAILABLE, so it cannot see that this value is in use. The next failed search hands it out again, even if the first process still holds it. The fix used here is to reset the global ASID space after a failed search, and to search again. If that search fails too, no global ASID is handed out. The test for !global_asid_available is removed, because use_global_asid() only calls allocate_global_asid() if global_asid_available > 0. The remaining test is for a condition that should be impossible. [riel: simplify code, massage changelog and comment]	7.8	More Details
CVE-2026-99591	In the Linux kernel, the following vulnerability has been resolved: iommu/amd: Force device table to maximum size IOMMU can perform out-of-bounds read of the device table entry (DTE) that is adjacent to the last DTE and falls outside the device table. For example, if the device table can hold 0 to N-1 entries, then the IOMMU can be made to fetch the Nth entry, which falls outside the device table. The device table is owned by the hypervisor (HV) and is assigned in 4KB pages that are contiguous in system physical address (SPA). If the out-of-bound DTE falls on a guest-owned page, then the guest can exploit this bug to setup malicious table walks to corrupt HV pages and/or bring down the system. Note that this bug does not manifest when the device table size is set to its maximum size of 2MB. Therefore, force the device table size to the maximum 2MB on affected processors. These are CPU family 0x17 (Zen1/Zen2) through family 0x1a (Zen5/BRH). This is a mitigation for CVE-2024-21963. kdump reuses the previous kernel's device table only when the size encoded in the IOMMU device-table register matches pci_seg->dev_table_size. A mixed rollout disagrees: the unpatched kernel uses the IVRS-derived size and the patched kernel uses 2MB, so __reuse_device_table() rejects the table. On SNP-capable IOMMUs that rejection panics the kdump kernel. Kernels that both have this patch agree on 2MB and reuse succeeds. Update the production kernel and the crash kernel together.	7.8	More Details
CVE-2026-99583	In the Linux kernel, the following vulnerability has been resolved: sctp: hold shkey across socket migration sctp_sock_migrate() transfers queued DATA skbs from the old socket to the new one. skb_orphan() invokes sctp_wfree() during that transfer and drops the skb-owned shared-key reference. If userspace has removed that key from the association, this can be the final reference. The following sctp_set_owner_w() then dereferences the freed chunk->shkey while trying to take the new owner reference. Take a temporary shared-key reference before orphaning the skb and release it after the new owner has taken its reference. This preserves the selected authentication key throughout the ownership transfer. Bug found and triaged by	7.8	More Details

	OpenAI Security Research and validated by Trail of Bits.		
CVE-2026-99582	In the Linux kernel, the following vulnerability has been resolved: sctp: re-point retained control chunks on association migration sctp_sock_migrate() transfers DATA chunk ownership but leaves retained control chunks pointing at the old socket. A later retransmission can therefore access the socket after it has been freed. Extend the migration walk to cover retained control chunks and use sctp_control_set_owner_w() to assign their new owner. Reuse the DATA ownership check so each chunk is processed once per pass, even when it is reachable through multiple lists or retained pointers. Preserve a control chunk's existing authentication key across migration so its final release can still generate SCTP_AUTH_FREE_KEY when needed.	7.8	More Details
CVE-2026-62964	Gortex is a high-performance code-intelligence engine for AI agents and integrated development environments. Prior to 0.55.0, resolveFilePath in internal/mcp/tools_fileops.go accepts absolute paths without checking indexed repository roots, while write_file, edit_file, and move_inline in internal/mcp/tools_move_inline.go omit guardSymlinkWithinRepo before reaching AtomicWriteFile and related write sinks. A caller that can influence MCP tool arguments can therefore create or overwrite files anywhere the daemon user can access, including startup or hook files that execute code on a later trigger. The flaw violates Gortex's repository-confinement invariant and can lead to arbitrary host code execution. This issue is fixed in version 0.55.0.	7.8	More Details
CVE-2026-132756	Geany through 2.1.0 contains an OS command injection vulnerability that allows attackers to run shell commands by embedding metacharacters in filenames substituted into build and execute commands. Attackers can craft a file named with backticks or $(...) that executes via /bin/sh when a user runs Compile, Build or Execute, gaining code execution as the user.	7.8	More Details
CVE-2026-99340	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Take SRCU in kvm_zap_gfn_range() kvm_zap_gfn_range() walks memslots and rmaps and may drop mmu_lock to reschedule. Callers that do not already hold kvm->srcu (or slots_lock) can race with memslot deletion: synchronize_srcu_expedited() does not wait, kvm_free_memslot() frees the old slot and its rmap, and the zap resumes on freed memory. The VFIO noncoherent-DMA path hits this by zapping the entire GPA space without SRCU. KASAN reported: BUG: KASAN: vmalloc-out-of-bounds in slot_rmap_walk_next+0x82/0x1c0 Read of size 8 at addr ffffc900005c1008 Call Trace: slot_rmap_walk_next+0x82/0x1c0 __kvm_rmap_zap_gfn_range+0x17a/0x280 kvm_zap_gfn_range+0x2a6/0x6a0 kvm_vfio_set_attr+0x576/0x770 kvm_device_ioctl+0x1ff/0x3b0 Take SRCU inside kvm_zap_gfn_range() so every caller is covered. Nesting with an existing kvm->srcu critical section is fine; the helper uses a local index. Drop the now-redundant SRCU pair from __kvm_set_or_clear_apicv_inhibit().	7.8	More Details
CVE-2026-99580	In the Linux kernel, the following vulnerability has been resolved: mm/userfaultfd: only allocate anon_vma in CoW mappings find_vma_and_prepare_anon() can end up allocating an anon_vma in VMAs that are not CoW mappings, which is wrong. If we have a MAP_SHARED R/O file mapped, we lack VM_SHARED (but do have VM_MAYSHARE) and lack VM_MAYWRITE. While uffd later rejects such mappings (uffd registration requires VM_MAYWRITE, see comment in validate_dst_vma) find_vma_and_prepare_anon() still wrongly allocates an anon vma before it validates the vma. This in turn later confuses copy_page_range()->vma_needs_copy() such that we end up calling copy_hugetlb_page_range() with ranges that actually do not contain *any* per-pagetable information that would require copying, which is the basis for PMD-table sharing in the first place. copy_hugetlb_page_range() should not have to worry about shared PMD tables at all: it doesn't make any sense to copy page tables that are shared, as they can be fully recreated later. Ending up in the function with shareable page tables is the real issue. If a destination PMD table is shared in copy_hugetlb_page_range(), it might get concurrently unshared as we are not holding relevant locks, resulting in us using a page table that is no longer mapped and possibly freed after calling: dst_pte = huge_pte_alloc(dst, dst_vma, addr, sz); To fix it, do not allocate an anon_vma in non-CoW mappings and add a safety guard to copy_hugetlb_page_range(), making sure that we never end up messing with shared PMD tables in there. Before commit 867a43a34ff8a ("userfaultfd: use per-vma locks in userfaultfd operations") we would always validate the VMA before allocating an anon_vma, and consequently never wrongly allocate an anon_vma. As uffd_lock_vma() is part of the same lookup machinery and also contains a misleading VM_SHARED check, change that one as well directly for consistency. A follow-up cleanup will stop taking the hugetlb_vma_lock_read(src_vma) lock in copy_hugetlb_page_range() entirely. Further, the usage of VM_SHARED in userfaultfd code should be cleaned up later to focus on CoW vs. non-CoW mappings.	7.8	More Details
CVE-2026-99574	In the Linux kernel, the following vulnerability has been resolved: f2fs: bound the SIT/NAT bitmaps by the checkpoint pack sit_ver_bitmap_bytesize and nat_ver_bitmap_bytesize are used as memcpy lengths into and out of sbi->ckpt, which is sized from cp_payload alone, and nothing compares either length against that allocation. sanity_check_ckpt() has checked these two fields since c77ec61ca0a4, which added them after a crafted image overran build_sit_info(). That check ties each length to a segment count; it never compares either against the allocation, which is the half this adds. Mounting over-reads in build_sit_info() and build_node_manager(), and the first checkpoint writes the same lengths back through the same pointers, so the over-read's contents reach the on-disk checkpoint. An unmount is enough; the filesystem does not have to be used. Bound both ends the way __bitmap_ptr() lays them out. The three branches differ, so a single sum of the two lengths would still admit a small overflow when cp_payload is non-zero.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: assoc_array: don't chain a new shortcut		

CVE-2026-99567	onto an existing one When a full node's leaves and the new key all share the node's segment, assoc_array_insert_into_terminal_node() replaces the node with a shortcut to a pair of new nodes. If the node is itself the target of a shortcut, the new shortcut is installed as that shortcut's next node, and a shortcut leading to a shortcut is something nothing expects: assoc_array_walk() indexes the second one as a node, and assoc_array_apply_edit() hits the BUG_ON() in the loop that adjusts nr_leaves_on_branch, on that same insertion. Deletion gets a node there. When the subtree below a shortcut shrinks enough, assoc_array_delete() collapses it into one node below the shortcut, and the leaves it collects can all share that node's segment. The next key that shares it too takes this path. Eighteen insertions and a deletion are enough: sixteen keys that agree on bits 0-11, one that leaves them at bit 8 and is deleted again, and one more that agrees with the sixteen. Have the new shortcut replace the one in front of the node as well. assoc_array_insert_shortcut() gives it that shortcut's back pointer and parent slot, so the caller installs it where that one was, and excises the old shortcut along with the node, which takes a second entry in excised_meta[]. The new shortcut's index key comes from the new key up to the new skip level, so it covers what the old one did. Found by a userspace test doing random insertions and deletions on keys that share long prefixes; a targeted case builds a shortcut, deletes down to the collapse and inserts one more key.	7.8	More Details
CVE-2026-84087	IBM Guardium Data Protection 12.2 is vulnerable to local privilege escalation via the SUID-root podman_wrapper binary on the Collector appliance. A local attacker with low-privileged access to the Collector can exploit insufficient argument validation in the SUID binary to execute arbitrary commands as root, resulting in full compromise of the Collector appliance.	7.8	More Details
CVE-2026-126826	A maliciously crafted SBSAR file, when imported or loaded through the Substance in 3ds Max plugin in Autodesk 3ds Max, can trigger an Out-of-Bounds Write vulnerability. A malicious actor may leverage this vulnerability to cause an application crash, data corruption, or potentially execute arbitrary code in the context of the current process.	7.8	More Details
CVE-2026-132858	A flaw was found in gdm. The CreateRemoteDisplay D-Bus method fails to perform proper authorization checks when initiating a remote display session. A local attacker with access to the login screen environment can invoke this method with an arbitrary username to bypass standard automatic login verification checks. This vulnerability allows the attacker to achieve privilege escalation and start an unauthorized session as any user on the system, including root.	7.8	More Details
CVE-2026-133168	OFFIS DCMTK 3.7.1 contains a heap buffer overflow vulnerability in togglePlanarConfiguration8 and togglePlanarConfiguration16 in djcodece.cc, which compute row offsets using rows instead of columns. Attackers can supply a crafted color-by-plane DICOM image with more rows than columns that is compressed with the true-lossless JPEG codec to corrupt heap memory, causing crashes and potentially code execution.	7.8	More Details
CVE-2026-133198	In JetBrains Junie before 252.819.35, 253.819.35 code execution was possible via terminal environment settings in project files	7.8	More Details
CVE-2026-99399	In the Linux kernel, the following vulnerability has been resolved: xfs: reject remote xattr entries with an out-of-range value length xfs_attr3_leaf_verify_entry() validates a remote attribute entry's name but never bounds its on-disk value length (xfs_attr_leaf_name_remote. valuelen, a __be32). A crafted leaf with valuelen = 0x80000000 passes the verifier and the CRC. That length is later assigned into the signed int args->rmtvaluelen (xfs_attr3_leaf_getvalue), becoming negative, which slips past the signed -ERANGE check in xfs_attr_copy_value(); a getxattr() with a small buffer then memcpy()s a full remote block into the small kvalue buffer (xfs_attr_rmtval_copyout) -- a heap out-of-bounds write with attacker-controlled content, from an unprivileged getxattr(2) on a mounted crafted image. Reject remote entries whose value length exceeds XFS_XATTR_SIZE_MAX in the leaf verifier, so the malicious block is rejected at read time.	7.8	More Details
CVE-2026-99728	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: fix list corruption when deferred_split_scan() races with deferred_split_folio() Since the deferred split queue became a list_lru, the isolate callback moves the folios it took a reference on to an on-stack list, linked through folio->_deferred_list, and deferred_split_scan() takes them off that list with list_del_init() without holding the list_lru lock. deferred_split_folio() does not know about the on-stack list. It only holds the list_lru lock when it looks at the same list_head, and it can run for such a folio: a THP is queued when it is mapped, as a candidate for the underused check, and is queued "again" when a part of it gets unmapped: CPU0 deferred_split_scan() CPU1 deferred_split_folio() list_del_init(&folio->_deferred_list) next->prev = prev prev->next = next entry->next = entry list_lru_lock_irqsave() folio_set_partially_mapped() __list_lru_add() list_empty(item) -> true list_add_tail(item, &l->list) item->next = &l->list item->prev = tail entry->prev = entry The folio now sits on the list_lru with ->prev pointing to itself. The next removal leaves the neighbours pointing to it, and they still do so after the folio has been freed. With CONFIG_DEBUG_LIST: list_del corruption. prev->next should be ffd4000000270090, but was ffd4000000278090. (prev=ffd4000000270090) Fix it by grabbing the lock around list_del_init() so deferred_split_folio() cannot race with it.	7.8	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: ALSA: hwdep: Serialize DSP load requests The check of the loaded DSP indices and the call of the driver's dsp_load callback in snd_hwdep_dsp_load() are not serialized. A hwdep file descriptor can be shared among threads, or inherited across fork, so two SNDRV_HWDEP_IOCTL_DSP_LOAD requests for the same index can both observe the index as unloaded and enter the driver's callback concurrently, even though the device itself is opened exclusively. This breaks the	7.8	More Details

99317	assumption behind hw->dsp_loaded that a firmware index is loaded only once. For example, snd-usb-usx2y hits WARN_ON(usx2y->as04.buffer) in usx2y_async_seq04_init() and may allocate the US-428 sequence URBs twice. Serialize the check, the callback invocation and the bookkeeping with a per-hwdep mutex, so that the second request returns -EBUSY.		
CVE-2026-131959	GNU Emacs 31.1 contains a code injection vulnerability in its handling of directory-local variables. A crafted .dir-locals.el can supply a (REGEXP FUNCTION t) entry in auto-mode-alist; when a matching file is opened, Emacs calls FUNCTION before set-auto-mode--dir-local-valid-p validates it, executing arbitrary Emacs Lisp with the privileges of the Emacs process. (Emacs 30.2 is unaffected.)	7.8	More Details
CVE-2026-131871	WLA DX 9.7 through 10.7 contains a stack-based buffer overflow vulnerability in parse_data_blocks() in wlalink/analyze.c that copies section namespaces without length checks. Attackers can supply a crafted object or library file with a namespace longer than 255 bytes or unterminated to overflow the 256-byte stack buffer when linked.	7.8	More Details
CVE-2026-132902	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 access keys of parent project AWS connections could be disclosed	7.7	More Details
CVE-2026-125829	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Rymera Web Co Product Feed PRO for WooCommerce woo-product-feed-pro allows Path Traversal.This issue affects Product Feed PRO for WooCommerce: from n/a through 13.5.8.	7.7	More Details
CVE-2026-133084	In JetBrains TeamCity before 2026.2.2, 2026.1.6 vCS root credentials could be exposed via change viewer links	7.7	More Details
CVE-2026-133235	Tandoor Recipes before 2.7.4 contains a server-side request forgery vulnerability that allows space admins to reach internal hosts because Nextcloud.get_client builds WebDAV clients from Storage.url without calling safe_request. Attackers can create a NEXTCLOUD storage via POST /api/storage/ pointing at internal hosts, import a file, then call GET /api/get_recipe_file/<id>/ to read internal service or cloud metadata responses.	7.7	More Details
CVE-2026-131266	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings edit permission could cause Splunk SOAR to send a stored Elasticsearch password to a host that the user controls. The vulnerability is possible because Splunk SOAR reuses a stored Elasticsearch password for a changed Search Settings endpoint without requiring a replacement password.	7.7	More Details
CVE-2026-132901	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 improper authorization allowed publishing packages to NuGet feeds of other projects	7.7	More Details
CVE-2026-132894	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 an SSRF in AWS connection settings could expose server credentials	7.7	More Details
CVE-2026-132893	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 an SSRF via pipeline repository branch listing was possible	7.7	More Details
CVE-2026-132240	In JetBrains TeamCity before 2026.1.3, 2025.11.7 XXE via Maven settings file upload was possible	7.7	More Details
CVE-2026-132897	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 improper authorization allowed access to credentials of other projects via build configuration settings	7.7	More Details
CVE-2026-132238	In JetBrains TeamCity before 2026.1.3, 2025.11.8 disclosure of password parameter values to project viewers was possible	7.7	More Details
CVE-2026-133094	In JetBrains TeamCity before 2026.2.2 missing authorization allowed project developers to read arbitrary HashiCorp Vault secrets	7.7	More Details
CVE-2026-133025	In JetBrains TeamCity before 2026.2.1 missing authorization allowed project viewers to download project settings	7.7	More Details
CVE-2026-130782	Incus before 7.6.0 contains an improper authorization vulnerability that allows restricted and view-only clients to read default project resources by requesting the default project directly. Attackers knowing resource names can read profiles, images, network zone peer keys and stored credentials, and export custom volume or bucket backups to obtain full contents or S3 admin keys.	7.7	More Details

CVE-2026-130787	Frappe ERPNext before 16.35.0 contains a missing authorization vulnerability in financial reporting endpoints that allows low-privilege authenticated users to generate reports outside their permitted scope. Attackers can send crafted requests to affected reporting endpoints to read other companies' tax and payment data, including personal and corporate identifiers and financial totals.	7.7	More Details
CVE-2026-68903	LLMKube is a Kubernetes operator for self-hosted large language model inference across heterogeneous hardware. Prior to 0.9.0, the ModelReconciler accepts arbitrary absolute paths and file:// URIs in Model.spec.source without restricting them to approved roots. A namespace user who can create or update Model and InferenceService resources can cause copyLocalModel to open an attacker-selected file in the controller pod or cause a HostPathVolumeSource to mount an attacker-selected node file when modelCache.enabled is enabled, which is the default. The selected file is copied into model storage and can expose controller service-account credentials or other node secrets outside the attacker's namespace. This issue is fixed in version 0.9.0.	7.7	More Details
CVE-2026-131801	Tanium addressed an improper input validation vulnerability in Threat Response.	7.7	More Details
CVE-2026-132918	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on build and queue pages was possible	7.6	More Details
CVE-2026-114585	IBM Engineering AI Hub 1.0.0 through 1.4.0 could allow a remote authenticated attacker to obtain sensitive information due to a server-side request forgery.	7.6	More Details
CVE-2026-130764	Obot from 0.8.0 before 0.27.0 contains a server-side request forgery vulnerability that allows Power User accounts to reach internal hosts by advertising a malicious OAuth token endpoint in remote MCP server metadata. Attackers can direct token exchange or refresh requests, including redirected GET requests, to cloud metadata services or private-network hosts and read error response bodies.	7.6	More Details
CVE-2026-132899	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the NPM Registry build feature was possible	7.6	More Details
CVE-2026-131179	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could add attacker-controlled Search Processing Language (SPL) to a shared dashboard. When another user opens the dashboard in Splunk Web, Splunk Enterprise could run the attacker-controlled SPL using the permissions of that dashboard viewer. The attacker-controlled SPL could expose stored credentials, modify indexed data, and reduce availability. The vulnerability is possible because Splunk Enterprise does not neutralize special elements in search filter values before it builds optimized search jobs. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The attacker should not be able to exploit the vulnerability at will.	7.6	More Details
CVE-2026-133085	In JetBrains TeamCity before 2026.2.2, 2026.1.6 an XXE was possible in the TFS helper process	7.6	More Details
CVE-2026-132243	In JetBrains TeamCity before 2026.1.3, 2025.11.7 stored XSS on the Docker Container Info tab was possible	7.6	More Details
CVE-2026-132242	In JetBrains TeamCity before 2026.1.3, 2025.11.7 stored XSS in the Promote build dialog was possible	7.6	More Details
CVE-2026-132239	In JetBrains TeamCity before 2026.1.3, 2025.11.7 missing authorization in the Kubernetes integration allowed project viewers to use stored credentials	7.6	More Details
CVE-2026-131175	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user who holds the "power" Splunk role could create and share a metric alert that causes Analytics Workspace to run a malicious Search Processing Language (SPL) when another authenticated user opens a link to the alert. The malicious SPL could expose data, modify data, or disrupt services available in the affected session. The vulnerability is possible because Analytics Workspace does not fully neutralize metric alert names before including the names in SPL. The vulnerability requires the attacker to phish the user by tricking them into opening a link that the attacker creates. The user who holds the "power" Splunk role should not be able to exploit the vulnerability at will. For more information see Alerts in the Analytics Workspace (https://help.splunk.com/en/splunk-enterprise/create-dashboards-and-reports/analytics-workspace/10.4/using-the-analytics-workspace/alerts-in-the-analytics-workspace) and Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities)	7.6	More Details

	in the Splunk documentation.		
CVE-2026-131178	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could craft a shared Classic Dashboard that runs arbitrary Search Processing Language (SPL) commands for another authenticated user. Exploitation requires the dashboard viewer to open the shared dashboard. The injected commands could access data available to the dashboard viewer, modify data available to the dashboard viewer, and partially disrupt Splunk Enterprise. For more information see Add sparklines to search results (https://help.splunk.com/en/splunk-enterprise/search/search-manual/10.6/calculate-statistics/add-sparklines-to-search-results), Configure dashboard permissions (https://help.splunk.com/en/splunk-enterprise/create-dashboards-and-reports/simple-xml-dashboards/10.6/manage-and-share-dashboards/configure-dashboard-permissions), and SPL safeguards for risky commands (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/best-practices-for-splunk-platform-security/spl-safeguards-for-risky-commands) in the Splunk documentation.	7.6	More Details
CVE-2026-132905	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the AWS Credentials build feature was possible	7.6	More Details
CVE-2026-131243	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to view an affected container could query the Representational State Transfer (REST) API for approvals to retrieve a live external prompt link belonging to another user. Using this link, the user could read the approval prompt without authentication and submit a response, potentially disrupting or altering the automated workflow governed by that approval. The vulnerability is possible because the REST API for approvals exposes the live external prompt link without verifying that the requester owns the approval.	7.6	More Details
CVE-2026-131182	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could store malicious Search Processing Language (SPL) in a shared dashboard. The malicious SPL could run using the permissions of the user who opened the dashboard and could access or modify data available to that user and disrupt Splunk Enterprise availability. The vulnerability is possible because Splunk Enterprise does not preserve quoted search argument boundaries when it processes a dashboard search. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The attacker, a user who holds the "power" Splunk role, should not be able to exploit the vulnerability at will.	7.6	More Details
CVE-2026-75522	OpenID4VP is an open-source credential parsing library for SD-JWT VC and mDOC/mDL formats. Prior to 0.8.1, src/trust/ChainBuilder.ts ChainBuilder.tryBuild() terminates chain construction when an attacker-controlled certificate Subject Distinguished Name equals a trusted anchor subject, without verifying cryptographic closure to the anchor public key. A remote unauthenticated attacker can supply a self-signed leaf or intermediate that reuses the trusted Subject Distinguished Name and have credentials signed by the attacker key accepted as issued by the trusted authority. The public src/trust/TrustEvaluator.ts TrustEvaluator.evaluate() path compounds the error by matching a terminus through serialNumber and falling back to anchors[0], which can attribute the forged chain to genuine trusted-authority metadata. This issue is fixed in version 0.8.1.	7.5	More Details
CVE-2026-75612	yt-mcp is a YouTrack MCP server for Claude Code. Prior to 1.18.1, add_attachment in src/yt_mcp/tools/attachments.py accepts a caller-controlled file_path, checks only that the file exists, reads it without directory confinement, and uploads the bytes to a YouTrack issue. A compromised MCP client or an LLM steered by indirect prompt injection can read any file accessible to the server process and exfiltrate secrets through an issue the caller can access. The same release also permits get_project_health to interpolate a caller-controlled project value into snapshot filenames, allowing parent-directory traversal to write JSON files outside the snapshots directory. This issue is fixed in version 1.18.1.	7.5	More Details
CVE-2026-76818	Axon is an agent-agnostic context engine that gives AI coding agents continuous memory across sessions, tools, and machines. Prior to 0.1.1, rtk-install in src/axon/context/rtk_bootstrap.py uses _http_download() to fetch a prebuilt rtkx release archive and _extract_binary() to mark it executable, after which compress_text_with_rtk() invokes it without verifying a checksum, signature, or build attestation. If the sammyjdev/rtkx release pipeline or maintainer release credentials are compromised, an attacker can supply a modified binary that executes with the Axon process privileges when a user installs or automatically bootstraps rtkx. This issue is fixed in version 0.1.1.	7.5	More Details
CVE-2026-132910	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 unauthenticated DoS via the agent protocol was possible	7.5	More Details
CVE-2026-127921	The AppPresser – Mobile App Framework plugin for WordPress is vulnerable to generic SQL Injection via the 'appp_cam_file[name]' parameter in all versions up to, and including, 5.0.0 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. Although the AJAX handler is registered for unauthenticated users, the only authorization gate is a nonce check, and the required nonce is only emitted to logged-in users viewing a page containing the [app-camera] shortcode, making Subscriber-level access the	7.5	More Details

	effective minimum requirement for exploitation.		
CVE-2026-121852	The AppPresser WordPress plugin before 5.0.0 does not properly escape the contents of its debug log when displaying it in the admin dashboard, and decodes escape sequences after sanitising input, allowing unauthenticated users to perform Stored XSS attacks against administrators when debug logging is enabled.	7.5	More Details
CVE-2026-129030	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in Mamunur Rashid ShopBuilder – Elementor WooCommerce Builder Addons shopbuilder allows Retrieve Embedded Sensitive Data.This issue affects ShopBuilder – Elementor WooCommerce Builder Addons: from n/a through 3.4.3.	7.5	More Details
CVE-2026-132908	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via Mercurial VCS root settings	7.5	More Details
CVE-2026-129013	Missing Authorization vulnerability in Deepen Bajracharya Video Conferencing with Zoom video-conferencing-with-zoom-api allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Video Conferencing with Zoom: from n/a through 4.6.12.	7.5	More Details
CVE-2026-14491	The WP Duplicate – WordPress Migration Plugin plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 1.1.11 via the 'file_name' parameter. This makes it possible for unauthenticated attackers to read the contents of arbitrary files on the server, which can contain sensitive information. The authentication bypass requires submitting the JSON boolean `true` as the `prod_key_random_id` value, which coerces against any non-empty stored secret so that the loose `!=` comparison evaluates to false, skipping the rejection branch entirely without knowledge of the secret.	7.5	More Details
CVE-2026-85776	The WP Duplicate – WordPress Migration Plugin plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 1.1.11. This is due to a missing capability check and missing nonce check on the start_zip_creation, start_zip_download, local_sync_get_init_root_files, and local_sync_get_files_by_key AJAX actions in the start_zip_creation_dev() and sibling functions. This makes it possible for authenticated attackers, with subscriber-level access and above, to overwrite the plugin's prod_site_url setting so it points at an attacker-controlled host and enumerate arbitrary server directories, which when combined with a subsequent administrator-triggered sync results in an unsigned attacker-supplied ZIP archive being fetched, extracted via a WordPress-independent bridge deployed under the web root, resulting in remote code execution on the affected site.	7.5	More Details
CVE-2026-121482	Issue summary: Verifying a maliciously crafted CMS or PKCS#7 SignedData message with CRL checking enabled may consume excessive CPU time, growing with the square of the message size. Impact summary: An attacker who can submit a CMS or PKCS#7 SignedData message to an application that verifies it with CRL checking enabled can occupy a CPU core for seconds to minutes, causing Denial of Service. Applications that leave CRL checking disabled are not affected. CWE: CWE-407: Inefficient Algorithmic Complexity Description: With CRL checking enabled, OpenSSL compares each CRL distribution point of a certificate against the issuing distribution point of every candidate CRL. For a nameRelativeToCRLIssuer distribution point, the full name is the CRL issuer name with the relative name appended. Since the change that deferred this computation, crl_crldp_check() rebuilds that full name for each candidate CRL and frees it again. CMS_verify() and PKCS7_verify() take both the distribution points and the candidate CRLs from the message, and this scoring runs before any signature is checked. An application is exposed only if it sets X509_V_FLAG_CRL_CHECK on the store and does not pass CMS_NOCRL or PKCS7_NOCRL; CRL checking is off by default. In TLS the CRLs come from the local store, so a peer controls only the distribution point count. The cost grows with the product of the distribution point count, the CRL count and the issuer name length, all chosen by the message author; a few hundred kilobytes can occupy a CPU core for tens of seconds. Verification still fails as it should, with no crash, no incorrect result and no memory growth. Requiring CRL checking is what limits the severity. FIPS impact: no The affected code is outside the FIPS module boundary.	7.5	More Details
CVE-2026-132900	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 secure parameter values were disclosed to project viewers in pipelines	7.5	More Details
CVE-2026-132895	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via Subversion VCS root settings	7.5	More Details
CVE-2026-76088	PGPainless is a simple-to-use OpenPGP API based on Bouncy Castle. Prior to 1.7.8 and from 2.0.0 until 2.0.4, pgpainless-sop/src/main/kotlin/org/pgpainless/sop/InlineDetachImpl.kt InlineDetachImpl and pgpainless-core/src/main/kotlin/org/pgpainless/decryption_verification/MessageInspector.kt MessageInspector manually descend PGPCompressedData packets without incrementing or enforcing MAX_LAYER_DEPTH. A remote unauthenticated attacker can supply an untrusted OpenPGP message containing deeply nested compressed packet structures to a downstream application using inline-detach or message inspection, bypass the intended 16-layer limit, and force continued decompression and packet processing. The resulting resource consumption can block processing threads, exhaust heap memory, and crash the consuming application. This issue is fixed in versions 1.7.8 and 2.0.4.	7.5	More Details
CVE-	Grav from 2.0.0 before 2.2.18 contains an information disclosure vulnerability that allows publishers and		

2026-130740	read-only user managers to read password hashes, 2FA secrets and reset tokens through unfiltered account serialization. Attackers can embed a Flex form for the user-accounts directory in a page or query the Flex API with admin.users.read to obtain super admin credentials and take over the account.	7.5	More Details
CVE-2026-59975	SharpIppNext is a .NET Standard library for building Internet Printing Protocol clients and servers. Prior to 4.1.0, server-side SharpIppNext implementations using SharpIppServer.ReceiveRequestAsync or SharpIppServer.ReceiveRawRequestAsync pass incoming data to IppProtocol.ReadIppRequestAsync in SharpIpp/Protocol/IppProtocol.cs, where ReadDocumentStream is enabled by default and the remaining document body is copied with CopyToAsync into a MemoryStream without a size limit. An unauthenticated remote client can send an arbitrarily large IPP document body, exhausting process memory and causing an out-of-memory failure or service termination. Client-only implementations using SharpIppClient are not affected. This issue is fixed in version 4.1.0.	7.5	More Details
CVE-2026-76417	A vulnerability in Object Linking and Embedding, 2nd Generation (OLE2) file parser for ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to logic error when parsing specifically formatted files with the OLE2 file parser. An attacker could exploit this vulnerability by by sending a crafted Microsoft Compound File Binary (CFB) file to be scanned. A successful exploit could allow the attacker to allocate all available memory, causing the ClamAV scanning process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-130686	vLLM from 0.6.0 before 0.33.0 lacks a request-wide media budget, allowing remote inference clients to exhaust frontend resources by submitting many remote image URLs in one request. Attackers can send a compact request to /v1/messages or /v1/responses that forces downloading and decoding of every image before prompt-length admission, exhausting memory, CPU, and connections.	7.5	More Details
CVE-2026-126336	Description: When BIG-IP APM is provisioned and a virtual server is configured with an APM Access profile and a SAML Service Provider (SP) agent with encrypted assertions enabled, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate. Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	7.5	More Details
CVE-2026-126721	The Magic Login Mail or QR Code plugin for WordPress is vulnerable to Privilege Escalation via Administrator Account Takeover in versions up to, and including, 2.07. This is due to the `curpageurl()` function falling back to the raw `$_SERVER['HTTP_HOST']` value — with no validation against the site's configured origin via `home_url()` or `site_url()` — when `get_the_permalink()` returns empty or false, allowing an attacker-controlled Host header to become the URL origin of the magic-login link generated by `generate_url()` and delivered to the target via `wp_mail()`. This makes it possible for unauthenticated attackers to intercept the one-time uid and token pair sent to a targeted Administrator's email and replay them against the real site, causing `autologin_via_url()` to call `wp_set_auth_cookie()` and establish a full Administrator session on the targeted account. Exploitation requires the `[magic_login]` shortcode to render on a path where `get_the_permalink()` returns empty or false (such as a 404 page or a block-theme header template part), the web server to forward non-canonical Host headers, the attacker to know the target administrator's email address, and the administrator to click the attacker-originated magic-login link. Successful exploitation gives the attacker control of a distinct, already-privileged Administrator account; it does not promote the attacker's original account. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	7.5	More Details
CVE-2026-126822	Description: When BIG-IP APM is provisioned, an Access Profile is configured as a SAML SP (Service Provider) and is applied to a virtual server, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate. Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	7.5	More Details
CVE-2026-130685	vLLM from 0.28.0 before 0.33.0 contains an improper input validation vulnerability that allows unauthenticated attackers to crash the shared EngineCore by sending out-of-range prompt token ids when served with --skip-tokenizer-init. Attackers can submit a large positive token id to /v1/completions, /v1/embeddings, or /pooling, bypassing the tokenizer-gated vocab check and causing a fatal embedding index error that denies service to all tenants.	7.5	More Details
CVE-2026-130677	Craft CMS from 5.0.0 before 5.11.4 contains an improper session management vulnerability that fails to invalidate existing sessions when 2FA is enabled and allows disabling 2FA without re-authentication. Attackers holding a hijacked or stale session can retain full account access and silently remove 2FA without the password, authenticator code, or recovery code.	7.5	More Details
CVE-2026-132245	In JetBrains TeamCity before 2026.1.3, 2025.11.7 DoS via agent command results was possible	7.5	More Details
	A vulnerability in the phishcheck functionality of ClamAV could allow an unauthenticated, remote attacker to		

CVE-2026-76411	cause a DoS condition on an affected device. This vulnerability is due to improper allocation of memory during processing of URLs. An attacker could exploit this vulnerability by sending a URL to be processed by the software. A successful exploit could allow the attacker to cause the ClamAV scanning process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76414	A vulnerability in Microsoft Compiled HTML Help (CHM) file scanning of ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to a logic error when scanning the CHM file and following the internal file pointers. An attacker could exploit this vulnerability by sending a crafted CHM file for scanning. A successful exploit could allow the attacker to cause the ClamAV decode process to terminate due to an infinite loop, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76415	A vulnerability in the Yoda'a Crytper (yC) Portable Executable (PE) unpacker of ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. . This vulnerability is due to lack of proper error checking for negative values when processing the PE file to be unpacked. An attacker could exploit this vulnerability by sending a PE file to be unpacked. A successful exploit could allow the attacker to cause the ClamAV scanning process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76416	A vulnerability in the InstallShield file parser of ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to incomplete error checking when the InstallShield File Parser is performing a scan. An attacker could exploit this vulnerability by sending a crafted InstallShield file to be scanned. A successful exploit could allow the attacker to cause the ClamAV scanning process to terminate due to an infinite loop, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76421	A vulnerability in the LZX data algorithm for Microsoft Compiled HTML Help (CHM) file scanning in ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to improper boundary checks when applying the LZX data algorithm to CHM files during the decoder stage, which may result in buffer overflow. An attacker could exploit this vulnerability by sending a crafted CHM file for scanning. A successful exploit could allow the attacker to cause the ClamAV scanning process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-99678	In the Linux kernel, the following vulnerability has been resolved: xfs: cap the block count of xfs_iomap_write_direct() at one extent xfs_iomap_write_direct() reserves blocks for the whole count in an unsigned int. For a pNFS layout the client chooses that count, and an RW LAYOUTGET for 16 TiB wraps the reservation and shuts XFS down. xfs_bmapi_write() maps at most XFS_MAX_BMBT_EXTLEN blocks here, so cap the count at that. nfsd asks again for the rest. The other callers get the same mapping as before.	7.5	More Details
CVE-2026-76422	A vulnerability in the debug logging functionality of ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to improper handling of logging values. An attacker could exploit this vulnerability by submitting a file to be scanned by the software that has debug logging enabled. A successful exploit could allow the attacker to cause the ClamAV scanning process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76430	A vulnerability in the Ultimate Packer for eXecutables (UPX) unpacker of ClamAV could allow an unauthenticated, remote attacker to cause a DoS condition on an affected device. This vulnerability is due to improper length checks, when unpacking the UPX files, which may result in an out-of-bounds read of heap memory. An attacker could exploit this vulnerability by submitting a crafted UPX for decode. A successful exploit could allow the attacker to cause the ClamAV decode process to terminate, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76515	A vulnerability in the Object Linking and Embedding, 2nd Generation (OLE2) file parser for ClamAV could allow an unauthenticated, remote attacker to exceed the configured MaxScanTime and cause a DoS condition on an affected device. This vulnerability is due to inefficiencies when parsing specifically formatted files with the OLE2 file parser. An attacker could exploit this vulnerability by by sending a crafted Microsoft Compound File Binary (CFB) file to be scanned. A successful exploit could allow the attacker to effectively cause the ClamAV OLE2 parser to get into an infinite loop, resulting in a DoS condition on the affected software.	7.5	More Details
CVE-2026-76516	A vulnerability in the eXtensible ARchive (XAR) GZIP file parser for ClamAV could allow an unauthenticated, remote attacker to exceed the configured MaxFileSize and cause a DoS condition on an affected device. This vulnerability is due to inefficiencies when parsing specifically formatted files with the XAR GZIP file parser. An attacker could exploit this vulnerability by by sending a crafted XAR GZIP file to be scanned. A successful exploit could allow the attacker to effectively cause the ClamAV to exceed the MaxFileSize and potentially get caught in an infinite loop, causing a DoS condition.	7.5	More Details
CVE-2026-128662	PHPGurukul Employee Leaves Management System V2.1 contains an information disclosure vulnerability because database backup files are stored within the publicly accessible web root. An unauthenticated remote attacker can download the SQL backup files and retrieve sensitive database contents, including employee personally identifiable information and password hashes.	7.5	More Details
CVE-2026-	Fleet before 4.86.2 contains a denial of service vulnerability that allows unauthenticated attackers to exhaust server memory via unbounded BER recursion in PKCS#7 parsing. Attackers can send deeply nested	7.5	More

130760	indefinite-length structures to Apple MDM check-in, enrollment, SCEP, or OTA endpoints to crash the entire Fleet instance.		Details
CVE-2026-130761	Fleet before 4.92.0 contains a denial of service vulnerability that allows enrolled hosts to trigger unbounded per-name database lookups by submitting osquery result logs. Attackers with a valid host node key can send one request to /api/osquery/log containing many distinct query names, exhausting database connections and read replica capacity.	7.5	More Details
CVE-2026-71481	ALTCHA PHP is a lightweight PHP library for creating and verifying ALTCHA challenges. From 2.0.0 until 2.0.3, the verifySolution method skips HMAC signature verification when a challenge payload omits the signature field, even when hmacSignatureSecret is configured. An attacker can remove the signature, supply attacker-selected low-difficulty challenge parameters, solve the resulting proof of work, and submit a payload that verifies successfully, bypassing the CAPTCHA control. This issue is fixed in version 2.0.3.	7.5	More Details
CVE-2026-73864	install-artifact-from-github installs prebuilt native artifacts hosted by GitHub without compiling them locally. Prior to 1.7.0, bin/install-from-cache.js downloads a native .node artifact, allows RE2_DOWNLOAD_MIRROR to select a source, follows HTTP redirects including HTTPS-to-HTTP downgrades, and writes the decompressed response to the artifact path without cryptographic integrity verification before writing or loading it. An attacker who controls a configured mirror, replaces a release asset, or can modify plaintext transport can supply native code that executes when the consumer loads the artifact. Version 1.7.0 adds SHA-256 verification for default-GitHub-source downloads when the consuming package publishes an artifactHashes bag; configured mirrors and packages without a bag remain outside that verification.	7.5	More Details
CVE-2026-127222	SGLang through 0.5.23, fixed in commit 7b5fb39, contains a missing authorization vulnerability in the POST /load_lora_adapter_from_tensors HTTP endpoint, which lacks the @auth_level(AuthLevel.ADMIN_OPTIONAL) decorator applied to its sibling admin endpoints, allowing unauthenticated network-reachable attackers to invoke it without an API key. Attackers can load a LoRA adapter from attacker-supplied serialized tensors to manipulate the served model, and when combined with a separately documented deserialization bypass, achieve unauthenticated remote code execution as root.	7.5	More Details
CVE-2026-132140	A security vulnerability has been detected in DrayTek Vigor2960 up to 1.5.1.6. Affected by this vulnerability is an unknown functionality of the file mainfunction.cgi of the component Apmcfg Upload. The manipulation leads to null pointer dereference. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. Upgrading to version 1.5.1.7 addresses this issue. Upgrading the affected component is advised. The vendor recommends: "1.5.1.6 isn't the latest version of 2960, the latest version fixes many security issues. And we have another one in our release plan which should solve the issues reported."	7.5	More Details
CVE-2026-132911	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via build access codes disclosed to project users	7.5	More Details
CVE-2026-133076	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via local repository URLs in Subversion VCS roots	7.5	More Details
CVE-2026-115677	Buffer Copy without Checking Size of Input ('Classic Buffer Overflow') vulnerability in AirTies Kablosuz İletişim Sanayi ve Dış Ticaret A.Ş. Air 5444 allows Overflow Buffers. This issue affects Air 5444: through 2026-11-19. NOTE: The vendor was contacted and it was learned that the product is not supported.	7.5	More Details
CVE-2026-133062	A flaw was found in Moodle's external database enrolment plugin. Insufficient sanitization of data passed to queries against the external database allows an authenticated user with course editing permissions to perform a Structured Query Language (SQL) injection attack. When this non-default plugin is enabled, an attacker can exploit this issue to view, modify, or delete sensitive records within the connected external database.	7.5	More Details
CVE-2026-118436	Allocation of resources without limits or throttling vulnerability in Apache Tomcat in the CGI Servlet when handling POST bodies can result in a DoS for some inputs. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.0.M1 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	7.5	More Details
CVE-2026-118233	Allocation of resources without limits or throttling vulnerability in Apache Tomcat during HTTP upgrade in some circumstances lead to possible DoS. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.41 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.61 through 8.5.100. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	7.5	More Details
CVE-2026-	Missing release of resource after effective lifetime vulnerability in Apache Tomcat's HTTP/2 initial connection handling leads to DoS. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.20 through 9.0.123. Users are recommended to upgrade to version 11.0.29,	7.5	More Details

116830	10.1.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released		
CVE-2026-116600	Missing release of memory after effective lifetime vulnerability in Apache Tomcat via WebSocket. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.22 through 9.0.123 The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.43 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.0.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	7.5	More Details
CVE-2026-131857	Supergateway 4.2.0 through 4.2.2 contains a denial of service vulnerability in stdio-to-SSE mode that allows unauthenticated attackers to crash the gateway via prototype-inherited sessionId values. Attackers can send a POST message request with sessionId set to toString or constructor, triggering an unhandled TypeError that terminates every attached MCP session.	7.5	More Details
CVE-2026-99451	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: unlink TWT flow if the MCU rejects the agreement mt7996_mac_add_twt_setup() links a new flow into dev->twt_list before asking the MCU to add the agreement, but exits without unlinking the flow if the request fails. The flow-ID and table masks are only committed after MCU success, so the still-linked flow is invisible to the mask-gated teardown paths. If the station retries the setup, the same slot is cleared and its embedded list node reinitialized while it is still linked, corrupting dev->twt_list. If the station disassociates instead, teardown skips the flow and station storage can be freed while the node remains linked. A later scheduling or debugfs twt_stats walk can then dereference the freed flow, resulting in a use-after-free. Unlink the flow when the MCU rejects the agreement. This restores the list to its pre-request state while leaving the unchanged mask and agreement accounting untouched.	7.5	More Details
CVE-2026-133201	A vulnerability was determined in smartalock wireguard-lwip up to cfba91666590e86bd67842c71cedf9a423be823a. Impacted is the function mul of the file src/crypto/refc/x25519.c of the component X25519 Ladder. Executing a manipulation can lead to out-of-bounds write. The attack can be executed remotely. This product implements a rolling release for ongoing delivery, which means version information for affected or updated releases is unavailable. The project was informed of the problem early through an issue report but has not responded yet.	7.5	More Details
CVE-2026-131299	In Splunk App for Stream versions below 8.1.7, an unauthenticated user who can reach the Splunk App for Stream Representational State Transfer (REST) API on the management port (8089) could upload an Extensible Markup Language (XML) vocabulary file, which lets the vocabulary upload endpoint read local files accessible to the Splunk process or send outbound requests. The vulnerability is possible because the vocabulary upload endpoint parses uploaded XML without an application-level policy that blocks external entity resolution. For more information see About Splunk Stream (https://help.splunk.com/en/splunk-enterprise/collect-stream-data/install-and-configure-splunk-stream/8.1/introduction/about-splunk-stream) and restmap.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/9.4/configuration-file-reference/9.4.3-configuration-file-reference/restmap.conf) in the Splunk documentation.	7.5	More Details
CVE-2026-126040	The Hide My WP Ghost WordPress plugin before 7.0.19 only enforces two-factor authentication on the interactive login flow and not on other authentication paths, allowing an attacker who already has a user's valid password to authenticate as that user without completing the second factor.	7.5	More Details
CVE-2026-99549	In the Linux kernel, the following vulnerability has been resolved: smb: client: do not pass a NULL buf to is_oplock_break() receive_encrypted_read() sets *num_mids = 1 and never writes the caller's bufs[]. Offload returns -1 with *num_mids = 0, so the demultiplex thread skips the frame. The inline decrypt path returns 0 even when smb2_find_mid() fails. cifs_demultiplex_thread() then treats mids[i] == NULL as an oplock break and calls is_oplock_break(bufs[i], server) with bufs[i] still NULL. smb2_is_valid_oplock_break() reads hdr.Command at offset 12 and oopses. A large encrypted READ whose MessageId does not match is enough to panic cifsd. Oops: general protection fault, probably for non-canonical address 0xdfffffc0000000001: 0000 [#1] SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f] CPU: 0 PID: 100 Comm: cifsd RIP: smb2_is_valid_oplock_break+0xc8/0x2270 [cifs] cifs_demultiplex_thread Kernel panic - not syncing: Fatal exception Return -1 when inline decrypt finds no mid, and skip is_oplock_break() plus the unknown-frame dump when bufs[i] is NULL. receive_encrypted_standard() already assigns bufs[*num_mids].	7.5	More Details
CVE-2026-131255	In Splunk SOAR versions below 8.8.0, an unauthenticated user could gain administrator privileges through a user account that receives an unintended role. A user that holds a role with permission to create or modify users and assign roles must assign the target account a role with a numeric name through the user-management Representational State Transfer (REST) API. The vulnerability is possible because the API resolves a numeric role name as a role identifier instead of preserving the requested role name.	7.5	More Details
	In the Linux kernel, the following vulnerability has been resolved: smb: smbdirect: release failed pending sockets of a listener A child of a listener only left the listen.pending or listen.ready list via smbdirect_socket_accept() or when the listener was destroyed. So every connection that failed before it was accepted (negotiate timeout, peer disconnect, invalid negotiate request, ...) stayed on the list and counted against the backlog until the listener was destroyed. After about backlog (10 for ksmbd) such connections smbdirect_listen_connect_request() rejected every new connection with -EBUSY, so an unauthenticated peer		

CVE-2026-99563	was able to permanently stop ksmbd from accepting SMB-Direct connections. Now a failed child moves itself to a new listen.orphaned list at the end of smbdirect_socket_cleanup_work(), after the disconnect was started, and queues listen.purge_orphaned_work, which releases it with smbdirect_socket_release(), without holding the listener's rdma_lock_handler() lock, in the same way smbdirect_socket_destroy() does for the remaining children. Orphaned sockets still count against the backlog until they are released. sc->accept.listener is now only changed under listener->listen.lock and whoever clears it is responsible for releasing the socket: smbdirect_socket_accept(), purge_orphaned_work or the destroy of the listener. In order to make that work: - a struct smbdirect_socket that is a listener is now freed via kfree_rcu(), so that a child can dereference sc->accept.listener under rcu_read_lock(), even if the listener is released concurrently. - smbdirect_listen_connect_request() adds nsc to the listen.pending list and sets nsc->accept.listener before it calls smbdirect_accept_connect_request(), and always returns 0 so the rdma_cm core keeps the connection id and nsc owns it. On a non-zero return it schedules nsc's teardown with smbdirect_socket_schedule_cleanup(), which orphans it off the listener so that purge_orphaned_work releases it; nsc destroys its own connection id with rdma_destroy_id() in that teardown. The teardown has to be deferred because the rdma_cm core holds nsc's id_priv->handler_mutex across the handler, so nsc->rdma.cm_id must not be destroyed synchronously from here. That same handler_mutex (and the listener's) also guarantees nsc can't go away under us, so setting the listener before the call is safe and lets a failure orphan nsc. - smbdirect_accept_negotiate_recv_work() no longer moves a socket to the ready list if it already failed or no longer belongs to the listener. Every accepting socket has a listener, so if sc->accept.listener is already cleared, it no longer sends a negotiate response, as the socket is about to be released. - smbdirect_socket_destroy() of the listener disables purge_orphaned_work and also releases the orphaned sockets.	7.5	More Details
CVE-2026-99564	In the Linux kernel, the following vulnerability has been resolved: smb: smbdirect: don't wait for RDMA_CM_EVENT_DISCONNECTED in smbdirect_socket_destroy_sync() smbdirect_socket_destroy_sync() waited for RDMA_CM_EVENT_DISCONNECTED after smbdirect_socket_cleanup_work() called rdma_disconnect(). That can take very long, e.g. if the peer just disappeared, until the rdma cm gives up, or forever if the event already happened, but the status was overwritten afterwards. This is not needed: smbdirect_socket_destroy() drains the qp under rdma_lock_handler(), destroys it and calls rdma_destroy_id(), which only waits for a currently running event handler and makes sure no further events are delivered. The rest of the disconnect protocol (DREQ/DREP and timewait for IB, abrupt close for iWarp) is handled by the rdma core asynchronously. drivers/nvme/host/rdma.c also just calls rdma_disconnect() and ib_drain_qp() before rdma_destroy_id(). So smbdirect_socket_destroy() now also accepts SMBDIRECT_SOCKET_DISCONNECTING and changes the status to SMBDIRECT_SOCKET_DISCONNECTED itself under rdma_lock_handler(). We could still get RDMA_CM_EVENT_DISCONNECTED in the small windows between rdma_unlock_handler() and rdma_destroy_id(), but in that case smbdirect_connection_rdma_event_handler() is basically a no-op.	7.5	More Details
CVE-2026-133234	Stoatchat before 0.16.5 contains an unbounded JXL image decoding vulnerability in revolt_files decode_image that allows unauthenticated attackers to exhaust memory via the january /proxy endpoint. Attackers can request a small hosted JXL file declaring 12000x12000 pixels through /proxy to trigger huge allocations, OOM-killing january and denying proxy and embed service.	7.5	More Details
CVE-2026-119220	The Dokan: AI Powered WooCommerce Multivendor Marketplace Solution – Build Your Own Amazon, eBay, Etsy plugin for WordPress is vulnerable to Authorization Bypass in all versions up to, and including, 5.3.3. The vulnerability exists because the `matches_mcp_endpoint()` function applies an unanchored regular expression (`#/(?:woocommerce|dokan)/mcp(?:/|$)#`) against the entire `REQUEST_URI` including the query string, meaning any request that contains the literal substring `/dokan/mcp/` or `/woocommerce/mcp/` anywhere in its URI — including inside a query parameter value — is incorrectly treated as a trusted internal MCP context, causing `scope_product_permissions()` to unconditionally return `true` on WooCommerce's `woocommerce_rest_check_permissions` filter regardless of whether the caller is authenticated. This makes it possible for unauthenticated attackers to read full authenticated WooCommerce REST product records for any published product — including private underscore-prefixed meta, purchase notes, downloadable file URLs for paid digital products, catalog-hidden products, and edit-context representations — across the `/wc/v1`, `/wc/v2`, `/wc/v3`, and `/wc-analytics` REST namespaces. The bypass is triggered by appending a query parameter whose value contains the literal string `/dokan/mcp/` to any WooCommerce REST product endpoint (e.g., `/wp-json/wc/v3/products/{id}?x=/dokan/mcp/`); note that percent-encoding the slashes prevents the bypass, as the unencoded literal slash characters are required for the regex to match.	7.5	More Details
CVE-2026-133245	parse5 7.0.0 before 8.1.1 contains an uncaught exception vulnerability in Preprocessor#_processSurrogate that allows remote attackers to crash applications by supplying two consecutive low surrogates. Attackers can submit input such as U+DC00 U+DC00, which decodes to out-of-range code point 0x110000 and throws a RangeError from parse, parseFragment, or streaming parsers, crashing Node.js processes.	7.5	More Details
CVE-2026-128139	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.6 to 3.10.3, XMLCorpusView._read_xml_fragment accumulates 1 KiB blocks and repeatedly applies whole-buffer XML validation while a comment or CDATA terminator remains ahead. A large comment in an attacker-controlled corpus file causes quadratic CPU consumption through documented XML corpus readers across the affected range. On 3.6 and 3.9.2, a large CDATA section incurs the quadratic scan before rejection, while on 3.10.3	7.5	More Details

	both comment and CDATA documents are read successfully. This vulnerability is fixed in 3.10.4.		
CVE-2026-129222	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, DependencyGraph.tree and DependencyGraph.triples walk parsed dependency edges through unbounded Python recursion and do not use a working cycle or depth guard. A crafted CoNLL or Malt-TAB dependency file containing a sufficiently long chain can raise RecursionError when the documented traversal APIs are called, aborting the current worker. This issue is fixed in version 3.10.4.	7.5	More Details
CVE-2026-129223	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, WordNet Synset.tree and acyclic_tree traverse relation graphs through recursive utility helpers without a depth bound. Exploitation requires the host application to load an attacker-influenced custom or foreign WordNet-format lexical resource; no relation in NLTK's bundled WordNet 3.0 corpus reaches the crash threshold. An excessively deep relation graph in such a resource can exceed the Python recursion limit and terminate the traversal or request worker. This issue is fixed in version 3.10.4.	7.5	More Details
CVE-2026-129229	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.10.0 until 3.10.4, TokenSearcher.findall and Text.findall accept caller-controlled regular expressions with an unbounded number of capturing groups throughout the affected range. Starting in 3.10.3, nltk.redos.compile and RegexpTokenizer expose the same compile-time vector. The third-party regex engine compiles group-dense patterns with roughly quadratic cost before the match-time timeout begins, allowing a small unauthenticated request to exhaust CPU and deny service. This issue is fixed in version 3.10.4.	7.5	More Details
CVE-2026-131196	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user who can connect to a Splunk-to-Splunk (S2S) data forwarding receiver could cause the Splunk daemon to terminate unexpectedly. The unexpected termination interrupts data indexing and access to the management API and Splunk Web until the Splunk Enterprise process restarts. The vulnerability is possible because the S2S data forwarding receiver does not validate an optional field in forwarded event data, which results in a null pointer dereference when the field is missing. For more information see Intermediate data routing using universal and heavy forwarders (https://help.splunk.com/en/splunk-enterprise/get-started/splunk-validated-architectures/getting-data-in-forwarding-and-preprocessing/intermediate-data-routing-using-universal-and-heavy-forwarders) in the Splunk documentation.	7.5	More Details
CVE-2026-131192	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could access live event data from the indexing pipeline through the Representational State Transfer (REST) API on a standalone instance or a search head that is not an indexer cluster member. The vulnerability is possible because Splunk Enterprise accepts an empty shared key during service authentication.	7.5	More Details
CVE-2026-129500	@fastify/http-proxy is a Fastify plugin that proxies HTTP and WebSocket traffic to an upstream server. In versions from 9.0.0 up to but not including 11.6.10, when the plugin is registered with websocket support enabled, it takes every upgraded socket and dispatches the request through the Fastify router before the WebSocket library takes ownership of the socket. Node.js removes its own error listener from the socket before emitting the upgrade event, and the plugin does not attach one of its own during routing, so a socket error such as a client connection reset that arrives before the WebSocket library takes over is unhandled and terminates the Node.js process. The error window is open whenever anything asynchronous runs between dispatch and the route handler, such as an async hook on the proxied route or any unrelated async route in the application. A single unauthenticated upgrade request followed by a connection reset is enough to crash the server, and it can be repeated to keep the process down. The issue is fixed in @fastify/http-proxy 11.6.10, and users should upgrade to 11.6.10 or later.	7.5	More Details
CVE-2026-129502	@fastify/http-proxy is a Fastify plugin that proxies HTTP and WebSocket traffic to an upstream server. In versions from 9.4.0 up to but not including 11.6.10, when the plugin is registered with websocket support enabled, a WebSocket upgrade request whose path carries a URL fragment is matched by the router and answered with a successful upgrade before the upstream link is created. Building the upstream target then throws, the error is only logged, and the accepted client socket is left open with no error listener, so the next frame the client sends becomes an unhandled error and terminates the Node.js process. A request target that is not a valid URL reference reaching a parameterized prefix triggers the same path. A single unauthenticated upgrade request followed by one frame is enough to crash the server, and it can be repeated to keep the process down. The issue is fixed in @fastify/http-proxy 11.6.10, and users should upgrade to 11.6.10 or later.	7.5	More Details
CVE-2026-130157	RepoManager is a self-hosted deb/rpm package repository. Prior to 5.13.3, an attacker who controls an upstream RPM repository mirrored by a RepoManager instance can provide a package filename through the primary.xml <location href> field; the mirror passes that filename into a shell-interpreted RPM query, allowing shell metacharacters to execute commands as the www-data service account. The target must mirror an upstream repository controlled by the attacker, and the default "Check GPG signatures" option triggers the vulnerable call. The command executes before a later signature-verification failure, so a valid signed package is not required. Fixed in 5.13.3.	7.5	More Details

CVE-2026-130647	@fastify/http-proxy is a Fastify plugin that proxies HTTP and WebSocket traffic to an upstream server. In versions from 9.0.0 up to but not including 11.6.10, when the plugin is registered with websocket support enabled, a WebSocket upgrade that is pipelined behind another HTTP request on the same cleartext connection arrives while that request's response is still attached to the socket. The shared upgrade listener sees the socket is already claimed and returns early without taking the upgrade, but it does not attach an error listener to the raw socket, and Node.js has already removed its own error listener before emitting the upgrade event. A later socket error on that connection, such as a client connection reset, is therefore unhandled and terminates the Node.js process. A single unauthenticated connection that pipelines an upgrade behind a normal request and then resets is enough to crash the server, and it can be repeated to keep the process down. Upgrades over TLS are not affected. The issue is fixed in @fastify/http-proxy 11.6.10, and users should upgrade to 11.6.10 or later.	7.5	More Details
CVE-2026-131190	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could access RapidDiag process metadata and log file inventory information through the Representational State Transfer (REST) API. The vulnerability is possible because certain RapidDiag REST API endpoints accept a factory-default shared key for node authentication, which lets requests bypass Splunk user authentication. For more information see Using RapidDiag (https://help.splunk.com/en/splunk-enterprise/administer/troubleshoot/9.0/contact-splunk-support/using-rapiddiag) and Secure Splunk Enterprise services with pass4SymmKey (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.0/secure-distributed-and-clustered-splunk-environments/secure-splunk-enterprise-services-with-pass4symmkey) in the Splunk documentation.	7.5	More Details
CVE-2026-115578	Go-Guerrilla SMTP Daemon is a lightweight SMTP server written in Go. In 1.6.7 and earlier, according to the repository advisory's structured affected range, server.go defines the seven-byte cmdXCLIENT token, but the XCLIENT handler slices input[8:] after matching only that seven-byte verb, so a bare command causes an out-of-bounds panic in handleClient and terminates the daemon because that connection goroutine does not recover the panic. The issue is reachable by an unauthenticated client when the non-default xclient_on setting is enabled; no HELO command is required. The repository fix is present in version 1.6.7-1.	7.5	More Details
CVE-2026-63124	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. Prior to 0.0.23.Final, BinaryHttpParser.readFieldLine in codec-bhttp/src/main/java/io/netty/incubator/codec/bhttp/BinaryHttpParser.java treats a complete field line that ends exactly at a known field-section boundary as incomplete and returns without consuming bytes. BinaryHttpParser.readFieldSection then leaves fieldSectionLength unchanged and repeats the loop because its forward-progress checks are Java assertions, which are normally disabled in production. A remote peer that sends a known-length Binary HTTP header or trailer section ending at that boundary can make the parser spin indefinitely and pin the Netty event-loop or worker thread responsible for the channel. This issue is fixed in version 0.0.23.Final.	7.5	More Details
CVE-2026-63202	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. Prior to 0.0.23.Final, BinaryHttpParser.readFieldSection in codec-bhttp/src/main/java/io/netty/incubator/codec/bhttp/BinaryHttpParser.java loops while fieldSectionLength is nonzero and relies on disabled-by-default assertions to require BinaryHttpParser.readFieldLine to return a value and consume bytes. A corrupt field line can consume more bytes than its understated declared section length and drive the counter negative, or a truncated field line can return without consuming data and leave the parser in the same state. An unauthenticated OHTTP peer can encrypt such Binary HTTP content under the gateway public configuration, reach OHttpRequestResponseContext.ContentDecoder.decodeChunk, and permanently pin a Netty event-loop thread at full CPU before application business logic. Multiple requests can exhaust the event-loop group and deny service until restart. This issue is fixed in version 0.0.23.Final.	7.5	More Details
CVE-2026-126540	Convert::ASN1 versions before 0.37 for Perl allow CPU exhaustion via quadratic bigint loop in os2ip. Every INTEGER or ENUMERATED longer than four content octets reaches os2ip, which iterates over each content octet with $result = ($result * 256) + $octet. Each multiplication operates on a bigint that grows by one octet per step, so decoding an n-octet value takes O(n^2) bigint work. Nothing limits content length before this loop runs. An attacker who controls the decoded data can exhaust CPU with a single oversized INTEGER or ENUMERATED value.	7.5	More Details
CVE-2026-126568	Convert::ASN1 versions before 0.37 for Perl allow memory exhaustion via unbounded MSG_PEEK buffer doubling in asn_recv. asn_recv peeks the socket in a loop, doubling the buffer size on each pass until a tag and length can be decoded. An inverted boolean guard ($tb || rather than $tb &&) prevents the need-more-data exit from firing once a tag byte has been decoded. When asn_decode_length returns no value because the length is absent or incomplete, no exit fires for that case either. MSG_PEEK leaves socket data unconsumed, so each re-peek returns at once while the buffer size doubles. A single truncated BER header sent to a process calling asn_recv doubles the peek buffer until memory is exhausted. Perl's realloc failure is fatal and not catchable with eval.	7.5	More Details
CVE-2026-	Insertion of Sensitive Information into Log File vulnerability in the Apache Iggy server. If Apache Iggy 0.9.0 started without IGGY_ROOT_USERNAME and IGGY_ROOT_PASSWORD, it generated a root password and wrote it to the server log file, which is enabled by default. The system snapshot command (GetSnapshot on the binary transports and POST /snapshot on HTTP) let a user with only the read_servers permission download the server log files. That user could read the generated root password and log in as root. This issue affects	7.5	More Details

118493	Apache Iggy: 0.9.0. Users are recommended to upgrade to version 0.10.0, which fixes the issue. Version 0.10.0 never generates a root password, and the snapshot requires the manage_servers permission. If a 0.9.0 server generated its root password, change that password.		
CVE-2026-133021	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via Mercurial subrepository paths	7.5	More Details
CVE-2026-133022	In JetBrains TeamCity before 2026.2.1 remote code execution was possible via the Notification Rules form	7.5	More Details
CVE-2026-22311	The failed-login limiter of the JTL-Wawi synchronisation interface of JTL-Shop derives the client identity from the client-controlled X-Forwarded-For header and counts requests without credentials as failed logins. An unauthenticated remote attacker can block the synchronisation between JTL-Shop and JTL-Wawi with a few requests and repeat this indefinitely, so that orders are no longer handed over to JTL-Wawi. No data is disclosed and no code is executed.	7.5	More Details
CVE-2026-121812	Description: When BIG-IP APM is provisioned, and a virtual server is configured with an APM access profile and a per-request policy, and the tmm.access.policytrace database key is enabled, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate. The key is disabled by default. Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system . There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	7.5	More Details
CVE-2026-133081	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via Git VCS root URLs	7.5	More Details
CVE-2026-133082	In JetBrains TeamCity before 2026.2.2, 2026.1.6 a DoS attack was possible via unauthenticated multipart requests	7.5	More Details
CVE-2026-107494	Issue summary: The QUIC stack queues a RETIRE_CONNECTION_ID frame also for retransmitted NEW_CONNECTION_ID frames. This allows a remote malicious QUIC stack to queue an unlimited number of RETIRE_CONNECTION_ID frames. Impact summary: A malicious remote QUIC peer can trigger an out-of-memory condition causing a Denial of Service. CWE: CWE-770: Allocation of Resources Without Limits or Throttling Description: The earlier fix for CVE-2026-84784 was not complete. It failed to ignore retransmissions of the NEW_CONNECTION_ID frame. The QUIC implementation created and inserted a RETIRE_CONNECTION_ID frame for every retransmitted NEW_CONNECTION_ID frame it received and there was no limit on how many RETIRE_CONNECTION_ID frames were queued. This issue allows a malicious remote QUIC stack to trigger an unbounded allocation of memory and eventually an out-of-memory (OOM) condition. FIPS impact: no The FIPS module is not affected as the QUIC implementation is outside the OpenSSL FIPS module boundary.	7.5	More Details
CVE-2026-129149	In Eclipse Mosquitto versions 2.1.0 to 2.1.2, using the mosquitto_ctrl utility in interactive shell mode to interact with a malicious server could result in a crash of the utility if the server sent a hostile response.	7.5	More Details
CVE-2026-126706	In Eclipse Mosquitto versions 2.1.0 to 2.1.2, the mosquitto_ctrl utility would crash if attempting to interact with a hostile broker and the broker sent a non-string "error" value.	7.5	More Details
CVE-2026-99618	In the Linux kernel, the following vulnerability has been resolved: nvmet: don't return memory contents beyond valid data when reading ANA log page The host side does not know how much valid data will be returned by the target so it will typically provide a buffer larger than the amount of valid data in the ANA log page. The memory allocation used by the target does not guarantee that prior kernel memory contents will not be returned past the end of the valid page data. Fix this to ensure that the SGL is zeroed before copying the data.	7.5	More Details
CVE-2026-133161	### Summary ui With `strip_ansi=True`, Pydantic AI Harness's `ToolOutputLimits` capability removes ANSI escape sequences from tool output with a pattern that takes time quadratic in the length of a line. A tool result of a few hundred kilobytes, crafted with many unterminated escape sequences, can occupy the event loop for tens of seconds to minutes, delaying every other request and agent run in the process. ### Details This affects applications that set `strip_ansi=True` on `ToolOutputLimits` (named `OverflowingToolOutput` before 0.13.0) and run tools whose output an attacker can influence, such as shell commands, file reads, fetched pages or MCP servers. Stripping runs on the event loop before the capability reduces the output, so its own size limits do not bound it. The assumed attacker can place text in content a tool returns. ### Mitigation Upgrade `pydantic-ai-harness` to 0.70.0 or later. Until then, leave `strip_ansi` at its default, `False`. On Pydantic AI v1, the harness cannot be upgraded past 0.4.0. Upgrade to Pydantic AI v2 and `pydantic-ai-harness` 0.70.0, or leave `strip_ansi` off.	7.5	More Details

CVE-2026-133195	In JetBrains Junie CLI before 26.6.29 files outside the project directory could be modified by agent file tools	7.5	More Details
CVE-2026-129891	IO::Socket::SSL::PublicSuffix versions through 2.099 for Perl let an on-path attacker choose the installed public suffix list because update_self_from_url fetches it over cleartext HTTP with no integrity check. Called with no argument, as the module's documentation shows, the updater fetches the list from the `http://` address in the URL constant and checks only that the request succeeded, with no signature or digest over the response. The lines that come back are written into the module's own source file as the built-in list. An attacker on the network path when the updater runs chooses the suffix boundary the module reports to every caller on the host, and that answer persists in the installed module until it is reinstalled. Dropping one rule makes the hostname check in IO::Socket::SSL accept a wildcard certificate that spans the whole public suffix that rule named.	7.5	More Details
CVE-2026-125695	The Lightbox with PhotoSwipe WordPress plugin before 5.19.0 does not escape image link URLs before outputting them in the attributes of the lightbox share menu, allowing unauthenticated users to perform Stored Cross-Site Scripting attacks via comments when the PhotoSwipe 4 frontend is used.	7.5	More Details
CVE-2026-133534	PowerLink developed by Neotroni Lighting has a SQL Injection vulnerability, allowing unauthenticated remote attackers to inject arbitrary SQL commands to read database contents.	7.5	More Details
CVE-2026-99383	In the Linux kernel, the following vulnerability has been resolved: smb: client: reclaim credits for decrypted reads with no matching mid When a decrypted SMB2 READ response carries a MessageId that matches no outstanding request, both the inline receive_encrypted_read() and the offloaded smb2_decrypt_offload() paths drop the frame without accounting the credits the server granted in it. The plaintext unknown-frame path in cifs_demultiplex_thread() already reclaims them with smb2_add_credits_from_hdr(), so encrypted reads leak credits that the plaintext path would not. This also happens for legitimate frames - e.g. a reply that arrives after the client has timed out and dropped the mid - so over a long-lived connection the client's credit count drifts down and can eventually starve it of the ability to issue requests. The frame is AEAD-authenticated, so the granted credits in the decrypted header come from the real server and are safe to apply. Reclaim the credits from the decrypted header before dropping the frame, and export smb2_add_credits_from_hdr() so the SMB2 receive paths can use it.	7.5	More Details
CVE-2026-99386	In the Linux kernel, the following vulnerability has been resolved: smb: client: don't resurrect an exiting tcon from cached dir workers cfids_laundromat_worker() bumps tcon->tc_count for every expired open cfid and queues cached_dir_offload_close() without checking whether the tcon is already being torn down. Racing with the final cifs_put_tcon(), the tcon comes back from zero and gets torn down twice: T1 T2 cifs_put_tcon() -- tc_count == 0, TID_EXITING tconInfoFree() free_cached_dirs() cancel_delayed_work_sync() cfids_laundromat_worker() ++tc_count == 1 queue_work(serverclose_wq) kfree(tcon->origin_fullpath) kfree(tcon) cached_dir_offload_close() close_cached_dir() [UAF] cifs_put_tcon() tconInfoFree() [double free] cached_dir_lease_break() also bumps tc_count without tc_lock, racing with the laundromat's locked update. Take the tcon reference through a helper that refuses once the tcon is TID_EXITING or its refcount has dropped to zero. Queued put_work and close_work then always hold a tcon reference, so tconInfoFree() can't run before they finish and free_cached_dirs() doesn't need to flush any workqueue. If the laundromat can't get a reference, put the cfid back on the dying list for free_cached_dirs() to free; its handle goes away with the tree connection. The lease break path finds the tcon on ses->tcon_list under cifs_tcp_ses_lock, so it can't see TID_EXITING, but it now takes tc_lock.	7.5	More Details
CVE-2026-130491	In Eclipse versions 1.5-1.6.13, 2.0.0-2.0.22 and 2.1.0-2.1.2, the client library libmosquitto does not reject Will messages with payloads greater than 65535 bytes when calling `mosquitto_will_set()` or `mosquitto_will_set_v5()`. This results in up to 65535 bytes of uninitialised heap memory being transmitted as the Will. This can be mitigated by only passing payloads with a maximum length of 65535 bytes.	7.5	More Details
CVE-2026-99389	In the Linux kernel, the following vulnerability has been resolved: smb: server: size directory name conversion buffer correctly A 2048-byte ASCII directory name needs 4098 bytes when converted to UTF-16, including the terminator. ksmbd_convert_dir_info_name() caps its allocation at PATH_MAX and writes past the end of that buffer. Allocate space for one UTF-16 code unit per input byte plus the terminator. Use kmalloc_array() to check the size calculation.	7.5	More Details
CVE-2026-131877	Convert::ASN1 versions before 0.37 for Perl allow CPU exhaustion via unbounded arc length in OBJECT IDENTIFIER decoding and encoding. The content of an OBJECT IDENTIFIER or RELATIVE-OID element is passed to unpack "w*" on decode, and the arcs of an OID string to pack "w*" on encode. Perl converts each arc to or from its decimal string in time quadratic in the arc's length, and nothing limits that length. An attacker who controls the decoded data, or an OID string passed to encode, can exhaust CPU with a single oversized arc.	7.5	More Details
CVE-2026-131868	Convert::ASN1 versions before 0.37 for Perl allow memory exhaustion via a declared element length passed unchecked to sysread in asn_read and asn_get. The length an element declares in its BER header is passed to sysread as the read size. Perl grows the target buffer to that size before reading, so memory is allocated from the declared length, not from the data received. A 6-byte header can declare a length of up to 4GB. A single header makes the reading process allocate the declared size. Where the allocation fails, perl exits with "Out	7.5	More Details

	of memory", which eval cannot catch.		
CVE-2026-126956	Convert::ASN1 versions before 0.37 for Perl allow memory exhaustion via unlimited recursion depth in constructed-form string decoding. Any constructed-form string element causes the decoder to recurse on the element's contents. The BER encoding controls nesting depth at four bytes per level, with no bound checked. Perl recursion is heap-allocated, so the process grows without a stack crash. An attacker who controls the decoded data can exhaust memory with a small deeply nested constructed string.	7.5	More Details
CVE-2026-129509	In the Apache Airflow Microsoft Azure provider 10.2.0 through 15.3.0, `S3ToAzureBlobStorageOperator` turned off TLS certificate verification for its connection to Amazon S3 by default (`s3_verify=False`), whatever the AWS connection's own verification settings were. An attacker positioned on the network between an Airflow worker and S3 could intercept that traffic, read the signed requests and alter the listed object keys and the copied file contents. Deployments are affected when a Dag uses `S3ToAzureBlobStorageOperator` without setting `s3_verify`. Users are advised to upgrade to apache-airflow-providers-microsoft-azure 15.3.1 or later, which uses the AWS connection's verification settings by default.	7.5	More Details
CVE-2026-126819	Convert::ASN1 versions before 0.37 for Perl allow memory exhaustion via zero-length BIT STRING content. A BIT STRING with zero content octets passes -1 to substr as the length. Perl reads a negative length as an offset from the end of the string, so the decoded value is nearly all of the input that follows the element. In the constructed form each zero-length segment copies the rest of the input again, so memory grows with the square of the input size. An attacker who controls the decoded data can exhaust memory with a short constructed BIT STRING. The resulting value also contains bytes from adjacent input.	7.5	More Details
CVE-2026-126682	Convert::ASN1 versions before 0.37 for Perl allow CPU exhaustion via exponent-notation string in INTEGER encoding. The encoder routes any value with abs >= 2**31 to i2osp without validating that it is an integer. A string in exponent notation such as "1e8000" passes that threshold. Math::BigInt expands it to an 8001-digit integer, and i2osp's per-octet conversion loop is quadratic in that size. An attacker who controls encoded INTEGER values can exhaust CPU with a short exponent-notation string.	7.5	More Details
CVE-2026-126576	Convert::ASN1 versions before 0.37 for Perl allow CPU exhaustion via unchecked exponent length in binary REAL decoding. Binary REAL decoding never bounds-checks the exponent length against the available content. For a minimal one-byte binary REAL, the exponent byte is read from outside the element, making the mantissa length -1. Perl reads a negative length as an offset from the end of the string, so the mantissa is nearly all of the input that follows the element. The long-form exponent path also omits the extra length octet when sizing the mantissa. An attacker who controls the decoded data can exhaust CPU with a list of such values, as each mantissa loop iterates over the remaining input. The decoded values are also computed from bytes outside each element.	7.5	More Details
CVE-2026-125272	Description: When BIG-IP APM is licensed and provisioned, and a virtual server with the APM access profile is configured with a per-request policy containing a form-based Logon Page, undisclosed requests can cause the Traffic Management Microkernel (TMM) to terminate. Impact: Traffic is disrupted while the TMM process restarts. This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) on the BIG-IP system. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	7.5	More Details
CVE-2026-133024	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via personal builds with versioned settings	7.5	More Details
CVE-2026-130176	aMQTT is an open source MQTT broker and client. From 0.11.0 until 0.12.1, the QoS 2 transaction handler in amqtt/mqtt/protocol/handler.py, _handle_qos2_message_flow(), can retain transaction state and packet identifiers when a client disconnects before the handshake completes, unlike QoS 1 flows that expire after the default five-second timeout. A remote client can repeatedly begin QoS 2 flows and disconnect without maintaining long-lived connections; anonymous clients can do this with the default configuration, while authenticated deployments remain affected when a client may initiate QoS 2 flows. The retained state grows broker memory until resource exhaustion can terminate the process; no confidentiality or integrity impact is identified. This issue is fixed in version 0.12.1.	7.5	More Details
CVE-2023-54420	LiveBOS builds through 2023-08-16 contain a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files via the imgName parameter in /feed/ShowImage.do using a ';.js.jsp' suffix bypass. Attackers can send crafted GET requests with directory traversal sequences in imgName to read system files such as /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-30.	7.5	More Details
CVE-2026-130749	Grav from 2.0.0 before 2.2.18 contains a resource exhaustion vulnerability that allows unauthenticated attackers to fill the disk by varying the HTTP Host header. Attackers can send requests with distinct Host values to create per-host compiled configuration, blueprint, language and page index caches, causing session, upload and logging failures.	7.5	More Details
CVE-	Traefik from 2.11.58 before 2.11.61 contains a resource leak vulnerability that allows unauthenticated attackers to exhaust resources by sending Authorization headers with NTLM or Negotiate schemes. Attackers		More

2026-130732	can open many short-lived connections that each allocate a dedicated backend transport never released, exhausting memory, file descriptors and backend connections.	7.5	Details
CVE-2026-84786	Issue summary: A DANE-enabled TLS client may crash by dereferencing a NULL pointer when processing an unusual combination of certificate and TLSA records. Impact summary: The impact is limited to a client crash (potential Denial of Service). CWE: CWE-476: NULL Pointer Dereference Description: A DANE-enabled TLS client that supports TLSA records with at least one of the PKIX-TA(0) or PKIX-EE(1) usages alongside the DANE-TA(2) usage may crash by dereferencing a NULL pointer if the server claims a trusted CA's certificate to be its end-entity certificate and has an unusual mix of associated TLSA records. DANE SMTP clients should not be affected since only the DANE-TA(2) and DANE-EE(3) TLSA record usages are specified applicable in SMTP. For example, the Postfix MTA is not affected, while the Exim MTA does not at present use OpenSSL's DANE implementation. Use of DANE in validation of RFC7250 raw public keys (RPKs) is also unaffected, as these only use DANE-EE(3) TLSA records. Only applications that support combining the DANE PKI with a traditional PKI are potentially affected. A DANE-enabled browser or similar, that uses OpenSSL's DANE API could be affected. FIPS impact: no The DANE API is outside the FIPS module boundary.	7.5	More Details
CVE-2026-130729	NocoDB before 2026.08.0 contains an authorization bypass in its public shared-view relation endpoints, which verify that a parent record exists but not that it passes the view's row filters. Unauthenticated attackers holding a public share link can supply guessable auto-incrementing ids of filtered-out records to has-many, many-to-many, one-many, or link endpoints to read their child records.	7.5	More Details
CVE-2023-54418	Nuxt 3.0.0-rc.1 before 3.0.0-rc.13 contains an arbitrary file read vulnerability caused by the Vite dev server default setting vite.server.fs.strict=false. Unauthenticated attackers can request /_nuxt/@fs/<absolute path> with a raw query to read arbitrary files like /etc/passwd from the system. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-18.	7.5	More Details
CVE-2023-54417	Nuxt 3.0.0 before 3.4.2 contains a semi-arbitrary file read vulnerability in the vite-node dev server middleware that allows unauthenticated attackers to read files by supplying an absolute path as the moduleId. Attackers can send GET requests to /__nuxt_vite_node__/module/ with an absolute file path to read sensitive system files. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	7.5	More Details
CVE-2026-130727	NocoDB before 2026.08.0 contains an authorization bypass vulnerability in public shared-view endpoints that fetch records by primary key without applying the view's row filter. Anonymous share link holders can enumerate sequential primary keys on link, has-many, and attachment download endpoints to read excluded records' linked data and download their attachments.	7.5	More Details
CVE-2022-51030	QiboCMS V7 contains an arbitrary file download vulnerability in inc/job/download.php that allows unauthenticated attackers to read files by supplying a base64-encoded url parameter to /do/job.php?job=download. Attackers can bypass the .php extension check using the Windows 'ph<' wildcard to download sensitive files such as configuration files containing database credentials. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-18.	7.5	More Details
CVE-2022-51029	Synway SMG Gateway Management Software builds through 2022-03-15 contain a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files via the downfile parameter in /down.php. Attackers can send a multipart POST request to /down.php with a traversal path in downfile to retrieve sensitive files like /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-20.	7.5	More Details
CVE-2021-48022	Rocket.Chat versions 0.61.0 before 3.11.4, 3.12.x before 3.12.4 and 3.13.x before 3.13.2 contain a missing authorization vulnerability allowing unauthenticated attackers to read any room's messages via the livechat:loadHistory Meteor method. Attackers can mint a guest token through livechat:registerGuest and send POST requests to /api/v1/method.callAnon with a target rid parameter to retrieve message history without membership checks. Exploitation evidence was first observed by the Shadowserver Foundation on 2025-02-13.	7.5	More Details
CVE-2019-25785	FineReport 8.0 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files through the resourcepath parameter of the /WebReport/ReportServer endpoint. Attackers can send GET requests with op=chart and cmd=get_geo_json to read sensitive files like privilege.xml and extract admin credentials. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	7.5	More Details
CVE-2026-130708	NocoDB before 2026.09.0 contains an improper authorization vulnerability that allows workspace Creator users to read plaintext notification webhook URLs by querying the webhook plugin listing endpoint. Attackers can replay exposed Slack, Teams, Discord and Mattermost webhook URLs directly against messaging providers to post arbitrary messages into organization chat channels.	7.5	More Details
CVE-2016-20106	NUUO NVRmini 2 through 3.0.8 contains a file disclosure vulnerability that allows unauthenticated attackers to read web-root files by passing a filename to the 'css' parameter of css_parser.php. Attackers can retrieve PHP source files confined to the document root, exposing hardcoded credentials and sensitive application logic. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-17.	7.5	More Details

CVE-2026-95823	tcp_pkt_linearize() in subsys/net/ip/tcp.c is used by th_get() whenever a received TCP header is not contiguous in a single net_buf. While pulling the header bytes out of the following fragments, the loop called net_buf_unref() on an emptied fragment without first detaching its frags pointer. net_buf_unref() walks the whole fragment list, so the payload buffers still linked behind that fragment were freed as well, and the stale pointer was then reinstalled into the packet as buf->frags, producing a use-after-free and a later double free during packet teardown. The same loop also advanced to second->frags unconditionally: when the fragment still held payload after the header remainder was pulled, it was spliced out of the packet without ever being unreferenced, permanently leaking that buffer and silently dropping the data it carried. The defect is reached from tcp_recv() via th_get() on every received TCP segment, before connection lookup or any header validation, so a single unauthenticated segment that matches a listening socket is sufficient. An attacker only needs the TCP header to straddle a net_buf boundary, which they can arrange from off-link: IPv6 extension headers are accepted on receive and their length is added to the header offset by th_get(), so a Destination Options header of a suitable length places the TCP header across the boundary of the fixed-size receive buffers (CONFIG_NET_BUF_DATA_SIZE). Packet chains assembled from several sources - IPv4/IPv6 reassembly, 6LoWPAN/802.15.4 reassembly, tunnel or virtual interfaces that strip an outer header, or drivers producing scatter-gather chains - can additionally present a short interior fragment that the loop empties exactly, which is the case that frees the remainder of the packet while it is still in use. The reliably reachable consequence is exhaustion of the receive buffer pool: each crafted segment leaks one net_buf from the fixed CONFIG_NET_BUF_RX_COUNT pool, the buffers are never reclaimed, and once the pool is empty the device can no longer receive network traffic until it is rebooted. Where a fragment is emptied exactly, the impact rises to memory corruption: freed pool buffers are written through by the remainder of the loop, their contents may be delivered to the application, and the packet is double-freed on teardown. The fix saves and clears second->frags before the net_buf_unref() so only the emptied header fragment is released, and advances to the next fragment only when second->len == 0, keeping any fragment that still holds payload in the packet. This restores the detach-before-unref ownership pattern used elsewhere in the stack, for example in net_pkt_trim_buffer().	7.5	More Details
CVE-2026-99715	In the Linux kernel, the following vulnerability has been resolved: nvme-rdma: fix NULL queue->ctrl dereference on reconnect Since commit 2a8513091d2f ("nvme-rdma: parallelize I/O queue allocation and startup"), a reconnect that gets more I/O queues than before grows the tag set before the new queues are set up. queue->ctrl is only set by nvme_rdma_setup_io_queues(), so I/O submitted in between reaches nvme_rdma_queue_rq() on a queue that was never allocated and dereferences a NULL queue->ctrl: BUG: kernel NULL pointer dereference, address: 000000000000034c RIP: 0010:nvme_rdma_queue_rq+0x52/0xb90 Set queue->ctrl for every queue when allocating the controller. Such requests are then requeued like on any other queue that is not live.	7.5	More Details
CVE-2026-115815	The rus to lat advanced plugin for WordPress is vulnerable to blind SQL Injection via the 'name' parameter in all versions up to, and including, 1.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for 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-2026-130718	NocoDB before 2026.08.0 contains an authorization bypass vulnerability that allows anonymous share link holders to access records excluded by a shared view's row filter, because the read-by-primary-key path ignores the view's filters. Attackers can request the public attachment-download or nested-link listing routes with a record id to mint signed download URLs for hidden files and read linked-record data.	7.5	More Details
CVE-2026-130723	NocoDB before 2026.08.0 contains an information disclosure vulnerability that allows unauthenticated attackers to access attachments of filtered-out records through the shared view attachment-download endpoint. Attackers holding a public share link can enumerate sequential primary keys to obtain signed download paths for files on records the view's row filter excludes.	7.5	More Details
CVE-2026-123123	An attacker that is able to send a crafted network packet to the affected component may trigger an integer overflow during module size validation. This may cause insufficient bounds checking and result in out-of-bounds memory reads, potentially leading to a denial of service condition due to process termination.	7.5	More Details
CVE-2026-131575	Tanium addressed an insertion of sensitive information into a log file vulnerability in TanOS.	7.5	More Details
CVE-2026-126763	In Eclipse Mosquitto versions 1.6 to 2.0.22 and 2.1.0 to 2.1.2, an MQTT v5.0 client connecting with a will-delay-interval that subsequently disconnects with a shorter interval causes all delayed wills to be suppressed.	7.5	More Details
CVE-2026-130748	Grav from 2.0.0 before 2.2.18 contains a password reset poisoning vulnerability that allows anonymous attackers to redirect reset links by forging the CloudFront-Forwarded-Proto header. Attackers can send a forgot-password request with a crafted scheme value so the emailed reset link points to their host, capturing the token to take over accounts including super admins.	7.5	More Details
CVE-2026-130014	The Team VCard generator plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 1.1.0 via the 'photo' parameter parameter. This makes it possible for unauthenticated attackers to read the contents of arbitrary files on the server, which can contain sensitive information.	7.5	More Details

CVE-2026-130746	Grav before 1.7.53.7 and 2.x before 2.2.18 serves page media files, and rendered pages when lazy sessions are enabled, with the visitor's session cookie and a shared-cacheable Cache-Control header. When a shared cache such as Apache mod_cache stores Set-Cookie responses, an attacker can lure a user signed in through the front-end login to open a media URL with a unique query string, then request that URL to receive the cached session cookie and take over the session.	7.5	More Details
CVE-2024-58403	AnyShare builds through 2024-11 contain a missing authentication vulnerability in the /api/ShareMgnt/Usrm_GetAllUsers endpoint that exposes user records including password fields. Attackers can send unauthenticated POST requests with a [start,limit] body to retrieve usernames, passwords, and department information for all users. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-12-25.	7.5	More Details
CVE-2026-118899	When using the listen directive with the ssl parameter in the http block with TLSv1.3 enabled, an attacker can use a specially crafted TLS connection to cause an infinite loop in the NGINX worker process while reading the HTTP request body. As a result, the worker process stops accepting new connections and processing existing connections. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated. Impact: This vulnerability allows a remote, unauthenticated attacker to cause a denial-of-service (DoS) of NGINX. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	7.5	More Details
CVE-2026-132825	Ruby before 3.3.14, 3.4.13, and 4.0.10 has a heap-based buffer overflow during case conversion. One byte can be written outside a buffer for, for example, ISO-8859-1 or Windows-1252.	7.4	More Details
CVE-2026-129938	An improper neutralization of delimiters in the MongoDB C# driver's DNS SRV connection-string resolution lets an unauthenticated party who can answer the client's DNS SRV query place a host outside the trusted domain into the client's server list. When an application resolves the connection string through the affected API and connects with the result, application data, authentication credentials and query results may pass through that host. This may result in unintended disclosure of data and credentials and in forged query results.	7.4	More Details
CVE-2026-131848	Snap7 through 1.4.3 contains an out-of-bounds write vulnerability in TSnap7MicroClient::opReadMultiVars() that allows malicious PLCs to overflow client buffers via unchecked DataLength fields. Attackers controlling or impersonating a PLC can return crafted ReadVar responses to Cli_ReadMultiVars() callers, writing up to 65535 bytes per item to corrupt memory and crash clients.	7.4	More Details
CVE-2026-133185	Goproxy before 1.9.4 contains a weak key generation vulnerability in NewCounterEncryptorRandFromKey that derives RSA and Ed25519 leaf certificate private keys from public data due to misuse of hash.Sum. Attackers with the CA certificate and one observed leaf certificate can regenerate that host's leaf private key and impersonate the host to clients trusting the proxy CA.	7.4	More Details
CVE-2026-115826	The Multi Step for Contact Form 7 plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 2.9.1 via the '_cf7mls_db_form_data_id' parameter due to missing validation on a user controlled key. This makes it possible for unauthenticated attackers to permanently delete arbitrary stored Contact Form 7 submissions from the database. Exploitation requires the companion CF7 Database plugin to be active and the per-form "Save every step, not just the final submission" option to be enabled.	7.4	More Details
CVE-2026-133097	In JetBrains TeamCity before 2026.2.2, 2026.1.6 agent Push authorization keys were predictable	7.4	More Details
CVE-2026-131849	Snap7 through 1.4.3 contains a heap-based buffer overflow in TSnap7MicroClient::opReadSZL() that allows malicious PLCs to overwrite client memory by sending oversized multi-slice SZL responses. A rogue or impersonated PLC can keep flagging more data with attacker-controlled DLen slices, overflowing the 65536-byte opData buffer via Cli_ReadSZL() and Cli_ReadSZLList(), crashing the client and corrupting heap memory.	7.4	More Details
CVE-2026-107644	The Bizuno Accounting – ERP/Accounting/CRM (for WordPress) plugin for WordPress is vulnerable to Authentication Bypass via Password Reset Token Disclosure in all versions up to, and including, 7.4.5 via the `bizUser` parameter of the `validateReset()` function on the `portal/api/lostPW` endpoint. This is due to `portalCtl::goAPI()` failing to call `loadBusinessCache()` before invoking the handler, which causes `isMailConfigured()` to read an empty business cache and return false regardless of the site's actual mail configuration, causing the fully-qualified reset URL — containing the victim's live reset UUID and numeric user ID — to be embedded directly in the HTML response rather than delivered by email. This makes it possible for unauthenticated attackers to obtain a valid password-reset token for any Bizuno account, including administrators, directly from the HTTP response and submit it to the `portal/api/lostVal` endpoint to rewrite the victim's password and achieve full account takeover. The administrator email address required to initiate the reset is frequently exposed by the plugin's own login error messages and the public `/bizuno/` portal page, making exploitation trivial against most installations.	7.4	More Details
CVE-2026-	Fleet before 4.92.0 contains an improper authentication vulnerability that allows attackers to reuse case-mutated security tokens because MySQL token columns used case-insensitive collation. Attackers holding a case-normalized invite or email-change token can redirect a victim's email for account takeover or redeem	7.4	More Details

130757	invites with roles up to Global Admin.		
CVE-2026-132128	@fastify/middie is the middleware engine that adds Express and Connect style middleware support to Fastify. In versions up to and including 9.3.6, middie does not honor the router's caseSensitive setting and matches path-scoped middleware mount paths with a case-insensitive regular expression that does not fold non-ASCII characters, while Fastify's router, when caseSensitive is set to false, lowercases the request path with a Unicode-aware conversion that maps the KELVIN SIGN character (U+212A) to the ASCII letter k. A request that spells a k in a protected mount path as the percent-encoded KELVIN SIGN is therefore routed to the protected handler while the middleware mounted on that path, such as authentication, authorization, or rate limiting, never matches and is skipped, letting an unauthenticated client reach the handler. Only applications that set caseSensitive to false and protect a mount path containing the letter k are affected. The issue is fixed in @fastify/middie 9.3.7, and users should upgrade to 9.3.7 or later. As a workaround, enforce authorization in a Fastify hook such as preHandler that runs after routing, or keep routing case-sensitive, which is the default.	7.4	More Details
CVE-2026-132131	@fastify/middie is the middleware engine that adds Express and Connect style middleware support to Fastify. In versions 8.3.0 through 9.3.6, the internal helper that normalizes a path-scoped middleware's mount path only handles string inputs, but the use method also accepts an array of mount paths and passes it through unchanged, so the normalization silently does nothing. As a result, a trailing slash on an array element survives into the generated route-matching regular expression, and an element that is exactly the root path, meaning a catch-all mount, is dropped entirely. Requests to subroutes of the affected mount paths then skip the path-scoped middleware, such as authentication, authorization, or rate limiting, while Fastify's router still dispatches them to the handler, allowing an unauthenticated client to reach protected routes. Applications are affected when they mount middleware on an array of paths in which an element ends with a slash, including one introduced by a plugin prefix that ends with a slash, or is exactly the root path. The issue is fixed in @fastify/middie 9.3.7, and users should upgrade to 9.3.7 or later. As a workaround, enforce authorization in a Fastify hook such as preHandler that runs after routing, and avoid array mount paths with a trailing slash element for security relevant middleware.	7.4	More Details
CVE-2026-129932	The MongoDB Java driver applies its DNS seedlist domain restriction to a record value that is later re-interpreted differently when the connection target is parsed, so the name that is checked and the name that is used can differ. An unauthenticated party who can set or forge the DNS service record for the seed domain in a connection string can cause the driver to direct its session to an unrelated host and to complete certificate verification against that host instead. This may result in unintended disclosure of application data and of configured database credentials, and in unintended content being returned to the application as query results.	7.4	More Details
CVE-2026-71861	The Chainguard SDK is a public Go SDK for integrating with the Chainguard platform. Prior to 0.1.107, auth/aws.VerifyToken trusts the Chainguard-Audience and Chainguard-Identity headers without requiring those headers to be covered by AWS SigV4 SignedHeaders. An attacker with a fresh STS GetCallerIdentity request signed by a victim can rewrite the unsigned binding headers and replay the proof to another relying party, causing the relying party to authenticate the victim’s AWS identity for an attacker-chosen audience and identity. Code that only uses auth/aws.GenerateToken is unaffected because the generator signs the binding headers, and exploitation requires a signed request within its validity window. This issue is fixed in version 0.1.107.	7.4	More Details
CVE-2026-133020	In JetBrains TeamCity before 2026.1.5 untrusted build approval for GitHub fork pull requests could be bypassed	7.4	More Details
CVE-2026-115463	Poweradmin is a web-based DNS administration tool for PowerDNS server. From 4.0.0 until 4.2.6, 4.3.5, and 4.4.1, lib/Domain/Service/MfaService.php verifies email second-factor codes through MfaService::verifyCode() without limiting failed attempts, and lib/Application/Controller/MfaVerifyController.php logs failures without throttling them. An attacker who knows the user's password can repeatedly guess the six-digit code during its ten-minute validity period when email MFA is enabled, potentially bypassing the second factor; the advisory notes MFA is disabled by default and that the practical request rate was not measured. This issue is fixed in versions 4.2.6, 4.3.5, and 4.4.1.	7.4	More Details
CVE-2026-133253	In JetBrains IntelliJ IDEA before 2026.3 untrusted projects could read files of other open projects via the built-in web server	7.4	More Details
CVE-2026-61581	The Datadog Agent collects events and metrics from hosts and sends them to Datadog. Prior to datadog-apm-inject 0.66.0, the Datadog Single Step Instrumentation injector stores helper tooling in a predictable host directory. On Linux or Kubernetes systems with SSI enabled and at least one Java service instrumented through SSI, a local unprivileged attacker with shell access to the same host can influence that tooling and execute code with the privileges of the instrumented Java service. Services that use dd-java-agent without SSI and non-Java language runtimes are not affected. This issue is fixed in datadog-apm-inject 0.66.0.	7.4	More Details
CVE-	In Splunk SOAR versions below 8.8.0, a user with access to the configured load balancer account could replace a load balancer helper that Splunk SOAR later runs with root privileges during cluster lifecycle		More

2026-131288	operations. If a Splunk SOAR administrator performs a cluster lifecycle operation, Splunk SOAR could execute arbitrary code with root privileges on the load balancer. The vulnerability is possible because Splunk SOAR does not verify the integrity or ownership of the load balancer helper before running it with root privileges.	7.3	Details
CVE-2026-132625	A vulnerability was determined in SourceCodester Student Project Allocation and Management System 1.0. Affected is an unknown function of the file /login.php of the component Login Handler. Executing a manipulation of the argument username/password can lead to sql injection. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2026-133415	A vulnerability has been found in sumiie24 School-Management-System up to d946d649c7c95e101fc5679ecd91999533bb10d7. The affected element is an unknown function of the file parent_Registration.php. Such manipulation of the argument student_reg leads to sql injection. The attack can be executed remotely. The exploit has been disclosed to the public and may be used. This product implements a rolling release for ongoing delivery, which means version information for affected or updated releases is unavailable.	7.3	More Details
CVE-2026-99379	In the Linux kernel, the following vulnerability has been resolved: smb: client: validate SMB1 EA list bounds before parsing CIFSSMBQAllEAs() checks the server-provided EA list length by adding it to the response pointer. On 32-bit systems, a positive length can wrap the pointer and pass the comparison. The parser then walks outside the received response while reading the next FEA header. A response with a list length of 0x40000000 reproduces this as a native i386 page fault and kernel panic from listxattr() against an SMB1 server. Check that the EA list starts within the SMB byte area before reading its header. Keep its length unsigned and compare it with the bytes left in the response without pointer addition. SMB1 is not negotiated by default and requires an explicit vers=1.0 mount.	7.3	More Details
CVE-2026-111513	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain an Improper Verification of Cryptographic Signature vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Arbitrary Code Execution and Elevation of Privilege.	7.3	More Details
CVE-2026-111514	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain an Incorrect Default Permissions vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Arbitrary Code Execution and Elevation of Privilege.	7.3	More Details
CVE-2026-111520	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain a Time-of-check Time-of-use (TOCTOU) Race Condition vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Code Execution and Elevation of Privilege.	7.3	More Details
CVE-2026-132627	A security flaw has been discovered in SourceCodester E-Commerce Shopping Cart Web Application 1.0. Affected by this issue is some unknown functionality of the file action.php. The manipulation of the argument keyword results in sql injection. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks.	7.3	More Details
CVE-2026-132626	A vulnerability was identified in SourceCodester Modern Ticket Reservation System 1.0. Affected by this vulnerability is an unknown functionality of the file /booking/validate.php. The manipulation of the argument ticket leads to sql injection. It is possible to initiate the attack remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2026-131211	In Splunk App for Payment Card Industry Compliance versions below 8.7.0, a user who holds the "ess_user" Splunk Enterprise Security role could add attacker-controlled values to lookup data used by the Primary Functions report. If another authenticated user later selects the value in the Primary Function Summary panel, the Splunk App for Payment Card Industry Compliance could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selected the value. With this access, the first user could expose or modify data available to the user who selected the value. The vulnerability is possible because the Primary Functions report does not safely escape report values before including them in drilldown searches. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) and Reports (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/user-manual/8.7/using-the-splunk-app-for-pci-compliance/reports) in the Splunk documentation.	7.3	More Details
CVE-2026-117517	The fix for CVE-2026-25854 was incomplete. This issue affects Apache Tomcat: from 11.0.19 through 11.0.27, from 10.1.53 through 10.1.61, from 9.0.116 through 9.0.123. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	7.3	More Details
CVE-2026-132611	A security vulnerability has been detected in SourceCodester Simple Form Builder Web App 1.0. This impacts the function save_response of the file /classes/Forms.php. The manipulation of the argument q[<index>] leads to unrestricted upload. Remote exploitation of the attack is possible. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2026-132610	A security flaw has been discovered in SourceCodester Simple Form Builder Web App 1.0. The impacted element is the function save_form of the file /classes/Forms.php. Performing a manipulation of the argument form_code/form_data results in code injection. The attack may be initiated remotely. The exploit has been released to the public and may be used for attacks.	7.3	More Details

CVE-2026-132609	A vulnerability was identified in SourceCodester Online Food Ordering System 1.0. The affected element is an unknown function of the file /web/admin/add_food.php. Such manipulation leads to unrestricted upload. The attack can be launched remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2026-132286	A security vulnerability has been detected in yangzongzhuan RuoYi up to 4.8.3. The impacted element is an unknown function of the file ruoyi-admin/src/main/java/com/ruoyi/web/controller/system/SysLoginController.java of the component User Registration Endpoint. The manipulation of the argument deptId/status/userId/createBy leads to improper input validation. Remote exploitation of the attack is possible. The exploit has been disclosed publicly and may be used. The identifier of the patch is d7ad977c9ca64e9b081c893ec56fe08f51f66698. It is recommended to apply a patch to fix this issue.	7.3	More Details
CVE-2026-131239	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in Updates data. If another user who can access the System Update Status report selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify data available to the user who selects the row. The vulnerability is possible because the System Update Status report does not validate destination, signature, and vendor-product values before inserting them into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and System Update Status (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/system-update-status) in the Splunk documentation.	7.3	More Details
CVE-2026-133478	A security vulnerability has been detected in rajesh2211 Doctor-Appointment-System up to ee974dcf4c50da85e6ea026e27f4370e9468b3cc. Affected is an unknown function of the file AppointmentSystem/AppointmentSystem/Connectivity.php of the component Login. The manipulation of the argument User leads to sql injection. It is possible to initiate the attack remotely. The exploit has been disclosed publicly 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.	7.3	More Details
CVE-2026-129234	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, Prover9Parent.config_prover9 and Mace accept a relative binary_location and pass it to find_binary as path_to_bin, disabling the bare-name absolute-path guard. A local attacker who can plant prover9 or mace4 in the resolved working directory can cause subprocess.Popen to execute that binary with the invoking user's privileges. This issue is fixed in version 3.10.4.	7.3	More Details
CVE-2026-131278	In Splunk SOAR versions below 8.8.0, a user with local access to the affected host could obtain a live Splunk SOAR session token or decrypted asset configuration when another user starts an App Editor debug action. The vulnerability is possible because App Editor debug actions expose sensitive runtime information to other local users on the affected host.	7.3	More Details
CVE-2026-130728	NocoDB before 2026.08.0 contains a stored cross-site scripting vulnerability that allows any authenticated user to register OAuth clients with javascript: redirect URIs. Attackers can send a consent link to a signed-in victim so either consent button runs script in the NocoDB origin, stealing the session token for account takeover.	7.3	More Details
CVE-2026-106012	Incorrect Authorization vulnerability in Apache Tomcat could occur during a request dispatch if the request depended on a security-role-ref provided by the dispatch target. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.0.M1 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124, which fixes the issue. Note: Tomcat 11.0.28 was not released	7.3	More Details
CVE-2026-126675	Incorrect default permissions in Nitro PDF Pro before 26.4 on Windows allow a local low-privileged user to execute code in the context of another user who launches the application.	7.3	More Details
CVE-2026-132241	In JetBrains TeamCity before 2026.1.3, 2025.11.7 XXE in the .NET build steps auto-discovery was possible	7.3	More Details
CVE-2026-133014	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via public reflection in the Kotlin DSL sandbox	7.2	More Details
CVE-2026-	AnsibleForms is a self-hosted web application that turns automation into self-service forms. Prior to 6.2.1, AnsibleForms allows a user with repository management privileges to inject shell syntax through repository uri, branch, name, or host values processed by server/src/models/repo.model.js. The Repo.clone and Repo.addKnownHosts paths interpolate these values into git and ssh-keyscan command strings, and	7.2	More

73968	server/src/lib/cmd.js executes the strings through Cmd.executeSilentCommand with shell interpretation enabled. Creating or updating a repository can therefore execute arbitrary commands as the AnsibleForms application process. This issue is fixed in version 6.2.1.		Details
CVE-2026-3232	The Easy Appointments plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'name' and 'phone' fields in all versions up to, and including, 3.12.23.1 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 an administrator accesses the Customers page.	7.2	More Details
CVE-2026-101113	Multiple ELECOM wireless LAN routers contain an OS command injection vulnerability that may cause an arbitrary OS command to be executed by an attacker who can access the administrative interface of the affected product.	7.2	More Details
CVE-2026-130333	The SNORDIAN's H5PxAPIkatchu plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'xapi' parameter in all versions up to, and including, 0.4.24 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 required to reach the unauthenticated insert endpoint is publicly exposed in the listener.js script URL on every page displaying H5P content, making the injection trivially exploitable by any anonymous visitor.	7.2	More Details
CVE-2026-62955	baserCMS is a website development framework. From 5.2.3 until 5.3.0, BcThemeFileService::getFullpath() in plugins/bc-theme-file/src/Service/BcThemeFileService.php skips its path containment check when realpath() fails for a nonexistent theme subdirectory selected by the type parameter. An authenticated administrator can supply a traversal path through the theme-file API to access locations outside the theme directory. The bypass of the CVE-2026-30940 fix can permit file writes, overwrites or deletion within the PHP process permissions and, where executable PHP can be written, remote code execution. This issue is fixed in 5.3.0.	7.2	More Details
CVE-2026-132748	Vtiger CRM through 8.4.0 contains a PHP code injection vulnerability that allows authenticated administrators to inject code via the updatedFields parameter into config.inc.php. Attackers can supply a crafted HELPDESK_SUPPORT_EMAIL_ID value that passes email validation but closes the PHP string literal to execute arbitrary operating system commands.	7.2	More Details
CVE-2026-120686	The AMP for WP – Accelerated Mobile Pages plugin for WordPress is vulnerable to Stored Cross-Site Scripting in versions up to, and including, 1.1.19 via comment content rendered on AMP post pages. The plugin registers an output-buffer callback ampforwp_the_content_filter_full() on the 'wp' action (priority 999) that applies a set of tag-agnostic and quote-agnostic regular expressions — notably /onclick=[^>]*/ — to the entire AMP page output after approved comments have already been echoed into it; these regexes do not understand HTML tag or attribute-value boundaries, so a kses-legal comment that places strings like 'onclick="' inside a title attribute value can trigger deletion of the title's closing quote, causing the browser to re-parse what was previously textual content inside the attribute value as live HTML attributes. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page, after an administrator has approved the comment.	7.2	More Details
CVE-2026-126131	The WP Instant Feeds plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'data[images][][url]' parameter in all versions up to, and including, 1.3.4 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. The cache-write AJAX endpoint is accessible to unauthenticated users because the only access control is a nonce check against a nonce minted for uid 0 and freely emitted on any front-end page rendering a feed via wp_localize_script, making it harvestable by any unauthenticated visitor.	7.2	More Details
CVE-2026-107516	Dell PowerStore contains an OS command injection vulnerability. An authenticated user with limited privileges could potentially exploit this vulnerability, potentially leading to arbitrary command execution with root privileges.	7.2	More Details
CVE-2026-119635	The Groundhogg plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'post_title' merge-tag of the [ghr] shortcode in versions up to, and including, 4.10 The post_title replacement is registered with 'nested' => true, so the nested branch of Replacements::do_replacement() returns the merge value through tackle_replacements() without passing it through escape_merge_tags()/escape_shortcodes(); the callback single_post() calls html_entity_decode(get_the_title(), ENT_QUOTES), converting the entity-encoded text that wp_filter_kses leaves intact on a Contributor's post title back into raw HTML, which is then emitted into the rendered post. 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.	7.2	More Details
CVE-2026-130758	Fleet before 4.92.0 contains an authorization bypass vulnerability in the deprecated POST /api/latest/fleet/mdm/apple/enqueue endpoint that allows team-scoped users to enqueue Apple MDM commands outside their teams. Attackers with team admin, maintainer, or GitOps roles can include one own-team host UUID alongside unresolved enrollment IDs to send commands like InstallProfile to deleted or user-channel enrollments without activity logging.	7.2	More Details
	The Slimbox plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment (anchor title		

CVE-2026-123968	within comment HTML)' parameter in all versions up to, and including, 1.0.8 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. WordPress core comment KSES permits unauthenticated commenters to post anchor elements with href and title attributes, so an entity-encoded payload stored in the title attribute survives sanitization verbatim; the browser then decodes the entities when populating the DOM .title property, and the Slimbox library's jQuery .html() caption sink renders it as executable HTML.	7.2	More Details
CVE-2026-131850	ILIAS 11.0 before 11.6 contains an sql injection vulnerability in RegistrationCodeRepository::getCodesData() that concatenates unvalidated sort fields into an ORDER BY clause. Authenticated users with read access to the Authentication and Registration administration node can supply SQL expressions as the sort parameter to extract database contents via boolean-based inference.	7.2	More Details
CVE-2026-119161	The HTML Forms – Simple WordPress Forms Plugin plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'arbitrary POST parameter NAME (the array key, not the value)' parameter in all versions up to, and including, 1.7.1 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Script execution does not require the Screen Options panel to be opened — injected attributes on the hidden-column checkboxes fire on page load as soon as the Submissions tab is rendered.	7.2	More Details
CVE-2026-131272	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings edit permission could use the Automation Dependency Wheel Representational State Transfer (REST) API to write arbitrary files outside the intended directory, which could lead to arbitrary code execution. The vulnerability is possible because Splunk SOAR does not restrict uploaded automation dependency wheel filenames to the intended storage directory.	7.2	More Details
CVE-2026-130121	The AI Puffer - MCP Server for Claude & ChatGPT, with AI Chatbot plugin for WordPress is vulnerable to Stored Cross-Site Scripting via AI-generated comment replies in all versions up to, and including, 2.5.28 due to insufficient sanitization of the AI model's output in the insert_comment_reply_logic() function, which saves replies with wp_insert_comment() and bypasses the kses filtering WordPress normally applies to comment content. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page, by submitting a comment that leads the AI model to reproduce an HTML payload in its automated reply. This is only exploitable when an administrator has configured a Comment Auto-Reply automation task with 'Action on reply' set to 'Approve' (the default is 'Hold for moderation'), and the attacker's comment has been approved.	7.2	More Details
CVE-2026-116536	The Custom Twitter Feeds – A Tweets Widget or X Feed Widget plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'atts (JSON object, specifically the inner `feed` key)' parameter in all versions up to, and including, 2.11.1 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 injected payload is first stored in the wp_ctf_feed_caches database table and later written into the ctf_cron_report WordPress option by the scheduled ctf_feed_update cron job, meaning execution is deferred until the cron runs and an administrator subsequently loads the Support page.	7.2	More Details
CVE-2026-133079	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via Maven dependencies of Kotlin DSL settings	7.2	More Details
CVE-2026-131268	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings edit permission could use a customer automation dependency to execute arbitrary code as the user account that runs Splunk SOAR. The vulnerability is possible because Splunk SOAR does not properly separate customer automation dependencies from trusted product code.	7.2	More Details
CVE-2026-125033	The Smash Balloon Social Post Feed – Simple Social Feeds for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'attachments.data[].description' parameter in all versions up to, and including, 4.15.1 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. This makes it possible for unauthenticated outside Facebook users — any party able to contribute an attachment description to a post ingested by the configured feed (visitor posts, shared posts, or attached content) — to plant payloads that execute against every anonymous visitor of any WordPress page carrying the shortcode, requiring no WordPress account whatsoever.	7.2	More Details
CVE-2026-133023	In JetBrains TeamCity before 2026.2.1, 2026.1.5 remote code execution was possible via Kotlin DSL compiler arguments	7.2	More Details
CVE-2026-113699	Improper neutralization of special elements used in an OS command ('OS command injection') vulnerability in Aidata Bilişim Kurumsal Tic. Ve San. A.Ş. Aidata WR854GVR allows OS Command Injection. This issue affects Aidata WR854GVR: before V4.0.18_261111.	7.2	More Details
CVE-2026-	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via security context spoofing in the Kotlin DSL sandbox	7.2	More Details

133078			
CVE-2026-131256	In Splunk SOAR versions below 8.8.0, an authenticated user that holds a role with the System Settings edit permission could execute arbitrary code on the host that runs Splunk SOAR. The vulnerability is possible because the Representational State Transfer (REST) API accepts package references from untrusted sources and passes them to the dependency installer.	7.2	More Details
CVE-2026-133077	In JetBrains TeamCity before 2026.2.2, 2026.1.6 privilege escalation to system administrator was possible via Project Import	7.2	More Details
CVE-2026-133026	In JetBrains TeamCity before 2026.2.1, 2026.1.5 code injection in generated Kotlin DSL was possible via matrix build parameters	7.2	More Details
CVE-2026-116530	The Check & Log Email – Easy Email Testing & Mail logging plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'author' parameter in all versions up to, and including, 2.0.17 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. Exploitation requires an administrator to approve the comment carrying the shortcode payload, and the target page must render the core Latest Comments block, which passes comment author names through do_shortcode via the_content.	7.2	More Details
CVE-2026-131248	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings edit permission and restricted tenant access could modify settings for a tenant outside the role scope through the Representational State Transfer (REST) API. With this access, the user could expose confidential tenant environment values, alter tenant settings and automation behavior, or take the tenant out of operation. The vulnerability is possible because the Tenant REST API checks the System Settings edit permission but does not enforce the tenant restrictions for the role before it updates the selected tenant. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises), Configure multiple tenants on your Splunk SOAR (On-premises) instance (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/configure-product-settings-for-your-splunk-soar-on-premises-instance/configure-multiple-tenants-on-your-splunk-soar-on-premises-instance), and REST Tenant (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/8.5.0/multi-tenancy-endpoints/rest-tenant) in the Splunk documentation.	7.2	More Details
CVE-2026-133246	modified eCommerce Shopsoftware through 3.3.0 contains a local file inclusion vulnerability that allows administrators to include arbitrary server files by storing ../ sequences in the content manager file name. Unauthenticated requests to shop_content.php then pass the stored path to include(), executing embedded PHP in uploaded content files or images as the web server user.	7.2	More Details
CVE-2026-132235	In JetBrains TeamCity before 2026.1.3, 2025.11.7 remote code execution was possible via path traversal in project import archives	7.2	More Details
CVE-2026-123972	The WP-Slimbox2 Plugin plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment' parameter in all versions up to, and including, 1.1.3.1 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. Successful exploitation requires a logged-in user, typically an administrator, to click the attacker's image link in the comment; however, because unmoderated comments appear in the WordPress moderation queue where administrators routinely interact with them, the moderation gate does not meaningfully prevent exploitation. Successful exploitation causes the attacker's own account to receive the complete Administrator role.	7.2	More Details
CVE-2026-125053	The WEB-Translation – eTranslation Multilingual plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment' parameter in all versions up to, and including, 2.0.7 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 exploit requires that at least one secondary language is configured in the plugin and that comments are open on the target post; because the stored payload contains no literal HTML brackets at save-time, WordPress core KSES filtering does not remove it, and the malicious script is only materialized when a visitor loads the post under a secondary-language URL slug.	7.2	More Details
CVE-2026-132896	In JetBrains TeamCity before 2026.2, 2026.1.4 remote code execution was possible via Kubernetes Executor settings	7.2	More Details
CVE-2026-130738	Grav from 2.0.0 before 2.2.18 contains a privilege escalation vulnerability that allows user managers with admin.users.update to take over super admin accounts by disabling the super-granting group. Attackers can set the group's enabled field to false via the Flex user-groups form, then reset the victim's password through the user-accounts form to gain super admin access.	7.2	More Details

CVE-2026-130737	Grav from 2.0.0 before 2.2.18 contains a privilege escalation vulnerability that allows user managers to grant super-admin rights by nesting access maps under a YAML merge key. Attackers with admin.users.create or api.users.write can submit '<<' keys, for example via PATCH /api/v1/users/<self>, that bypass super-grant filtering and gain full site control.	7.2	More Details
CVE-2026-116758	The Advanced Country Blocker plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'REQUEST_URI' parameter in all versions up to, and including, 2.3.2 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The injected script is stored via the advcb_live_activity transient on the init hook and executed when an administrator views the Live Monitor page, as the stored URI is unsafely concatenated into a double-quoted HTML title attribute and inserted into the DOM via jQuery .html().	7.2	More Details
CVE-2026-123892	SDDC Manager contains an authenticated command-injection vulnerability due to improper sanitization of user-supplied input. Affected versions: 9.1.x.x (fixed in 9.1.1.0200); 9.0.x.x (fixed in 9.0.2.0100); 5.x (fixed in 5.2.5).	7.2	More Details
CVE-2026-96571	The FancyBox for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Comment Anchor 'href' Attribute 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 will execute whenever a user accesses an injected page. Exploitation requires the attacker to wrap the payload anchor around a smilie shortcode (e.g., :mrgreen:), which WordPress renders as an <img> element, satisfying the plugin's a:has(img) enrollment selector and disguising the malicious link as an ordinary comment emoticon.	7.2	More Details
CVE-2026-107739	The Search & Replace WordPress plugin before 3.2.4 does not correctly restrict which database tables an administrator of a single site may operate on in a multisite network, allowing subsite administrators to read and modify the tables of other sites and of the network, and to make themselves Super Admin.	7.2	More Details
CVE-2026-84998	The NextScripts: Social Networks Auto-Poster plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'comment_author' and 'comment_agent' Comment Fields in all versions up to, and including, 4.4.8 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 an administrator's browser session when the administrator views the WordPress comment moderation queue, including pending and unapproved comments, meaning no prior approval of the comment is required for exploitation.	7.2	More Details
CVE-2026-96665	The SendPress Newsletters plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'sp_firstname' parameter in all versions up to, and including, 1.26.1.20 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 public subscriber-management route uses an unsigned, forgeable base64-encoded JSON token, allowing an unauthenticated attacker to impersonate any subscriber and obtain a valid nonce — bypassing the POST handler's only access control check.	7.2	More Details
CVE-2026-132912	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 remote code execution was possible via class injection in the Kotlin DSL sandbox	7.2	More Details
CVE-2026-116519	The Geo Mashup plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'author (comment form) / display_name (profile form)' parameter in all versions up to, and including, 1.13.23 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 comment-author path requires the crafted comment to be approved before the payload is served to visitors; the subscriber-level path requires the plugin's user-location mapping feature to be enabled.	7.2	More Details
CVE-2026-133090	In JetBrains TeamCity before 2026.2.2, 2026.1.6 improper authorization allowed overwriting versioned settings of other projects	7.2	More Details
CVE-2026-107740	The Search & Replace WordPress plugin before 3.2.4 does not restrict one of its database tools to network administrators on multisite installations, allowing administrators of a single subsite to run arbitrary SQL against the whole network's database and make themselves Super Admin.	7.2	More Details
CVE-2026-95789	Command injection vulnerability in the affected interface of AOS-10 or AOS-8 could allow an authenticated remote attacker to perform command injection only on certain impacted devices. Successful exploitation could allow an attacker to execute unauthorized arbitrary commands with the privileges of the affected mechanism on the underlying operating system.	7.2	More Details
CVE-2026-133088	In JetBrains TeamCity before 2026.2.2, 2026.1.6 remote code execution was possible via FreeMarker notification templates	7.2	More Details

CVE-2026-133093	In JetBrains TeamCity before 2026.2.2, 2026.1.6 code injection in parent project Kotlin DSL was possible via subproject settings	7.2	More Details
CVE-2026-111825	The Event Tickets with Ticket Scanner plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'code' parameter in all versions up to, and including, 3.3.1 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 payload persists through WooCommerce session data into order-item metadata and executes when rendered in the ticket scanner interface.	7.2	More Details
CVE-2026-132898	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 arbitrary file write on the server was possible via Kotlin DSL settings	7.2	More Details
CVE-2026-132916	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 arbitrary file deletion on the server was possible via the Kotlin DSL sandbox	7.1	More Details
CVE-2026-99546	In the Linux kernel, the following vulnerability has been resolved: net/sched: cls_flower: avoid stale mask references after delete Flower releases a mask reference when a filter is removed from the classifier. The filter object can remain alive while dump, walk, or hardware reoffload paths hold a reference to it. Those paths can then dereference f->mask after the mask has already been released. Keep the mask reference until the filter reference count reaches zero. Release the mask synchronously from the final __fl_put() while the caller still holds the tcf_proto reference, and leave the existing RCU work to free the mask. This keeps the mask available to every reader that still holds the filter. The filter is still removed from the classifier indexes immediately. Therefore, a deleted filter can temporarily retain an otherwise empty mask; fl_delete() uses the handle IDR, rather than the mask list, to determine whether the classifier has become empty. Failed mask assignment can leave fnew->mask as an error or borrowed pointer without an owned mask reference. Clear it before the common failure path so the final put cannot release an unowned reference. This fix does not add fields, snapshots, or per-filter memory.	7.1	More Details
CVE-2026-126631	A flaw was found in OpenSCAP. This vulnerability allows for information disclosure when performing offline scans of untrusted filesystem images, such as container images or virtual machine (VM) disks. Because the scanner does not properly restrict symbolic link (symlink) resolution to the target image's root directory, an attacker can supply a crafted image containing links pointing to files on the host system. When the scan runs, typically with elevated privileges, OpenSCAP follows these links to read sensitive host files and expose their contents in the scan report or manipulate compliance results.	7.1	More Details
CVE-2026-129034	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Chill Strong Testimonials strong-testimonials allows Stored XSS.This issue affects Strong Testimonials: from n/a through 3.3.18.	7.1	More Details
CVE-2026-132915	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 an LDAP injection allowed bypassing login restrictions	7.1	More Details
CVE-2026-131262	In Splunk SOAR versions below 8.8.0, an authenticated user who can view and edit containers in at least one container label could create artifacts with parent relationships to containers and artifacts outside those labels. This could expose fields from restricted artifacts and modify incident lineage. The vulnerability is possible because Splunk SOAR does not verify access to the selected parent container and parent artifact when it creates an artifact.	7.1	More Details
CVE-2026-111515	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain an Improper Authorization in Handler for Custom URL Scheme vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to Command Execution and Launch of Phishing Attacks.	7.1	More Details
CVE-2026-99343	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/mmu: Bug the VM if KVM attempts to walk more levels than the MMU has Extend the "EFER.LMA && !CR4.PAE" check, which exists largely to guard against KVM configuring a paging32 MMU with more than 2 levels of paging, with a very explicit check for exactly that: that KVM isn't trying to walk more levels of paging than the MMU template provides. I.e. harden KVM against all bugs that would cause KVM to generates accesses beyond the bounds of guest_walker's arrays, regardless of how KVM ended up with the misconfigured MMU. Note, don't use w->cpu_role.base.level directly as the paging64 template only provides two levels of page tables for PAE paging on 32-bit hosts, and handles the third level by manually emulating the PDPTR access.	7.1	More Details
CVE-2026-131245	In Splunk SOAR versions below 8.8.0, a user with edit permission for the container label of a pending external prompt could obtain the external prompt link through the Representational State Transfer (REST) API and use the link to view and resolve the associated approval. The vulnerability is possible because the approval detail REST API includes the external prompt link in responses outside the assigned approval scope. For more information see REST Approval (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/6.3.1/approval-endpoints/rest-approval), REST External Prompt (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/7.1.0/external-prompt-endpoints/rest-external-prompt), and Manage	7.1	More Details

	roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.5.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.		
CVE-2026-99602	In the Linux kernel, the following vulnerability has been resolved: udf: don't let the extent type hide an exhausted free-space table extent udf_table_new_block() takes the first block of the free-space table extent closest to the goal. It keeps that extent's type and length together in goal_elen, subtracts one block, and deletes the extent only if goal_elen then reaches zero. UDF 2.60 section 2.3.7.1 requires the free-space extents of an Unallocated Space Entry to be of type 1 (allocated but not recorded), and mkudffs writes them that way. For such an extent the type bits keep goal_elen non-zero, so taking its last block leaves a type-1 extent of length zero behind, starting at the block after the extent, which is in use. The next allocation that picks this empty extent returns that in-use block. Subtracting a block from 0x40000000 then borrows from the type bits, and the extent becomes type 0 with a length of 2^30 - blocksize: about a gigabyte of "free" space overlapping live metadata, the other table extents and whatever lies behind the partition. From then on blocks are handed out twice, or past the end of the partition. Extents that udf_table_free_blocks() adds are type 0, where the check works, so only tables written by mkudffs or by another implementation are affected. Nothing more than filling such a filesystem is needed: on a fresh 1 MiB "mkudffs --space=unalloctable" image, creating empty files until the partition is full writes file entries over the reserve volume descriptor sequence and the backup anchor, and df then reports a negative amount of used space. On syzbot's images the same double allocation is what the two reports below trip over. In the first, ftruncate() extends a new file with enough hole extents to need a chain of allocation extent descriptors; one AED is placed on another inode's file entry, and a later one is placed on the block of the AED currently being filled, which udf_setup_indirect_aext() zeroes, so __udf_add_aext() finds lengthAllocDescs out of step with its cursor. In the second, an AED is placed past the end of the device, sb_getblk() fails, and the error path of udf_do_extend_file() calls udf_truncate_extents() on an extent list that no longer covers i_size. Keep the type separately from the length, as udf_table_prealloc_blocks() already does.	7.1	More Details
CVE-2026-129001	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Gabe Livan Asset CleanUp: Page Speed Booster wp-asset-clean-up allows Reflected XSS.This issue affects Asset CleanUp: Page Speed Booster: from n/a through 1.4.0.6.	7.1	More Details
CVE-2026-76101	keystore-go is a Go implementation of the Java KeyStore encoder and decoder. In 2.1.0 and earlier, keystore.Decode in decoder.go does not check the error returned when reading the trailing SHA-1 integrity digest and bounds its comparison by the length of the digest actually read. If an attacker with write access removes the final 20-byte digest from a modified JKS file, the empty actual digest causes the comparison loop to execute zero times and Decode returns the parsed keystore without an error. For certificate-only stores using TrustedCertificateEntry values, this bypasses both tamper detection and password verification, allowing trusted CA certificates to be added or replaced without knowing the keystore password and potentially enabling interception or modification of TLS connections that trust the forged CA. This issue is fixed in version 3.0.0.	7.1	More Details
CVE-2026-132773	Attendize through 2.8.0 contains an authorization bypass vulnerability in EventOrdersController::postCancelOrder() that allows authenticated organisers to cancel attendees belonging to other accounts' events. Attackers can cancel one of their own orders while supplying foreign attendee IDs, which Attendee::findFromSelection() loads unfiltered, invalidating tickets for any event on the installation.	7.1	More Details
CVE-2026-77313	Flyto2 Core is an execution kernel for automation and AI-agent workflows. Prior to 2.26.8, validate_url_ssrf in src/core/utils.py resolves and checks a hostname but returns the hostname URL rather than the validated address, and session.get in http.get performs an independent DNS lookup when connecting. A low-privileged caller with control of the hostname's DNS can make the first lookup return a public address and the connection-time lookup return an internal address, bypassing the SSRF guard for all modules that call it and exposing readable internal or metadata-service responses. Successful exploitation requires attacker-controlled DNS and a successful timing-dependent answer change between the two lookups. This issue is fixed in version 2.26.8.	7.1	More Details
CVE-2026-95788	An unauthenticated buffer overflow vulnerability exists in the underlying operating system of AOS-10 or AOS-8 and could allow an unauthenticated adjacent attacker to conduct a denial of service attack. Successful exploitation could allow an attacker to disrupt the availability of the affected system, with a limited impact to the integrity of the affected component.	7.1	More Details
CVE-2026-99627	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/gfx12: disallow GFXOFF while updating the SPM VMID gfx_v12_0_update_spm_vmid() does a read-modify-write of RLC_SPM_MC_CNTL over MMIO without keeping GFX powered, unlike gfx11. When GFX is in GFXOFF at that point the update doesn't stick, and with the SPM VMID programmed on submission again, the job of a process with a reserved VMID can still run with the SPM writing through VMID 0. On navi48, after a few other submissions and a VM fault on the same VM, a COPY_DATA of RLC_SPM_MC_CNTL from such a job reads 0 although the update was called for it: $ amd_uapi_regression --run-subtest cs-secure-ib-compute,gem-mapping-info-high-va,gpuvm-fault-high-va,reserved-vmid-spm-vmid amdgpu 0000:03:00.0: [gfxhub] page fault (src_id:0 ring:40 vmid:7 pasid:185) amdgpu 0000:03:00.0: Faulty UTCL2 client ID: SDMA1 (0xe) RLC_SPM_MC_CNTL in the job: 0x00000000 With GFXOFF disallowed around the update it reads the reserved VMID 15. (cherry picked from commit 907aebc5578ae09fa5c0e2f0784201295837e60b)	7.1	More Details

CVE-2026-133096	In JetBrains TeamCity before 2026.2.2, 2026.1.6 improper authorization allowed queuing builds with personal changes of other users	7.1	More Details
CVE-2026-71535	Mareforma is a local verification layer for AI-assisted research. From 0.3.0 until 0.3.7, the _maybe_update_replicated_unlocked convergence path in mareforma/db/core.py uses generated_by labels without requiring a valid signature_bundle from a distinct enrolled signer, so an unsigned claim written to a shared graph can converge with a signed peer and be promoted to REPLICATED. The forged trust state is returned by queries requiring REPLICATED support, allowing false multi-party endorsement signals to reach downstream consumers. This issue is fixed in version 0.3.7.	7.1	More Details
CVE-2026-108761	The UpdraftPlus: WP Backup & Migration Plugin for WordPress is vulnerable to second-order SQL Injection via the user profile description field in all Free versions up to, and including, 1.26.15 and all Premium versions up to, and including, 2.26.15. This makes it possible for authenticated attackers, with subscriber-level access and above, to store a malicious payload that is written into a database backup and executed as SQL when an administrator later restores that backup. Successful exploitation can modify sensitive WordPress data and, on sites where plugin installation is available, lead to remote code execution through the installation and activation of an attacker-controlled plugin.	7.1	More Details
CVE-2026-128999	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Saphali Saphali Woocommerce Lite saphali-woocommerce-lite allows Stored XSS.This issue affects Saphali Woocommerce Lite: from n/a through 2.0.2.	7.1	More Details
CVE-2026-131303	In Splunk AI Assistant versions below 2.0.0, a user who does not hold the "admin" or "power" Splunk roles could obtain limited information about indexes and data sources outside their authorized scope through the assistant. The vulnerability is possible because Splunk AI Assistant does not fully apply authorization controls to all information that it retrieves.	7.1	More Details
CVE-2026-120391	The WPvivid — Backup, Migration & Staging WordPress plugin before 0.9.140 does not prevent serialized data it cannot safely unserialize from being re-serialized as a PHP object during a site migration, allowing users with a role as low as Subscriber to plant a PHP object that WordPress instantiates after an administrator restores a backup onto a different URL.	7.1	More Details
CVE-2026-99516	In the Linux kernel, the following vulnerability has been resolved: nvmet-fc: fix use-after-free in nvmet_fc_ls_create_connection() nvmet_fc_ls_create_connection() looks up the association with nvmet_fc_find_target_assoc(), which returns it with a reference held, and drops that reference as soon as the target queue has been allocated. iod->assoc is then dereferenced again by nvmet_fc_makeconnid() while building the connection id. The queue only borrows the pointer and takes no reference of its own. Once the lookup reference is dropped the association is kept alive solely by the reference taken when it was created, so a disconnect LS or a remote port unregister running concurrently can free it before nvmet_fc_makeconnid() reads it: BUG: KASAN: slab-use-after-free in nvmet_fc_handle_ls_rqst_work (drivers/nvme/target/fc.c:193 drivers/nvme/target/fc.c:1841 drivers/nvme/target/fc.c:1995 drivers/nvme/target/fc.c:2021) Read of size 8 at addr ffff888179f1b000 by task kworker/u9:2/1167 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: nvmet-wq nvmet_fc_handle_ls_rqst_work Call Trace: <TASK> dump_stack_lvl (lib/dump_stack.c:94 lib/dump_stack.c:120) print_report (mm/kasan/report.c:378 mm/kasan/report.c:482) kasan_report (mm/kasan/report.c:595) nvmet_fc_handle_ls_rqst_work (drivers/nvme/target/fc.c:193 drivers/nvme/target/fc.c:1841 drivers/nvme/target/fc.c:1995 drivers/nvme/target/fc.c:2021) process_scheduled_works (kernel/workqueue.c:3322 kernel/workqueue.c:3405) worker_thread (kernel/workqueue.c:3486) kthread (kernel/kthread.c:436) ret_from_fork (archkernel/process.c:158) ret_from_fork_asm (arch/x86/entry/entry_64.S:245) </TASK> Freed by task 1167 on cpu 0 at 405.998150s: kasan_save_track (mm/kasan/common.c:57 mm/kasan/common.c:78) kasan_save_free_info (mm/kasan/generic.c:584) __kasan_slab_free (mm/kasan/common.c:253 mm/kasan/common.c:285) kfree (include/linux/kasan.h:235 mm/slub.c:2705 mm/slub.c:6405 mm/slub.c:6720) nvmet_fc_handle_ls_rqst_work (drivers/nvme/target/fc.c:1808 drivers/nvme/target/fc.c:1995 drivers/nvme/target/fc.c:2021) process_scheduled_works (kernel/workqueue.c:3322 kernel/workqueue.c:3405) worker_thread (kernel/workqueue.c:3486) kthread (kernel/kthread.c:436) ret_from_fork (archkernel/process.c:158) ret_from_fork_asm (arch/x86/entry/entry_64.S:245) The buggy address belongs to the object at ffff888179f1b000 which belongs to the cache kmalloc-2k of size 2048 The buggy address is located 0 bytes inside of freed 2048-byte region [ffff888179f1b000, ffff888179f1b800) Hold the lookup reference until the response has been built, and drop it on the error path as well.	7.1	More Details
CVE-2026-131286	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to author and run playbooks could view action data outside the authorized scope for that user and change the status and message of another playbook run. The vulnerability is possible because the playbook API does not verify that the requested playbook run is within the authorized scope of the authenticated user.	7.1	More Details
CVE-	Flyto2 Core is an execution kernel for automation and AI-agent workflows. Prior to 2.26.8, validate_url_ssrf() in src/core/utils.py validates one DNS answer and then returns the hostname, while guarded_aiohttp_request(), enforce_outbound_url(), and validate_url_with_env_config() allow the outbound connection to perform an independent second lookup. A caller with control of the hostname's DNS can make		More

2026-77316	validation see a public address and make the connection see an internal address, bypassing the guard across its outbound consumers and exposing readable internal or metadata-service responses. The attack requires low or zero DNS TTL behavior and a successful timing-dependent answer change; static multiple-address answers are checked and do not provide the same bypass. This issue is fixed in version 2.26.8.	7.1	Details
CVE-2026-99536	In the Linux kernel, the following vulnerability has been resolved: ksmbd: prevent session binding in compound requests A failed multichannel binding can retain its target session on the work item so that the error response can be signed. Subsequent compound commands can then use that session together with a tree connection left over from an earlier command. Limit binding SESSION_SETUP requests to standalone messages during message validation, before smb2_sess_setup() can replace work->sess. Check both NextCommand and the current request offset to cover binding in the first, middle and last positions of a compound. The existing validation error path returns STATUS_INVALID_PARAMETER and stops processing the remaining commands. This intentionally limits support for compounded binding. It is not a protocol requirement that binding be standalone. Standalone binding and non-binding compounds retain their existing behavior, without changing session, tree or response encryption handling.	7.1	More Details
CVE-2026-128669	A Session Fixation vulnerability exists in the admin login endpoint (/admin/index.php) of PHPGurukul News Portal Project V4.1. The application fails to regenerate the session identifier (PHPSESSID) upon successful authentication. A remote attacker can exploit this by setting a known session ID in the victim's browser prior to login. Once the victim authenticates, the attacker can reuse the fixed session ID to hijack the victim's authenticated session.	7.1	More Details
CVE-2026-99430	In the Linux kernel, the following vulnerability has been resolved: ocfs2: fix chunk number of the first chunk in a local quota file Mounting a crafted OCFS2 image can corrupt kernel memory. When the local quota file header of such an image claims zero chunks, the first quota entry allocated during the mount gets a chunk number taken from a kernel pointer, and releasing that entry clears a bit far outside the chunk bitmap. The write will corrupt the data stored at that address, and the corruption may lead to a system crash or damage unrelated data. Mounting requires CAP_SYS_ADMIN, so it takes an untrusted image, such as removable media or a loop mount of a file from elsewhere. Since the chunk number comes from a kernel pointer, the address of the bit that gets cleared differs from boot to boot, and the issue has been reported under several titles: - BUG: unable to handle kernel paging request in ocfs2_local_release_dquot - KASAN: use-after-free Write in ocfs2_local_release_dquot - KASAN: slab-out-of-bounds Write in ocfs2_local_release_dquot - KASAN: slab-use-after-free Write in ocfs2_local_release_dquot - KFENCE: use-after-free write in ocfs2_local_release_dquot Note that none of these is a use-after-free in the quota code: the chunk and its buffer head are alive. The bit that gets cleared lies far outside the bitmap. KASAN and KFENCE name each report based on the object occupying that address, which explains why the same issue is reported under so many different titles. In the report I sent, the address fell in a free page, which KASAN labels use-after-free. There are known reproducers. Both a syzkaller and a C reproducer are available from the Google Drive link [1] in my report thread. The issue was found by our modified syzkaller, not a theoretical thing found by an LLM. I have not seen it outside fuzzing, so it is not affecting me in the real world. The local quota file in OCFS2 is divided into chunks, and each chunk begins with a header block holding a bitmap of the quota entries that chunk has handed out. Chunks are numbered from zero, and that number is used to convert the file offset of an entry back into a bit position in the bitmap. ocfs2_local_quota_add_chunk() appends a new chunk to the in-memory list and numbers it one past the chunk that was last: list_add_tail(&chunk->qc_chunk, &oinfo->dqi_chunk); chunk->qc_num = list_entry(chunk->qc_chunk.prev, struct ocfs2_quota_chunk, qc_chunk)->qc_num + 1; The predecessor is looked up after the new chunk is added to the list, so if the list was empty, the prev pointer is the list head itself. The head is the dqi_chunk member of struct ocfs2_mem_dqinfo and is not a chunk, so reading qc_num through it lands 16 bytes past the start of the head, on the dqi_gqinode pointer that follows it. The first chunk of the file is then numbered with the lower half of a kernel pointer instead of 0. The list is empty when the local quota file header claims the file has no chunks. ocfs2_local_read_info() takes dqi_chunks from that header without validating it, so an image with dqi_chunks == 0 takes this path when the first quota entry is allocated. ocfs2_create_local_dquot() turns the bad number into a file offset with ol_dqblk_off(), which shifts a 32-bit block number left by the block size bits, so the top bits of such a large block number are lost. ocfs2_local_release_dquot() turns the offset back into a bit index with ol_dqblk_chunk_off(), using the full chunk number, so the lost bits push that index far outside the bitmap, and clearing it corrupts unrelated memory. Compute the chunk number before putting the chunk on the list, and use 0 when the list is empty.	7.1	More Details
CVE-2026-120390	The WPvivid — Backup, Migration & Staging WordPress plugin before 0.9.140 does not limit its table prefix rewrite to table names when restoring a database backup, which corrupts the length of stored serialized values and allows users with a role as low as Subscriber to plant a PHP object that WordPress instantiates after an administrator restores a backup.	7.1	More Details
CVE-2026-130776	Incus before 7.6.0 contains an information disclosure vulnerability that allows clients with project event-view access to read migration websocket secrets broadcast in operation metadata. Attackers can connect to in-flight migration control and fs websockets before the legitimate peer, hijacking the transfer to obtain instance configuration and filesystem data.	7.1	More Details
	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could use Data Model Acceleration to run malicious Search Processing Language (SPL) commands and perform actions through the Representational State Transfer (REST) API that the user cannot perform directly. The vulnerability is possible because Splunk Enterprise does not neutralize user-		

CVE-2026-131181	provided data before it adds that data to a generated search. For more information, see Accelerate data models (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.2/use-data-summaries-to-accelerate-searches/accelerate-data-models), Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities), and About the search language (https://help.splunk.com/en/splunk-enterprise/search/search-manual/10.6/search-overview/about-the-search-language).	7.1	More Details
CVE-2026-130697	NocoDB from 0.301.4 before 2026.10.0 contains an information disclosure vulnerability that allows authenticated workspace viewers to read pending invitation tokens through the workspace member list. Attackers can redeem an exposed higher-privileged invitation token at signup to claim the invited account and escalate to roles such as workspace owner.	7.1	More Details
CVE-2026-131177	In Splunk Enterprise versions below 10.4.4, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could run arbitrary Search Processing Language (SPL) searches with elevated privileges through saved search dispatch and access stored credentials and indexed data or modify lookup data. The vulnerability is possible because Splunk Enterprise does not correctly neutralize caller-supplied saved search dispatch values before it places the values into an SPL search. For more information see Search endpoint descriptions (https://help.splunk.com/en/splunk-enterprise/leverage-rest-apis/rest-api-reference/10.4/search-endpoints/search-endpoint-descriptions) in the Splunk documentation.	7.1	More Details
CVE-2026-131188	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, an unauthenticated user could craft a malicious link that, if opened by an authenticated user, executes arbitrary Search Processing Language (SPL) searches with the victim's permissions. Successful exploitation could allow the unauthenticated user to access or modify data, alter Splunk platform configurations, and disrupt operations available to the affected user. The vulnerability is possible because the Portable Document Format (PDF) generation component processes request data before it rejects deprecated GET requests and does not apply SPL safeguards for risky commands to subsearches. For more information see Generate PDFs of your reports and dashboards (https://help.splunk.com/en/splunk-enterprise/create-dashboards-and-reports/reporting-manual/10.4/report-management/generate-pdfs-of-your-reports-and-dashboards) and SPL safeguards for risky commands (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/best-practices-for-splunk-platform-security/spl-safeguards-for-risky-commands) in the Splunk documentation.	7.1	More Details
CVE-2026-131198	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the mc_view_im_data capability could cause the Mission Control incident mapping Representational State Transfer (REST) API to run arbitrary Search Processing Language (SPL) commands with elevated permissions. With this access, the user could expose all data in Splunk indexes and all stored credentials and modify protected indexes and lookup data. The vulnerability is possible because the Mission Control incident mapping REST API does not validate incident identifiers before including them in SPL searches. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) and Overview of Mission Control in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/mission-control/overview-of-mission-control-in-splunk-enterprise-security) in the Splunk documentation.	7.1	More Details
CVE-2026-128135	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.9.4 to 3.10.3, NLTK ZIP package extraction can write through an existing in-root hardlink whose other link points outside the NLTK data root. A local attacker who can prepare the extraction target can cause a package member write to overwrite an out-of-root file, violating the pathsec containment boundary and affecting integrity and availability. This issues is fixed in version 3.10.4.	7.1	More Details
CVE-2026-45142	Chamilo LMS is an open-source learning management system. Prior to 1.11.40 and 2.0.1, main/lp/lp_ajax_save_objectives.php and public/main/lp/lp_ajax_save_objectives.php trust the request uid parameter instead of the authenticated user identity when saving SCORM learning objectives. Any authenticated student can replace uid with another user's identifier and overwrite that user's objectives because the endpoint performs no ownership check. This issue is fixed in versions 1.11.40 and 2.0.1.	7.1	More Details
CVE-2026-130673	Craft CMS from 5.0.0 before 5.11.5 contains an sql injection vulnerability that allows authenticated low-privilege users to inject raw SQL by passing nested array parameters into Element::findOne() criteria. Attackers can set the query's where property through endpoints like users/delete-address and users/save-address to perform blind boolean or time-based injection and read database contents.	7.1	More Details
	In the Linux kernel, the following vulnerability has been resolved: media: imx-jpeg: cancel timeout worker when streaming stops Each per-fd context ctx owns a delayed_work (ctx->task_timer, callback mxc_jpeg_device_run_timeout) armed via schedule_delayed_work() at the end of mxc_jpeg_device_run() to recover a stalled encode/decode job. The only existing cancellation is cancel_delayed_work() in the frame-done IRQ handler, which de-queues a pending work item but does not wait for a callback that has already started, and it only runs when a frame completes. When the fd is closed while a job is in flight (the frame-done IRQ has not fired yet), nothing syncs the worker before mxc_jpeg_release() frees ctx with kfree() after v4l2_m2m_ctx_release(). A queued or executing mxc_jpeg_device_run_timeout() can then recover ctx through container_of(&ctx->task_timer) and dereference it (ctx->mxc_jpeg, slot_data, dev_warn) after ctx		

CVE-2026-99468	has been freed. Cancel the worker from mxc_jpeg_stop_streaming(). The cancel cannot live in mxc_jpeg_release(): mxc_jpeg_device_run() arms the timer while holding only hw_lock, not the mxc_jpeg->lock mutex that release holds, so a cancel in release could still race a concurrent mxc_jpeg_device_run() that re-arms the timer afterwards. mxc_jpeg_stop_streaming() instead runs inside v4l2_m2m_ctx_release() -> vb2_queue_release(), i.e. after v4l2_m2m_cancel_job() has set TRANS_ABORT and waited for any in-flight job to finish (so __v4l2_m2m_try_queue() will not queue and v4l2_m2m_try_run() will not run any further job for this context, which prevents mxc_jpeg_device_run() from re-arming the timer) and before the m2m context is freed. cancel_delayed_work_sync() removes a pending work item and waits for a running callback, so the worker can no longer race with the subsequent kfree(). The cancel is placed before the buffer-release loop so a concurrently running timeout callback cannot race with it over the same buffers. If the frame-done IRQ canceled a still-pending timer, this cancel is a no-op; if the timeout callback has already started, it waits for the callback to finish. The same mxc_jpeg_stop_streaming() call is also reached from VIDIOC_STREAMOFF, which drains the worker early, although STREAMOFF itself does not free ctx -- the use-after-free arises only when the fd is later closed. This bug was found by static analysis.	7.0	More Details
CVE-2026-99625	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix duplicate dependency fence releases drm_sched_job_add_dependency() consumes its fence reference on success and failure. Do not release the fence again if drm_sched_job_add_dependency() fails. (cherry picked from commit e04f3dfb2071451633e6440a3911a9bb65df9144)	7.0	More Details
CVE-2026-99659	In the Linux kernel, the following vulnerability has been resolved: mm/mmu_gather: fix flushing ranges with 4 GiB or larger __tlb_adjust_range() takes the size of the range which has to be flushed as an unsigned int. Its callers, tlb_flush_{pte,pmd,pud,p4d}_range(), get and pass an unsigned long. Consequently, a range that is 4 GiB or larger gets truncated, possibly to 0. On arm64 with hugetlb, a 64k kernel can have 16 GiB hugetlb folios mapped contiguously at the PMD level and flush them when unmapping, and a 16k kernel can flush 64 GiB PUDs through tlb_unshare_pmd_ptdesc(). So at least these two scenarios are affected. Fix it by flushing the proper range instead of the truncated one by properly using "unsigned long" for range_size. Fix generated by an LLM and sent by Josef Bacik in private. [david: use two-tab indent, simplify description]	7.0	More Details
CVE-2026-99320	In the Linux kernel, the following vulnerability has been resolved: ipv6: raw: fix use-after-free in rawv6_send_hdrinc() rawv6_send_hdrinc() caches the route before handing the skb to l3mdev_ip6_out(). The VRF output path may replace skb->dst and release the original route, leaving the cached pointer stale. The function later dereferences that pointer for statistics accounting and as the output device passed to NF_HOOK(). Its RCU read-side critical section starts after l3mdev_ip6_out(), so it cannot prevent the old route from being reclaimed in between. Reload the route from the current skb destination after l3mdev_ip6_out() returns. KASAN reports a slab-use-after-free in rawv6_sendmsg(), with the route allocated by ip6_dst_alloc(), freed by dst_destroy() from rcu_core(), and the related RCU callback queued from vrf_l3_out().	7.0	More Details
CVE-2026-99473	In the Linux kernel, the following vulnerability has been resolved: KVM: guest_memfd: Gracefully handle xarray errors when binding a memslot If inserting a memslot into a guest_memfd's bindings xarray fails, propagate the error back to the caller, i.e. fail memslot creation as well. Signalling success and continuing on with memslot creation results in use-after-free, as the guest_memfd instance will remain reachable via the memslot after the file is freed (kvm_gmem_release() won't nullify the file pointer due to lack of a valid binding). Opportunistically WARN and reject binding if KVM_MEMSLOT_GMEM_ONLY is already set, partly to guard against goofs elsewhere, but mostly so that KVM doesn't need to worry about clobbering flags when unwinding on failure. Regarding the unwind, the slot must be fully prepared before inserting it into the bindings, at which point the slot becomes reachable. I.e. waiting to update the slot in order to avoid the ugly unwind isn't an option. And as part of the unwind, explicitly nullify the relevant bindings, as xarray can store a subset of entries when populating a range.	7.0	More Details
CVE-2026-99316	In the Linux kernel, the following vulnerability has been resolved: ALSA: usx2y: Initialize idle URBs before submitting streams usb_stream_start() submits the first two capture and playback URBs and only afterwards assigns sk->idle_inurb, sk->idle_outurb, sk->completed_inurb and sk->completed_outurb to the remaining two URBs of each direction. A transfer can complete as soon as usb_submit_urb() has been called, so the completion handler i_capture_start() or i_playback_start() may run before those pointers are set. While the stream is still being synchronized (state below usb_stream_sync1), stream_start() then calls playback_prep_freqn(sk, sk->idle_outurb) with sk->idle_outurb still NULL, and the packet descriptors are written through that NULL pointer: Oops: general protection fault, probably for non-canonical address 0xdffffc0000000018: 0000 [#1] SMP KASAN KASAN: null-ptr-deref in range [0x00000000000000c0-0x00000000000000c7] RIP: 0010:playback_prep_freqn+0x11c/0x4a0 sound/usb/usx2y/usb_stream.c:33 Call Trace: <IRQ> stream_start+0x5dc/0x8d0 sound/usb/usx2y/usb_stream.c:596 i_capture_start+0x3da/0x710 sound/usb/usx2y/usb_stream.c:654 __usb_hcd_giveback_urb+0x238/0x550 drivers/usb/core/hcd.c:1657 usb_giveback_urb_bh+0x20b/0x530 drivers/usb/core/hcd.c:1691 tasklet_hi_action+0xa/0x20 kernel/softirq.c:994 handle_softirqs+0x16f/0x440 kernel/softirq.c:645 </IRQ> The four URBs referenced by these pointers are allocated long before and do not depend on any transfer state, so assign them before the submit loop. Keeping the assignment inside the retry path also restores a well-defined URB rotation when the start frame mismatch handling kills and resubmits the streams.	7.0	More Details
	In the Linux kernel, the following vulnerability has been resolved: accel/amdxdna: clear the mailbox channel pointer when starting it fails aie2_create_context() allocates a mailbox channel, starts it, and on a failed start frees it again: hwctx->priv->mbox_chann = xdna_mailbox_alloc_channel(ndev->mbox); ... ret =		

CVE-2026-99485	xdna_mailbox_start_channel(hwctx->priv->mbox_chann, ...); if (ret) goto free_channel; ... free_channel: xdna_mailbox_free_channel(hwctx->priv->mbox_chann); hwctx->priv->mbox_chann keeps pointing at the freed channel. That pointer is the driver's own test for whether a hardware context has a usable channel: aie2_destroy_context() returns early on NULL, aie2_sched_job_run() refuses to run a job on NULL, and the message helpers return -ENODEV on NULL. A stale pointer passes all of them. It matters on the restart path. aie2_hwctx_restart() calls aie2_create_context() again on a hardware context that is already live, after aie2_hwctx_stop() has torn its channel down, and the context survives a failure there: aie2_sched_job_timedout() discards the return value entirely, and aie2_hwctx_resume() only propagates it. What is left is a live hardware context holding a freed channel, and every later user of it takes the non-NULL branch: - the next job submitted reaches xdna_mailbox_send_msg() on the freed channel, via the !mbox_chann guard in aie2_sched_job_run(), - a second command timeout reaches aie2_hwctx_stop() -> aie2_destroy_context(), - closing the device reaches aie2_hwctx_fini() -> aie2_release_resource() -> xrs_release_resource() -> aie2_xrs_unload() -> aie2_destroy_context(), and aie2_destroy_context() then calls xdna_mailbox_stop_channel() followed by xdna_mailbox_free_channel() on memory that was already freed: a use-after-free, and a second free of the same channel. The AIE4 management channel already does this correctly - aie4_mailbox_start() clears ndev->aie.mgmt_chann right after freeing it. Do the same here. The create path itself is not affected: when aie2_xrs_load() fails, xrs_allocate_resource() removes the solver node without calling ->unload, and aie2_hwctx_init() frees hwctx->priv, so the stale pointer never outlives the structure holding it. Starting a channel can fail today without this patch: the ring buffer sizes firmware reports are rejected unless both are powers of two, and request_irq() can fail.	7.0	More Details
CVE-2026-99667	In the Linux kernel, the following vulnerability has been resolved: mm/khugepaged: always recheck the pmd in try_collapse_pte_mapped_thp() try_collapse_pte_mapped_thp() drops the PTE lock after it has cleared the page table, and takes the PMD lock and the PTE lock to remove it. In between, MADV_DONTNEED with CONFIG_PT_RECLAIM can reclaim the page table. The pmd_same() check is supposed to handle this, but it is only done with CONFIG_SPLIT_PTE_PTLOCKS. Fix it by performing the check unconditionally. Fix generated by an LLM and sent by Josef Bacik in private. [david: adjust description]	7.0	More Details
CVE-2026-99476	In the Linux kernel, the following vulnerability has been resolved: fsverity: Set S_VERITY only after rollback is no longer possible Commit f77f281b6118 ("fsverity: use a hashtable to find the fsverity_info") changed the implementation of FS_IOC_ENABLE_VERITY to make concurrent reads start going through fsverity verification slightly earlier than they did before. Instead of it happening when enable_verity() calls fsverity_set_info(), it now happens when ->end_enable_verity() sets S_VERITY. However, ext4 and btrfs can roll back verity after they set S_VERITY, causing a use-after-free. Fix that by making ext4 and btrfs set S_VERITY only after rollback is no longer possible. On failure, don't roll back S_VERITY, but do roll back the filesystem-specific verity flag. Also update enable_verity() to handle any future instances of this bug more safely by warning and *not* freeing the fsverity_info. Finally, tweak some comments. Clarify one, and move information from another to its proper place in the comment for ->end_enable_verity().	7.0	More Details
CVE-2026-99510	In the Linux kernel, the following vulnerability has been resolved: media: hackrf: fix use-after-free in hackrf_alloc_urbs() error path hackrf_alloc_urbs() frees the URBs it allocated so far when one allocation fails, but leaves the entries in dev->urb_list[] and dev->urbs_initialized untouched. The caller, hackrf_start_streaming(), then calls hackrf_free_urbs() on the error path, which walks dev->urbs_initialized entries and calls usb_free_urb() a second time on the already-freed URBs, causing a use-after-free (slab-use-after-free Write in usb_free_urb()). Drop the redundant cleanup loop inside hackrf_alloc_urbs() and let hackrf_free_urbs(), which the caller already invokes on error, own the cleanup of the successfully allocated URBs.	7.0	More Details
CVE-2026-99630	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/userq: validate the moved individual BOs on restore An individual BO of a user queue VM, i.e. one which doesn't share the root reservation, goes to the individual needs_update list when it is evicted. The user queue restore looked at vm->invalidated for them, which the VM state machine rework turned into the always_valid evicted list and commit 18decfc9e8cc ("drm/amdgpu/userq: actually validate userq BOs") into the individual evicted one. That list only holds imported DMA-bufs of compute VMs though, so for a user queue VM nothing is validated. The restore then neither locks such a BO nor moves it back. Its PTEs are updated to wherever it is now, cleared if it went to system memory, and it doesn't get the new eviction fence, so it can move again under the running queue. On navi48, once the BOs of an SDMA user queue are evicted from GTT (debugfs amdgpu_evict_gtt) the queue faults on its ring and rptr: amdgpu 0000:03:00.0: in page starting at address 0x0000000100000000 from client 10 amdgpu 0000:03:00.0: Faulty UTCL2 client ID: SDMA0 (0xd) amdgpu 0000:03:00.0: in page starting at address 0x0000000100101000 from client 10 Validate the BOs on the individual needs_update list like before. (cherry picked from commit eb0d94765d4be18d41e93ad09947dd9cc06ed31a)	7.0	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: KVM: guest_memfd: take the invalidate lock when unbinding a dying file kvm_gmem_unbind() takes mapping->invalidate_lock only if the get_file_active() succeeds. When the guest_memfd file is dying, its reference count is already zero so get_file_active fails and the binding is removed without holding the invalidate lock. kvm_gmem_invalidate_{start,end}() checks f->bindings independently to decide whether to begin or end KVM MMU invalidation. So, the bindings must therefore remain stable between the two calls. Otherwise, unbind can remove a binding after start increments mmu_invalidate_in_progress, but before end finds the binding and decrements it. Example,		

2026-99336	unbind race with memory failure: CPU 0: memory failure CPU 1: memslot delete ---------- (guest_memfd file is dying) kvm_gmem_error_folio() kvm_gmem_invalidate_start() finds binding mmu_invalidate_in_progress++ kvm_gmem_unbind() get_file_active() fails removes binding kvm_gmem_invalidate_end() no binding found counter stays elevated mmu_invalidate_retry() then returns 1 forever, so guest page faults retry without ever installing a mapping and the guest hangs. Remove the distinction between live and dying files in the unbind path. Always use slot->gmem.file and take the invalidate lock when it's non-NULL. Normal unbind callers hold slots_lock, which prevents release from clearing the pointer or freeing the file state until unbind completes. Final VM teardown can only see a NULL pointer because guest_memfd pins the KVM until release has cleared all bindings.	7.0	More Details
CVE-2026-99645	In the Linux kernel, the following vulnerability has been resolved: mm/cma: don't release CMA pages still in use When a driver calls dma_free_contiguous() before quiescing DMA, the page might still have a reference from the device. put_page_testzero() there returns false, WARN fires, but the code proceeds to free_contig_frozen_range() putting a live page onto buddy and clearing the bitmap. A subsequent allocation can hand the same PFN to a new owner while the original holder still references it. A concurrent put_page() that drops the last reference between the testzero loop and free_contig_frozen_range() can double-queue the page via page->lru, corrupting buddy lists. Avoid the corruption by not freeing a CMA region if any of the pages inside are still in use. Make it explicit in the warning that the driver allowed a leak. While at it, use page_ref_freeze() to make the semantics clearer and switch to pr_warn_ratelimited(). [david: use page_ref_freeze() + pr_warn_ratelimited()]	7.0	More Details
CVE-2026-133306	VideoLAN libaacs 0.7.0 through 0.12.0 on non-macOS, non-Windows builds contains stack-based buffer overflows in aacs_resolve_path() in src/file/path.c from length checks against the wrong buffer. Local attackers can place a symlink with a long target on a disc path a user opens to overflow resolved_path[] and link[], crashing the process or potentially executing code if mitigations are bypassed.	7.0	More Details
CVE-2026-99401	In the Linux kernel, the following vulnerability has been resolved: rtnetlink: wait for RCU readers in rtnl_link_unregister() rtnl_link_ops_get() finds ops on link_ops under rcu_read_lock() and only then takes ops->srcu. rtnl_link_unregister() unlinks ops, then calls synchronize_srcu() and cleanup_srcu_struct() without first waiting for RCU readers. A reader that found ops before list_del_rcu() and has not yet taken ops->srcu is not waited for. Such a reader then locks a cleaned-up srcu_struct and keeps using ops after the module that owns it is gone. rtnl_newlink() and rtnl_dump_ifinfo() both look up ops without RTNL, the latter for any user that dumps links filtered by kind. With delays added between the lookup and srcu_read_lock() in rtnl_link_ops_get(), and before rtnl_link_ops_put() in rtnl_dump_ifinfo(), running "ip link show type dummy" alongside "rmmod dummy" on a KASAN kernel gives: BUG: KASAN: vmalloc-out-of-bounds in __srcu_read_unlock+0x57/0x70 Read of size 8 at addr ffffffffc02ff318 by task ip/147 Call Trace: __srcu_read_unlock+0x57/0x70 rtnl_dump_ifinfo.cold+0x30/0x3b rtnl_dumpit+0x7a/0x1b0 netlink_dump+0x486/0x11b0 __netlink_dump_start+0x661/0x910 rtnetlink_rcv_msg+0x611/0x9e0 netlink_rcv_skb+0x126/0x380 netlink_unicast+0x605/0x830 netlink_sendmsg+0x6d0/0xbb0 __sys_sendto+0x335/0x3a0 Call synchronize_rcu() before synchronize_srcu(), as rtnl_af_unregister() already does. This should be enough because rtnl_link_ops_get() is the only RCU walker of link_ops, and it takes ops->srcu before rcu_read_unlock(), so every reader that found ops is covered by synchronize_srcu() once the RCU grace period ends. The extra grace period only costs time at module unload. Readers and the writer were serialized by RTNL until link_ops got its own mutex, which is why the missing RCU wait did not matter before.	7.0	More Details
CVE-2026-99725	In the Linux kernel, the following vulnerability has been resolved: cipso: fix optlen corruption in cipso_v4_delopt() cipso_v4_delopt() recomputes the option length by walking the option buffer after the CIPSO option has been moved out. The walk then reads the stale bytes at the tail of the buffer as option data, and a crafted setsockopt(IP_OPTIONS) buffer can make it treat one of them as an option length, increasing opt->opt.optlen beyond the valid option data. The corrupted length stays on the socket and is later used as the copy length by getsockopt(IP_OPTIONS) and by the udp/raw/ping sendmsg paths, where it exceeds the size of the buffers they copy into. Compute the new length from the intact options before the move and subtract the CIPSO length, and fill the padding with IPOPT_END.	7.0	More Details
CVE-2026-132975	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in issue tracker links was possible	6.8	More Details
CVE-2026-132974	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Promote build dialog for dependent builds was possible	6.8	More Details
CVE-2026-117492	The Flatsome WordPress theme before 3.20.13 does not sanitise post meta saved through an AJAX action, nor escape it when rendering the product page, allowing users with a role as low as Contributor (who lack the unfiltered_html capability) to store JavaScript that executes in the browser of anyone who views the product, including administrators.	6.8	More Details
CVE-2026-83631	On Intel-based macOS systems, the Okta Device Access (ODA) component of Okta Verify for macOS configures a LaunchDaemon service whose program path resolves via a symbolic link in /usr/local/bin. In environments where non-standard directory permissions grant local non-root users write access to this path, the symbolic link target can be modified, causing the LaunchDaemon to execute the replacement binary	6.8	More Details

	upon service restart.		
CVE-2026-132621	A vulnerability has been found in GearUP Hyperev up to 20261020. This affects an unknown function. Such manipulation leads to on-chip debug and test interface with improper access control. The attack can be executed directly on the physical device. The vendor was contacted early about this disclosure but did not respond in any way.	6.8	More Details
CVE-2026-132976	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Inspections (IntelliJ IDEA) build step editor was possible	6.8	More Details
CVE-2026-132977	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Agent Statistics tooltip was possible	6.8	More Details
CVE-2026-132978	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in agent requirements was possible	6.8	More Details
CVE-2026-97329	The Team Members WordPress plugin before 5.4.2 does not properly escape a team member field when it is rendered back in the administration area, allowing users with the contributor role and above to perform Stored Cross-Site Scripting attacks against higher-privileged users who edit the team.	6.8	More Details
CVE-2026-132979	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the checkout rules editor was possible	6.8	More Details
CVE-2026-97330	The Team Members WordPress plugin before 5.4.2 does not properly escape team member social link URLs when they are rendered back in the administration area, allowing users with the Contributor role and above to perform Stored Cross-Site Scripting attacks against higher privileged users who edit the team.	6.8	More Details
CVE-2026-104773	CKAN is an open-source data management system for data hubs and data portals. Prior to 2.10.11 and 2.11.6, the web user registration flow authenticates a newly registered account without first regenerating the session identifier when server-side session storage is configured. An attacker who can set a session cookie in a victim's browser or obtain the victim's current valid session identifier can reuse that identifier after the victim registers, gaining access to the new account. CKAN uses cookie-based sessions by default, and web registration is disabled by default through ckan.auth.create_user_via_web. This issue is fixed in CKAN 2.10.11, 2.11.6, and 2.12.0.	6.8	More Details
CVE-2026-132971	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Schedule Trigger dependency settings was possible	6.8	More Details
CVE-2026-132957	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the configuration health popup was possible	6.8	More Details
CVE-2026-86403	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to check whether the calling session originated from an OAuth authorization grant, which allows a third-party application holding a delegated access token to register a new, independent OAuth application under the identity of the delegating user via the OAuth application registration endpoint.. Mattermost Advisory ID: MMSA-2026-00780	6.8	More Details
CVE-2026-132970	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Finish Build Trigger editor was possible	6.8	More Details
CVE-2026-130717	NocoDB before 2026.04.1 fails to clear reset_password_token and reset_password_expires in users.service.ts when a password is changed, leaving earlier reset links valid. Attackers holding an unexpired reset link, such as from prior mailbox access, can set a new password and regain control of the account.	6.8	More Details
CVE-2026-133228	A flaw was found in Keycloak. When a client application enforces step-up authentication or multi-factor authentication (MFA) via a minimum authentication level, Keycloak fails to apply this requirement to login sessions initiated through the credential reset endpoint. An attacker who possesses a user's initial credentials (such as a password) can navigate from the reset interface back to the standard login form and authenticate without completing the required second factor. This flaw allows an attacker to obtain valid access tokens and gain unauthorized access to protected services without satisfying mandatory security controls.	6.8	More Details
CVE-2026-133537	A flaw was found in Keycloak. When evaluating access permissions, Keycloak Authorization Services improperly matches resource paths containing wildcard file extensions by evaluating only the file extension suffix and ignoring configured path prefixes. An authenticated remote attacker could exploit this vulnerability by requesting a restricted resource that shares a file extension with a more permissive resource, leading to an authorization bypass and unintended access to protected data.	6.8	More Details

CVE-2026-132921	In JetBrains TeamCity before 2026.2 two-factor authentication bypass was possible via access token management	6.8	More Details
CVE-2026-130233	The Okta .NET SDK Abstractions library improperly synchronizes authentication credentials set on a shared HTTP client during request execution. When concurrent requests are dispatched for different principals, one principal's credential may be transmitted on an outbound request intended for another principal, which can result in credential exposure and lead to unauthorized actions performed as the affected principal at the Okta Management API.	6.8	More Details
CVE-2026-132980	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Schedule Trigger branch filter settings was possible	6.8	More Details
CVE-2026-132967	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the build Issues tab was possible	6.8	More Details
CVE-2026-132939	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Maven settings health report was possible	6.8	More Details
CVE-2026-132969	In JetBrains TeamCity before 2026.2, 2026.1.4 stored XSS in the Xcode runner configuration selection was possible	6.8	More Details
CVE-2026-132973	In JetBrains TeamCity before 2026.2, 2026.1.5 stored XSS on the external Kotlin DSL editor status page was possible	6.8	More Details
CVE-2026-130638	Relative path traversal in the object detection dataset preparation tool in AutoGluon Multimodal before 1.6.0 might allow adversary-in-the-middle attackers to write arbitrary files with the privileges of the user running the tool via a crafted dataset archive substituted for a dataset downloaded over plaintext HTTP. To remediate this issue, users should upgrade autogluon.multimodal to version 1.6.0 or later.	6.8	More Details
CVE-2026-133037	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS in the GitHub App repository selector was possible	6.8	More Details
CVE-2026-133047	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS on the thread dump page was possible	6.8	More Details
CVE-2026-133035	In JetBrains TeamCity before 2026.2.1 stored XSS on the change Files tab was possible	6.8	More Details
CVE-2026-133049	In JetBrains TeamCity before 2026.2.1 stored XSS in the Pull Requests health report for VCS roots was possible	6.8	More Details
CVE-2026-133033	In JetBrains TeamCity before 2026.2.1 stored XSS in Maven Build Info module links was possible	6.8	More Details
CVE-2026-133031	In JetBrains TeamCity before 2026.2.1 stored XSS in Git VCS root settings was possible	6.8	More Details
CVE-2026-133256	In JetBrains YouTrack before 2026.2.19840 improper authorization allowed project administrators to obtain GitLab tokens of other projects	6.8	More Details
CVE-2026-133038	In JetBrains TeamCity before 2026.2.1 stored XSS in the Maven Build Info project header was possible	6.8	More Details
CVE-2026-133046	In JetBrains TeamCity before 2026.2.1 stored XSS in recipe export errors was possible	6.8	More Details
CVE-2026-	In JetBrains TeamCity before 2026.2.1 stored XSS in VMware vSphere cloud profile settings was possible	6.8	More Details

133045			
CVE-2026-133042	In JetBrains TeamCity before 2026.2.1 stored XSS in the Amazon EC2 cloud profile editor was possible	6.8	More Details
CVE-2026-133050	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored and reflected XSS in Space project connection dialogs was possible	6.8	More Details
CVE-2026-133051	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS in the JaCoCo coverage settings was possible	6.8	More Details
CVE-2026-133041	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS in the Commit Status Publisher settings was possible	6.8	More Details
CVE-2026-130731	NocoDB before 2026.05.3 contains an account takeover vulnerability that allows invited users to complete other invitees' signups because batch invitations share one token across all recipients. Attackers holding a valid batch invite can submit another invitee's email with the shared non-expiring token to set its password and take over that account, including higher-role ones.	6.8	More Details
CVE-2026-133052	In JetBrains TeamCity before 2026.2.1, 2026.1.5 a DOM-based XSS in the plugin installation dialog was possible	6.8	More Details
CVE-2026-127003	The RSVP and Event Management WordPress plugin before 2.7.19 does not sanitise and escape attendee data imported from a spreadsheet before outputting it on the frontend RSVP form, which could allow users with the Author role and above to perform Stored Cross-Site Scripting attacks against anyone, including administrators, viewing the affected attendee, even when the unfiltered_html capability is disallowed.	6.8	More Details
CVE-2026-133098	In JetBrains TeamCity before 2026.2.2 stored XSS in GitHub issue tracker suggestions was possible	6.8	More Details
CVE-2026-133054	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS in the Create Project dialog was possible	6.8	More Details
CVE-2026-133100	In JetBrains TeamCity before 2026.2.2, 2026.1.6 command injection in Project Import artifact copy scripts was possible	6.8	More Details
CVE-2026-133039	In JetBrains TeamCity before 2026.2.1 stored XSS in the Pull Requests health report for Space connections was possible	6.8	More Details
CVE-2026-133103	In JetBrains TeamCity before 2026.2.2, 2026.1.6 stored XSS in the enforced settings restriction tooltip was possible	6.8	More Details
CVE-2026-133106	In JetBrains TeamCity before 2026.2.2, 2026.1.6 stored XSS on the Projects Import page was possible	6.8	More Details
CVE-2026-132983	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Report Tabs settings page was possible	6.8	More Details
CVE-2026-132982	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the build Report Tab dialog was possible	6.8	More Details
CVE-2026-133036	In JetBrains TeamCity before 2026.2.1 stored XSS in the Issue Trackers health report was possible	6.8	More Details
CVE-2026-123893	SDDC Manager contains a local privilege-escalation vulnerability. Affected versions: 9.1.x.x (fixed in 9.1.1.0200); 9.0.x.x (fixed in 9.0.2.0100); 5.x (fixed in 5.2.5).	6.7	More Details
CVE-	VMware Tools for Windows contains a SAML token signature bypass vulnerability. Affected versions: 13.x.x.x		More

2026-123894	(fixed in 13.1.5.0); 12.3.x (fixed in 12.5.5).	6.7	Details
CVE-2026-130663	ImageMagick before 7.1.2-35 contains a code injection vulnerability in the MAGICK, XBM, XPM and UIL encoders that allows local attackers to inject content via crafted filenames. Attackers with access to the output directory and control over the filename can inject additional headers or code into the generated image source output.	6.7	More Details
CVE-2026-131213	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_analyst" Splunk Enterprise Security role could store attacker-controlled data in finding exclusion rule metadata. Another authenticated user could later open the Finding Exclusion Rules page. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who opened the page. The first user could access data and perform actions available to the later user. The vulnerability is possible because Splunk Enterprise Security does not escape backslashes in finding exclusion rule metadata before adding the values to a page search. For more information see Finding exclusions in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/user-and-entity-behavior-analytics/finding-exclusions-in-splunk-enterprise-security) and Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) in the Splunk documentation.	6.7	More Details
CVE-2026-132759	GPAC MP4Box before 26.07.0 contains a use-after-free vulnerability because gf_node_free() frees nodes without detaching routes that still reference them via FromNode. Attackers can supply a crafted XMT scene with nested PROTO graphs so gf_sg_route_del() accesses freed memory, crashing MP4Box and potentially corrupting the heap.	6.6	More Details
CVE-2026-63483	X-Road is open-source data exchange layer software. Prior to 7.8.2 and 7.7.3, the Configuration Proxy uses attacker-controlled Content-location metadata from a validly signed federated global configuration part when constructing an output path without verifying that the resulting path remains inside the intended configuration output directory. A malicious or compromised federated ecosystem can provide an absolute path or directory traversal sequence during the normal refresh cycle, causing the Configuration Proxy to create a new file with attacker-controlled content under the permissions of the non-privileged xroad process. The vulnerability can poison downstream global configuration or create files in other writable locations, but it does not overwrite existing files. This issue is fixed in versions 7.8.2 and 7.7.3.	6.6	More Details
CVE-2026-132764	GPAC before 26.07.0 contains a use-after-free vulnerability in the XMT loader that leaves ROUTE commands of a discarded node queued after being freed by gf_sg_command_del(). Attackers can supply a crafted XMT scene file imported with MP4Box so xmt_resolve_routes() writes to freed memory, corrupting the heap and crashing the process.	6.6	More Details
CVE-2026-132765	GPAC MP4Box through 26.07.0 contains a use-after-free vulnerability in gf_node_try_destroy because Conditional node commands keep raw pointers to target nodes freed by gf_sg_reset(). Attackers can supply a crafted XMT scene processed with -info or -add to trigger heap use-after-free, crashing the process and corrupting heap memory.	6.6	More Details
CVE-2026-132766	GPAC before 26.07.0 contains a use-after-free vulnerability in MP4Box because gf_sg_reset frees DEF'd proto sub-graph nodes without removing them from the GF_Proto node_code list. Attackers can supply a crafted XMT scene with nested PROTOs so gf_sg_proto_del unregisters freed nodes during teardown, crashing MP4Box and corrupting the heap.	6.6	More Details
CVE-2026-131694	Tanium addressed an improper check for certificate revocation vulnerability.	6.6	More Details
CVE-2026-132772	GPAC MP4Box before 26.07.0 contains a use-after-free vulnerability in the XMT parser that leaves parser->command_buffer pointing into a freed InputSensor node. Attackers can supply a crafted XMT scene file triggering a parse error so xmt_node_start() writes NULL into freed memory, corrupting the heap and crashing MP4Box.	6.6	More Details
CVE-2026-132775	GPAC before 26.07.0 contains a use-after-free vulnerability in xmt_node_end() that allows attackers to corrupt memory by attaching nodes from unrelated proto sub-graphs to field-replace commands. Attackers can supply a crafted XMT scene buffering such a command in a Conditional node, causing gf_node_try_destroy() to access the freed node, crashing MP4Box and corrupting the heap.	6.6	More Details
CVE-2026-132763	GPAC before 26.07.0 contains an integer underflow vulnerability in gf_xml_sax_peek_node() that allows attackers to write a zero byte outside the heap line buffer via crafted XML attribute values. Attackers can supply a crafted XMT file with a ROUTE or command referencing an undefined node, imported with MP4Box, to crash the process and corrupt memory.	6.6	More Details
CVE-	OliveTin gives safe and simple access to predefined shell commands from a web interface. From 3000.0.0 until 3000.18.1, OliveTin passes attacker-influenceable output and ot_username values into shellAfterCompleted processing in service/internal/executor/executor.go without applying checkShellArgumentSafety, the shell-safety validation used for primary action execution. The same post-execution template path also expands exitCode without that validation, but the advisory does not establish		More

2026-75506	attacker control of the exit-code value. When an action configures shellAfterCompleted with output or ot_username, shell metacharacters from command output, external data processed by an action, or a trusted-header username can become syntax in the command executed by the stepExecAfter path. This can execute arbitrary commands as the OliveTin process user, including when the primary action uses the safer exec-array mode. This issue is fixed in version 3000.18.1.	6.6	Details
CVE-2026-132933	In JetBrains TeamCity before 2026.2 privilege escalation of project-scoped access tokens via concurrent project deletion was possible	6.6	More Details
CVE-2026-111516	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain a Least Privilege Violation vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Elevation of Privilege and Denial of Service.	6.6	More Details
CVE-2026-132246	In JetBrains TeamCity before 2026.1.3, 2025.11.7 privilege escalation to system administrator was possible after disabling per-project permissions	6.6	More Details
CVE-2026-132777	GPAC MP4Box before 26.07.0 contains a use-after-free vulnerability in Conditional_execute() because SetupConditional() stores a BIFS decoder pointer that gf_sm_load_run_isom() later frees without clearing. Attackers can supply a crafted MP4 file with a BIFS Conditional node processed with -nstatx to crash MP4Box and potentially corrupt the heap.	6.6	More Details
CVE-2026-77378	Bolt is a fully open source, adaptable platform for building and running modern websites. Prior to 6.1.4, SQL injection in ListFormatHelper allows authenticated users with upload permission to execute arbitrary SQL queries through /async/select-options and exfiltrate sensitive database contents when general/caching/list_format is enabled. This issue is fixed in version 6.1.4.	6.5	More Details
CVE-2026-127132	The GamiPress – Gamification plugin to reward points, badges & ranks in WordPress, now with AI plugin for WordPress is vulnerable to generic SQL Injection via the 'name-1[compare]' parameter in all versions up to, and including, 8.1.3 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with 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 that an administrator has published at least one GamiPress achievement step triggered by a Forminator field-value submission event, as the GamiPress Forminator integration listener is the only path that routes a form submission's raw field array — including the attacker-controlled compare operator — into gamipress_query_logs().	6.5	More Details
CVE-2026-88902	Mattermost versions 11.7.x <= 11.7.11, 11.10.x <= 11.10.2, 11.11.x <= 11.11.1 fail to bound the work performed while reading EXIF orientation metadata, which allows an authenticated user to cause disproportionate server-side CPU usage by uploading or linking a small image with specially structured EXIF metadata.. Mattermost Advisory ID: MMSA-2026-00784	6.5	More Details
CVE-2026-133029	In JetBrains TeamCity before 2026.2.1 missing authorization allowed guests to push VCS labels using stored VCS credentials	6.5	More Details
CVE-2026-133028	In JetBrains TeamCity before 2026.2.1 missing authorization allowed running modified build steps via pipeline debug runs	6.5	More Details
CVE-2026-121591	Mattermost Desktop App versions <=6.3 6.0.3 6.3.2.0 fail to bound the Calls widget resize dimensions before applying them to the window, which allows a configured Mattermost server to terminate the Desktop App process via a crafted resizeCallsWidget IPC message sent from the Calls widget window. Mattermost Advisory ID: MMSA-2026-00787	6.5	More Details
CVE-2026-131185	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who does not hold the "admin" or "power" Splunk roles could use the apps backup and restore Representational State Transfer (REST) API to prevent subsequent apps backup or restore operations until Splunk Enterprise restarts. This could result in a Denial of Service (DoS). The vulnerability is possible because the apps backup and restore REST API handlers do not correctly enforce authorization.	6.5	More Details
CVE-2026-131186	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user that holds a role with the list_deployment_server capability could access sensitive information from inputs.conf files outside the intended deployment app directory through the Representational State Transfer (REST) API. The vulnerability is possible because the REST API does not restrict a deployment app name to an existing deployment app before reading input configurations. For more information see authorize.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.4/configuration-file-reference/10.4.4-configuration-file-reference/authorize.conf), Deployment server architecture (https://help.splunk.com/en/splunk-enterprise/administer/update-your-deployment/9.0/deployment-server-and-forwarder-management/deployment-server-architecture), and inputs.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/9.4/configuration-file-	6.5	More Details

	reference/9.4.16-configuration-file-reference/inputs.conf) in the Splunk documentation.		
CVE-2026-120564	The Search Everything plugin for WordPress is vulnerable to Sensitive Information Exposure in versions up to, and including, 8.1.9 via the search_everything AJAX action. The handler registered at `wp_ajax_search_everything` performs no capability check, no ownership check, and no nonce verification before running a WP_Query over `$_GET['s']` and returning the raw WP_Post objects — including the `post_password` field — as JSON. This makes it possible for authenticated attackers, with Subscriber-level access and above, to extract cleartext passwords for every password-protected post on the site and, when the plugin's draft/attachment/metadata search options are enabled by an administrator, to read draft and pending post contents they otherwise have no capability to view.	6.5	More Details
CVE-2026-131876	fqterm 0.9.9 before 0.9.11 contains an out-of-bounds memory access vulnerability in its SSH client that trusts unchecked server-supplied length fields for allocations, copies and read positions. Malicious servers or man-in-the-middle attackers can send crafted lengths in getString(), KEXINIT name-lists, key parsing or packet framing before host key verification to crash the client.	6.5	More Details
CVE-2026-130159	RepoManager is a self-hosted deb/rpm package repository. Prior to 5.13.3, the authenticated get-unit-log action accepts a caller-controlled logfile value. Validate::string() escapes HTML but leaves path traversal components unchanged. The action joins the caller-controlled logfile to SERVICE_LOGS_DIR without enforcing containment, so traversal can resolve the file read outside that directory. www/controllers/ajax/status/service.php reads and returns the resolved file, and any authenticated user can invoke the action to retrieve files readable by the www-data service account, including application configuration, source code, and the GPG signing-key passphrase. Fixed in 5.13.3.	6.5	More Details
CVE-2026-2307	The Elementor Website Builder – More Than Just a Page Builder plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 4.0.5 via `/wp-json/elementor/v1/user` REST endpoint. This makes it possible for authenticated attackers, with Contributor-level access and above, to extract sensitive data including emails, password hashes and usernames for all WordPress users.	6.5	More Details
CVE-2026-127060	@fastify/http-proxy is a Fastify plugin that proxies HTTP and WebSocket traffic to an upstream server. In versions before 11.6.10, it rejects destinations containing parent-directory path segments by an exact string match, so a segment that carries a semicolon path parameter, such as a dot-dot segment followed by a semicolon and an optional parameter, is not equal to the plain dot-dot form and passes the check, and the WHATWG URL parser keeps it as an ordinary segment so the rewrite-prefix containment check also passes. Servlet containers such as Apache Tomcat and Jetty strip semicolon path parameters from each segment before resolving dot-segments, so a request routed through the configured prefix is rewritten and reaches the upstream where it resolves above the intended subtree, escaping the path the proxy is meant to expose and, on some containers, crossing application contexts. Both the HTTP and WebSocket proxy paths are affected. The issue is fixed in @fastify/http-proxy 11.6.10, and users should upgrade to 11.6.10 or later.	6.5	More Details
CVE-2026-133099	In JetBrains TeamCity before 2026.2.2 vCS root passwords could be extracted via the REST API property locator	6.5	More Details
CVE-2026-133101	In JetBrains TeamCity before 2026.2.2 a CSRF attack was possible on the diagnostic page	6.5	More Details
CVE-2026-133030	In JetBrains TeamCity before 2026.2.1 the build parameter editor allowed project settings viewers to recover protected credentials	6.5	More Details
CVE-2026-131844	Poppler through 26.11.0 contains an uncontrolled memory allocation vulnerability in the libcurl-based remote file loader that trusts unvalidated HEAD Content-Length values in CurlCachedFileLoader::init(). Malicious HTTP servers can return an enormous Content-Length so CachedFile resizes its chunk cache accordingly, aborting pdfinfo or PDFDocFactory-based applications with std::bad_alloc for denial of service.	6.5	More Details
CVE-2026-132937	In JetBrains TeamCity before 2026.2, 2026.1.4 arbitrary file write via read-write file channels in the Kotlin DSL sandbox was possible	6.5	More Details
CVE-2026-95790	A vulnerability in the web-based management interface of network device management systems of AOS-8 Mobility Controllers could allow an authenticated remote attacker to bypass authorization controls. Successful exploitation could allow a low-privilege attacker to read sensitive configuration data potentially leading to unauthorized network access and further compromise of network security.	6.5	More Details
CVE-2026-131652	Tanium addressed an information disclosure vulnerability.	6.5	More Details
	Missing Authorization vulnerability in the Apache Iggy server. The system snapshot command (GetSnapshot on the binary transports and POST /snapshot on HTTP) checked only that the caller was logged in, not the		

CVE-2026-58320	read_servers or manage_servers permission. Any authenticated user, including a user with no permissions, could download a diagnostic archive with the server log files, the host process list, a file system listing and the runtime server configuration. The runtime configuration contains the HTTP JWT signing secret and the message encryption key when they are set in the configuration. If HTTP is enabled, the user could use the signing secret to create tokens for any account, including root. This issue affects Apache Iggy: before 0.9.0. Users are recommended to upgrade to version 0.10.0, which fixes the issue.	6.5	More Details
CVE-2026-16021	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to bound the resources consumed while processing an uploaded custom emoji image which allows an authenticated user with default permissions to make the server unresponsive for all users via repeated custom emoji uploads.. Mattermost Advisory ID: MMSA-2026-00720	6.5	More Details
CVE-2026-120545	The PowerPress Podcasting plugin by Blubrry plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 11.18.2 via the 'podcast-images[] (actually `${feed_slug}-images[]`)' parameter parameter. This makes it possible for authenticated attackers, with contributor-level access and above, to read the contents of arbitrary image files on the server, which can contain sensitive information.	6.5	More Details
CVE-2026-129049	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Go Maps WP Go Maps wp-google-maps allows Stored XSS.This issue affects WP Go Maps: from n/a through 10.1.12.	6.5	More Details
CVE-2026-133173	WatchAlert 2.0.0 through 3.9.20 contains an insecure direct object reference vulnerability in the /api/w8t/datasource/dataSourceGet endpoint that lets authenticated users read other tenants' datasources because lookups omit tenant_id filtering. Tenant members with datasource view permission can request another tenant's datasource ID to obtain basic auth passwords, AWS and Alibaba Cloud keys, and Kubernetes kubeconfigs.	6.5	More Details
CVE-2026-128140	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.6.7 until 3.10.4, nltk.downloader follows collection membership without tracking already visited collection nodes. A self-referencing or mutually referencing collection index can therefore make the package-closure walk loop indefinitely while its queue and memory usage grow, denying service when a victim processes the attacker-influenced index. This issue is fixed in version 3.10.4.	6.5	More Details
CVE-2026-133113	Gogs through 0.14.4 contains an authorization bypass vulnerability that allows collaborators with write access to push to protected branches by merging pull requests. Because MergePullRequest skips branch protection checks and PullRequest.Merge omits hook environment variables, attackers can merge commits into branches restricted to allowlisted users.	6.5	More Details
CVE-2026-133043	In JetBrains TeamCity before 2026.2.1 missing authorization allowed reading VCS file contents and patches of other projects	6.5	More Details
CVE-2026-118521	The Toolbox for Asgaros Forum plugin for WordPress is vulnerable to arbitrary file deletion due to insufficient file path validation in the replace_temp_pics function in all versions up to, and including, 1.2.4. This makes it possible for authenticated attackers, with subscriber-level access and above, to delete arbitrary files on the server, which can easily lead to remote code execution when the right file is deleted (such as wp-config.php). This vulnerability is only reachable when the site administrator has enabled the Inline Image Upload feature (disabled by default).	6.5	More Details
CVE-2026-130696	NocoDB before 2026.10.1 contains a credential exposure vulnerability in the connection-test endpoint that allows authorized integration users to redirect stored credentials by overriding connection host, port, and user fields. Attackers can point the test at a server they control to capture offline-crackable PostgreSQL md5 or MySQL handshakes and recover the stored database password.	6.5	More Details
CVE-2026-101167	Missing Authorization vulnerability in Dataprom Bilişim Teknolojileri LTD. ŞTİ. DataPer Access Control and Security Systems allows Privilege Escalation. This issue affects DataPer Access Control and Security Systems: before v14.15(2026/11).	6.5	More Details
CVE-2026-129014	Missing Authorization vulnerability in DynamiApps Frontend Admin by DynamiApps acf-frontend-form-element allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Frontend Admin by DynamiApps: from n/a through 3.30.2.	6.5	More Details
CVE-2026-121584	Mattermost versions 11.7.x <= 11.7.11, 11.10.x <= 11.10.2, 11.11.x <= 11.11.1 fail to bound the number of records a single GraphQL request may hydrate which allows an authenticated user to exhaust server memory and cause a denial of service via a small query that repeats the runs field under many aliases. Mattermost Advisory ID: MMSA-2026-00767	6.5	More Details
CVE-2026-108323	Tanium addressed an uncontrolled resource consumption vulnerability in Tanium Client.	6.5	More Details
	OpnForm before 2.8.2 contains an improper authorization vulnerability that allows unauthenticated attackers		

CVE-2026-133239	to move other forms' shared assets by supplying their URLs in file upload fields. Attackers can submit a public form with another form's cover or logo URL, causing Storage::move() to relocate it into submissions and break that asset.	6.5	More Details
CVE-2026-130762	Fleet before 4.92.0 contains an SQL injection vulnerability that allows authenticated observer-level users to sort software listings by arbitrary columns via the order_key parameter. Attackers can order by joined hosts table columns such as node keys and use pagination as a binary-search oracle to extract sensitive column values.	6.5	More Details
CVE-2026-132747	Vtiger CRM through 8.4.0 contains an SQL injection vulnerability that allows authenticated users with export access to inject SQL through unvalidated selected_ids values in ExportData.php. Attackers can send a UNION SELECT payload in selected_ids[] to the ExportSelectedRecords mode and read password hashes and access keys from the exported CSV.	6.5	More Details
CVE-2026-131222	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the edit_modinput_identity_manager capability could use Asset and Identity Management to modify protected saved searches and role capabilities. With this access, the user could obtain full administrator access to the Splunk platform. The vulnerability is possible because Asset and Identity Management does not validate the app namespace before sending a Representational State Transfer (REST) API request with a system authentication token. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) and Add asset and identity data to Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/asset-and-identity-management/add-asset-and-identity-data-to-splunk-enterprise-security) in the Splunk documentation.	6.5	More Details
CVE-2026-14879	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 does not invalidate session after a logout which could allow an authenticated user to impersonate another user on the system.	6.5	More Details
CVE-2026-125806	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Media WP Rocket wp-rocket allows Stored XSS.This issue affects WP Rocket: from n/a before 3.24.1.	6.5	More Details
CVE-2026-131294	In Splunk SOAR versions below 8.8.0, an authenticated user who can edit and run playbooks on a container label could associate a custom-function callback with another playbook run. The callback could modify the state of a container outside the authorized scope for that user. The vulnerability is possible because Splunk SOAR does not bind playbook API requests to the trusted playbook run and container context.	6.5	More Details
CVE-2026-130726	NocoDB before 2026.07.0 contains an SQL injection vulnerability in its MySQL and SQLite clients that inserts the column precision value into generated data-definition statements without numeric validation. Authenticated users with the Creator or Owner base role can supply a crafted precision value when creating or updating a column to close the type declaration and execute arbitrary SQL against the base's data source.	6.5	More Details
CVE-2026-130781	Incus before 7.6.0 contains a resource exhaustion vulnerability in the storage bucket S3 listener, which buffers entire signed request bodies in memory before verifying the signature or credential role. Attackers knowing only a valid access key, or holding a read-only credential, can send many concurrent large requests to exhaust daemon memory.	6.5	More Details
CVE-2026-86866	Tanium addressed a SQL injection vulnerability in multiple products.	6.5	More Details
CVE-2026-132750	Meilisearch before 1.55.1 and 1.56.x before 1.56.1 contains a quadratic memory consumption vulnerability in filter-parser that allows attackers with a search API key to exhaust server memory. Attackers can send a roughly 100 KB filter, such as a _geoPolygon with thousands of coordinates, to the search endpoint, crashing the server and taking all indexes offline.	6.5	More Details
CVE-2026-132749	Vtiger CRM through 8.4.0 contains a blind SQL injection vulnerability in ReportRun::generateAdvFilterSql() that allows authenticated low-privileged users to inject SQL via the bw comparator. Attackers can send a crafted advanced_filter parameter containing a single quote to the Reports getRecordsCount action to break out of BETWEEN literals and extract database contents.	6.5	More Details
CVE-2026-131280	In Splunk SOAR versions below 8.8.0, an authenticated user whose role grants system_settings:view but not assets:view could access asset and application inventory, tenant identifiers, connector run messages, and stored configuration snapshots through the ingestion status Representational State Transfer (REST) API. The vulnerability is possible because the ingestion status REST API checks the system_settings:view permission but does not enforce assets:view or tenant restrictions before it returns asset-derived ingestion records. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.3.1/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises), Add and configure apps and assets to provide actions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.4.0/manage-your-	6.5	More Details

	splunk-soar-on-premises-apps-and-assets/add-and-configure-apps-and-assets-to-provide-actions-in-splunk-soar-on-premises), and Configure multiple tenants on your Splunk SOAR (On-premises) instance (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.4.1/configure-product-settings-for-your-splunk-soar-on-premises-instance/configure-multiple-tenants-on-your-splunk-soar-on-premises-instance) in the Splunk documentation.		
CVE-2026-131285	In Splunk SOAR versions below 8.8.0, an authenticated user who can edit and run playbooks could cause automation batches for other events to stop running. The vulnerability occurs because Splunk SOAR does not bind playbook API requests to the server-provided execution context.	6.5	More Details
CVE-2026-130775	Incus before 7.6.0 contains an authorization bypass vulnerability in the /dev/incus guest-API image export handler that fails to pin the project parameter to the guest's own project. Container root users with the guest-API image entitlement can supply another project in the project query parameter to export public or cached images, including non-public images pulled from authenticated remotes.	6.5	More Details
CVE-2026-133274	THC-Hydra through 9.7 contains a stack buffer overflow vulnerability in analyze_server_response() of hydra-http-form.c that allows malicious HTTP servers to corrupt memory via crafted Set-Cookie headers. Attackers can send a Set-Cookie header with an extra leading space followed by a repeated cookie name, causing a negative strncpy length that overflows tmpcookie and crashes Hydra.	6.5	More Details
CVE-2026-132746	Vtiger CRM through 8.4.0 contains a blind SQL injection vulnerability that allows authenticated low-privileged users to inject SQL via unescaped clientid values in the wsapp_map webservice operation. Attackers can register a sync application with wsapp_register, then send delete-mode entries with quoted clientid values processed by SyncServer.php to extract database contents including administrator access keys.	6.5	More Details
CVE-2026-76545	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could determine whether a username exists through the Representational State Transfer (REST) API or Splunk Web and lock out existing user accounts through repeated login attempts. The vulnerability is possible because Splunk Enterprise returns different login failure responses for locked existing user accounts and nonexistent usernames.	6.5	More Details
CVE-2026-131291	In Splunk SOAR versions below 8.8.0, a Splunk SOAR user who can edit and run playbooks for a permitted container label could associate action and callback state with a running playbook outside that container label. This could affect playbook automation for containers that the user cannot view. The vulnerability is possible because the Playbook Automation API accepts a supplied playbook run identifier without verifying that it is bound to the server-provided playbook execution context. The vulnerability requires a running playbook on a container label that the user cannot view.	6.5	More Details
CVE-2026-132745	Vtiger CRM through 8.4.0 contains a blind SQL injection vulnerability that allows authenticated non-administrative users to inject SQL through the k (does not contain) comparator in list view search_params. Attackers can send crafted search_params to the ListAjax getRecordsCount mode, breaking out of the unescaped NOT LIKE clause to extract database contents via boolean or time-based techniques.	6.5	More Details
CVE-2026-129993	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, the settings.getWebServerSettings protected procedure returns the complete global web-server settings object to any authenticated member, including the stored sshPrivateKey, even though settings.saveSSHPrivateKey requires an administrator. A default member can therefore retrieve infrastructure, deployment, or server-management credentials that only an administrator can configure. This issue is fixed in version 0.31.1.	6.5	More Details
CVE-2026-19865	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to limit the aggregate memory allocated when processing ZIP archive entries during document content extraction and Slack workspace import, which allows an authenticated attacker to cause a server-side denial of service via a crafted ZIP archive.. Mattermost Advisory ID: MMSA-2026-00772	6.5	More Details
CVE-2026-76539	In Splunk Enterprise versions below 10.4.4, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could use the search jobs Representational State Transfer (REST) API to access indexed events that the srchFilter setting for the role should hide. The same user could also run internal Search Processing Language (SPL) commands and bypass SPL safeguards for risky commands within indexes that the role can search. The vulnerability is possible because Splunk Enterprise does not correctly escape required field values before it includes them in an internally generated search. For more information, see Search endpoint descriptions (https://help.splunk.com/en/splunk-enterprise/rest-api-reference/9.4/search-endpoints/search-endpoint-descriptions), authorize.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.4/configuration-file-reference/10.4.5-configuration-file-reference/authorize.conf), and SPL safeguards for risky commands (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/best-practices-for-splunk-platform-security/spl-safeguards-for-risky-commands).	6.5	More Details
CVE-2026-130769	Incus before 7.6.0 contains a denial of service vulnerability that allows clients permitted to create instances to crash the daemon by importing a backup with a null snapshot entry in index.yaml. Attackers can trigger an unrecovered nil pointer dereference during configuration rewrite, terminating the daemon and all containers and virtual machines on the node.	6.5	More Details

CVE-2026-131300	In Splunk App for Stream versions below 8.1.7, a user who can read splunkd_access.log or the _internal Splunk index could recover the Stream forwarder authentication key and use it to impersonate a trusted Stream forwarder. The vulnerability is possible because Splunk App for Stream sends the authentication key in an internal request URL, and Splunk records that URL in splunkd_access.log. For more information see About Splunk Stream (https://help.splunk.com/en/splunk-enterprise/collect-stream-data/install-and-configure-splunk-stream/8.1/introduction/about-splunk-stream) and restmap.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/9.4/configuration-file-reference/9.4.3-configuration-file-reference/restmap.conf) in the Splunk documentation.	6.5	More Details
CVE-2026-69138	djangocms-versioning is a general-purpose versioning package for django CMS 4 and later. Prior to 2.6.1, the check_discard and check_edit_redirect Conditions on djangocms_versioning.models.Version omit user_can_change, while VersionAdmin.discard_view and VersionAdmin.edit_redirect_view obtain a version by its primary key and rely on those conditions for object-level authorization. With CMS_PERMISSION set to True, the affected views are gated only by is_staff, allowing any staff user to select content without object-level change permission. With the default CMS_PERMISSION value of False, a staff user must also hold the global Django change_ permission, but that model-level permission does not enforce the intended per-object authorization, so the user can delete a draft version and its content object through discard or create an editable draft from a published version through edit_redirect even when the user cannot change that specific content object. This issue is fixed in version 2.6.1.	6.5	More Details
CVE-2026-131292	In Splunk SOAR versions below 8.8.0, a user who can edit and run a playbook on a permitted label could cause Splunk SOAR to return custom-function status and callback state from another playbook run. The issue results from insufficient binding between a custom-function-status request and the playbook run that initiates it.	6.5	More Details
CVE-2026-130715	NocoDB before 2026.08.0 contains an sql injection vulnerability that allows authenticated Creator or Owner users to execute database functions by supplying zero-argument function calls as column default values. Attackers can plant functions like version() or inet_server_addr() in PostgreSQL or MySQL column definitions and read results through record APIs, disclosing server addresses, role names, and engine details.	6.5	More Details
CVE-2026-130772	Incus before 7.6.0 contains an improper authorization vulnerability that allows restricted TLS clients to read other projects' bucket S3 keys and profile configuration through recursive and all-projects responses. Attackers can query bucket listings with recursion=2 or profiles with all-projects=true to obtain S3 secret keys and sensitive environment, cloud-init and user values across projects.	6.5	More Details
CVE-2026-130711	NocoDB before 2026.09.0 contains an authorization bypass vulnerability in extract-ids.middleware.ts that allows authenticated users to resolve roles against another base by supplying an MCP token id in the query string. Attackers with Viewer access on one base can mint an MCP token and pass its id to read another base's metadata, including table lists and member email addresses.	6.5	More Details
CVE-2026-131854	Ettercap before 0.8.5.2 contains a heap out-of-bounds read vulnerability in dissector_mysql() in src/dissectors/ec_mysql.c that copies seed and password bytes without bounds checks. A malicious MySQL server or client can send truncated greeting or login packets to leak adjacent heap memory through hex-encoded output in the interface and log.	6.5	More Details
CVE-2026-63012	baserCMS is a website development framework. Prior to 5.3.0, getAuthenticationService permits an Api/Admin session authentication fallback based on the Referer header, although operator accounts have an Admin auth prefix and cannot obtain an Api/Admin JWT. An authenticated operator can therefore access /baser/api/admin/ endpoints whose permission.php entries have auth => true, including /baser/api/admin/baser-core/site_configs/view.json, even when the corresponding administrative functionality is restricted. This can disclose administrator-restricted configuration information. This issue is fixed in 5.3.0.	6.5	More Details
CVE-2026-71425	Outstatic is a static site CMS for Markdown and MDX content stored in GitHub repositories. Prior to 2.2.1, the authenticated GitHub REST proxy accepts an unvalidated endpoint value, constructs an upstream URL by string concatenation, and attaches the current dashboard user's GitHub OAuth token without ensuring the destination remains the exact GitHub API origin. An authenticated dashboard user can cause the server to send that user's token to a non-GitHub host, potentially exposing private repositories and other resources allowed by the token's scopes. The issue does not let an unauthenticated attacker access private data and does not directly expose another user's token without an additional same-origin execution or request-forging primitive. This issue is fixed in version 2.2.1.	6.5	More Details
CVE-2026-133251	In JetBrains Datalore before 2026.4.1 improper authorization allowed notebook viewers to read attached environment variable values	6.5	More Details
CVE-2026-108312	Tanium addressed an improper access controls vulnerability in Tanium Server.	6.5	More Details
CVE-2026-107949	A SQL injection vulnerability affects an administrative API in HCL Commerce versions 8, 9.0.1 and 9.1.0 - 9.1.21. A remote attacker with an administrative role could use a crafted URL to infer unauthorized data from the database.	6.5	More Details

CVE-2026-62245	X-Road is open-source data exchange layer software. Prior to 7.8.2, 7.7.3, and 7.6.4, HashChainVerifier in the common-message module resolves attacker-controlled batch-signature hash chains before verifying the message signature and signing certificate, without cycle detection, recursion-depth or total-step limits, or per-step memoization. A registered ecosystem member can send a small crafted hashchain MIME part with repeated or cyclic StepRef values so that verifyHashChain() performs exponential work or unbounded recursion before authentication. Repeated requests can exhaust Security Server proxy threads and deny data-exchange service, and a self-reference can raise an uncaught error. This issue is fixed in versions 7.8.2, 7.7.3, and 7.6.4.	6.5	More Details
CVE-2026-123343	The BlockArt Blocks – Gutenberg Blocks, Page Builder Blocks ,WordPress Block Plugin, Sections & Template Library plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.4.3. 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 contributor-level access and above, to view the title, excerpt, and full content of other users' draft, pending, and private posts, exposing that content to unauthenticated visitors on the published permalink. This is exploitable by an Author-level user because WordPress's wp_filter_post_kses preserves block-comment JSON in HTML comments, allowing the malicious queryArgs and query[post_status] attributes to survive post saving and execute on every unauthenticated request to the published URL. Exploitation by attackers with contributor-level access and above requires that attacker only needs to see the data themselves (not expose it to unauthenticated visitors) — satisfied by viewing their own preview URL. On a default installation, where that is not configured, the vulnerability requires attackers with author-level access and above.	6.5	More Details
CVE-2026-121592	Mattermost Plugins versions <=12.-1 11.0.11 11.11.7 11.2.10.0 fail to check for a nil result after decoding the HTTP request body which allows an authenticated user to crash the Zoom plugin process, causing denial of service, via sending a JSON null body to the plugin's action endpoints. Mattermost Advisory ID: MMSA-2026-00722	6.5	More Details
CVE-2026-84065	IBM Guardium Data Protection 12.2 could allow a remote attacker to obtain sensitive information or cause a denial of service due to an out-of-bounds heap read.	6.5	More Details
CVE-2026-133260	In JetBrains YouTrack before 2026.2.19840 privilege escalation from project administrator to system administrator was possible via group merging	6.5	More Details
CVE-2026-131242	In Splunk SOAR versions below 8.8.0, a user that holds a role with the containers:view permission could access automation history for containers outside the label or tenant scope assigned to the user. This could expose action parameters, connector result summaries, and other investigation and response activity. The vulnerability is possible because the Splunk SOAR Representational State Transfer (REST) API checks the permission to view containers but does not restrict automation history to containers within the label or tenant scope assigned to the user. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) and Configure multiple tenants on your Splunk SOAR (On-premises) instance (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/configure-product-settings-for-your-splunk-soar-on-premises-instance/configure-multiple-tenants-on-your-splunk-soar-on-premises-instance) in the Splunk documentation.	6.5	More Details
CVE-2026-133060	A flaw was found in Moodle's Blackboard question import format. Insufficient validation of file paths during the question import process allows an authenticated user with permissions to import questions into a question bank to read arbitrary files accessible to the web server process. An attacker could leverage this vulnerability to view sensitive files, resulting in unauthorized information disclosure.	6.5	More Details
CVE-2026-131610	In ConnectWise Automate versions prior to 2026.11, a condition in Web Control Center (WCC2) authorization handling could allow an authenticated contact user to access information associated with other clients on the same Automate server through API requests. This may result in disclosure of information outside the user’s assigned access scope. The issue is limited to information disclosure and does not extend existing modification or command execution permissions. This issue is resolved in ConnectWise Automate 2026.11.	6.5	More Details
CVE-2026-129323	The Masteriyo LMS – LMS Course Builder, Quizzes & Certificates plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'content' parameter in all versions up to, and including, 3.5.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. On a default install, any unauthenticated visitor can self-register as a Masteriyo Student through the plugin's public /account/signup/ form without administrator approval, granting them the publish_course_reviews capability required to submit reviews and exploit this vulnerability.	6.4	More Details
CVE-2026-	The Flying Images: Optimize and Lazy Load Images for Faster Page Speed plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'post_content' parameter in all versions up to, and including, 2.4.16 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. Although wp_kses_post sanitizes post content on save, the payload is kses-legal	6.4	More Details

125012	at save time; the plugin's output-buffer rewrite transforms it into an executable form at render time by overwriting the valueless data-src attribute with the unescaped src value, meaning save-time sanitization provides no effective mitigation.		
CVE-2026-24205	NVIDIA TensorRT-LLM for any platform contains a vulnerability where a user may cause a race condition by running concurrent requests. A successful exploit of this vulnerability may lead to data tampering.	6.4	More Details
CVE-2026-131697	Tanium addressed an improper access controls vulnerability.	6.4	More Details
CVE-2026-133034	In JetBrains TeamCity before 2026.2.1 stored XSS on the Changes page was possible	6.4	More Details
CVE-2026-131685	Tanium addressed an improper signature validation vulnerability.	6.4	More Details
CVE-2026-125058	The Image Source Control Lite – Show Image Credits and Captions plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'meta[isc_image_licence]' parameter in all versions up to, and including, 3.14.0 due to insufficient input sanitization and output 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. The licence display feature must be enabled in the plugin settings (enable_licences option), which is a standard configuration for sites using the licence-attribution functionality.	6.4	More Details
CVE-2026-131226	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in Intrusion Detection data. If another user who can access the Intrusion Search dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify data available to the user who selects the row. The vulnerability is possible because the Intrusion Search dashboard encodes data-model field values for the drilldown link but does not escape the values as SPL search literals before inserting the values into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Network dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/network-dashboards) in the Splunk documentation.	6.4	More Details
CVE-2026-133104	In JetBrains TeamCity before 2026.2.2 stored XSS in external change viewer links was possible	6.4	More Details
CVE-2026-132940	In JetBrains TeamCity before 2026.2 credential leakage via Mercurial subrepository URLs was possible	6.4	More Details
CVE-2026-131197	In Splunk Enterprise versions below 10.4.5, 10.2.9, and 10.0.12, a low-privileged user who holds the "power" Splunk role could overwrite shared application dashboard views with arbitrary Search Processing Language (SPL) commands. When another authenticated user opens an affected dashboard in Splunk Web, Splunk Enterprise executes the injected SPL using the viewing user's permissions, potentially allowing the low-privileged user to access or modify data available to that user. The vulnerability is possible because the affected app grants the "power" Splunk role write access to shared dashboard views that should be writable only by the "admin" Splunk role. The vulnerability is possible because the Splunk App for Splunk Observability Cloud grants the "power" Splunk role write access to shared dashboard views that should be writable only by the "admin" Splunk role	6.4	More Details
CVE-2026-85413	The Foxtool All-in-One: Contact chat button, Custom login, Media optimize images plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'foxtool_image_id' parameter in all versions up to, and including, 2.5.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. A subscriber-level attacker can store a payload in their own profile's foxtoolpic user-meta that executes in an administrator's browser session when the administrator views the user via the Users → Edit screen.	6.4	More Details
CVE-2026-	The Jeg Kit for Elementor – Powerful Addons for Elementor, Widgets & Templates for WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'description' parameter in all versions up to, and including, 3.3.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Successful execution requires a published page	6.4	More Details

116594	containing both a Post Author widget in user_bio mode (rendering the attacker's biography) and any legitimate jkit_client_logo widget, so that client-logo.js is enqueued and the loadCarousel handler fires against the forged widget markup from the bio field.		
CVE-2026-123971	The Slimbox Plugin plugin for WordPress is vulnerable to Stored DOM-Based Cross-Site Scripting via the 'el.title (HTML-decoded title attribute of a stored <a rel^=lightbox> element)' parameter in all versions up to, and including, 1.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Exploitation requires a companion plugin such as bbPress that permits low-privileged users to store anchor elements with arbitrary title and rel attributes, as WordPress core comment handling does not preserve rel on anchor tags by default.	6.4	More Details
CVE-2026-105328	The Delivery & Pickup Date Time for WooCommerce plugin for WordPress is vulnerable to Stored Cross-Site Scripting in versions up to, and including, 1.3.90. This is due to a nonce-only settings AJAX handler that stores an attribute-breaking value and an administrator settings template that outputs it without HTML attribute escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-118068	The JobBoardWP – Job Board Listings and Submissions plugin for WordPress is vulnerable to Stored DOM-Based Cross-Site Scripting via the 'job_description' parameter in all versions up to, and including, 1.3.6 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with custom-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The plugin's default public submission workflow auto-provisions a jb_employer account for any visitor, meaning the effective entry point is unauthenticated — the custom role requirement imposes no real barrier on a stock install.	6.4	More Details
CVE-2026-54747	Chamilo LMS is an open-source learning management system. Prior to 2.0.3, public/main/course_description/course_description_controller.php assigns contentDescription directly from a teacher's request in the Doctrine-based add and edit paths without cross-site scripting sanitization. The stored value is rendered by src/CoreBundle/Resources/views/Course/about.html.twig through description with raw output, and the /course/{id}/about route is publicly accessible. A teacher can store active content that executes when an authenticated or unauthenticated visitor views the course description, enabling session theft, credential phishing, defacement, or actions in an authenticated victim's session. This issue is fixed in version 2.0.3.	6.4	More Details
CVE-2026-120557	The WP Delicious – Recipe Plugin for Food Bloggers (formerly Delicious Recipes) plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'details[].icon' parameter in all versions up to, and including, 1.10.11 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. The payload bypasses wp_kses_post on save because Gutenberg serializes block attribute JSON with JSON_HEX_QUOT, encoding double-quotes as \u0022 so no raw quote or angle bracket appears in post_content; parse_blocks() restores the literal quote at render time, producing the HTML attribute breakout.	6.4	More Details
CVE-2026-131238	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled certificate values. If another user who can access the Secure Sockets Layer Search dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify data available to the user who selects the row. The vulnerability is possible because the Secure Sockets Layer Search dashboard encodes certificate values for the drilldown link but does not escape the values as SPL search literals before inserting the values into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Protocol intelligence dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/protocol-intelligence-dashboards) in the Splunk documentation.	6.4	More Details
CVE-2026-129322	The Orbit Fox: Duplicate Page, Menu Icons, SVG Support, Cookie Notice, Custom Fonts & More plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'title (post_title)' parameter in all versions up to, and including, 3.0.13 due to insufficient input sanitization and output 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. The exploit is further enabled by wptexturize deliberately skipping texturization of content inside <code> elements and the default WordPress kses allow-list preserving <code> tags and raw double quotes on save, allowing a crafted title attribute breakout payload to survive storage and the the_title filter chain intact.	6.4	More Details
CVE-2026-120566	The WP Delicious – Recipe Plugin for Food Bloggers (formerly Delicious Recipes) plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'recipeTitle' parameter in all versions up to, and including, 1.10.10 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. WordPress core's block-attribute kses pass strips HTML tags but	6.4	More Details

	preserves bare double-quote characters, allowing a payload such as x" onerror="alert(1) to survive both the save pipeline and the render-time strip_tags() call, breaking out of the alt attribute of the rendered img element.		
CVE-2026-126009	The (Simply) Guest Author Name plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the user Biographical Info (description) field in versions up to, and including, 4.42. The plugin unconditionally calls remove_filter('pre_user_description', 'wp_filter_kses') at plugin load (sfly-guest-author.php:17), which strips WordPress core's sanitization of the user description field for users who lack the unfiltered_html capability (Subscribers, Contributors, Authors). This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in their own profile Biographical Info that will execute whenever a user accesses a page rendering that author's description (such as an author archive or a single-post author bio in default WordPress themes like Twenty Seventeen and Twenty Fifteen).	6.4	More Details
CVE-2026-131184	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a low-privileged user who does not hold the "admin" or "power" Splunk roles could inject a malicious script into the name of a private saved search. If another user subsequently enables, disables, or edits the corresponding platform alert in the Monitoring Console, the injected script executes within that user's browser session, potentially allowing the low-privileged user to view sensitive information or modify data in Splunk Enterprise. The vulnerability is possible because the Monitoring Console does not safely handle saved-search names in the platform alert workflow. For more information see Enable and configure platform alerts (https://help.splunk.com/en/splunk-enterprise/administer/monitor/10.4/configure-the-monitoring-console/enable-and-configure-platform-alerts) and About configuring role-based user access (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/about-configuring-role-based-user-access) in the Splunk documentation.	6.4	More Details
CVE-2026-131232	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in vulnerability scan data. If another user who can access the Vulnerability Scan Details dashboard selects an affected value, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the value. The user who stores the values could then access sensitive data and modify data available to the user who selects the affected value. The vulnerability is possible because the Vulnerability Scan Details dashboard does not validate destination values before inserting them into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Vulnerability Scan Details (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/vulnerability-scan-details) in the Splunk documentation.	6.4	More Details
CVE-2026-12040	The FS Poster - WordPress Social media Auto Poster & Scheduler [Facebook, Instagram, Twitter, Pinterest] plugin for WordPress is vulnerable to Server-Side Request Forgery in all versions up to, and including, 8.0.1 via the 'site_url' parameter. This makes it possible for authenticated attackers, with subscriber-level access and above, to make web requests to arbitrary locations originating from the web application and can be used to query and modify information from internal services. The Guzzle HTTP client is instantiated with TLS certificate verification disabled and redirect following enabled by default, allowing redirect-based SSRF bypasses and enabling internal port scanning via distinguishable server responses.	6.4	More Details
CVE-2026-129218	The WP Coder – Insert & Manage Code Snippets plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'alt_text (WP media REST) → _wp_attachment_image_alt post meta' parameter in all versions up to, and including, 5.0 due to insufficient input sanitization and output 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. This requires that an admin-authored image-type snippet exists on the site, as the attacker uploads an image via the standard WordPress media REST endpoint, sets its Alt Text to a double-quote-based payload that survives WordPress core's sanitize_text_field(), and then publishes a post embedding the shortcode.	6.4	More Details
CVE-2026-111371	The SchemaEngine AI – AI Schema Markup, Reviews & Rich Snippets for SEO plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'content/excerpt (REST wp/v2/posts) → stored in post_content/post_excerpt' parameter in all versions up to, and including, 3.1.3 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. A Contributor's entity-encoded payload survives save-time wp_kses filtering verbatim and is only decoded into raw script tags at render time via html_entity_decode(), meaning the injected script executes for every visitor — including authenticated Administrators — once an Editor publishes the post through the standard Contributor workflow.	6.4	More Details
CVE-2026-120404	The built-in regular expression name finder for latitude and longitude coordinates uses a pattern that is vulnerable to excessive backtracking. An attacker who can supply text that an application analyzes with this name finder can cause high CPU consumption, leading to denial of service. Users are recommended to upgrade to version 2.5.13 or 3.0.0, which fix the issue.	6.4	More Details
	A flaw was found in CRI-O. When the optional Linux LinkLogs feature is enabled, the host pod log directory is bind-mounted into a pod emptyDir and is intended to be read-only, but Linux ignores MS_RDONLY on the		

CVE-2026-106096	initial bind and the CRI mount stays writable. A non-privileged UID 0 container can place a symlink at a future kubelet log path so that privileged conmon later appends attacker-controlled log data to an existing host file. This is an append-only host integrity issue, not universal remote code execution; LinkLogs is not enabled by default.	6.4	More Details
CVE-2026-125999	The Widget Manager Light plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'posts' parameter in all versions up to, and including, 1.18 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-130767	Obot from 0.9.0 before 0.27.0 contains a command injection vulnerability that allows authenticated users to execute shell commands by embedding an EOF line in file-type MCP server configuration values. Attackers can terminate the heredoc in the generated init container script to run commands, modify the MCP server files volume, and reach services on Obot's Docker network.	6.4	More Details
CVE-2026-96841	The ITRO Popup Plugin plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'custom_html' parameter in all versions up to, and including, 5.2.6 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This is only exploitable by low-privileged roles such as Subscriber when an administrator has explicitly granted that role access to the plugin settings page via the plugin's role selection option.	6.4	More Details
CVE-2026-129992	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.0, failed Railpack builds and database backups propagate an execution error containing the full command, exposing application environment variables, database passwords, and S3 access keys embedded in those commands. When the corresponding App Build Error or Database Backup notification setting is enabled, the disclosure can send secrets to email, Slack, Discord, Telegram, or webhooks; the same secrets are written separately to persistent Dokploy container logs, which may be accessible to more principals than the affected service. This issue is fixed in version 0.31.0.	6.3	More Details
CVE-2026-132279	A flaw has been found in elunez eladmin up to 2.7. Affected by this vulnerability is the function SqlUtils.sanitizeJdcUrl of the component JDBC URL Sanitization. Executing a manipulation can lead to improper input validation. The attack may be performed from remote. The exploit has been published and may be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-102884	Connaisseur is an admission controller that integrates container image signature verification into a Kubernetes cluster. Prior to 3.11.1, the Cosign validator counts valid signature layers rather than distinct trust roots for trustRoot: "*" policies using an explicit threshold or an implicit require-all policy. An attacker who controls one configured trust root's signing key and can attach multiple signatures to an image can satisfy a threshold intended to require multiple independent roots, causing admission without the intended signer quorum. Policies using a single trust root or the required option are not affected. This issue is fixed in 3.11.1.	6.3	More Details
CVE-2026-131180	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could create and share a dashboard that contains attacker-controlled Search Processing Language (SPL). The SPL could run using the permissions of the user who opens the dashboard. With this access, the user who created the dashboard could expose stored credentials and write data to an index. The vulnerability is possible because Splunk Enterprise does not correctly escape field names in searches that use the rename command. For more information see rename (https://help.splunk.com/en/splunk-enterprise/spl-search-reference/10.4/search-commands/rename) and SPL safeguards for risky commands (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/best-practices-for-splunk-platform-security/spl-safeguards-for-risky-commands) in the Splunk documentation.	6.3	More Details
CVE-2026-133490	A vulnerability was detected in declare-lab tango up to 34ecd38898508017ca937508a047b42817d2b070. Impacted is the function Tango.__init__ of the file tango.py of the component Model Loading. Performing a manipulation of the argument scheduler_name results in improper input validation. The attack is possible to be carried out remotely. The exploit is now public and may 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.	6.3	More Details
CVE-2026-133488	A weakness has been identified in XiangLi1999 Diffusion-LM up to 759889d58ef38e2eed41a8c34db8032e072826f4. This vulnerability affects the function os.system of the file improved-diffusion/scripts/infill.py. This manipulation of the argument notes causes os command injection. Remote exploitation of the attack is possible. The exploit has been made available to the public and could be used for attacks. 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.	6.3	More Details
	A security vulnerability has been detected in SEO Panel 6.0.0. This affects the function mime_content_type of		

CVE-2026-132601	the file controllers/website.ctrl.php of the component Website Import. Such manipulation of the argument website_csv_file leads to unrestricted upload. It is possible to launch the attack remotely. The exploit has been disclosed publicly and may be used. Upgrading to version 7.0.0 mitigates this issue. The name of the patch is 2a2d4aeb718177327b0cd8ce17e2f40328772127. You should upgrade the affected component.	6.3	More Details
CVE-2026-132602	A vulnerability was detected in SEO Panel 6.0.0. This impacts an unknown function of the file controllers/siteauditor.ctrl.php of the component Site Auditor. Performing a manipulation of the argument order_col/order_val results in sql injection. The attack can be initiated remotely. The exploit is now public and may be used. Upgrading to version 7.0.0 will fix this issue. The patch is named 389783b5d8235f5637f4e4587f524b5fef07864d. The affected component should be upgraded.	6.3	More Details
CVE-2026-132603	A flaw has been found in SEO Panel 6.0.0. Affected is an unknown function of the file controllers/siteauditor.ctrl.php of the component Site Auditor. Executing a manipulation of the argument order_col/order_val can lead to sql injection. The attack can be launched remotely. The exploit has been published and may be used. Upgrading to version 7.0.0 is able to address this issue. This patch is called 2a2d4aeb718177327b0cd8ce17e2f40328772127. It is suggested to upgrade the affected component.	6.3	More Details
CVE-2026-77314	Flyto2 Core is an execution kernel for automation and AI-agent workflows. Prior to 2.26.8, the MCP run_recipe tool in src/core/mcp_handler.py accepts a caller-controlled recipe_name that load_recipe() in src/cli/recipe.py appends to the src/recipes path without resolved-path containment. Parent-directory segments can therefore make Flyto load and execute a YAML workflow outside the bundled recipe directory, bypassing the boundary represented by list_recipes(). This does not necessarily provide unrestricted code execution because available workflow modules remain subject to Flyto's module policy, but it permits unintended execution of allowed browser, network, file, and data-processing modules. This issue is fixed in version 2.26.8.	6.3	More Details
CVE-2026-130695	NocoDB before 2026.10.1 contains an authentication bypass vulnerability that allows holders of revoked OAuth access tokens to authenticate as first-party sessions by presenting them through the session header. Attackers can bypass revocation, client-existence, route allowlist and granted_resources scope checks to modify victim profiles, upload to storage, and force global sign-outs.	6.3	More Details
CVE-2026-131264	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings edit permission could cause Splunk SOAR to send the stored root credential for an external credential store to a host that the user controls. Depending on the configuration, the exposed credential could include an access token, application secret, account password, or client certificate. The vulnerability is possible because Splunk SOAR reuses the stored credential when the user changes the external credential store endpoint without supplying replacement credentials.	6.3	More Details
CVE-2026-131267	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings view and edit permissions could test a Lightweight Directory Access Protocol authentication provider and cause Splunk SOAR to send a stored directory service account bind password to a host that the user controls. The vulnerability is possible because Splunk SOAR reuses a stored directory service provider credential when an authentication test uses a different directory service connection destination.	6.3	More Details
CVE-2026-132137	A vulnerability was identified in Relaticle up to 3.10.0. This affects the function mapToField of the file packages/ImportWizard/src/Livewire/Steps/MappingStep.php of the component Import Wizard. Such manipulation of the argument Target leads to authorization bypass. The attack can be executed remotely. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure.	6.3	More Details
CVE-2026-132222	BizPrint (print-google-cloud-print-gcp-woocommerce) before 4.10.0 contains a missing authorization vulnerability in its Debug\Core component that allows any authenticated user, including the lowest-privilege Subscriber role, to invoke the getPrinters and sendPrintRequest AJAX actions without a nonce or capability check. Attackers can send a crafted admin-ajax.php request with an arbitrary WooCommerce order_id to dispatch print jobs for other customers' orders and disclose printer configuration, exposing cross-customer PII and abusing printer resources.	6.3	More Details
CVE-2026-131277	In Splunk SOAR versions below 8.8.0, a user that holds a role with permission to edit playbooks and repositories could store a modified utility block in a playbook. If a user that holds a role with permission to edit playbook code later opens and saves the playbook in the Visual Playbook Editor, a subsequent playbook run could send a Representational State Transfer (REST) API delete request to a resource outside the intended data list. The user account that runs the playbook must have permission to delete the target resource. The vulnerability is possible because Splunk SOAR does not prevent a utility block from changing the intended REST API request path. For more information see Data management automation API (https://help.splunk.com/en/splunk-soar/soar-on-premises/python-playbook-api-reference/6.3.0/automation-api/data-management-automation-api) and Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	6.3	More Details
CVE-2026-	A vulnerability was identified in Dolibarr ERP CRM up to 23.0.4. This affects the function dol_convert_file in the library htdocs/core/lib/files.lib.php of the component Document Upload. Such manipulation leads to code injection. The attack can be launched remotely. The exploit is publicly available and might be used.	6.3	More Details

133327	Upgrading to version 24.0.0 is able to mitigate this issue. The name of the patch is 9b6368b2a9eae3eec690b2ce22c4f291f0051540. Upgrading the affected component is advised.		
CVE-2026-119634	The User Frontend – Membership, User Registration, User Profile, User Directory & Content Restriction with Frontend Post Submission plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 4.3.14 via the `draft_post()` function. This is due to the function storing the `post_content` parameter via `wp_insert_post()` without the protected-shortcode allow-list check that was introduced in 4.3.14 for the sibling `submit_post()` handler — a gap that `wp_kses` filtering does not close, as it strips only HTML tags while leaving shortcode text such as `[wpuf-registration role="editor"]` intact. This makes it possible for authenticated attackers, with Author-level access and above, to embed a registration shortcode in a post and cause any site visitor who subsequently registers through the rendered form to receive Editor-level access. The render-time `wpuf_modify_shortcodes` filter that is designed to neutralize protected shortcodes provides no protection here because posts created via the `wpuf_draft_post` AJAX endpoint carry no `_edit_last` post meta, causing the filter to return the content unchanged and allowing the shortcode to render and execute fully.	6.3	More Details
CVE-2026-131227	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in Network Resolution data. If another user who can access the Domain Name System Search dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify some data available to the user who selects the row. The vulnerability is possible because the Domain Name System Search dashboard encodes data-model field values for the drilldown link but does not escape the values as SPL search literals before inserting the values into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Protocol intelligence dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/protocol-intelligence-dashboards) in the Splunk documentation.	6.3	More Details
CVE-2026-133486	A vulnerability was identified in hasan-py Hayroo up to fd0fbe763153f72bfaa04ed1e1e4a1bb6f7c61df. Affected by this issue is the function Product.deleteImages of the file server/controller/products.js of the component Image Deletion Handler. The manipulation leads to path traversal. The attack may be initiated remotely. The exploit is publicly available and might be used. 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 9714947bb43ec8559287aa3962b70a7be8c8d6b1. It is recommended to apply a patch to fix this issue.	6.3	More Details
CVE-2026-133487	A security flaw has been discovered in nodejs undici up to 7.24.0. This affects the function Object.assign of the file balanced-pool.js of the component BalancedPool. The manipulation results in improperly controlled modification of object prototype attributes. The attack may be launched remotely. The exploit has been released to the public and may be used for attacks. Upgrading to version 7.24.1 is able to mitigate this issue. The patch is identified as 0d7ec33ff37563d3e7c98d11d7bca736f330d156. You should upgrade the affected component.	6.3	More Details
CVE-2026-132282	A vulnerability was determined in elunez eladmin up to 2.7. This vulnerability affects the function FileUtil.upload of the file /api/localStorage of the component File Upload Handler. This manipulation causes path traversal. The attack can be initiated 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.	6.3	More Details
CVE-2026-132281	A vulnerability was found in elunez eladmin up to 2.7. This affects the function SqlUtils.batchExecute of the file /api/database/upload of the component Database Maintenance Module. The manipulation results in sql injection. It is possible to launch the attack remotely. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-132280	A vulnerability has been found in elunez eladmin up to 2.7. Affected by this issue is the function CheckBean of the file /api/jobs/exec of the component Quartz Handler. The manipulation of the argument methodName leads to improper input validation. It is possible to initiate the attack 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.	6.3	More Details
CVE-2025-36441	IBM Financial Transaction Manager (FTM) 3.2.13.0 through 3.2.13.0 Interim Fix 006 stores potentially sensitive information in environment variables that could be read by a local user.	6.2	More Details
CVE-2026-130750	Network-AI from 5.11.0 before 5.17.2 contains a CORS misconfiguration in McpStreamableServer that reflects localhost Origins into Access-Control-Allow-Origin regardless of bind interface. Attackers controlling a localhost-origin page or DNS-rebinding name can read unauthenticated /tools and /health responses, exposing tool catalogs and bridge metadata on non-loopback binds.	6.2	More Details
CVE-2026-75801	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 could allow a local attacker to obtain sensitive credentials due to the use of a hard-coded cryptographic key.	6.2	More Details

CVE-2026-132949	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS on the Clean-up Rules page was possible	6.1	More Details
CVE-2026-131873	Gargoyle 2023.1 through 2026.1.1 contains a stack-based buffer overflow vulnerability in the bundled Plus interpreter's ReadAction() function in terps/plus/loaddatabase.c, which never bounds-checks the index into the 1024-entry conditions array. Attackers can supply a crafted or truncated .plus game file with unterminated action conditions to write file-controlled values past the array and crash the player.	6.1	More Details
CVE-2024-58401	Bagisto 2.1.0 before 2.2.1 contains a client-side template injection vulnerability in the search page where the 'query' parameter is rendered unescaped inside a Vue app root. Unauthenticated attackers can send crafted VueJS template expressions to the /bagisto-common/search endpoint via the 'query' parameter to execute arbitrary JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-07-30.	6.1	More Details
CVE-2026-132754	MediaInfoLib before 26.10 contains an out-of-bounds read vulnerability in File_Mpeg4::meta_iloc that allows attackers to trigger a global buffer over-read via oversized iloc field sizes. Attackers can supply a crafted MP4 file with invalid offset_size or length_size values to overrun the BitStream_Fast mask table and crash applications analysing the file.	6.1	More Details
CVE-2026-130788	SemanticMediaWiki from 3.0.0 before 7.4.0 contains a reflected cross-site scripting vulnerability in CsvFileExportPrinter and DsvResultPrinter via unescaped valuesep and separator parameters. Attackers can craft Special:Ask links using request_type=special_page so export output renders as text/html, executing JavaScript in victims' browsers under the wiki origin.	6.1	More Details
CVE-2019-25783	Tiny Tiny RSS before commit c68ac04 (2019-03-03) contains an open redirect vulnerability that allows unauthenticated attackers to redirect users through the unvalidated 'return' parameter. Attackers can craft requests to /public.php?op=login with an absolute external URL in 'return', which is placed in the Location header to send victims to attacker-controlled sites. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-02-26.	6.1	More Details
CVE-2022-51031	PMB 7.3.1 before 7.3.15 and 7.4.x before 7.4.5 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject script through unescaped POST parameter names. Attackers can send crafted parameter names to opac_css/index.php extended_search, which empr.inc.php and languages.inc.php reflect into hidden form inputs to execute JavaScript. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2026-132922	In JetBrains TeamCity before 2026.2, 2026.1.4 stored XSS in the unified diff page syntax highlighting was possible	6.1	More Details
CVE-2022-51032	Reddit Top RSS through commit 5e38bfd contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the unescaped score GET parameter. Attackers can craft links with a score value that closes the input value attribute in html.php, executing JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	6.1	More Details
CVE-2026-87956	When Private data source connect (PDC) or a secure SOCKS proxy is used, the MQTT data source plugin does not apply TLS to broker connections for the tls, ssl, tcps, mqtts and mqtt+ssl schemes. The configured TLS settings are ignored, and the broker username and password are sent unencrypted on the hop between the proxy and the broker. The server certificate is not verified, and a configured client certificate is not sent.	6.1	More Details
CVE-2026-82951	Latvijas Valsts radio un televīzijas centrs (LVRTC) eParakstītājs 3.0 version 1.10.0-resolute allows an unauthenticated remote attacker to cause HTTP GET requests from a user's workstation by convincing the user to open a crafted signed eDOC file. Automatic certificate validation retrieves a signer-controlled Authority Information Access (AIA) caIssuers URL and follows redirects without restricting the destination, enabling blind SSRF to loopback and private-network services. Additionally, unrestricted loading of local file URIs (such as file:///dev/zero) causes unbounded memory buffering resulting in Denial of Service (OutOfMemoryError). In automated server verification pipelines with UI:N, impact scales up to CVSS 8.6 High.	6.1	More Details
CVE-2026-133105	In JetBrains TeamCity before 2026.2.2, 2026.1.6 reflected XSS on the Edit VCS Root page was possible	6.1	More Details
CVE-2026-128136	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.10.0 until 3.10.4, allowlisted_pickle_load keys dangerous NumPy globals by export name, allowing submodule-local numpy.lib file input/output callables to evade the denylist under the broad module allowance used by TransitionParser. A crafted pickle can invoke recfromtxt to read a local file or open_memmap to create or truncate an attacker-selected file when a victim loads an untrusted model. This issue is fixed in version 3.10.4.	6.1	More Details
	Retool versions 3.253.0 before 3.253.19, 3.254.0 before 3.284.12, and 3.285.0 before 3.300.4 contain a		

CVE-2026-126751	DOM-based cross-site scripting vulnerability in the custom-component-collections.html postMessage handler that imports received code without origin validation. Unauthenticated attackers can send crafted postMessage events that load attacker-controlled code via data:text/javascript URIs to execute arbitrary JavaScript in the victim's context. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-04-22.	6.1	More Details
CVE-2026-128666	A Reflected Cross-Site Scripting (XSS) vulnerability exists in PHPGurukul News Portal Project V4.1. The $currenturl variable in news-details.php is constructed using unencoded REQUEST_URI and HTTP_HOST server variables and echoed directly into the href attributes of social share links. An unauthenticated remote attacker can exploit this by crafting a URL containing HTML metacharacters to break out of the href attribute and execute arbitrary HTML or JavaScript in the context of the victim's browser.	6.1	More Details
CVE-2026-77376	Bolt is a fully open source, adaptable platform for building and running modern websites. Prior to 6.1.4, LoginFormAuthenticator accepts unsafe destinations for the login success redirect and login failure redirect. An attacker can direct a user to an attacker-controlled domain after authentication succeeds or fails, enabling phishing and potential token theft. This issue is fixed in version 6.1.4.	6.1	More Details
CVE-2026-123417	The Check & Log Email – Easy Email Testing & Mail logging plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment' parameter in all versions up to, and including, 2.0.17 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. Successful exploitation requires an administrator to open the relevant log entry and click the planted anchor link in the Preview panel, at which point the ThickBox gadget parses the title attribute as live HTML in the wp-admin origin. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	6.1	More Details
CVE-2023-54421	phpLDAPadmin 1.1.0 before 1.2.5 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the dn parameter of cmd.php?cmd=template_engine. Attackers can supply an unparseable dn value that pretty_print_dn returns unescaped in the 'No such entry' error, executing JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2022-51035	SLiMS Senayan Library Management System 9.4.0 before 9.4.2 contains a reflected cross-site scripting vulnerability in the destination parameter of /index.php?p=member, which is echoed unmodified into the login form action attribute. Unauthenticated attackers can craft a malicious destination value to inject script and hijack PHPSESSID sessions of victims who access the link. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2026-128665	A stored cross-site scripting (XSS) vulnerability in the Public News Article Rendering Module of PHPGurukul News Portal Project v4.1 allows authenticated attackers with administrative or sub-administrative privileges to inject arbitrary HTML and JavaScript through the postdescription parameter when creating or editing news articles. The application renders the stored post content without proper output encoding in news-details.php, allowing injected JavaScript to execute in the browsers of users who visit the affected article.	6.1	More Details
CVE-2023-54416	Tiki Wiki CMS Groupware 12.12 through 12.14, 15.5 before 18.10, 19.0 before 21.9, 22.0 before 24.3, and 25.0 contain a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the objectId parameter of tiki-ajax_services.php. Attackers can send a crafted GET request with the comment controller and list action to execute arbitrary JavaScript and steal session tokens or credentials. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	6.1	More Details
CVE-2026-126753	LimeSurvey 6.5.0 before 6.16.11 contains an open redirect vulnerability that allows unauthenticated attackers to redirect users by supplying arbitrary values to the url parameter. Attackers can send crafted GET requests to index.php?r=editorLink, where EditorLinkController concatenates url and route unvalidated, redirecting victims to attacker-controlled sites for phishing. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-03-20.	6.1	More Details
CVE-2026-132945	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS in the inspection profile selector popup was possible	6.1	More Details
CVE-2026-126754	cPanel & WHM 11.134.0.13 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the mpidentity parameter of /mailman/listinfo. Attackers can send crafted requests with a script-context breakout payload in mpidentity to execute arbitrary JavaScript in the victim's browser. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-06-08.	6.1	More Details
CVE-2026-130888	A flaw was found in nbdkit. The qcow2dec filter's zstd_compressed_cluster() function incorrectly sets the ZSTD input buffer size to the uncompressed cluster size instead of the compressed buffer size. A crafted QCOW2 image with ZSTD-compressed clusters can trigger a heap over-read, leading to a denial of service and potentially exposing heap memory.	6.1	More Details
CVE-2026-	GPAC before 26.07.0 contains a type confusion vulnerability in xmt_parse_script_field() that allows attackers to trigger a heap out-of-bounds read via duplicate Script field names with conflicting types. Attackers can	6.1	More

132769	craft an XMT file that, when imported with MP4Box, causes gf_sg_vrml_mf_reset() to treat a single-value field as multi-value, crashing the application.		Details
CVE-2025-71474	lwext4 through 1.0.0 contains an out-of-bounds write vulnerability in ext4_fs_init_block_bitmap() that allows attackers to corrupt heap memory via crafted superblock geometry fields. Attackers can supply an ext4 image with oversized s_reserved_gdt_blocks and BLOCK_UNINIT groups to write past the bitmap buffer, crashing applications that mount it.	6.1	More Details
CVE-2026-131843	Poppler pdftohtml through 26.11.0 contains a cross-site scripting vulnerability in HtmlLink::getLinkStart() in utils/HtmlLinks.cc that writes PDF link annotation URIs into href attributes without escaping in HTML output modes. Attackers can craft a PDF whose /URI entry contains a double quote followed by HTML markup or event handlers to execute script in browsers viewing the generated HTML.	6.1	More Details
CVE-2026-132929	In JetBrains TeamCity before 2026.2, 2026.1.4 reflected XSS on the notification rules page was possible	6.1	More Details
CVE-2026-132755	MediaInfoLib before 26.10 contains a heap-based buffer over-read vulnerability in tfsxml_leave() caused by an off-by-one error when handling generic <! markup in truncated XML. Attackers can supply a crafted media file with malformed ADM XML parsed by File_Adm.cpp to crash MediaInfo or applications using MediaInfoLib.	6.1	More Details
CVE-2026-132946	In JetBrains TeamCity before 2026.2, 2026.1.5 reflected XSS on the Test History page was possible	6.1	More Details
CVE-2026-132947	In JetBrains TeamCity before 2026.2 reflected XSS in Kubernetes cloud profile namespace and deployment choosers was possible	6.1	More Details
CVE-2026-132948	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS on the Build Chains tab was possible	6.1	More Details
CVE-2026-132943	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS on the Build Log page was possible	6.1	More Details
CVE-2026-132951	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS in the notification rule form was possible	6.1	More Details
CVE-2026-124977	The Flexible PDF Invoices for WooCommerce & WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'billing_address.first_name / extensions.wpdesk-flexible-invoices-checkout._billing_vat_number' parameter in all versions up to, and including, 7.1.2 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Execution requires an administrator to open the invoice edit page for the attacker's order and interact with (e.g., type into or trigger a change/keyup event on) a field in the customer metabox, at which point the payload stored in billing fields such as first_name or _billing_vat_number is decoded from its entity-escaped input value and passed unsanitized to jQuery .html(). Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	6.1	More Details
CVE-2025-71473	lwext4 through 1.0.0 contains a heap buffer overflow in ext4_fs_init_inode_bitmap() in src/ext4_fs.c because it does not validate s_inodes_per_group against the inode bitmap size. Attackers can supply a crafted ext4 image with an oversized s_inodes_per_group so memset writes past the heap buffer, corrupting memory and crashing the mounting application.	6.1	More Details
CVE-2026-92236	The Asset CleanUp: Page Speed Booster plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via 'wpacu_unload_css' and 'wpacu_unload_js' Query Parameters in all versions up to, and including, 1.4.0.6 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 that the victim has sufficient privileges to see the plugin's admin-bar entry, typically an administrator.	6.1	More Details
CVE-2019-25786	Gnuboard5 5.0.1 before 5.3.3.3 and 5.4.0.1 before 5.4.0.2 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the arr_ajax_msg parameter. Attackers can craft GET requests to /plugin/sms5/ajax.sms_emoticon.php, where common.php's @extract($_GET) injects the value into the echoed JSON output, executing script in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2010-	Virtual Keyboard plugin for SquirrelMail through 0.9.1 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject script by supplying unsanitized input in the passformname parameter. Attackers can craft GET requests to /plugins/vkeyboard/vkeyboard.php with script payloads in	6.1	More

20126	passformname to execute code in victims' browsers and steal session cookies. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.		Details
CVE-2026-133032	In JetBrains TeamCity before 2026.2.1 reflected XSS on the Cloud Profiles page was possible	6.1	More Details
CVE-2026-133247	modified eCommerce Shopsoftware through 3.3.0 contains a reflected cross-site scripting vulnerability in admin/support.php that allows unauthenticated attackers to inject script via the module parameter. Attackers can trick a logged-in administrator into opening a crafted link to run script in the admin panel, reading CSRF tokens and changing shop settings.	6.1	More Details
CVE-2026-110521	IBM Engineering AI Hub 1.0.0 through 1.4.0 is vulnerable to cross-site scripting. This vulnerability allows an unauthenticated attacker to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session.	6.1	More Details
CVE-2026-123969	The Fancy Lightbox plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment' parameter in all versions up to, and including, 1.0.0 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. Successful exploitation without administrator interaction requires that the attacker first have one comment approved so that a subsequent payload comment is auto-published; the injected script only executes when a user clicks the malicious image link inside the comment. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	6.1	More Details
CVE-2026-132931	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS on the build feature parameters page was possible	6.1	More Details
CVE-2022-51033	JoomBri Careers 3.3.0 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the 'keyword' parameter of the /for-jobseekers/search-jobs endpoint in com_jbcareers. Attackers can craft malicious URLs using an attribute-breakout payload with onfocus and autofocus to execute arbitrary JavaScript and steal session tokens or credentials. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2026-132930	In JetBrains TeamCity before 2026.2, 2026.1.4 reflected XSS in build problem details was possible	6.1	More Details
CVE-2026-132962	In JetBrains TeamCity before 2026.2, 2026.1.4 stored XSS in the JVM update warning on the Agent Summary page was possible	6.1	More Details
CVE-2021-48028	Open Journal Systems 2.1.0 before 3.0.0 contains an open redirect vulnerability that allows unauthenticated attackers to redirect users to external domains by abusing the source parameter. Attackers can send GET requests to /index.php/index/user/setLocale/<locale> with a source value beginning with '@', which UserHandler.inc.php appends to the server host, redirecting victims to attacker-controlled sites for phishing. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-12-05.	6.1	More Details
CVE-2026-109909	A out of bounds read after using negative integers during offset calculation during PCF decoding in X.Org libXfont before 2.0.10 could be used by authenticated X clients to crash the X server or expose sensitive information in its memory.	6.1	More Details
CVE-2026-133244	parse5-html-rewriting-stream 7.0.0 before 8.1.1 contains a cross-site scripting vulnerability in RewritingStream#emitText, which skips escaping based on a stale lastStartTagName. Attackers can place entity-encoded markup after a closed script or style element to bypass RewritingStream-based filters and execute script in users' browsers.	6.1	More Details
CVE-2021-48027	OTOBO 10.0.1 before 10.0.13 contains an open redirect vulnerability that allows unauthenticated attackers to redirect users by supplying arbitrary URLs to the ExternalURLJump action. Attackers can craft GET requests to /otobo/index.pl?Action=ExternalURLJump with a malicious URL parameter, which the ExternalURLJump.pm module redirects to without an allow-list check, enabling phishing attacks. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-23.	6.1	More Details
CVE-2026-126699	The Majestic Support – The Leading-Edge Help Desk & Customer Support Plugin plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'mjsupport_message' parameter in all versions up to, and including, 1.2.4 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. Successful exploitation requires that a staff member or administrator opens the affected ticket and clicks the attacker-planted link on the ticket-detail screen, as the VenoBox payload is triggered on click rather than on page load. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	6.1	More Details

CVE-2026-132932	In JetBrains TeamCity before 2026.2, 2026.1.4 reflected XSS in the build log tree view was possible	6.1	More Details
CVE-2026-131657	Tanium addressed an incorrect privilege assignment vulnerability.	6.0	More Details
CVE-2026-55507	Wavelog is web-based amateur radio logging software. In 3.2.5 and earlier, the Hams of Note feature stores the administrator-supplied hon_url configuration value through application/controllers/Options.php hon_save() and application/models/Update_model.php update_hams_of_note() fetches that URL without restricting the scheme, host, loopback, link-local, or private network destinations. An administrator can cause the Wavelog server to request arbitrary HTTP or HTTPS destinations from its network position and have parsed response content rendered in the administrator interface, providing a read primitive against otherwise unreachable internal HTTP services. The HTTP status requirement prevents file scheme responses from succeeding, so this issue does not provide local file disclosure. No fixed version is available as of this review.	6.0	More Details
CVE-2026-133200	A flaw was found in quarkus-vertx-utils. When a client disconnects while receiving a streaming response, the server fails to detect the closed connection and continues processing the data indefinitely. An unauthenticated remote attacker can exploit this issue by repeatedly initiating streaming requests and disconnecting, exhausting worker threads and resulting in a Denial of Service (DoS).	5.9	More Details
CVE-2026-18079	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 Fail to revoke a user's active sessions when a bulk import links an existing user account to a bot record or changes the account's auth_service/auth_data, which allows an attacker in possession of a session token issued before the import to retain full authenticated access to the affected account via a system administrator-triggered bulk import job.. Mattermost Advisory ID: MMSA-2026-00746	5.9	More Details
CVE-2026-118519	The DataPress (Dataverse Integration) plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 2.97 via the 'src' parameter parameter. This makes it possible for unauthenticated attackers to read the contents of arbitrary files on the server, which can contain sensitive information. Exploitation requires WP_DEBUG to be defined as true on the target site — not the default on production, but commonly enabled on staging, development, and misconfigured production environments.	5.9	More Details
CVE-2026-133111	@fastify/http-proxy is a Fastify plugin that proxies HTTP and WebSocket traffic to an upstream server. In versions from 7.1.0 up to but not including 11.6.10, when the plugin is registered with websocket support enabled, it opens the upstream WebSocket connection with a header set built by the default header rewriter, which forwards the client cookie but never sets the upstream origin option, so the browser origin header is dropped while the authenticating cookie is forwarded. Many WebSocket servers defend against cross-site hijacking by checking the origin and treat a missing origin as a trusted non-browser client, so behind the proxy that check never sees the foreign origin. A malicious page visited by a logged-in user can therefore open a WebSocket through the proxy to an upstream that authenticates the handshake with a cookie, and read or send messages on the user's behalf. The issue is fixed in @fastify/http-proxy 11.6.10, and users should upgrade to 11.6.10 or later. As a workaround, provide a header rewriter that sets a correct origin, or validate the origin before the proxy.	5.9	More Details
CVE-2026-59964	uftpd is a UNIX FTP and TFTP server. Prior to 2.16, uftpd's compose_path function in src/common.c can overflow the fixed dir[PATH_MAX] stack buffer when the server runs without chroot and prepends the configured home directory to a near-PATH_MAX attacker-supplied path using memmove. A logged-in FTP client, including an anonymous or guest client where enabled, can trigger the out-of-bounds write through commands such as CWD, RETR, STOR, or LIST, causing memory corruption and a reliable service crash. Remote code execution has not been demonstrated. This issue is fixed in version 2.16.	5.9	More Details
CVE-2026-131208	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user could create and deliver a crafted Investigation Workbench link to an authenticated Splunk Enterprise Security user. If the authenticated user opens the link, Investigation Workbench could run arbitrary Search Processing Language (SPL) commands using the permissions of the authenticated user. With this access, the unauthenticated user could expose or modify data available to the authenticated user. The vulnerability is possible because Investigation Workbench does not safely neutralize special characters in artifact values before adding the values to searches. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see Investigate a potential security incident on the investigation workbench in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-7/user-guide/7.3/investigations/investigate-a-potential-security-incident-on-the-investigation-workbench-in-splunk-enterprise-security) and Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) in the Splunk documentation.	5.9	More Details
CVE-2026-127039	The certificate authentication info handling in gnutls from 2.5.9 up to 3.8.13 does not reset the certificate count and list pointer when re-allocation of raw_certificate_list fails, which allows remote attackers to cause a denial of service for a peer under heap pressure.	5.9	More Details

CVE-2026-132952	In JetBrains TeamCity before 2026.2 privilege escalation to system administrator was possible via project administrator roles in simple authorization mode	5.9	More Details
CVE-2026-132938	In JetBrains TeamCity before 2026.2 improper authorization allowed project-limited access tokens to act with full permissions via XML-RPC	5.9	More Details
CVE-2026-119631	The Kirki - Freeform Page Builder, Website Builder & Customizer plugin for WordPress is vulnerable to Unauthorized Data Access in versions up to, and including, 6.4.1 via the `kirki_post_apis` AJAX action's `get-dynamic-content-batch` and `get-collection-batch` endpoints. This is due to the plugin opening these resolvers behind a site-wide, resource-unbound preview token (`EditorPreview::has_valid_token()`) that is NOT bound to the specific page the administrator intended to share, combined with the absence of any `list_users` or `can_user_read_post()` capability check inside `HelperFunctions::get_user_dynamic_content()` and `Collection::get_collection_batch()` — the equivalent checks ARE present on the plugin's own REST route (`FrontendRESTController::get_item_permissions_check()`). This makes it possible for unauthenticated attackers holding any valid read-only share link to extract from the entire user table every user's `user_email`, `first_name`, `last_name`, `nickname`, `registered_date`, arbitrary string-valued `user_meta` (including custom keys such as `billing_phone`), and the e-mail address of every approved commenter — information that WordPress itself withholds from anonymous readers (served only in the `edit` context behind `list_users`).	5.9	More Details
CVE-2026-131220	In Splunk Enterprise Security versions below 8.7.0, a user who can influence event values that populate an affected data model based on the Common Information Model could cause attacker-controlled values to appear in an affected Dashboard Studio view. Another authenticated user could later click a corresponding row in that view. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who clicked the row. The commands could access or modify some data available to the user who clicked the row. The vulnerability is possible because affected Dashboard Studio drilldowns do not safely quote data model values before including them in searches. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) and Create a dashboard in Dashboard Studio (https://help.splunk.com/en/splunk-enterprise/create-dashboards-and-reports/dashboard-studio/10.2/create-a-dashboard-in-dashboard-studio/create-a-dashboard-in-dashboard-studio) in the Splunk documentation.	5.9	More Details
CVE-2026-132959	In JetBrains TeamCity before 2026.2 remote code execution on build agents was possible via service message injection in build configuration names	5.9	More Details
CVE-2026-132934	In JetBrains TeamCity before 2026.2, 2026.1.4 authentication bypass via polling agent name rebinding was possible	5.9	More Details
CVE-2026-84211	IBM Guardium Data Protection 12.2 could allow a remote attacker to inject forged audit records due to improper output neutralization for logs.	5.9	More Details
CVE-2026-131223	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user could trick an authenticated user into opening a link that the unauthenticated user creates. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the authenticated user. The commands could access data and perform actions available to that user. The vulnerability is possible because the Lookup Editor does not validate lookup identifiers from URL parameters before adding them to searches. The vulnerability requires the attacker to phish the user by tricking them into opening a link that the attacker creates. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see Create and manage lookups in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.0/managing-security-content/create-and-manage-lookups-in-splunk-enterprise-security) in the Splunk documentation.	5.9	More Details
CVE-2026-76542	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could trick a user who has an active Splunk Web session into selecting the Top values by time field report on the Search page. The selection could run arbitrary Search Processing Language (SPL) searches and potentially expose stored credentials and indexed data. The vulnerability is possible because Splunk Web does not correctly neutralize user-controlled data when it generates an SPL search for a field report. The vulnerability requires the attacker to phish the user by tricking them into opening a link that the attacker creates. The unauthenticated user should not be able to exploit the vulnerability at will.	5.9	More Details
CVE-2026-131584	Tanium addressed an information disclosure vulnerability.	5.9	More Details
	A vulnerability in the Audio, Web, and Video Conferencing (AWV) component of Mitel MiCollab through 10.3 (10.3.0.18) could allow an unauthenticated attacker, with knowledge of a specific system-generated random		

CVE-2026-133311	identifier or pre-staged content, to conduct an unauthorized access attack due to improper authentication mechanisms. A successful exploit could allow an attacker to obtain sensitive information. The vulnerability is limited to the affected functionality and does not provide unrestricted access to arbitrary files elsewhere on the system. This is MTLVULN-1863.	5.9	More Details
CVE-2026-108318	Tanium addressed an uncontrolled recursion vulnerability in Tanium Server.	5.9	More Details
CVE-2026-77382	falcon-rust is a Rust implementation of the Falcon post-quantum signature scheme. Prior to 0.2.0, the discrete-Gaussian sampler in src/samplerz.rs evaluates approx_exp, ber_exp, and sampler_z acceptance probabilities with FixedPoint64 values having only 32 fractional bits. The scaling in approx_exp quantizes its input and limits relative precision to about 2^-33, below Falcon's required bound, causing reference known-answer vectors to diverge and invalidating the proof assumption that secret-dependent signature and key distributions are negligibly close to the ideal discrete Gaussian. This is a cryptographic-soundness defect and no practical key-recovery attack has been demonstrated. This issue is fixed in version 0.2.0.	5.9	More Details
CVE-2026-130710	NocoDB before 2026.09.0 contains a missing authorization vulnerability in the job status polling endpoint of jobs.controller.ts that lacks an @Acl decorator. Unauthenticated attackers with a job id can read buffered status messages and completion payloads, exposing signed export download URLs or new base identifiers.	5.9	More Details
CVE-2026-131201	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user with network-path interception capability could cause attacker-controlled JavaScript to run in the browser of a security analyst who previews a Portable Document Format (PDF) attachment in Mission Control. The JavaScript could expose data and perform actions available to the security analyst. The vulnerability is possible because Mission Control loads the PDF rendering worker from an external content delivery network. The vulnerability requires the unauthenticated user to intercept the network path to the external content delivery network and a security analyst to preview a PDF attachment in Mission Control. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see Collaborate on investigations in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/investigations/collaborate-on-investigations-in-splunk-enterprise-security) in the Splunk documentation.	5.9	More Details
CVE-2026-127040	The RSA-PSK ClientKeyExchange handler in gnutls before 3.8.14 does not return after detecting a memory allocation failure when copying the PSK username, which allows remote attackers to cause a denial of service for a peer under heap pressure.	5.9	More Details
CVE-2026-131195	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could send repeated HTTP PubSub channel subscription requests that cause a Denial of Service (DoS) by permanently increasing memory usage until Splunk Enterprise exhausts available memory. The vulnerability affects instances that turn on the [pubsubsvr-http] stanza in server.conf because the HTTP PubSub broker does not release memory associated with a channel after the client disconnects.	5.9	More Details
CVE-2026-127231	A double free in the X.509 name constraints enforcement of GnuTLS 3.8.13 on all platforms allows a remote attacker to cause a denial of service (application crash) via a certificate chain containing a NameConstraints extension processed while the application is under memory pressure. When an already populated name constraints object is reinitialized, previously freed list pointers are not cleared, so a memory allocation failure precisely at reinitialization leaves a dangling pointer that is freed again during a later cleanup.	5.9	More Details
CVE-2026-120472	A flaw was found in OpenShift Hive. The ClusterClaim controller assigns hive-claim-owner RBAC based on a user-controlled ClusterClaim.spec.namespace and only verifies ClusterPoolRef.ClaimName. A tenant can point a same-named claim at another tenant's ClusterDeployment and obtain get access to that cluster's admin kubeconfig and admin password secrets.	5.9	More Details
CVE-2026-130745	Grav from 2.0.0 before 2.2.18 contains an access control bypass vulnerability that allows anonymous visitors to read protected pages because admin and site scopes share a page cache key. Routine editor activity through the API plugin rebuilds the page index without frontmatter.yaml access rules, exposing members-only page content until the cache is rebuilt.	5.9	More Details
CVE-2026-121481	Issue summary: Calling EVP_EncodeBlock() with an input length close to INT_MAX may overflow the calculation of the encoded output size and write past the end of the caller-supplied output buffer. Impact summary: No path by which an attacker can trigger this has been identified. An application that passes an input length near INT_MAX to EVP_EncodeBlock() could suffer a heap buffer overflow, crashing or corrupting its own memory, but it must itself supply such a length. CWE: CWE-190: Integer Overflow or Wraparound Description: EVP_EncodeBlock() base64-encodes a buffer of a caller-supplied length into an output buffer that the caller allocates. The caller is expected to reserve 4*((len+2)/3)+1 bytes. For input lengths close to INT_MAX, the arithmetic that computes the encoded length overflows a signed integer, so the function writes more data than the documented output size. The result is a heap buffer overflow in the calling process. Reaching this requires an application to pass an input length near INT_MAX to EVP_EncodeBlock(). OpenSSL's TLS, QUIC and CMS implementations never encode buffers of that size, and the command line tools cannot practically be made to do so: the limits on reading a proxy password from a file and on bignum sizes prevent it. An application that accepts an attacker-controlled length and forwards it to EVP_EncodeBlock() without its	5.9	More Details

	own bound check is exposed. The severity is Low because no attacker-controlled path to the overflow is known. The caller must already hold a buffer of about two gigabytes before the defect can be reached. FIPS impact: no The affected code is outside the FIPS module boundary.		
CVE-2026-131846	Poppler 0.76.1 through 26.11.0 contains a heap-based buffer overflow in SplashOutputDev::iccTransform() that allows attackers to write out of bounds when rendering in splashModeDeviceN8. Attackers can supply a PDF with an ICCBased image whose alternate is DeviceN, causing getDeviceNLine() to overflow the per-row buffer, crashing the renderer and corrupting heap memory.	5.8	More Details
CVE-2026-53526	Reposilite is lightweight and easy-to-use repository management software dedicated to Maven-based artifacts in the JVM ecosystem. From 3.5.16 until 3.6.0, as stated by the advisory, the Maven directory-index page interpolates the request URI without HTML escaping in its title, heading, base URL, and parent-directory link. The advisory says the content-embedded ctx.uri() is percent-encoded, preventing tag injection, but that a single quote passes through and breaks the quoted Parent-Directory link; it also says Location.of strips single quotes, so the reachability of the described quote-based trigger remains unresolved. The listing shares the Admin UI origin and does not set the sandbox Content-Security-Policy used for artifact and Javadoc responses, so the reported script execution could access a user's local-storage token. A manager token could enable full instance compromise. The advisory reports version 3.6.0 as fixed.	5.7	More Details
CVE-2026-115565	storyblok-backup is an npx CLI tool to create backups of a Storyblok space and restore individual resources. Prior to 0.7.1, the writeJson function used the editable Storyblok component name field as a filename without path containment. A person with component-management access in a Storyblok space could store a component name containing path traversal sequences; when another user later ran a backup for that space, the stored name escaped the configured backup directory and caused a JSON file to be written outside it on the backup host. Depending on the path chosen and the backup process's permissions, this could overwrite existing files at attacker-chosen locations. The vulnerability is specific to components: stories and datasources use numeric identifiers, and asset names are reduced to a final path segment. This issue is fixed in version 0.7.1.	5.7	More Details
CVE-2026-131216	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the manage_all_investigations capability could store attacker-controlled script content in a legacy investigation workbench panel title. Another authenticated user could later open a legacy investigation workbench that renders the panel. The script could then run in the browser of the user who opened the workbench and expose data or actions available through Splunk Web to the user who opened the workbench. The vulnerability is possible because Splunk Enterprise Security does not escape legacy investigation workbench panel titles before rendering them in Splunk Web. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) and Review and finish existing legacy investigations in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.4/investigations/review-and-finish-existing-legacy-investigations-in-splunk-enterprise-security) in the Splunk documentation.	5.7	More Details
CVE-2026-133193	In JetBrains Junie GitHub Action before 1.10.0 code execution was possible via path traversal in feedback session markers	5.6	More Details
CVE-2026-133194	In JetBrains Junie GitHub Action before 1.10.0 command injection was possible via the action version input	5.6	More Details
CVE-2026-109920	Use of uninitialized memory when processing certain font files in X.Org libXfont before 2.0.10 could be used by authenticated attackers to crash the X server.	5.5	More Details
CVE-2026-133182	In JetBrains IntelliJ IDEA before 2026.3 the GitHub plugin exposed GitHub access tokens via its public API	5.5	More Details
CVE-2025-71472	littlefs 2.8.0 through 2.11.3 contains a divide-by-zero vulnerability in lfs_mount_() that allows attackers to crash processes mounting crafted filesystem images when block_count is 0. Attackers can supply an image lacking a superblock entry or recording block_count 0, causing lfs->seed % lfs->block_count to trigger SIGFPE and denial of service.	5.5	More Details
CVE-2026-133519	FFmpeg 7.0 through 9.0.2 contains an out-of-bounds read vulnerability in dvdvideo_menu_next_ps_block() in libavformat/dvdvideodec.c when demuxing DVD menu domains. Attackers can craft a DVD image whose last cell sector and NAV next-VOBU pointer bypass the end-of-cell check, overreading pgc->cell_playback and crashing the process.	5.5	More Details
CVE-2026-133167	OFFIS DCMTK 3.5.4 through 3.7.1 contains an integer overflow in DJCodecEncoder::encodeTrueLossless size checks that allows attackers to trigger a heap out-of-bounds read. Attackers can supply a crafted DICOM file with large Columns, Rows and SamplesPerPixel values that, when compressed with true-lossless JPEG,	5.5	More Details

	crashes the application.		
CVE-2026-126827	A maliciously crafted SBSAR file, when imported or loaded through the Substance in 3ds Max plugin in Autodesk 3ds Max, can trigger a Stack Overflow vulnerability. A malicious actor may leverage this vulnerability to cause an application crash.	5.5	More Details
CVE-2026-126825	A maliciously crafted SBSAR file, when imported or loaded through the Substance in 3ds Max plugin in Autodesk 3ds Max, can trigger an Out-of-Bounds Read (table-index lookup) vulnerability. A malicious actor may leverage this vulnerability to cause an application crash or potentially read sensitive data from the current process.	5.5	More Details
CVE-2026-129927	A credential-handling weakness in the MongoDB Python Driver allows a configuration option value that protects private key material to be written in clear text into the embedding application's own output and log records when a driver object is converted to text. A party able to read that output may obtain the credential. Obtaining it requires read access to the application's local output or log files.	5.5	More Details
CVE-2026-92906	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 could allow a remote authenticated attacker to cause a denial of service due to client-side enforcement of business rules.	5.5	More Details
CVE-2026-126824	A maliciously crafted SBSAR file, when imported or loaded through the Substance in 3ds Max plugin in Autodesk 3ds Max, can trigger an Out-of-Bounds Read (unvalidated element begin-pointer) vulnerability. A malicious actor may leverage this vulnerability to cause an application crash or potentially read sensitive data from the current process.	5.5	More Details
CVE-2026-131874	Gargoyle 2023.1 through 2026.1.1 contains an infinite loop vulnerability in the bundled Plus interpreter's ReadString() function in terps/plus/loaddatabase.c that fails to handle end of file. Attackers can supply an empty, truncated, or crafted .plus game file that hangs Gargoyle at full CPU until the user kills it.	5.5	More Details
CVE-2026-129928	The MongoDB PHP extension includes the full connection string it was given in the error it raises when that string cannot be parsed. A party able to influence any part of the connection string, including a part that holds no secret, can cause the resulting message to carry credentials into the application's standard error output and log files. This may result in unintended disclosure of authentication material to anyone able to read that output.	5.5	More Details
CVE-2026-131206	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the edit_modinput_data_migrator capability could cause the data migration input to run arbitrary Search Processing Language (SPL) commands with elevated permissions. With this access, the user could read data outside the assigned role scope and write data to indexes and lookup tables. The vulnerability is possible because the data migration input does not restrict configured searches to product-supplied migration tasks before it runs them with elevated permissions. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) in the Splunk documentation.	5.5	More Details
CVE-2026-128137	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.9.4 until 3.10.4, NLTK corpus-reader path pointers and view methods validate a requested relative path without consistently resolving intermediate directory symlinks against the reader's scoped root. A local attacker who can place such a symlink can make abspaths or words reach a readable file outside the corpus root and disclose its contents. On releases with pathsec enforcement, the file must remain inside a global NLTK allowed root; version 3.9.4 has enforcement disabled and permits a broader escape. This issue is fixed in version 3.10.4.	5.5	More Details
CVE-2026-133007	A flaw was found in GIMP. Opening a crafted XCF with a truncated uncompressed tile can expose uninitialized main-process stack memory through the imported layer's pixel data, which may include process addresses that weaken ASLR.	5.5	More Details
CVE-2026-129230	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, WordNet Synset.max_depth, min_depth, hypernym_paths, and hypernym_distances recursively follow attacker-controlled corpus pointers without a visited-node set or depth cap. A cyclic or excessively deep hypernym graph can raise RecursionError and terminate traversal APIs and similarity operations that depend on them. This issue is fixed in version 3.10.4.	5.5	More Details
CVE-2026-129939	The MongoDB C Driver places a submitted connection-string option value into a diagnostic message without redaction when that option fails validation. The message is returned to the calling application and is also written to the process's standard error by the driver's default logging behavior. A party that can influence the connection string given to an application using the driver may cause that field's contents to be recorded, which may result in unintended disclosure of information the application placed in that field.	5.5	More Details
CVE-2026-129924	The MongoDB Ruby driver can include unmasked connection string content in the error it raises when a connection string cannot be parsed. If that connection string carries a secret, the secret may be disclosed in cleartext through the application's error output or logs. No privileges on the driver and no network access are	5.5	More Details

	needed to produce the error; a party able to read the application's error output may obtain the secret.		
CVE-2026-120260	InfoScan does not enforce authentication on an endpoint used to retrieve sales-automation templates. A local attacker can submit an HTTP request to the affected endpoint without a valid session and retrieve the returned template data.	5.5	More Details
CVE-2026-131872	SOIL2 before 1.4.2 contains a NULL pointer dereference vulnerability in its DDS, PVR and PKM image loaders that write results through unchecked caller-supplied output pointers. Attackers can supply a valid DDS, PVR or PKM file to applications passing NULL output pointers to SOIL_load_image or stbi_info, causing a crash.	5.5	More Details
CVE-2026-130694	NocoDB from 2026.07.0 before 2026.10.0 contains an sql injection vulnerability that allows users with Creator role to inject SQL through unescaped backslashes in the DATETIME_FORMAT formula format argument on MySQL-backed bases. Attackers can place a backslash before a quote to terminate the DATE_FORMAT string literal and read any data reachable by the data source's database user, including tables outside the base.	5.5	More Details
CVE-2026-73504	cassandra-jdbc-wrapper provides a JDBC-compliant API over the Java driver for Apache Cassandra. Prior to 4.16.4 and 5.0.2, DriverUtil.toStringWithoutSensitiveValues() and DriverUtil.redactSensitiveValuesInJdbcUrl() redact the password property but not the equally sensitive token property used for DataStax Astra and other token-based authentication. When DEBUG logging is enabled, SessionHolder passes connection properties and JDBC URLs containing token through these helpers, causing the live token to be written in plaintext when connections are acquired or released. An attacker who can read those logs can use the token to access Cassandra data and operations authorized to that credential. This issue is fixed in versions 4.16.4 and 5.0.2.	5.5	More Details
CVE-2026-132760	GPAC before 26.07.0 contains a NULL pointer dereference vulnerability in gf_isom_get_meta() in src/isomedia/meta.c that reads file->moov->trackList without checking the moov box exists. Attackers can supply a crafted ISOBMFF file with malformed AV1 configuration data that leaves moov unset, crashing MP4Box -infon and causing denial of service.	5.5	More Details
CVE-2026-132776	GPAC MP4Box before 26.07.0 contains a use-after-free vulnerability in the XMT loader that allows attackers to crash the process by supplying crafted XMT files. Attackers can assign a second node to an SFNode command field so xmt_node_end() frees a forward-referenced child still cached in peeked_nodes, triggering a crash during DEF lookup in xmt_find_node().	5.5	More Details
CVE-2026-132778	GPAC before 26.07.0 contains an uncontrolled recursion vulnerability in gf_sg_proto_field_is_sftime_offset() in src/scenegraph/vrml_proto.c that allows attackers to crash MP4Box via crafted files. Attackers can supply an MP4 file whose BIFS scene contains cyclic PROTO IS relations, causing unbounded recursion that exhausts the stack and causes denial of service.	5.5	More Details
CVE-2026-130688	NocoDB before 2026.10.1 contains a stored HTML injection vulnerability that allows authenticated users with record edit rights to inject forms, images, and styled overlays into rich text cells, which are sanitized only with DOMPurify's default policy. Attackers can place a page-sized spoofed sign-in form or remote image in a cell to harvest passwords and disclose viewer IP addresses in expanded records, gallery and kanban cards, and shared views.	5.5	More Details
CVE-2026-132768	GPAC MP4Box before 26.07.0 contains a use-after-free vulnerability in xmt_find_node() in loader_xmt.c that leaves freed forward-referenced nodes cached in parser->peeked_nodes. Attackers can supply a crafted XMT file with a node targeted by an SFCommandBuffer command, triggering a later DEF lookup on the freed node and crashing MP4Box.	5.5	More Details
CVE-2026-132767	GPAC before 26.07.0 contains a NULL pointer dereference vulnerability in xmt_parse_command() in loader_xmt.c that allows attackers to crash MP4Box via unchecked field lookups. Attackers can supply a crafted XMT scene whose atIndexField names a non-existent node field, causing a NULL read during import and denial of service.	5.5	More Details
CVE-2026-129829	The rsyslog's`imuxsock` special parser can end up writing into the received message raw buffer when trying to replace the message timestamp with the system's timestamp. A local attacker with the ability to write to the `imuxsock` socket can leverage that by crafting a malformed PRI value with a short timestamp to lead the rsyslog to perform the out-of-bounds write.	5.5	More Details
CVE-2026-132762	GPAC MP4Box before 26.07.0 contains a divide by zero vulnerability in DumpTrackInfo() in applications/mp4box/filedump.c that fails to check the media timescale for zero. Attackers can supply a crafted MP4 file with a zero media timescale to make MP4Box -infon raise SIGFPE and crash, causing denial of service.	5.5	More Details
CVE-2026-132770	GPAC MP4Box before 26.07.0 contains a NULL pointer dereference vulnerability in xmt_parse_proto() in loader_xmt.c that allows attackers to crash the process via a redefined ProtoDeclare name. Attackers can supply a crafted XMT file with a duplicate ProtoDeclare name, causing gf_sg_proto_new() to return NULL and crashing MP4Box during import.	5.5	More Details
CVE-2026-132761	GPAC before 26.07.0 contains a use-after-free vulnerability in gf_node_try_destroy that allows attackers to crash MP4Box by supplying crafted XMT scene files. Commands created inside a PROTO scope keep a raw pointer to a freed prototype sub-scene graph, which is read during teardown, causing denial of service.	5.5	More Details

CVE-2026-132771	GPAC before 26.07.0 contains a use-after-free vulnerability in ReplaceDEFNode that allows attackers to crash MP4Box by triggering access to a partly freed child list. Attackers can supply a crafted XMT scene file whose InputSensor command buffer cascades prototype sub-graph deletion during gf_sg_vrml_parent_destroy, causing denial of service.	5.5	More Details
CVE-2026-132758	GPAC MP4Box before 26.07.0 contains a NULL pointer dereference vulnerability in xmt_node_end() that allows attackers to crash the application through an unmatched closing ProtoDeclare element. Attackers can supply a crafted XMT-A scene file with a mismatched closing ProtoDeclare or ExternProtoDeclare element, causing a dereference of a NULL parsing_proto pointer and denial of service.	5.5	More Details
CVE-2026-132757	GPAC before 26.07.0 contains a NULL pointer dereference vulnerability in DoWriteMeta() in src/isomedia/isom_store.c that allows attackers to crash MP4Box via crafted files. Attackers can supply an ISOBMFF file with a meta box containing an iloc entry but no dref box, crashing MP4Box during rewrite operations like -unhint.	5.5	More Details
CVE-2026-132753	MP4v2 through 2.1.0 contains a NULL pointer dereference vulnerability in MP4RtpHintTrack::GetPayload() in src/rtphint.cpp that passes an unchecked rtpmap value to strchr(). Attackers can supply a crafted MP4 file whose RTP hint track has a corrupted hinf atom to crash mp4info or applications calling MP4GetHintTrackRtpPayload().	5.5	More Details
CVE-2026-115622	A flaw was found in ansible-core. When a task or its parameters are marked no_log: true, sensitive values passed to a module's warn() or deprecate() calls are not redacted and are printed in cleartext to stdout/stderr in warning and deprecation messages. The 'warnings' and 'deprecations' keys are retained in the PRESERVE allow-list in lib/ansible/_internal/_task.py when the task result is censored, so sensitive parameters bypass masking and can be disclosed in output captured in terminal sessions, CI logs, and any automation wrapping ansible-playbook, even though every task touching the secrets is marked no_log: true.	5.5	More Details
CVE-2026-132956	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Build Features page was possible	5.4	More Details
CVE-2026-133430	A flaw was found in Keycloak. This vulnerability allows an authenticated remote attacker to cause a denial of service by terminating active sessions of arbitrary users. When processing Security Assertion Markup Language (SAML) documents that lack a root signature, Keycloak accepts the request as long as any assertion within the document is signed, without verifying its context or placement. An attacker holding a validly signed assertion can embed it into an unsigned logout request for another user, causing Keycloak to accept the forged request and terminate the victim's session.	5.4	More Details
CVE-2026-132958	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the artifact dependency tree was possible	5.4	More Details
CVE-2026-132248	In JetBrains TeamCity before 2026.1.3 improper authorization allowed project viewers to change project IDs	5.4	More Details
CVE-2026-133221	A flaw was found in Keycloak. An authenticated remote user with a pending required security action, such as configuring multi-factor authentication (MFA), can bypass the mandatory action by submitting a specially crafted request to the consent endpoint without a valid session code. Keycloak incorrectly interprets the submission as a page refresh and completes the authentication process, allowing the user to access the application without fulfilling the required security policies.	5.4	More Details
CVE-2026-131233	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in endpoint summary data. If another user who can access the Time Center dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify some data available to the user who selects the row. The vulnerability is possible because the Time Center dashboard does not safely escape endpoint summary values before inserting them into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Endpoint dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/endpoint-dashboards) in the Splunk documentation.	5.4	More Details
CVE-2026-132963	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the checkout rules path preview was possible	5.4	More Details
	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_user" Splunk Enterprise Security role could store attacker-controlled asset data in an asset source lookup. A user that holds a role with the mc_assets_read capability could later open the Asset Center dashboard and select that value in the Asset Information table. Splunk Enterprise Security could then run attacker-controlled Search Processing		

CVE-2026-131205	Language (SPL) commands using the permissions of that user. With this access, the user who stored the asset data could expose or modify data available to the user who selected the value. The vulnerability is possible because Asset Center does not escape asset values before adding them to drilldown searches. The vulnerability requires Asset Center to process the asset data and another user to select the value. The user who stored the data should not be able to exploit the vulnerability at will. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security), Add asset and identity data to Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/asset-and-identity-management/add-asset-and-identity-data-to-splunk-enterprise-security), and Asset and identity dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/asset-and-identity-dashboards) in the Splunk documentation.	5.4	More Details
CVE-2026-132752	ScreenTinker before 3.2.4 contains a stored cross-site scripting vulnerability in the GET /api/kiosk/:id/render endpoint that allows users with kiosk page edit rights to store javascript: URLs in page buttons. Attackers can trigger script execution in the dashboard origin when an administrator activates the button, hijacking the victim's session and taking full control of the instance.	5.4	More Details
CVE-2026-131207	In Splunk Enterprise Security versions below 8.7.0, an authenticated user who does not hold the edit_correlationsearches capability could set Status to Testing for a versioned detection or clear that setting. This could affect the integrity and availability of detection testing. The vulnerability is possible because the Test Mode Representational State Transfer (REST) API does not require the edit_correlationsearches capability before it creates or deletes scheduled test copies of detections. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/capability-reference-for-splunk-enterprise-security) and Test detections on production data using Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/detections/test-detections-on-production-data-using-splunk-enterprise-security) in the Splunk documentation.	5.4	More Details
CVE-2026-131237	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in data used by the Weak Encrypted Communication report. If another user who can access the report selects an affected drilldown, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the drilldown. The user who stores the values could then access sensitive data and modify some data available to the user who selects the drilldown. The vulnerability is possible because the Weak Encrypted Communication report does not validate source, destination, and Secure Sockets Layer version values before inserting them into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Weak Encrypted Communication (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/weak-encrypted-communication) in the Splunk documentation.	5.4	More Details
CVE-2026-131275	In Splunk SOAR versions below 8.8.0, a user with container edit permission could supply timestamps when creating containers and artifacts through the Representational State Transfer (REST) API. This could alter the timeline shown to analysts and cause configured data retention strategies to delete affected records earlier than intended. The vulnerability is possible because the REST API accepts client-supplied values for timestamps that Splunk SOAR uses to order records and apply data retention. The early deletion impact requires an applicable data retention strategy; the timeline impact does not. For more information see Manage your PostgreSQL database with data retention strategies (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.4.0/manage-your-splunk-soar-on-premises-system-health-and-performance/manage-your-postgresql-database-with-data-retention-strategies) in the Splunk documentation.	5.4	More Details
CVE-2026-131231	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in managed-lookup audit data. If another user who can access the Managed Lookups Audit dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify some data available to the user who selects the row. The vulnerability is possible because the Managed Lookups Audit dashboard does not safely escape audit data before inserting it into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Audit dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/audit-dashboards) in the Splunk documentation.	5.4	More Details
CVE-2026-133116	Termix 1.7.0 before 26.11.2 contains a path traversal vulnerability that allows authenticated users to write files outside the export directory via crafted usernames. Attackers can self-register with a username containing ../ segments and request POST /database/export to plant attacker-controlled SQLite files in writable directories, enabling disk exhaustion.	5.4	More Details
CVE-			More

2026-132954	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Triggers page was possible	5.4	Details
CVE-2026-132924	In JetBrains TeamCity before 2026.2, 2026.1.4 improper authorization allowed users with the Run build permission to delete Perforce workspaces	5.4	More Details
CVE-2026-120478	The HivePress – Business Directory, Listings & Classified Ads Plugin plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'description' parameter in all versions up to, and including, 1.8.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	5.4	More Details
CVE-2026-133308	ClassiCube before 1.4.0 contains a heap-based out-of-bounds write vulnerability in Png_Reconstruct in src/Bitmap.c that occurs when AVERAGE and PAETH filters process zero-width PNG images. A malicious multiplayer server can supply a crafted texture pack URL or player skin to write past the heap buffer, corrupting memory and crashing connecting clients.	5.4	More Details
CVE-2026-132928	In JetBrains TeamCity before 2026.2, 2026.1.4 stored XSS in Perforce job details was possible	5.4	More Details
CVE-2026-132936	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the npm connection health report was possible	5.4	More Details
CVE-2026-133225	A flaw was found in Keycloak. When processing token requests using the User-Managed Access (UMA) grant, Keycloak fails to verify whether a target client or resource server is enabled before issuing an authorization token. An authenticated remote user can exploit this flaw to obtain valid Requesting Party Tokens (RPTs) for disabled clients, bypassing administrative access controls and permissions restrictions.	5.4	More Details
CVE-2026-133224	A flaw was found in Keycloak. When an administrator revokes tokens by setting a not-before revocation cutoff, Keycloak fails to enforce this boundary when processing User-Managed Access (UMA) grant requests using ID token claims. An authenticated remote user with an active session can submit a previously revoked ID token to receive a fresh Requesting Party Token (RPT). This flaw allows an attacker to bypass token revocation and retain unauthorized access to protected resources.	5.4	More Details
CVE-2026-132944	In JetBrains TeamCity before 2026.2, 2026.1.4 improper authorization allowed project agent administrators to manage agents of other projects	5.4	More Details
CVE-2026-129221	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, nltk.classify.weka.ARFF_Formatter writes attacker-controlled class labels into the Weka ARFF @ATTRIBUTE '-label-' declaration without safely formatting the label tokens. Crafted labels can alter the generated ARFF schema or records, corrupting classifier input and causing integrity loss or processing failure. This issue is fixed in version 3.10.4.	5.4	More Details
CVE-2026-107517	Dell PowerStore contains a stored cross-site scripting (XSS) vulnerability. A low-privileged authenticated attacker could exploit this vulnerability, potentially leading to execution of malicious script in a higher-privileged user's session.	5.4	More Details
CVE-2026-133521	Bar Assistant 3.0.0 before 6.11.0 contains an authorization bypass vulnerability in BarController::join that allows authenticated users to join bars without a valid invite code. Attackers can send POST requests to /api/bars/join omitting invite_code, matching invite-disabled bars with NULL codes, to read cocktails, ingredients, members, settings and search tokens.	5.4	More Details
CVE-2026-131225	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in the Malware data model. Another authenticated user with access to the Malware Search dashboard could later select a table row containing those values. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selected the row. The commands could expose data available to the user who selected the row and add attacker-controlled records to an index. The vulnerability is possible because the Malware Search dashboard does not safely quote row values before adding them to drilldown searches. For more information see Endpoint dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.4/analytics/endpoint-dashboards) and Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) in the Splunk documentation.	5.4	More Details
CVE-2026-128717	An issue in PHPGurukul Employee Leaves Management System (ELMS) Version 2.1 allows a remote attacker to escalate privileges via the PHP session management mechanism	5.4	More Details

CVE-2026-131228	In Splunk Enterprise Security versions below 8.7.0, a user who holds the "ess_user" Splunk Enterprise Security role could run a search that causes attacker-controlled Search Processing Language (SPL) text to appear in the Search Audit dashboard. If another user who can access the dashboard selects the affected search, Splunk Enterprise Security could run attacker-controlled SPL commands using the permissions of the user who selects the search. The user who runs the first search could then access sensitive data and modify some data available to the user who selects the affected search. The vulnerability is possible because the Search Audit dashboard does not safely escape logged search text before inserting it into drilldown searches. For more information see Users and roles for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.7/installation/users-and-roles-for-splunk-enterprise-security) and Audit dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/audit-dashboards) in the Splunk documentation.	5.4	More Details
CVE-2026-131229	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in vulnerability data. If another user who can access the Vulnerability Search dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data and modify some data available to the user who selects the row. The vulnerability is possible because the Vulnerability Search dashboard encodes data-model field values for the drilldown link but does not escape the values as SPL search literals before inserting the values into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Network dashboards (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/analytics/network-dashboards) in the Splunk documentation.	5.4	More Details
CVE-2026-108315	Tanium addressed an improper access controls vulnerability in Tanium Server.	5.4	More Details
CVE-2026-108314	Tanium addressed an incorrect default permissions vulnerability in Tanium Server.	5.4	More Details
CVE-2026-132964	In JetBrains TeamCity before 2026.2 stored XSS on the Duplicates report tab was possible	5.4	More Details
CVE-2026-132960	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the diff view page was possible	5.4	More Details
CVE-2026-132925	In JetBrains TeamCity before 2026.2 stored XSS in the duplicate remote parameter health report was possible	5.4	More Details
CVE-2026-131254	In Splunk SOAR versions below 8.8.0, a user with permission to create containers on one label could use the Representational State Transfer (REST) API to create a container with a parent relationship to a container that the user cannot access. This could let the user alter incident lineage across authorization boundaries and identify existing restricted containers. The vulnerability is possible because the REST API accepts a customer-supplied parent relationship when creating a container without authorizing access to the referenced container. For more information see REST Containers (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/8.6.0/container-endpoints/rest-containers) and Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	5.4	More Details
CVE-2026-130676	Craft CMS from 5.0.0 before 5.11.4 contains a stored cross-site scripting vulnerability that allows low-privilege control panel users to inject JavaScript by poisoning cached license data. Attackers can store a malicious licensed domain via the app/process-api-response-headers endpoint, which executes unescaped in administrators' dashboard license-mismatch warnings.	5.4	More Details
CVE-2026-133242	diaspora* before 0.10.0 contains a stored cross-site scripting vulnerability that allows authenticated users to inject script by crafting og:video:secure_url values. Attackers can post links to pages whose allowlisted-prefix video URL contains a quote and onload handler, executing script in readers' sessions when they click the video preview.	5.4	More Details
CVE-2026-130684	vLLM before 0.33.0 contains a path traversal vulnerability that allows malicious model publishers to redirect the served model and tokenizer by setting speculators_config verifier.name_or_path in config.json. Attackers can point this field at another local checkpoint directory to leak its tokenizer vocabulary via unauthenticated /detokenize requests and gain query access to private models.	5.4	More Details
CVE-	A flaw was found in Moodle. Insufficient authorization checks in the quiz editing functionality allow an		

2026-133072	authenticated user with teacher privileges to view slot details and modify sections of quizzes they do not have permission to manage. An attacker could exploit this flaw to alter quiz configurations or disclose quiz structure information without proper authorization.	5.4	More Details
CVE-2026-133071	A flaw was found in Moodle. Due to a missing authorization check when editing Lesson activity pages, authenticated users with teacher privileges can view and move lesson pages from courses they are not permitted to access into their own lessons. This issue can lead to unauthorized access and modification of restricted course content.	5.4	More Details
CVE-2026-133069	A flaw was found in Moodle. Missing authorization checks in the forum discussion web services allow authenticated users to view discussion information in courses they do not have permission to access. In addition, users with teacher privileges may lock forum discussions in courses where they lack proper authorization. This issue can lead to unauthorized information disclosure and unauthorized discussion state modification.	5.4	More Details
CVE-2026-133048	In JetBrains TeamCity before 2026.2.1 improper authorization allowed using AWS connections of other projects in cloud profiles	5.4	More Details
CVE-2026-133067	A flaw was found in Moodle. Due to missing ownership checks in badge alignments, an authenticated user with badge management privileges can view and modify alignments associated with badges they do not own. This vulnerability can lead to unauthorized data disclosure and tampering with badge configurations.	5.4	More Details
CVE-2026-130649	WWBN AVideo from 11.1 contains a cross-site scripting vulnerability that allows users with streaming permission to inject HTML by supplying an unvalidated url parameter to plugin/LiveLinks/play.php. Attackers can embed a --> sequence to escape the HTML comment in liveVideo.php, creating a shareable link that executes JavaScript in any visitor's session.	5.4	More Details
CVE-2026-131174	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could make unauthorized changes when restoring archived indexed data or cause a Denial of Service (DoS) on the Splunk Enterprise instance. The vulnerability is possible because the Representational State Transfer (REST) API does not require the indexes_edit capability for archived-data restoration actions. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation. For more information see Restore archived indexed data (https://help.splunk.com/en/splunk-enterprise/administer/manage-indexers-and-indexer-clusters/10.6/back-up-and-archive-your-indexes/restore-archived-indexed-data) in the Splunk documentation.	5.4	More Details
CVE-2026-133066	A flaw was found in Moodle. Missing ownership verification in Feedback activity templates allows an authenticated teacher to view, use, or delete templates associated with courses other than their own. This can result in unauthorized access, modification, or deletion of course materials.	5.4	More Details
CVE-2026-100181	The Asset CleanUp: Page Speed Booster plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment' parameter in all versions up to, and including, 1.4.0.6 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 exploit succeeds even though wp_kses is applied on save, because the XSS is manufactured post-sanitization by the plugin's regex pass: kses-permitted tags such as <a title> and <code> carry the attacker's marker and terminator strings intact into the rendered HTML buffer, where the plugin's own replacement then corrupts the document structure.	5.4	More Details
CVE-2026-76543	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could store attacker-controlled Search Processing Language (SPL) in a shared data model. The vulnerability is possible because the Data Model Editor does not properly neutralize data model field names when it builds a preview search. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The attacker should not be able to exploit the vulnerability at will. For more information see Add a regular expression field (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.4/define-data-model-dataset-fields/add-a-regular-expression-field), Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities), and About the search language (https://help.splunk.com/en/splunk-enterprise/search/search-manual/10.4/search-overview/about-the-search-language) in the Splunk documentation.	5.4	More Details
CVE-2026-120049	The Link Library plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'linkfile[name]' parameter in all versions up to, and including, 7.10.0 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. This requires an administrator to have enabled the "Link File" option within the plugin's settings, which activates the vulnerable file upload field on the public add-link form.	5.4	More Details

CVE-2026-131296	In Splunk SOAR versions below 8.8.0, a user with permission to author and run playbooks could add messages to a different playbook run log and cause that run to fail, including runs outside the labels that the user could access. The vulnerability is possible because Splunk SOAR does not verify that playbook API requests use the trusted runtime context for the current playbook run.	5.4	More Details
CVE-2026-131250	In Splunk SOAR versions below 8.8.0, a user with permission to edit containers for an allowed label could forge the author attribution of container pins and access limited user details outside the expected authorization boundary. The Container Pin Representational State Transfer (REST) API does not prevent caller-provided author fields from reaching generic model binding, which lets the request reassign a pin author or query user-related data outside the expected authorization scope. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.3.1/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	5.4	More Details
CVE-2026-1212	The BuddyPress plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 14.4.0. This is due to the `get_item_permissions_check` function in the Messages REST API endpoint trusting a client-supplied `user_id` parameter instead of using the authenticated user's ID. This makes it possible for authenticated attackers with Subscriber-level access and above to read, update, or delete other users' private message threads by specifying the victim's user ID in the request.	5.4	More Details
CVE-2026-130734	Grav from 2.0.0 before 2.2.18 contains a stored cross-site scripting vulnerability that allows users with media upload rights to bypass the upload sanitizer. Attackers can upload SVG files with inner extensions like .svg.defaults, payloads that crash the fail-open sanitizer, or unblocked .xml/.xhtm files to run script when administrators open the raw file URL.	5.4	More Details
CVE-2026-133257	In JetBrains YouTrack before 2026.2.19840 improper authorization allowed read-only users to modify and delete articles	5.4	More Details
CVE-2026-132966	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in build statistics charts was possible	5.4	More Details
CVE-2026-133262	In JetBrains YouTrack before 2026.2.19766 improper authorization allowed read-only users to change the visibility of restricted issues via agile boards	5.4	More Details
CVE-2026-133261	In JetBrains YouTrack before 2026.2.19840 an HTML injection in email notifications was possible via issue descriptions	5.4	More Details
CVE-2026-131257	In Splunk SOAR versions below 8.8.0, a user with permission to view and edit containers and at least one container label could set misleading user and app provenance on containers that the user creates and learn whether user IDs and installed app identifiers exist. The vulnerability is possible because the Representational State Transfer (REST) API accepts client-supplied values for server-managed container fields.	5.4	More Details
CVE-2026-55504	Chamilo LMS is an open-source learning management system. Prior to 2.0.3, POST /api/social_posts accepts attacker-controlled content that SocialPost::setContent stores without cross-site scripting sanitization, and assets/vue/components/social/SocialWallPost.vue renders post.content through v-html without sanitization. Any authenticated user can create a social-wall post that executes active content when another authenticated user views the wall. The script can run in an administrator's browser session and expose session data or perform actions with the victim's privileges. This issue is fixed in version 2.0.3.	5.4	More Details
CVE-2026-131259	In Splunk SOAR versions below 8.8.0, a user with permission to create and view containers for a label could use the Representational State Transfer (REST) API to create a container that references a workbook phase in a case the user cannot view, disclosing the phase name and creating an inconsistent phase association. The vulnerability exists because the Container REST API does not enforce the same case-scoped authorization when it accepts a workbook-phase reference during container creation.	5.4	More Details
CVE-2026-132984	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the project Report Tab dialog was possible	5.4	More Details
CVE-2026-132968	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Large Build Log health report was possible	5.4	More Details
CVE-2026-81499	Incus is a system container and virtual machine manager. Prior to 7.3.0, the six network address set handlers in cmd/incusd/network_address_sets.go resolved the effective project with project.NetworkProject(), which returns default when features.networks is unset. The expandProject authorization switch in cmd/incusd/daemon.go omitted auth.ObjectTypeNetworkAddressSet, so authorization used the requested project instead of the effective project. This issue is fixed in version 7.3.0.	5.4	More Details

CVE-2026-55503	Chamilo LMS is an open-source learning management system. Prior to 1.11.40 and 2.0.3, main/lp/scormItem.class.php and public/main/lp/scormItem.class.php store an attacker-controlled title element from imsmanifest.xml without cross-site scripting sanitization when a teacher imports a SCORM package. The title reaches c_lp_item.title and is rendered by src/CoreBundle/Resources/views/LearnPath/scorm_list.html.twig through item.title with raw output while Twig auto-escaping is disabled; the default configuration does not apply the optional AI-disclosure sanitization path. Enrolled students or administrators who view the learning-path table of contents can execute the stored content in their browser sessions. This issue is fixed in versions 1.11.40 and 2.0.3.	5.4	More Details
CVE-2026-131183	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a low-privileged user who holds the "power" Splunk role could inject arbitrary Search Processing Language (SPL) commands into a Classic Dashboard search. When another user opens the affected dashboard in Splunk Web, those injected commands execute with the viewing user's permissions, potentially allowing the low-privileged user to access stored credentials and write data to indexes. The vulnerability is possible because Splunk Enterprise does not correctly escape user-supplied content when it rewrites a search, which lets the rewritten search include additional SPL commands. The vulnerability requires another user to open the dashboard in Splunk Web. Consequently, a user who holds the "power" Splunk role should not be able to exploit the vulnerability at will. For more information see SPL safeguards for risky commands (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/best-practices-for-splunk-platform-security/spl-safeguards-for-risky-commands) and Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.	5.4	More Details
CVE-2026-130752	OpenWrt LuCI luci-app-uhttpd before commit d38658bedbdd811b798da16b49301478c6cf21af grants '/*' file read in its rpcd read ACL, which rpcd's fnmatch matching applies to every path. Users delegated only luci-app-uhttpd read access can read any file as root, including /etc/shadow, /etc/config/* credentials, /etc/uhttpd.key, and /etc/dropbear keys.	5.3	More Details
CVE-2026-132138	A flaw was found in cyrus-sasl. When processing authentication requests using the Salted Challenge Response Authentication Mechanism (SCRAM), a remote attacker can trigger a null pointer dereference by submitting unexpected channel-binding data. This causes the server process to crash before validating authentication proofs, leading to a Denial of Service (DoS).	5.3	More Details
CVE-2026-1214	The BuddyPress plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 14.5.0. This is due to missing authorization checks in the `bp_nouveau_ajax_get_single_activity_content` AJAX handler. The function only verifies a nonce that is publicly exposed to any visitor who loads a BuddyPress page, but does not call `bp_activity_user_can_read()` to verify the user has permission to view the activity. This makes it possible for unauthenticated attackers to read private group activity content by iterating activity IDs.	5.3	More Details
CVE-2026-92535	The Make Connector plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 1.6.6 via the include_custom_fields. This makes it possible for unauthenticated attackers to extract arbitrary admin-selected user, post, comment, and term metadata exposed via core REST endpoints — including active session tokens with source IP, user-agent, login and expiry timestamps, capability sets, and any other admin-selected private meta keys including protected underscore-prefixed entries. This is exploitable because the plugin's meta picker queries all meta keys from the database without filtering, offering protected underscore-prefixed keys such as session_tokens and wp_capabilities for administrator selection; if any such keys are selected, their values become readable without any credentials.	5.3	More Details
CVE-2026-102434	IBM Engineering AI Hub 1.0.0 through 1.4.0 could allow a remote attacker to obtain sensitive information due to a Cross Origin Resource Sharing (CORS) misconfiguration.	5.3	More Details
CVE-2026-130666	ImageMagick before 7.1.2-35 contains a policy bypass vulnerability that allows attackers to evade resource limits because a policy check is missing when reading XMP profiles. Attackers can supply an image with a crafted XMP profile to trigger excessive memory usage and cause denial of service.	5.3	More Details
CVE-2026-131847	Poppler through 26.11.0 contains a use-after-free vulnerability in NSSCryptoSignBackend.cc that allows attackers to trigger reads of freed CMS arena memory during repeated signature verification. Attackers can supply a signed PDF embedding the signer certificate to crash applications or pdfsig runs validating multiple signatures, or cause incorrect certificate validation results.	5.3	More Details
CVE-2026-126041	The Hide My WP Ghost WordPress plugin before 7.0.19 does not remove or restrict a debug branch in its output rewriting, allowing unauthenticated attackers to reveal the hidden login page, the hidden admin path and the renamed content folders the Hide My WP Ghost WordPress plugin before 7.0.19 is meant to conceal.	5.3	More Details
CVE-2026-77336	Aperi'Solve is an open-source steganalysis web platform. Prior to 3.4.2, an unauthenticated user can submit a PNG to POST /upload, which reaches aperisolve/utils/png.py through analyze_image, PCRTAnalyzer, PNG.repair(), and check_idat(). A malformed IDAT chunk can make attacker-controlled length and newline data drive an unbounded combinations search in PNG._fix_dos2unix(), causing repeated CRC checks that occupy a worker and delay or block image analysis for other users. Per-IP upload limits and eventual job	5.3	More Details

	termination bound the denial of service, and the issue has no confidentiality or integrity impact. This issue is fixed in version 3.4.2.		
CVE-2020-37290	Kiwi TCMS before 8.6 contains a missing authorization vulnerability in the User.filter JSON-RPC method that allows unauthenticated attackers to retrieve user account information. Attackers can send a POST request to /json-rpc/ invoking User.filter with a query parameter to disclose sensitive user data without authentication. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	5.3	More Details
CVE-2026-133484	A vulnerability was found in ByteDance Lark Desktop 8.0.0.14. This vulnerability affects the function FUN_18001b770/FUN_18001ba40 in the library backend.dll. Performing a manipulation of the argument day results in signed to unsigned conversion error. The attack needs to be approached locally. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2026-133324	A vulnerability in the Audio, Web, and Video Conferencing (AWV) component of Mitel MiCollab through 10.3 (10.3.0.18) could allow an unauthenticated attacker to conduct an authentication bypass attack due to improper handling of authentication. A successful exploit could allow an attacker to bypass authentication and perform unauthorized actions within the context of the application. This is MTLVULN-1877.	5.3	More Details
CVE-2026-120735	The Appointment Booking Plugin – LatePoint | Calendar & Scheduling for WordPress plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 5.7.7 via the 'key' parameter. This makes it possible for unauthenticated attackers to extract full invoice and receipt records including customer full name, email address, phone number, invoice line items and totals, and the invoice's real access key. This is exploitable via the unauthenticated public routes invoices__view_by_key and transactions__view_receipt_by_key, which apply no ownership, capability, or nonce check before passing the attacker-supplied key into the ORM lookup.	5.3	More Details
CVE-2026-130655	ImageMagick before 7.1.2-35 contains a denial of service vulnerability in the BGR, CMYK, GRAY, RGB, YCBCR and YUV decoders that allows attackers to crash the application. Attackers can supply a crafted image to these decoders to trigger a crash and disrupt image processing services.	5.3	More Details
CVE-2026-11602	The FS Poster - WordPress Social media Auto Poster & Scheduler [Facebook, Instagram, Twitter, Pinterest] plugin for WordPress is vulnerable to missing authorization to Unauthorized Social App Data Modification in all versions up to, and including, 8.0.1 This is due to missing authentication, nonce, and capability checks in the `standardFSApp()` function, which is registered on the `wp` action hook and executes on every frontend page load without any login or permission verification. This makes it possible for unauthenticated attackers to insert arbitrary social network app records into the database, or overwrite the OAuth `client_id` and `client_secret` of existing social network connections, by supplying attacker-controlled values via the `fsp_app_redirect` and `app` GET parameters. An attacker can target a specific existing social app record for credential overwrite by supplying matching `social_network` and `slug` values, requiring no knowledge of internal record IDs, which disrupts scheduled social media posts and forces administrators to re-authenticate their connected accounts.	5.3	More Details
CVE-2026-71427	WPGraphQL Content Blocks is a WordPress plugin that extends WPGraphQL to support querying Gutenberg blocks as data. Prior to 4.8.6, the editorBlocks renderedHtml field in includes/Data/ContentBlocksResolver.php uses resolve_content_blocks() to read a post with get_post() without applying post_password_required(). An unauthenticated caller who knows a supported post identifier can therefore retrieve the complete body of a password-protected published post without a post password, nonce, or valid cookie, exposing paid, member-only, preview, or internal content. This issue is fixed in version 4.8.6.	5.3	More Details
CVE-2026-132283	A vulnerability was identified in Directus up to 12.7.0. This issue affects the function AuthenticationService.refresh of the file api/src/services/authentication.ts of the component Authentication Service. Such manipulation leads to session expiration. The attack can be launched remotely. The vendor was contacted early about this disclosure.	5.3	More Details
CVE-2026-130719	NocoDB before 2026.08.0 contains an information disclosure vulnerability that allows unauthenticated shared-base link holders to read the base's stored plaintext password through the general base metadata endpoint. Attackers with a share link can query the base read or list responses in bases.controller.ts, which skip password stripping in sanitizeBase.ts, to obtain reusable plaintext passwords.	5.3	More Details
CVE-2026-69133	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. From 0.0.20.Final until 0.0.27.Final, OHttpClientCodec in codec-ohttp/src/main/java/io/netty/incubator/codec/ohttp/OHttpClientCodec.java and OHttpServerCodec in codec-ohttp/src/main/java/io/netty/incubator/codec/ohttp/OHttpServerCodec.java retain the consumed outer cumulationBuffer after a valid OHTTP exchange completes, even though OHttpRequestResponseContext.destroy in codec-ohttp/src/main/java/io/netty/incubator/codec/ohttp/OHttpRequestResponseContext.java releases the inner Binary HTTP cumulation and the channel remains active. A peer can complete a large request or response and leave the keep-alive connection idle, causing the client and server handlers to retain the current high-water pooled or direct buffer capacity until later inbound content replaces it or the handler is removed. Many long-lived idle connections can pin substantial memory and create availability pressure. The issue does not	5.3	More Details

	cause unbounded growth for repeated exchanges on a single channel. This issue is fixed in version 0.0.27.Final.		
CVE-2026-130725	NocoDB before 2026.10.0 contains an insecure direct object reference vulnerability that allows authenticated users to obtain signed download URLs for other bases' attachments by storing foreign attachment paths in their own rows via prepareNocoData. Attackers with write access to an attachment column can have _convertAttachmentType or the attachment download endpoint sign the foreign path, reading cross-base files despite secure attachment mode.	5.3	More Details
CVE-2026-77379	Bolt is a fully open source, adaptable platform for building and running modern websites. Prior to 6.1.4, the public Relation API lacks an access-control filter. Unauthenticated users can query its REST endpoints and GraphQL endpoints to discover metadata about unpublished, draft, and viewless content. This issue is fixed in version 6.1.4.	5.3	More Details
CVE-2026-14880	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 could allow unauthorized network access due to insufficient network segmentation controls.	5.3	More Details
CVE-2026-130659	ImageMagick before 7.1.2-35 contains a policy bypass vulnerability in various coders that allows attackers to allocate excessive memory by supplying images with many profiles. Attackers can craft images containing a high number of embedded profiles to exhaust memory and cause denial of service, bypassing security policy limits.	5.3	More Details
CVE-2026-132607	Insufficient verification of data authenticity in the decryption component of Amazon S3 Encryption Client for PHP in AWS SDK for PHP before 3.383.1 might allow an authenticated remote user to forge ciphertext that decrypts to arbitrary plaintext via manipulation of the GCM authentication tag length metadata. To remediate this issue, users should upgrade to AWS SDK for PHP version 3.383.1 or later.	5.3	More Details
CVE-2026-133255	In JetBrains Ktor before 3.6.1 a DoS attack was possible via the CIO engine	5.3	More Details
CVE-2026-131853	Ettercap before 0.8.5.2 contains a heap out-of-bounds read vulnerability in dissector_ssh() in src/dissectors/ec_ssh.c that copies a peer-supplied data_len without checking remaining packet bytes. A malicious SSHv1 client or server in an intercepted session can send a short packet with a large data_len to disclose up to 64 bytes of adjacent heap memory through captured credentials in the interface and log.	5.3	More Details
CVE-2026-61799	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. Prior to 0.0.23.Final, BinaryHttpParser in codec-bhttp/src/main/java/io/netty/incubator/codec/bhttp/BinaryHttpParser.java accumulates attacker-controlled Binary HTTP variable-length integers into the int offset sumBytes. BinaryHttpParser.readRequestHead, BinaryHttpParser.readFieldLine, and BinaryHttpParser.getIndeterminateLength can narrow long lengths into that offset, causing integer wraparound before buffer indices are calculated. A remote peer that supplies a valid large encoded length to BinaryHttpParser or BinaryHttpDecoder can trigger an unchecked bounds exception from a small malformed message, fail or close the channel, and repeatedly cause denial of service. Java and Netty bounds checks prevent memory corruption, and no information disclosure was observed. This issue is fixed in version 0.0.23.Final.	5.3	More Details
CVE-2026-133237	Manyfold 0.124.0 before 0.154.2 contains a server-side request forgery vulnerability that allows unauthenticated remote attackers to make the instance fetch internal URLs by supplying malicious ActivityPub actor icon.url and image.url fields. Attackers can deliver Create or Update activities to the inbox, causing the Shrine remote_url plugin to request internal or cloud metadata hosts and store responses as creator avatar or banner attachments.	5.3	More Details
CVE-2026-110728	The AIKTP plugin for WordPress is vulnerable to Authorization Bypass via Predictable Token in all versions up to, and including, 5.0.7. The vulnerability exists because the plugin's sync token (`aiktpz_token`) is generated using PHP's non-cryptographic `uniqid()` function — producing a wall-clock-derived value with approximately one million microsecond candidates per uncertain generation second — while the unauthenticated, unthrottled REST endpoint `/wp-json/aiktp/checkToken` (registered with `permission_callback => __return_true`) acts as a public brute-force oracle for the token; simultaneously, the capability checks in `check_valid_token` and its wrappers (`verify_aiktp_token_with_edit_capability`, `verify_aiktp_token_with_upload_capability`) inspect only a server-side stored Administrator account (option `aiktp_author`) rather than the actual REST caller, meaning possession of the token alone grants Administrator-equivalent authority to any external request. This makes it possible for unauthenticated attackers to recover the sync token and then invoke `POST /wp-json/aiktp/createpost` to publish arbitrary posts attributed to a site Administrator, `POST /wp-json/aiktp/doUpdatePostWithImageURL` to overwrite the body content of any existing post, and `POST /wp-json/aiktp/doUploadImageToWP` to import remote media and attach it as the featured image of any post. Additionally, when the `aiktpz_token` option has never been initialized, the `checkToken` oracle returns success for an empty token submission due to a loose comparison in `check_valid_token`, providing a zero-effort bypass that requires no brute-forcing whatsoever.	5.3	More Details

CVE-2026-130789	SemanticMediaWiki from 7.3.0 before 7.4.0 contains a missing authorization vulnerability in the smwtask API check-query task that allows unauthenticated attackers to bypass object-level authorization by relying on the wiki-wide edit right. Attackers can supply arbitrary subject titles, including protected pages, to run synchronous Semantic MediaWiki queries with caller-controlled conditions, limit, offset, and querymode parameters.	5.3	More Details
CVE-2026-103033	ALTCHA is a Go library for creating and verifying ALTCHA challenges. Prior to 2.1.0, parseServerSignature can leave ServerSignatureVerificationData.Extra uninitialized when url.ParseQuery returns partial values and an error for a malformed percent escape in VerificationData. A later write of an unknown parsed key to the nil map panics before VerifyServerSignature or VerifyServerSignatureSafe completes signature verification, so an unauthenticated client can abort a request. Standard net/http recovers the request goroutine, while application goroutines or frameworks without recovery can crash the process. This issue is fixed in version 2.1.0.	5.3	More Details
CVE-2026-129232	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, XMLCorpusView.read_block rebuilds the full root-to-node context path with "/".join(context) for every start tag while searching for a matching element. Deeply nested attacker-influenced corpus XML therefore causes quadratic pure-Python work and can pin a processing worker, resulting in denial of service. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-133275	A vulnerability in the Audio, Web, and Video Conferencing (AWV) component of Mitel MiCollab through 10.3 (10.3.0.18) could allow an unauthenticated attacker to conduct an arbitrary file upload attack due to insufficient file type validation. A successful exploit could allow an attacker to upload arbitrary files to the target system. This is MTLVULN-1861.	5.3	More Details
CVE-2026-108320	Tanium addressed an improper access controls vulnerability in Tanium Server.	5.3	More Details
CVE-2026-132136	A weakness has been identified in Directus up to 12.6.1. Affected by this issue is the function getSecret of the file api/src/services/shares.ts of the component Share Login Service. This manipulation causes incorrect authorization. It is possible to initiate the attack remotely. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure.	5.3	More Details
CVE-2026-131301	In Splunk App for Stream versions below 8.1.7, an unauthenticated user could cause Splunk App for Stream to reset its in-memory indexer cache. The vulnerability exists because the Splunk Web handler accepts the cache-reset request without requiring a login. For more information see Splunk Stream installation package overview (https://help.splunk.com/en/data-management/collect-stream-data/install-and-configure-splunk-stream/8.0/introduction/splunk-stream-installation-package-overview) and Launch Splunk Web (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.6/administer-splunk-enterprise-with-splunk-web/launch-splunk-web) in the Splunk documentation.	5.3	More Details
CVE-2026-121594	Mattermost Plugins versions <=12.-1 11.11.7 11.2.10 11.1.11.0 fail to enforce the GitLab project permission check for confidential issues when the project is publicly visible which allows an authenticated user with no GitLab membership in the project to obtain the contents of confidential issues, including title, description, author, and URL via subscribing a channel to the parent public GitLab group with the confidential_issues feature and waiting for a confidential issue webhook event. Mattermost Advisory ID: MMSA-2026-00738	5.3	More Details
CVE-2026-130661	ImageMagick before 7.1.2-35 contains a policy bypass vulnerability in the ASE decoder caused by a missing security policy check. Attackers can supply crafted ASE files that evade configured resource policies to cause excessive memory usage and denial of service.	5.3	More Details
CVE-2026-133414	A security flaw has been discovered in sofastack sofa-bolt up to 1.6.12. Affected by this issue is the function HessianSerializer.deserialize of the file src/main/java/com/alipay/remoting/serialization/HessianSerializer.java of the component Hessian2 Deserialization. Performing a manipulation results in allocation of resources. 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
CVE-2026-95690	Issue summary: Receiving crafted out-of-sequence DTLS handshake fragments may cause OpenSSL to allocate a full message-length reassembly buffer for each buffered message, so a small amount of wire data consumes a large amount of memory. Impact summary: A remote attacker can send a few small crafted DTLS handshake fragments to make a DTLS server or client allocate well over a megabyte of memory per connection, which may lead to memory exhaustion and Denial of Service if many connections are made. CWE: CWE-770: Allocation of Resources Without Limits or Throttling Description: DTLS handshake messages can arrive out of order and fragmented, so OpenSSL buffers fragments of up to ten out-of-sequence messages until reassembly is possible. When the first fragment of a new message arrives, OpenSSL allocates a buffer sized to the full message length declared in the fragment header, plus a reassembly bitmap, instead of storing only the bytes actually received. The declared length is capped by the maximum certificate list size, 128KB by default, and that cap applies to every handshake message type rather than only to Certificate messages. Only DTLS endpoints reach this code. TLS, QUIC and applications that do not use DTLS are unaffected. An attacker sends fragments that declare a large message length but carry almost no data, using up to ten distinct message sequence numbers. A single connection can then hold roughly 1.4MB of buffers	5.3	More Details

	from about 156 bytes on the wire. Because the number of buffered messages is capped, memory exhaustion requires many connections or many peers, which is why the severity is Low. FIPS impact: no The affected code is outside the FIPS module boundary.		
CVE-2021-48026	Zhongqing Naboo Education Cloud Platform builds through 2021-07-09 contain a missing authentication vulnerability in the /api/TeacherQuery/SearchTeacherInSiteWithPagerRecords endpoint that allows unauthenticated attackers to access sensitive teacher data. Attackers can send unauthenticated GET requests to this endpoint to retrieve user names, login names, and admin ids from the JSON response. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-11-15.	5.3	More Details
CVE-2026-132942	In JetBrains TeamCity before 2026.2, 2026.1.4 unauthenticated extraction of build parameter values via a timing attack on build status icons was possible	5.3	More Details
CVE-2026-129233	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. From 3.10.0 until 3.10.4, StreamBackedCorpusView performs os.stat on a caller-supplied plain string path during construction before pathsec containment is enforced. The PathPointer branch used by in-tree corpus readers is validated, so exploitation requires application code to pass an attacker-influenced string directly to the public constructor. That string can identify an out-of-root path and disclose its existence, size, and modification time even though opening the file would be rejected. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-130722	NocoDB before 2026.08.0 contains an authorization bypass in shared-view relation endpoints because parent-record lookups apply only the primary key, not the view's row filter. Anonymous share-link holders can enumerate auto-increment primary keys against the has-many, many-to-many, one-to-many and link-to-another-record routes to read linked records of filtered-out rows.	5.3	More Details
CVE-2026-133196	In JetBrains Junie CLI before 26.11.1 files outside the project could be overwritten via shell output redirection	5.3	More Details
CVE-2026-129040	Missing Authorization vulnerability in WP User Manager WP User Manager wp-user-manager allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP User Manager: from n/a through 2.10.4.	5.3	More Details
CVE-2026-129231	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, SeekableUnicodeStreamReader.readline appends each new block to a growing character buffer and reruns splitlines over the entire buffer on every iteration. An attacker-controlled corpus file containing one very long unterminated line therefore causes quadratic CPU consumption in line-oriented corpus readers and can deny service. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-133053	In JetBrains TeamCity before 2026.2.1 arbitrary file read was possible via symlinks in Mercurial subrepositories	5.3	More Details
CVE-2026-120271	InfoScan contains an unauthenticated information disclosure vulnerability in its mobile-enumeration and version-reporting endpoints. A remote attacker can send HTTP requests without authentication to retrieve system configuration information, including usernames and display names, multifactor authentication configuration status, session settings, application version information, server details, hardware module firmware versions, and the BACnet/IP port used by the installation.	5.3	More Details
CVE-2026-133171	Tiptap @tiptap/core 3.7.0 before 3.30.5 contains a regular expression denial of service vulnerability in parseAttributes and parseShortcodeAttributes that allows attackers to exhaust CPU through quadratic backtracking. Attackers can supply Markdown with long crafted attribute strings parsed by createBlockMarkdownSpec, createAtomBlockMarkdownSpec or createInlineMarkdownSpec to stall browser pages or server processes.	5.3	More Details
CVE-2026-129069	Missing Authorization vulnerability in John-Michael L'Allier Create mediavine-create allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Create: from n/a through 3.1.3.	5.3	More Details
CVE-2026-129227	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, Tree.fromlist recursively converts nested Python list input into Tree objects without enforcing a maximum tree depth. An attacker-controlled list nested beyond the interpreter recursion limit can raise RecursionError and terminate the current parse operation or request worker. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-129226	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, the semantic-logic read_type parser recursively processes both halves of nested complex type strings without applying LogicParser.MAX_PARSE_DEPTH. An attacker-controlled type string or signature with excessive nesting can	5.3	More Details

	raise RecursionError through Type.fromstring or type-checking entry points and abort the current worker. This issue is fixed in version 3.10.4.		
CVE-2026-129225	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. In 3.10.3 and earlier, DependencyGraph._tree and triples recurse once per dependency edge without a depth bound, and get_cycle_path also lacks a visited-set termination guard. A deeply chained or cyclic attacker-controlled CoNLL or Malt-TAB dependency structure can raise RecursionError and abort a parse worker. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-130720	NocoDB before 2026.08.0 contains an insecure direct object reference vulnerability that allows unauthenticated attackers to access records hidden by a shared view's row filter by supplying arbitrary primary keys. Attackers with the share link can request attachment downloads and many-to-many or has-many link listings for excluded records to obtain their attachment files and linked record identities.	5.3	More Details
CVE-2026-130703	NocoDB before 2026.10.0 contains an information disclosure vulnerability that allows unauthenticated attackers to read hidden columns by supplying comma-separated group-by column lists. Attackers with a shared view link can pair a visible and hidden column to bypass restrictSharedViewColumnReferences and retrieve hidden values tied to individual rows.	5.3	More Details
CVE-2026-120993	When NGINX Plus or NGINX Open Source is configured to use the ngx_stream_ssl_module module with dynamic certificate loading and the ngx_stream_proxy_module module with the proxy_protocol v2 directive, an abbreviated SSL handshake can cause a worker process to terminate. Impact: This vulnerability may allow remote attackers to cause a denial-of-service (DoS) on the NGINX system. This is a data plane vulnerability only; there is no control plane exposure. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	5.3	More Details
CVE-2026-120298	InfoScan fails to properly handle certain malformed or incomplete requests to its login and password-reset endpoints. These requests can trigger unhandled server exceptions that return detailed Node.js stack traces to the requester. The responses disclose internal filesystem paths, application directory names, source file locations, and line numbers.	5.3	More Details
CVE-2026-130811	@fastify/static is the Fastify plugin for serving static files. In versions up to and including 10.1.11, when it is configured to deny or ignore dotfiles together with directory listing enabled, a request for a denied dot-directory returns its directory listing instead of being refused. The underlying send step answers with a forbidden response, but the listing fallback misreads that response on a path ending in a slash as a directory without an index file and serves the listing anyway. An unauthenticated client can enumerate the names of files and subdirectories, recursively, inside dot-directories the operator explicitly denied, such as version control, backup, and key file names. File contents remain protected and only names are disclosed. The issue is fixed in @fastify/static 10.1.12, and users should upgrade to 10.1.12 or later. As a workaround, disable directory listing or do not serve a root that contains dot-directories.	5.3	More Details
CVE-2026-130705	NocoDB before 2026.10.1 contains an information disclosure vulnerability that allows unauthenticated shared view link holders to read hidden columns by grouping on physical column names. Attackers can supply a hidden column's physical name or the unchecked sub-group parameter to the public group-by endpoint to reveal hidden values, counts, and column existence.	5.3	More Details
CVE-2026-115515	@fastify/static is the Fastify plugin for serving static files. In versions before 10.1.8, the allowedPath guard receives only the requested pathname, but the underlying send step selects the file to serve afterwards, so a request for a directory resolves to its default index file and an extensionless request resolves through the configured extensions to a matching file. The selected file is streamed without being re-checked against allowedPath, so a guard that denies a specific file is bypassed by requesting its directory or extensionless form. An unauthenticated client can read files inside the static root that the application meant to keep private whenever the denied file is reachable as a directory index or an extension match. The issue is fixed in @fastify/static 10.1.8, and users should upgrade to 10.1.8 or later. As a workaround, also deny the directory and extensionless forms in allowedPath, or disable the index and extensions options.	5.3	More Details
CVE-2026-125283	When NGINX Plus or NGINX Open Source is configured to use the try_files and alias directives with a map variable, certain combinations of regular expression captures in map and try_files can cause a segmentation fault in an NGINX worker process. Impact: An attacker could leverage this vulnerability to cause an NGINX worker process to terminate unexpectedly. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	5.3	More Details
CVE-2026-132986	In JetBrains TeamCity before 2026.2 password parameters containing a caret were not masked in build logs on Windows	5.3	More Details
CVE-2026-123640	Description When NGINX Plus or NGINX Open Source is configured with the ngx_stream_ssl_preread_module module, a specially crafted ClientHello message from an unauthenticated attacker causes forged values of the ngx_stream_ssl_preread_module variables, affecting virtual server selection. In addition, it may also cause a disagreement with a proxied server terminating the SSL/TLS connection without Transport Layer Security (TLS) 1.3 support. Impact This vulnerability may allow an attacker to forge values of the ngx_stream_ssl_preread_module variables and bypass the NGINX systems stream routing segmentation. As a	5.3	More Details

	result, the attacker may be able to route the connection to a privileged or internal backend that should be unreachable and gain access to sensitive information. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.		
CVE-2026-130161	Tamari is a fully-featured recipe manager web application built using Python and the Flask Framework. Prior to 1.7.2, an authenticated user can bypass the recipe-autofill SSRF check by submitting a hostname under the user's control through the Autofill from URL feature; the hostname can resolve to a public address during validation and a private or internal address during the later connection. In app/myrecipes/routes.py, is_safe_autofill_url validates the first DNS response, but recipe autofill passes the original URL to requests.get, which independently resolves the hostname and can connect to the internal address. The fetched response is parsed into recipe fields, enabling internal-service reconnaissance and potential data exposure; no redirect is required, and disabling redirects does not stop this DNS-rebinding bypass. This issue is fixed in 1.7.2.	5.3	More Details
CVE-2026-131189	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, an unauthenticated user could inject arbitrary fields into Splunk Enterprise audit records via the Representational State Transfer (REST) API. Successful exploitation could allow the unauthenticated user to forge audit evidence, compromising the integrity of audit reviews and security detections that rely on legacy audit-field parsing. The vulnerability exists because Splunk Enterprise does not neutralize field delimiters in user-supplied data before writing the data to a legacy audit record. The affected workflow also does not create a corresponding structured audit record. For more information see Auditing activities in a Splunk platform instance (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.0/audit-activity-in-splunk-enterprise/auditing-activities-in-a-splunk-platform-instance) and About structured audit trail logs (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/audit-activity-in-splunk-enterprise/about-structured-audit-trail-logs) in the Splunk documentation.	5.3	More Details
CVE-2026-129228	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, ToolboxSettings.parse builds a deeply nested attacker-controlled settings tree iteratively, but the later public helpers to_settings_string, remove_blanks, add_default_fields, sort_fields, and add_blank_lines recursively traverse that tree without a depth bound. Invoking any of those helpers on input nested beyond the interpreter recursion limit can raise RecursionError and abort the operation or request worker. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-120157	The publicly distributed InfoScan Electron installer contains configuration files identifying hostnames associated with customer installations and internal development, testing, and database infrastructure. An individual who obtains and inspects the installer can extract these hostnames without authenticating to the affected systems.	5.3	More Details
CVE-2026-129224	NLTK (Natural Language Toolkit) is a suite of open source Python modules, data sets, and tutorials supporting research and development in Natural Language Processing. Prior to 3.10.4, the CCG lexicon parser's matchBrackets and augParseCategory functions recursively process nested parenthesized categories without a depth bound. An attacker-controlled lexicon string with excessive nesting can raise RecursionError through nltk.ccg.lexicon.fromstring and terminate the parse operation or request worker. This issue is fixed in version 3.10.4.	5.3	More Details
CVE-2026-107099	The Royal Addons for Elementor WordPress plugin before 1.7.1071 does not restrict or validate the post meta keys used when counting or rendering a filtered grid, and does not reject private (hidden) meta keys, allowing unauthenticated users to disclose the value of arbitrary private post meta one value at a time from the filtered result count. This is an incomplete fix of CVE-2026-17585, whose 1.7.1067 patch covered only a separate search-based path and left the grid filtering path affected.	5.3	More Details
CVE-2026-77377	Bolt is a fully open source, adaptable platform for building and running modern websites. Prior to 6.1.4, the authenticated URL-upload endpoint and authenticated embed endpoint fetch user-supplied URLs server-side without protection against internal targets, allowing an attacker with the applicable backend permission to cause the Bolt server to request those targets. This issue is fixed in version 6.1.4.	5.2	More Details
CVE-2026-77380	Bolt is a fully open source, adaptable platform for building and running modern websites. Prior to 6.1.5, the /bolt/async/upload-url and /bolt/async/embed endpoints validate only a submitted URL before fetching it server-side. The fetchers then follow redirects without revalidating each destination, allowing an authenticated user to bypass UrlSafetyChecker and cause the server to request internal or cloud-metadata addresses. This issue is fixed in version 6.1.5.	5.2	More Details
CVE-2026-131230	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user could trick a user that holds a role with the ueba_can_view_ui capability into opening a link that the unauthenticated user creates for the User and Entity Behavior Analytics (UEBA) Overview dashboard. When the user opens the link, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who opens the link. The commands could expose or modify some data available to the user who opens the link and partially disrupt search processing. The vulnerability is possible because the User and Entity Behavior Analytics Overview dashboard does not escape dashboard filter values from the link before substituting the values into SPL searches. The vulnerability requires the attacker to phish the user by tricking them into opening a link that the attacker creates. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see UEBA dashboards	5.0	More Details

	(https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.8/analytics/ueba-dashboards) in the Splunk documentation.		
CVE-2026-133197	In JetBrains Junie CLI before 26.9.22 command injection was possible via file names in code review	5.0	More Details
CVE-2026-130777	Incus before 7.6.0 contains an information disclosure vulnerability that allows project-restricted authenticated clients to receive other projects' OVN ACL logs because these events are broadcast with an empty project. Attackers can subscribe to network-acl events on one accessible project to observe internal IPs, ports, protocols, and allow/drop decisions across all projects.	5.0	More Details
CVE-2026-133520	Dokploy 0.28.0 before 0.31.1 contains a path traversal vulnerability in the patch.readRepoDirectories tRPC procedure that allows authenticated organization members to list tracked files of arbitrary git repositories. Attackers with service read permission can supply a crafted repoPath, which is passed to git ls-tree unconfined, to enumerate other organizations' application and compose source trees.	5.0	More Details
CVE-2026-121593	Mattermost Plugins versions <=12.-1 11.0.11 11.11.7 11.2.10.0 fail to verify that the caller can post in a run's channel when executing a checklist slash command which allows an authenticated run participant to post arbitrary messages into a private channel they are not a member of via the checklist item run endpoint. Mattermost Advisory ID: MMSA-2026-00730	5.0	More Details
CVE-2026-102945	Aether Project AMF supports termination of NAS signalling, NAS ciphering & integrity protection, registration management, connection management, mobility management, access authentication and authorization, security context management. Prior to 3.1.6, a malformed NGAP PathSwitchRequest can contain a UESecurityCapabilities information element whose extensible APER bit-string representation has no bytes even though decoding reports no error. In ngap/handler.go, HandlePathSwitchRequest checks only whether the element is present before six index-zero reads of its encryption and integrity algorithm fields, causing an unrecovered panic in the NGAP receive goroutine and terminating the AMF. The described attacker is a rogue or compromised gNB on the N2 interface and additionally needs a registered UE's AMF-UE-NGAP-ID; a bare over-the-air UE cannot trigger the issue. This issue is fixed in version 3.1.6.	4.9	More Details
CVE-2026-131856	Pika before 1.5.0 contains an uncontrolled resource allocation vulnerability that allows malicious brokers to exhaust client memory by declaring oversized frame payloads. Attackers can send frame headers declaring up to 4 GiB payloads, unchecked against frame_max, or very large message bodies to terminate consuming applications.	4.9	More Details
CVE-2026-109235	The WP Maps – Google Maps, OpenStreetMap, Mapbox, Store Locator with Search, Filters & Listings plugin for WordPress is vulnerable to generic SQL Injection via the 'csv_columns' parameter in all versions up to, and including, 5.0.2 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-132985	In JetBrains TeamCity before 2026.2 project administrators could read OAuth connection secrets of ancestor projects	4.9	More Details
CVE-2026-120672	The Meow Gallery plugin for WordPress is vulnerable to generic SQL Injection via the 'limit' parameter in all versions up to, and including, 5.5.9 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-133102	In JetBrains TeamCity before 2026.2.2, 2026.1.6 out-of-band data exfiltration via DNS resolution was possible in the Kotlin DSL sandbox	4.9	More Details
CVE-2026-55505	Wavelog is web-based amateur radio logging software. Prior to 3.2.4, the application/controllers/Lotw.php import() workflow passes the lotw_download_url configuration value to curl without scheme validation, accepts file:// content, and writes the response to a predictable path under ./uploads/ that can be retrieved through an unauthenticated HTTP request. A standard user can trigger the LotW import, although changing lotw_download_url normally requires installer administrator access, and related vulnerabilities can remove that prerequisite. The same missing status-validation pattern affects the LotW lotw_users fetch helper and per-user download helper. This exposes files readable by the www-data user, including environment data, mounted secrets, and application configuration. This issue is fixed in version 3.2.4.	4.9	More Details
CVE-2026-132927	In JetBrains TeamCity before 2026.2, 2026.1.4 the Perforce Administrator Access password could be written to the server log	4.9	More Details
CVE-	The Flatsome WordPress theme before 3.20.13 does not verify that the acting user is authorised to modify		

2026-117491	the target post when saving a template, allowing users with a role as low as Contributor to overwrite the content and type of an arbitrary post, including another user's published page, removing it from the front end.	4.9	More Details
CVE-2026-95791	A vulnerability in the web management and file transfer component of AOS web could allow an authenticated remote attacker to perform arbitrary file reads. Successful exploitation could allow an attacker to access sensitive files on the underlying system, potentially exposing configuration data or credentials.	4.9	More Details
CVE-2026-132133	@fastify/middie is the middleware engine that adds Express and Connect style middleware support to Fastify. In versions up to and including 9.3.6, middie builds the match regular expression for each path-scoped middleware from its mount path without collapsing duplicate slashes, even when the application enables the ignoreDuplicateSlashes router option. Under that option Fastify's router collapses duplicate slashes in both the registered route and the request path, so a mount path containing a doubled slash is served by the router at its single-slash form while the middleware's compiled match still requires the doubled slash and never matches. A path-scoped middleware such as authentication, authorization, or rate limiting is therefore skipped while the router dispatches the request to the protected handler, without any authentication required. Applications are affected when they enable ignoreDuplicateSlashes and a middleware mount path reaches middie with a duplicate slash, most commonly when a plugin prefix that ends with a slash is joined to a mount path that starts with one. The issue is fixed in @fastify/middie 9.3.7, and users should upgrade to 9.3.7 or later. As a workaround, enforce authorization in a Fastify hook such as preHandler that runs after routing, and avoid mount paths and plugin prefixes that introduce duplicate slashes.	4.8	More Details
CVE-2026-131214	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the edit_timeline capability could store attacker-controlled Search Processing Language (SPL) in a legacy investigation timeline entry. Another user that holds a role with the manage_all_investigations capability could later open the investigation and select Run Search from the timeline. Splunk Enterprise Security could then run the SPL stored by the first user using the permissions of the user who selected Run Search. The SPL stored by the first user could access data or perform actions available to the user who selected Run Search. The vulnerability is possible because Splunk Enterprise Security does not treat stored investigation search content as a literal search before building the Search page link. For more information see Capability reference for Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/install/8.3/installation/capability-reference-for-splunk-enterprise-security) and Add details to an investigation in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-7/user-guide/7.3/investigations/add-details-to-an-investigation-in-splunk-enterprise-security) in the Splunk documentation.	4.8	More Details
CVE-2026-132955	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the Pull Requests branch filter health report was possible	4.8	More Details
CVE-2026-132249	In JetBrains TeamCity before 2026.1.3, 2025.11.7 stored XSS in the Shared Resources build feature was possible	4.8	More Details
CVE-2026-76541	In Splunk Enterprise versions below 10.6.0, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who holds the "power" Splunk role could store attacker-controlled Search Processing Language (SPL) in a shared data model calculation. When a user who can run a search against the shared data model runs one, the injected SPL can run using the permissions of the user who ran the search. This could expose sensitive credentials to the user who created the shared data model. The vulnerability is possible because Splunk Enterprise does not safely neutralize data model calculation fields before it includes them in generated searches. For more information see About data models (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.4/build-a-data-model/about-data-models) in the Splunk documentation. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.	4.8	More Details
CVE-2026-131235	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in data used by the Credit Card Data Found dashboard. If another user who can access the dashboard selects an affected row, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selects the row. The user who stores the values could then access sensitive data available to the user who selects the row. The vulnerability is possible because the Credit Card Data Found dashboard does not validate source and destination values before inserting them into drilldown searches. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) and Credit Card Data Found (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/credit-card-data-found) in the Splunk documentation.	4.8	More Details
CVE-2026-	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in SSH key management was possible	4.8	More

132923			Details
CVE-2026-130790	SemanticMediaWiki from 2.5.0 before 7.4.0 contains a stored cross-site scripting vulnerability that allows users with the smw-patternedit right to inject script by seeding markup reference names in MediaWiki:Smw_allows_pattern. Once an editor annotates the matching Allows pattern value, AllowsPatternValue emits it unescaped on Special:Browse, executing JavaScript in every viewer's session.	4.8	More Details
CVE-2026-132124	modified eCommerce Shopsoftware through 3.3.0 contains a stored cross-site scripting vulnerability that allows authenticated administrators to inject scripts via the orders_status_name parameter in admin/orders_status.php. Attackers with per-page order status access can store payloads that execute in other administrators' browsers when the unencoded status name is displayed.	4.8	More Details
CVE-2026-133169	modified eCommerce Shopsoftware through 3.3.0 contains a stored cross-site scripting vulnerability that allows authenticated administrators to inject scripts via the unsanitized products_model field. Attackers with product editing permission can store JavaScript that executes in browsers of other administrators, storefront visitors, and customers viewing carts and order history.	4.8	More Details
CVE-2026-132965	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the unknown VCS root page was possible	4.8	More Details
CVE-2026-132972	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Slack Notifier settings page was possible	4.8	More Details
CVE-2026-128660	PHPGurukul Beauty Parlour Management System V 1.1 is vulnerable to Cross Site Scripting (XSS) in admin/search-invoices.php via the searchdata parameter.	4.8	More Details
CVE-2026-132123	modified eCommerce Shopsoftware through 3.3.0 contains a stored cross-site scripting vulnerability that allows authenticated administrators to inject scripts via the cross_sell_group_name parameter in admin/cross_sell_groups.php. Attackers with per-page access to cross-selling groups can store a script payload that executes in other administrators' browsers, including full administrators, when they view the group listing.	4.8	More Details
CVE-2026-132981	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Versioned Settings page was possible	4.8	More Details
CVE-2026-14885	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 could allow a remote attacker to obtain sensitive information due to improper access control.	4.8	More Details
CVE-2026-133258	In JetBrains YouTrack before 2026.2.19840 app widgets could perform API requests on behalf of viewing users without their consent	4.8	More Details
CVE-2026-120967	When NGINX Plus or NGINX Open Source is configured to require client certificates ("ssl_verify_client on") in a non-default stream virtual server, a remote attacker could bypass verification of client certificates. Impact: An attacker could leverage this vulnerability to bypass verification of client certificates. There is no control plane exposure; this is a data plane issue only. Note: Software versions which have reached End of Technical Support (EoTS) are not evaluated.	4.8	More Details
CVE-2026-133248	modified eCommerce Shopsoftware through 3.3.0 contains a stored cross-site scripting vulnerability that allows administrators with product editing permission to inject script via the products_name field. The unescaped name is rendered in the h1 heading of product detail and reviews pages, executing script in the browsers of shoppers and logged-in customers.	4.8	More Details
CVE-2026-132987	In JetBrains TeamCity before 2026.2 server certificates were not verified in the SVN integration	4.8	More Details
CVE-2026-132953	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the queued build Overview tab was possible	4.8	More Details
CVE-2026-104357	A flaw was found in libtirpc. The SunRPC UDP/datagram server transport reuses a single receive buffer across requests and resets the XDR decode position without rebinding the decode length to the size of the datagram actually received. As a result, when a short UDP RPC request is received after a longer one, the XDR decoder can read stale bytes left in the reused buffer and interpret them as part of the current request, including credential, verifier, or procedure-argument fields. The highest threat from this vulnerability is to data confidentiality and integrity for services that use a UDP-based SunRPC server.	4.8	More Details

CVE-2026-107511	Dell PowerStore contains a CSV Injection vulnerability. A high-privileged attacker could exploit this vulnerability, potentially leading to code execution on another user's workstation.	4.8	More Details
CVE-2026-18553	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to sanitize terminal escape sequences when rendering server logs via {{mmctl logs}}, which allows an unauthenticated attacker to hijack a system administrator's clipboard or spoof terminal output via a crafted HTTP request whose logged content is later viewed with {{mmctl logs}}. Mattermost Advisory ID: MMSA-2026-00753	4.8	More Details
CVE-2026-133170	Tiptap @tiptap/core 2.0.0-alpha.0 before 3.30.4 contains a cross-site scripting vulnerability that allows attackers to inject DOM attributes via a __proto__ key passed to mergeAttributes(). Attackers can supply untrusted JSON attribute data whose __proto__ replaces the merged object's prototype, causing DOMSerializer.renderSpec to render injected event handlers that execute JavaScript.	4.7	More Details
CVE-2026-112163	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Kayaport Ödeme Kuruluşu A.Ş. BulutTahsilat – Virtual POS Infrastructure allows Cross-Site Scripting (XSS). This issue affects BulutTahsilat – Virtual POS Infrastructure: through 2026-11-05.	4.7	More Details
CVE-2026-129936	A MongoDB C# Driver component that interprets connection configuration supplied by the hosting application can include the caller-supplied configuration value in the text of the error it raises when that value cannot be interpreted. The redaction applied to that text does not cover all forms in which secret material may appear, so authentication material embedded in the configuration value can be preserved in the resulting error text and in anything that records it. Exploitation requires only the low privilege needed to observe the application's own error output or logs on the system where the driver runs; no database credentials or server connectivity are required. The result is disclosure of sensitive information to parties whose access to error output is broader than their intended access to the configuration secret.	4.7	More Details
CVE-2026-132950	In JetBrains TeamCity before 2026.2, 2026.1.4 stored XSS in the Xcode runner SDK selection was possible	4.7	More Details
CVE-2026-133479	A vulnerability was detected in WebKul QloApps up to 1.8.0. Affected by this vulnerability is the function postProcessDownload/installTheme of the file controllers/admin/AdminModulesController.php of the component Archive Import. The manipulation results in path traversal. It is possible to launch 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.	4.7	More Details
CVE-2026-129434	The OneLogin SAML SSO WordPress plugin before 3.7.1 does not validate a user-supplied parameter before using it to redirect an already-authenticated user, allowing unauthenticated attackers to redirect a logged-in user to an arbitrary external domain via a crafted link (e.g. for phishing).	4.7	More Details
CVE-2026-129931	A component of the MongoDB Java Driver that processes connection configuration places credential values taken from that configuration into generated text and error output without replacing them with a placeholder, unlike other parts of the driver. A party able to read output produced by an application that embeds the driver may obtain database or proxy credentials in clear text.	4.7	More Details
CVE-2026-129934	The MongoDB Go driver's connection string parser includes the raw, undecoded value of a query option in the error it returns when that value contains invalid percent encodings. When the option holds a secret, such as a TLS key file password or an AWS session token, the secret appears in plain text in the error. A party with low-privileged read access to wherever the application records or displays that error may learn the secret.	4.7	More Details
CVE-2026-115814	The Kirki – Freeform Page Builder, Website Builder & Customizer plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'comment (title attribute of <abbr>)' parameter in all versions up to, and including, 6.4.1 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. Successful exploitation requires an administrator to run the Kirki Convert (migration) workflow on a page containing the attacker's approved comment, as the browser-side DOM parser decodes the kses-allowed abbr title attribute value into a raw XSS payload that is then persisted without sanitization and later rendered unescaped. Successful exploitation additionally requires an action by an administrator — the attacker cannot trigger the vulnerability on their own.	4.7	More Details
CVE-2026-131842	Poppler pdfdetach 24.03.0 through 26.11.0 contains a path traversal vulnerability in utils/pdfdetach.cc that allows attackers to write files outside the output directory by bypassing a string-prefix containment check. Attackers can craft PDFs with embedded file names containing '..' sequences or absolute paths, so -saveall or -savefile extraction creates or overwrites files with the invoking user's privileges.	4.7	More Details
CVE-2026-76534	In Splunk Enterprise for Windows versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user that holds a role with the edit_user capability could create multiple user accounts that share the same user directory. This could expose user-specific configuration between those accounts, let one account modify configuration used by another account, or disrupt features that depend on separate user directories. The vulnerability is caused by improper validation of usernames before Splunk Enterprise performs Windows filesystem operations. Windows removes trailing periods from directory names, so distinct usernames can resolve to the same user	4.7	More Details

	directory. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.		
CVE-2026-132122	agentic-radar 0.7.0 before 0.14.2 contains an untrusted configuration loading vulnerability that allows malicious repositories to inject environment variables by having load_dotenv() in agentic_radar/cli.py load the scanned repository's .env file. Attackers can set OPENAI_BASE_URL to an attacker-controlled host so that scan openai-agents or --harden-prompts send requests carrying the user's OpenAI API key to the attacker.	4.7	More Details
CVE-2026-130153	The Okta Client SDK for Python (okta-client-python) includes client authentication values (client_secret or JWT assertions) as query parameters in the front-channel authorization URL when executing the authorization code flow without Pushed Authorization Requests (PAR). When PAR is disabled, unsupported by the authorization server, or fails, the SDK falls back to standard authorization URL construction, transmitting these parameters via user-agent requests.	4.7	More Details
CVE-2026-133241	diaspora* before 0.10.0 contains a cross-site scripting vulnerability that allows unauthenticated attackers to inject iframe attributes via the maxheight parameter of the GET /oembed endpoint. Attackers can supply a single quote in maxheight to inject an onload handler into the returned html, executing script in third-party oEmbed consumers that embed it unsanitized.	4.7	More Details
CVE-2026-130765	Obot before 0.27.0 contains an open redirect vulnerability that allows unauthenticated attackers to bypass same-site validation of the rd parameter by inserting a tab character. Attackers can craft sign-in or sign-out links on the Obot domain that browsers resolve to external hosts, redirecting victims to phishing pages.	4.7	More Details
CVE-2026-129930	A weakness in the MongoDB Node.js driver can place sensitive information taken from an application-supplied connection string into the text of an error produced while that connection string is being parsed. The error is generated inside the application process before any database connection is made, so no database account and no network access are needed for it to occur. A party able to read the application's error output or log files may thereby obtain database credentials that the application's configuration source would otherwise keep to a narrower audience.	4.7	More Details
CVE-2026-131217	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in data that populates the System Misconfigurations report. Another user with access to the report could later select a related row. The System Misconfigurations report could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selected the row. The commands could access or modify data available to the user who selected the row. The vulnerability is possible because the System Misconfigurations report does not treat values from the Vulnerabilities data model as literal search values when it builds drilldown searches. For more information see System Misconfigurations (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/system-misconfigurations) and Configure users and roles (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-the-app/configure-users-and-roles) in the Splunk documentation.	4.6	More Details
CVE-2026-131221	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in intrusion detection data. Another user who holds the "pci_user" role could later select the corresponding row in the Intrusion Detection and Prevention System Alert Activity Details panel. The report could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who selected the row. The commands could access data and perform actions available to the user who selected the row. The vulnerability is possible because the Intrusion Detection and Prevention System Alert Activity report does not treat intrusion detection values as literal search values when it builds the details drilldown search. For more information see Reports (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/user-manual/8.7/using-the-splunk-app-for-pci-compliance/reports) and Configure users and roles (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-the-app/configure-users-and-roles) in the Splunk documentation.	4.6	More Details
CVE-2026-131212	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in data that populates the Malware Activity report. Another user with access to the report could later click a related row. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who clicked the row. The commands could access data and perform actions available to that user. The vulnerability is possible because the Malware Activity report does not treat values from the Malware data model as literal search values when it builds drilldown searches. For more information see Malware Activity (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/malware-activity) in the Splunk documentation.	4.6	More Details
CVE-2026-133064	A flaw was found in Moodle. Insufficient sanitization of objective identifiers in the SCORM (Sharable Content Object Reference Model) objectives report allows an authenticated user with content upload permissions to inject malicious scripts. When another user views the report, the injected script executes in their browser session, potentially allowing the attacker to access sensitive information or perform unauthorized actions on	4.6	More Details

	their behalf.		
CVE-2026-130687	NocoDB before 2026.10.1 contains a stored CSS injection vulnerability that allows authenticated Commenter or Editor users to inject inline styles by placing single quotes in mention text. Attackers can break out of the single-quoted data-id attribute in comments or rich text cells to overlay or rearrange the interface for other viewers.	4.6	More Details
CVE-2026-131194	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a low-privileged user who holds the "power" Splunk role could inject malicious content into a Splunk Web view. When another authenticated user opens the affected view in Splunk Web, the injected content executes unauthorized JavaScript code within the viewing user's browser session. The vulnerability is possible because Splunk Web does not properly validate content in a view definition before rendering the view. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.	4.6	More Details
CVE-2026-113458	Mattermost versions 11.10.x <= 11.10.2, 11.11.x <= 11.11.1, 11.7.x <= 11.7.12 fail to validate the format of configured database connection strings before redacting them, which allows a party with access to a generated Support Packet or server log files to recover plaintext database credentials when the connection string is configured in a non-URL format.. Mattermost Advisory ID: MMSA-2026-00783	4.5	More Details
CVE-2026-130763	Fleet before 4.92.0 contains a CSV formula injection vulnerability in the host inventory export endpoint that allows attackers controlling an enrolled host or enroll secret to plant spreadsheet formulas. Attackers can set osquery-reported fields like hostname or hardware serial to formula values, enabling data exfiltration or DDE command execution when users open the exported CSV.	4.4	More Details
CVE-2026-131845	Poppler 21.01.0 through 26.11.0 contains a time-of-check time-of-use race condition in FormWidgetSignature::signDocument() that allows local attackers to obtain valid signatures over altered content. Attackers with write access to the output file or directory can overwrite or swap the PDF between writing and hashing, causing the signer's key to sign attacker-controlled content.	4.4	More Details
CVE-2026-111517	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain a Deserialization of Untrusted Data vulnerability. A low privileged attacker with local access could potentially exploit this vulnerability, leading to Arbitrary Code Execution.	4.4	More Details
CVE-2026-130667	ImageMagick before 7.1.2-35 contains a heap buffer over-read vulnerability in the TIFF decoder that allows attackers to corrupt heap memory by supplying a crafted TIFF image. Attackers can provide a malicious TIFF file that, when encoded in a specific output format, triggers a heap buffer over-write causing memory corruption.	4.4	More Details
CVE-2026-132961	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS on the Test Details page was possible	4.4	More Details
CVE-2026-130657	ImageMagick before 7.1.2-35 contains a policy bypass vulnerability on Windows that allows attackers to evade path-based security policies by mixing forward slash and backslash separators. Attackers can supply crafted image file paths using alternate separators to read or write image files that the configured policy is intended to block.	4.4	More Details
CVE-2026-120541	The Product Options and Price Calculation Formulas for WooCommerce – Uni CPO plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'data' parameter in all versions up to, and including, 4.9.69 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with shop manager-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	4.4	More Details
CVE-2026-131298	In Splunk SOAR versions below 8.8.0, an authenticated user with edit permission on a workbook could configure a service-level agreement at only the workbook phase or task level. Splunk SOAR then flags the task as overdue immediately and does not display the configured deadline in the work queue. The incorrect deadline status could obscure genuine deadline breaches. The vulnerability is possible because Splunk SOAR treats an absent phase-level or task-level service-level agreement as immediately due when it compares the configured due times. For more information see Define tasks using workbooks (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/7.0.0/configure-product-settings-for-your-splunk-soar-on-premises-instance/define-tasks-using-workbooks) in the Splunk documentation.	4.3	More Details
CVE-2026-133075	A flaw was found in Moodle. Due to insufficient authorization checks in the quiz random question filter web service, an authenticated user with teacher privileges can modify random question settings in quizzes they do not have permission to manage. This flaw allows unauthorized users to alter assessment configurations across courses they do not oversee.	4.3	More Details
	In Splunk SOAR versions below 8.8.0, a user with edit permission for containers on only a subset of container labels could modify container pins outside the container labels that the role permits. The vulnerability is		

CVE-2026-131241	possible because the Container Pin Representational State Transfer (REST) API does not enforce the container label authorization boundary when handling requests. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.4.1/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	4.3	More Details
CVE-2026-133073	A flaw was found in Moodle's single view grade report web service. A missing authorization check allows authenticated students to view the names of all grade items in a course, including those configured to be hidden. This vulnerability leads to unauthorized information disclosure.	4.3	More Details
CVE-2026-131247	In Splunk SOAR versions below 8.8.0, an authenticated Splunk SOAR user with permission to edit and run playbooks could read action status and callback state from a playbook run outside the authorized container-label or tenant scope for that user. The vulnerability is possible because the playbook API uses mutable runtime context to identify a playbook run instead of preserving the server-provided context.	4.3	More Details
CVE-2026-133074	A flaw was found in Moodle. Insufficient authorization checks in the quiz overrides web service allow an authenticated user with teacher privileges to modify user and group overrides for quizzes they do not have permission to manage. An attacker could exploit this vulnerability to improperly alter exam access, deadlines, or testing conditions across unauthorized courses.	4.3	More Details
CVE-2026-131283	In Splunk SOAR versions below 8.8.0, an authenticated user with playbook edit permission could cause later playbook conversions to include spurious branch conditions. This can prevent downstream blocks from running in affected converted playbooks. The vulnerability is possible because the playbook converter retains branch conditions from previous conversions.	4.3	More Details
CVE-2026-131251	In Splunk SOAR versions below 8.8.0, a user with permission to view and modify containers with an allowed label could add artifacts to containers that the user cannot view or modify. The vulnerability is possible because the Representational State Transfer (REST) API artifact-creation workflow does not preserve the authorization boundary for the selected container.	4.3	More Details
CVE-2026-76532	In Splunk Enterprise versions below 10.6.1, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who does not hold the "admin" or "power" Splunk roles could view credentials configured for email notifications in a saved search shared with that user. The vulnerability is possible when email notifications use authentication credentials because the Saved Searches Representational State Transfer (REST) API returns those credentials in clear text. For more information see Search endpoint descriptions (https://help.splunk.com/en/splunk-enterprise/leverage-rest-apis/rest-api-reference/10.6/search-endpoints/search-endpoint-descriptions), Manage knowledge object permissions (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-manual/10.6/get-started-with-knowledge-objects/manage-knowledge-object-permissions), and Email notification action (https://help.splunk.com/en/splunk-enterprise/alert-and-respond/alerting-manual/10.6/configure-alert-actions/email-notification-action) in the Splunk documentation.	4.3	More Details
CVE-2026-131276	In Splunk SOAR versions below 8.8.0, an authenticated user that holds a role with the System Settings view permission but not the Assets view permission could view asset and tenant metadata in the ingestion history view. The ingestion history view does not restrict summary results to assets and tenants that the user can view.	4.3	More Details
CVE-2026-131284	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to view and edit containers could import a container archive that makes artifacts appear to be part of a case or have notes when no corresponding case associations or notes exist. This could cause Representational State Transfer (REST) API clients and filters to use incorrect artifact state. The vulnerability is possible because container archive import accepts artifact workflow state from the archive instead of deriving that state from associated cases and notes. For more information see Splunk SOAR REST API reference (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference) in the Splunk documentation.	4.3	More Details
CVE-2026-131244	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to edit and run playbooks could modify a child playbook callback handle associated with another playbook run, which could alter the callback behavior of that run. The vulnerability exists because the Playbook API does not verify that callback-handle updates belong to the related parent and child playbook runs.	4.3	More Details
CVE-2026-131302	In Splunk AI Assistant versions below 2.0.0, a user who does not hold the "admin" or "power" Splunk roles could turn on the AI chat training-data consent setting through the Representational State Transfer (REST) API. The missing authorization is possible because the affected configuration endpoint does not enforce authorization checks.	4.3	More Details
CVE-2026-2841	The Advanced Custom Fields (ACF) plugin for WordPress is vulnerable to incorrect authorization in all versions up to, and including, 6.7.0. This is due to the 'ACF_Ajax_Query::get_args' method returning the user-supplied ''query' array directly without any key whitelisting or validation. This makes it possible for authenticated attackers, with Subscriber-level access and above, to inject arbitrary query parameters to bypass field-configured role restrictions and exfiltrate sensitive user data including administrator email addresses, login names, and private user metadata via blind oracle techniques.	4.3	More Details
	A flaw was found in Moodle. The grade feedback web service does not properly verify group and course		

CVE-2026-133059	enrolment permissions. An authenticated user with non-editing teacher privileges can exploit this vulnerability to access grade feedback and user profile details across groups or courses they are not authorized to view, leading to unauthorized information disclosure.	4.3	More Details
CVE-2026-133068	A flaw was found in Moodle. An improper authorization check when displaying individual glossary entries allows an authenticated user to view approved glossary items from courses they do not have permission to access. This vulnerability leads to unauthorized information disclosure of restricted educational materials.	4.3	More Details
CVE-2026-131281	In Splunk SOAR versions below 8.8.0, a user with asset edit permission could change the approval requirements and effective automation user for an asset without Splunk SOAR recording field-level audit details for those changes. This could weaken investigation and accountability for changes to security-sensitive asset settings. The vulnerability is possible because Splunk SOAR does not include the asset approval voting fields and effective automation user field in its field-level audit coverage.	4.3	More Details
CVE-2026-69131	django-filer is a file and image management application for Django. Prior to 3.5.0, the delete_files_or_folders action in filer/admin/folderadmin.py checks global Django model-level deletion permission but omits check_folder_edit_permissions and check_files_edit_permissions when FILER_ENABLE_PERMISSIONS is enabled. An authenticated staff user who has a global delete_file, delete_folder, or delete_image permission but is limited to read-only access by FolderPermission can therefore delete files, folders, and descendants from the restricted folder. Deployments using the default disabled per-folder permission system are not affected, and the issue does not create a read-confidentiality impact. This issue is fixed in version 3.5.0.	4.3	More Details
CVE-2026-131279	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to edit artifacts could cause indicator correlation data to retain stale field associations, which can produce incorrect correlation results and event counts. The vulnerability is possible because Splunk SOAR does not remove stale indicator associations when an artifact changes.	4.3	More Details
CVE-2026-131176	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.17, a user that holds a role with the edit_own_objects capability but does not hold the edit_view_html capability could create or modify a private Hypertext Markup Language (HTML) view through the Representational State Transfer (REST) API. The vulnerability is possible because the REST API applies the edit_own_objects capability when it creates or modifies private HTML-based views instead of requiring the edit_view_html capability. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.	4.3	More Details
CVE-2026-131249	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to view, edit, and run playbooks and access to at least one container label could retrieve a callback handle from a playbook run outside the container-label scope granted to that account. The vulnerability is possible because Splunk SOAR does not verify that the requested callback handle belongs to the parent playbook run that launched the child playbook. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.8.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) and Data access automation API (https://help.splunk.com/en/splunk-soar/soar-on-premises/python-playbook-api-reference/8.8.0/automation-api/data-access-automation-api) in the Splunk documentation.	4.3	More Details
CVE-2026-76537	In Splunk Enterprise versions below 10.6.0, 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who does not hold a role with either the list_indexer_cluster capability or the edit_indexer_cluster capability could download the cluster manager configuration bundle through the Representational State Transfer (REST) API and access configuration and app files that the cluster manager stages for indexer distribution. The vulnerability is possible because the REST API does not require the list_indexer_cluster or edit_indexer_cluster capability before it sends the configuration bundle. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.4/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation. For more information see Cluster endpoint descriptions (https://help.splunk.com/en/splunk-enterprise/leverage-rest-apis/rest-api-reference/10.6/cluster-endpoints/cluster-endpoint-descriptions) in the Splunk documentation.	4.3	More Details
CVE-2026-131290	In Splunk SOAR versions below 8.8.0, an authenticated user with permission to edit and run playbooks could view child playbook status and loop-state metadata for playbook runs outside the authorized scope for that user. The vulnerability is possible because the playbook status API does not bind the requested run to the current execution context before it returns metadata.	4.3	More Details
CVE-2026-133063	A flaw was found in Moodle. Due to insufficient validation when saving activity settings, an authenticated user with teacher privileges can modify activities in courses they are not authorized to access. This flaw allows unauthorized changes to course materials, compromising the integrity of course content.	4.3	More Details
CVE-2026-133070	A flaw was found in Moodle. Due to a missing authorization check when viewing and restoring wiki page versions, an authenticated user can access content from courses they are not permitted to view. This issue allows an attacker to inspect restricted course material and copy it into their own wiki pages, leading to unauthorized information disclosure.	4.3	More Details

CVE-2026-131265	In Splunk SOAR versions below 8.8.0, a user that holds a role with the System Settings view permission but not the Assets view permission could view asset names, asset identifiers, app names, tenant names, and ingestion schedule settings for polling assets in the ingestion schedule view. The vulnerability is possible because the ingestion schedule view does not verify the Assets view permission before it returns data derived from assets.	4.3	More Details
CVE-2026-1209	The BuddyPress plugin for WordPress is vulnerable to unauthorized modification of data due to a misconfigured capability check on the 'update_item_permissions_check' function in all versions up to, and including, 14.4.0. This makes it possible for authenticated attackers, with Author-level access and above, to enable and disable plugin components.	4.3	More Details
CVE-2026-131922	In OpenStack keystonemiddleware before 13.0.3, by supplying a malicious regular expression when creating an application credential (aka AppCred) access_rules entry, an authenticated user may consume resources in the consuming process making a service unresponsive.	4.3	More Details
CVE-2026-133065	A flaw was found in Moodle. Due to insufficient permission checks and input validation in the quiz question bank callback, an authenticated user can access and view questions from question banks they do not have permission to access. Furthermore, this issue allows users to trigger arbitrary class instantiation (creating unintended application objects), potentially leading to further unauthorized actions.	4.3	More Details
CVE-2026-133061	A flaw was found in Moodle. Due to a missing capability check in the grade items web service, an authenticated user with a student role can view the names of all grade items in a course, including those intended to be hidden. This flaw leads to unauthorized information disclosure of hidden course assessment details.	4.3	More Details
CVE-2026-131193	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user holding a role with the edit_token_http capability could create a persistent HTTP Event Collector (HEC) token that cannot be disabled or deleted via the Representational State Transfer (REST) API or Splunk Web. As a result, the token remains active and can continue ingesting events even after administrative attempts to revoke it. The vulnerability is possible because HTTP Event Collector token creation accepts names that later management operations cannot process. For more information see authorize.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.4/configuration-file-reference/10.4.5-configuration-file-reference/authorize.conf) and Set up and use HTTP Event Collector in Splunk Web (https://help.splunk.com/en/splunk-enterprise/get-data-in/get-started-with-getting-data-in/10.6/get-data-with-http-event-collector/set-up-and-use-http-event-collector-in-splunk-web) in the Splunk documentation.	4.3	More Details
CVE-2026-69139	djangocms-moderation provides an approval workflow mechanism for organizations that require content approval before publication. Prior to 2.4.0, CollectionCommentAdmin and RequestCommentAdmin in the Django admin use changelist_view and changeform_view to scope comment access through caller-controlled object identifiers and model-level view_collectioncomment or view_requestcomment permissions without invoking ModerationRequest.user_can_view_comments() or an equivalent collection participant check. An authenticated staff user with either view permission can enumerate collection or moderation-request identifiers and read full comment messages for workflows they do not participate in. The exposure affects confidentiality of internal editorial discussions; comment creation, modification, and deletion remain protected, and integrity and availability are unaffected. This issue is fixed in version 2.4.0.	4.3	More Details
CVE-2026-133305	sqlmap before 1.10.11.89 contains a credential exposure vulnerability that allows attackers to capture user-supplied Cookie and Authorization headers by redirecting requests to another host. Attackers controlling the scanned target or abusing an open redirect can make RedirectHandler.http_error_302 forward session cookies or bearer tokens to an attacker-controlled host.	4.3	More Details
CVE-2026-17575	Mattermost versions 11.8.x <= 11.8.5, 11.7.x <= 11.7.10, 11.10.x <= 11.10.2 fails to validate the type of card property fields which allows a board Editor to make a board or card permanently unopenable for all users via a crafted PATCH request storing a non-string property name or value.. Mattermost Advisory ID: MMSA-2026-00743	4.3	More Details
CVE-2026-18045	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to enforce the MFA requirement on the update and delete user endpoints, which allows an authenticated user who has not configured MFA to modify or delete their own account via the API despite system-wide MFA enforcement being enabled. Mattermost Advisory ID: MMSA-2026-00744	4.3	More Details
CVE-2026-132774	Attendize through 2.8.0 contains a cross-account authorization bypass vulnerability in EventAttendeesController::postMessageAttendees() that allows authenticated organisers to message other accounts' attendees by supplying a foreign event_id. Attackers, including self-registered organisers, can post to the attendees message endpoint for another event to send arbitrary emails to all attendees, enabling phishing and impersonation.	4.3	More Details
CVE-2026-18396	Mattermost versions 11.7.x <= 11.7.11, 11.10.x <= 11.10.2, 11.11.x <= 11.11.1 fail to enforce playbook membership management permissions when creating a new playbook which allows a low-privileged user to grant the playbook_admin role to an arbitrary user via a crafted playbook creation request. Mattermost Advisory ID: MMSA-2026-00749	4.3	More Details

CVE-2026-130672	Craft CMS from 5.0.0 before 5.11.5 contains a missing authorization vulnerability in ChartsController::actionGetNewUsersData() that lacks a viewUsers permission check. Authenticated control panel users can POST to the charts/get-new-users-data action to read aggregate user-registration time series and totals for the site or any user group.	4.3	More Details
CVE-2026-132608	A vulnerability was identified in SECUDOCX DMP up to 9.0.0 Build 1027. This affects an unknown function of the file /DMP/public/login/forward_login.jsp of the component Login. The manipulation of the argument LoginErrorMsg leads to cross site scripting. The attack may be initiated 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.	4.3	More Details
CVE-2026-131653	Tanium addressed an improper input validation vulnerability in Discover.	4.3	More Details
CVE-2026-133174	HotGo 2.18.6 through 2.19.0 contains an information disclosure vulnerability that allows authenticated backend users to read other administrators' private letters because the messageIds helper omits the receiver filter. Low-privilege users can call GET /admin/notice/pullMessages, which is exempt from role permission checks via exceptAuth, to retrieve titles, contents and senders of private letters.	4.3	More Details
CVE-2026-133226	A flaw was found in Keycloak. The Account API endpoint for sharing resources fails to verify whether owner-managed access is enabled and accepts permission scopes not assigned to the resource. An authenticated user can exploit this vulnerability to generate pre-granted permission tickets, causing Keycloak to bypass policy evaluation and issue access tokens with unauthorized scopes.	4.3	More Details
CVE-2026-131273	In Splunk SOAR versions below 8.8.0, an authenticated user who holds a role with permission to view system settings but not assets could view asset, app, and container-label metadata in Ingestion filters. The vulnerability is possible because Splunk SOAR returns Ingestion filter metadata without applying asset authorization and tenant restrictions. For more information see Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.5.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	4.3	More Details
CVE-2026-130755	Froxlor before 2.3.16 contains an authorization bypass vulnerability that allows authenticated customers to read hidden traffic statistics by calling the Traffic.listing API command. Attackers can authenticate to api.php with their API key and secret to retrieve HTTP, FTP, and mail traffic volumes despite the panel.customer_hide_options restriction.	4.3	More Details
CVE-2026-132941	In JetBrains TeamCity before 2026.2, 2026.1.4 missing authorization allowed users to remove all build configuration assignments from an agent	4.3	More Details
CVE-2026-133494	Greenbone gvmd before 26.43.0 does not restrict certificate creation. Even the Observer role, documented to be read-only, can execute create_tls_certificate.	4.3	More Details
CVE-2026-133535	A flaw was found in Hawtio. An authenticated remote attacker can cause a Denial of Service (DoS) by sending a specially crafted request URI to monitored endpoints. The FlightRecordingDownloadFacade servlet filter uses an inefficient regular expression pattern that causes catastrophic backtracking when processing certain long requests. This issue leads to excessive CPU utilization, causing handling threads to hang and degrading service availability.	4.3	More Details
CVE-2026-133480	A flaw has been found in hrshadhin CloudSchool up to 3.0.1. Affected by this issue is the function PublicController.checkSanity4Student of the file app/Http/Controllers/Backend/PublicController.php of the component Student Data Access Check. This manipulation of the argument student_id causes authorization bypass. The attack can be initiated remotely. The exploit has been published and may be used. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-133536	A flaw was found in Keycloak. The OpenID Connect (OIDC) logout endpoint continues to honor post-logout redirect rules for disabled clients rather than rejecting the requests. If a disabled client retains permissive redirect configurations, a remote attacker can exploit this issue by supplying an arbitrary redirection target in a logout request, leading to an open redirect that can be used in phishing attacks.	4.3	More Details
CVE-2026-131798	Tanium addressed a server-side request forgery vulnerability.	4.3	More Details
CVE-2026-132935	In JetBrains TeamCity before 2026.2 missing authorization in statistics chart export allowed viewing data of project	4.3	More Details
CVE-	The Flatsome WordPress theme before 3.20.13 does not perform a capability or nonce check on an AJAX		More

2026-117489	action available to authenticated users, allowing users with a role as low as Subscriber to disclose the name and content of non-public templates stored on the site.	4.3	Details
CVE-2026-120255	InfoScan uses hard-coded cryptographic keys and initialization-vector values in its database-encryption implementation. The application uses these values to encrypt the MongoDB password before storing it in the Windows registry. Because the cryptographic material is embedded in the distributed application, an individual with access to the encrypted registry value and the application code may be able to recover the database password.	4.3	More Details
CVE-2026-130600	Insertion of Sensitive Information Into Sent Data vulnerability in BdThemes Element Pack Elementor Addons bdthemes-element-pack-lite allows Retrieve Embedded Sensitive Data.This issue affects Element Pack Elementor Addons: from n/a through 10.0.9.	4.3	More Details
CVE-2026-133489	A security vulnerability has been detected in XiangLi1999 Diffusion-LM up to 759889d58ef38e2eed41a8c34db8032e072826f4. This issue affects the function main of the file improved-diffusion/scripts/text_sample.py. Such manipulation of the argument out_dir leads to path traversal. The attack can be executed remotely. The exploit has been disclosed publicly and may be used. This product implements a rolling release for ongoing delivery, which means version information for affected or updated releases is unavailable. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-132919	In JetBrains TeamCity before 2026.2 missing authorization allowed reading Qodana reports of other projects	4.3	More Details
CVE-2026-132920	In JetBrains TeamCity before 2026.2, 2026.1.4 an SSRF via connection tests was possible for users without project edit permissions	4.3	More Details
CVE-2026-17537	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to enforce the channel Access Control Policy's file download permission check when serving file metadata and thumbnails through the drafts, scheduled posts, and post edit history endpoints, which allows an authenticated user who has lost eligibility under the policy to view metadata and thumbnails of file attachments via those endpoints, even though direct file access is correctly denied. Mattermost Advisory ID: MMSA-2026-00740	4.3	More Details
CVE-2026-18399	Mattermost versions 11.7.x <= 11.7.11, 11.10.x <= 11.10.2, 11.11.x <= 11.11.1 fail to enforce the configured GitLab group/namespace restriction on the permalink code-preview feature and the {{get_projects}} MCP tool, which allows an authenticated user with a connected GitLab account to retrieve repository file content or project metadata from projects outside the restricted group via a crafted GitLab permalink or MCP tool call. Mattermost Advisory ID: MMSA-2026-00750	4.3	More Details
CVE-2026-133259	In JetBrains YouTrack before 2026.2.19840 improper authorization allowed read-only users to change issue fields via agile boards	4.3	More Details
CVE-2026-16319	IBM API Connect 10.0.8.0 through 10.0.8.13 could allow a remote authenticated attacker to delete API keys belonging to other users due to improper authorization checks.	4.3	More Details
CVE-2026-92907	IBM watsonx.data 2.4.0 through 2.4.0 patch 8 could allow a remote authenticated attacker to conduct a server-side request forgery due to a server-side request forgery (SSRF) vulnerability.	4.3	More Details
CVE-2026-133114	Caddy – WooCommerce Side Cart & Free Shipping Bar plugin for WordPress 3.0.0 before 3.2.1 contains a missing authorization vulnerability in the saved-items/add REST endpoint that exposes non-public product data. Authenticated attackers with Subscriber-level access can enumerate product_id values to read names, prices, images and stock status of private, draft, hidden and password-protected products.	4.3	More Details
CVE-2026-133115	Redmine 6.1.0 before 6.1.6 and 7.0.0 before 7.0.3 contains an incorrect authorization vulnerability that lets OAuth clients exceed token scope during project creation. A client with only project-creation scope can set is_public, enabled_module_names and inherit_members, making projects public and granting parent-project members subproject access.	4.3	More Details
CVE-2026-131191	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17, a user who does not hold the "admin" or "power" Splunk roles could read private index configuration values for other users through the deployment server Representational State Transfer (REST) API. The deployment server REST API does not limit a supplied repository location to configured deployment repositories before it reads indexes.conf stanzas.	4.3	More Details
CVE-2026-133481	In Dolibarr before 24.0.3, /api/index.php/eventattendees discloses attendee and project data regardless of the caller's permissions. Only 22.0.0 and later (where api_eventattendees.class.php exists) are affected. A remote authenticated user can obtain column values one character at a time.	4.3	More Details
CVE-	Weblate 2026.6 through 2026.11.2 contains an information disclosure vulnerability that allows authenticated		

2026-132126	users to enumerate all workspaces through an unfiltered ProjectSerializer.workspace field. Low-privilege users can request the browsable /api/projects/ page to render the project creation form, revealing every workspace UUID and name they cannot access.	4.3	More Details
CVE-2026-89218	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to enforce the Access Control Policy's download_file_attachment check across several post- and file-serving code paths (bulk post retrieval, channel bookmarks, and persistent notification delivery), which allows an authenticated channel member denied a specific file's download to obtain that file's metadata and a rendered preview thumbnail via those endpoints and events.. Mattermost Advisory ID: MMSA-2026-00785	4.3	More Details
CVE-2026-130034	The Nexter Blocks – Gutenberg Blocks, Page Builder & AI Website Builder plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 5.0.10 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 create, modify, and delete entries in the site-wide global button-preset store, and inject arbitrary CSS rules into the generated plus-button-presets stylesheet served to all site visitors.	4.3	More Details
CVE-2026-131855	Ettercap before 0.8.5.2 contains a heap out-of-bounds read vulnerability in dissector_vnc() in src/dissectors/ec_vnc.c that checks total packet length instead of remaining bytes. A malicious VNC server can send the authentication type and challenge in one short packet to leak up to 3 bytes of adjacent heap memory via the logged challenge.	4.3	More Details
CVE-2026-131187	In Splunk Enterprise versions below 10.6.1, 10.4.5, and 10.2.9, a low-privileged user who does not hold the "admin" or "power" Splunk roles could access file path and metadata listings from search head knowledge bundles via the Representational State Transfer (REST) API. This unauthorized access could expose sensitive metadata, including Splunk platform usernames, private knowledge-object configuration file names, installed application directories, and the file path names of security-sensitive configuration and key files. The vulnerability is possible because the REST API actions that list bundle contents and denied files use a capability available to the "user" role instead of requiring distributed-search administration capabilities. For more information see Knowledge bundle replication overview (https://help.splunk.com/en/splunk-enterprise/administer/distributed-search/10.6/knowledge-bundle-replication/knowledge-bundle-replication-overview) in the Splunk documentation. For more information see Define roles on the Splunk platform with capabilities (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.6/manage-splunk-platform-users-and-roles/define-roles-on-the-splunk-platform-with-capabilities) in the Splunk documentation.	4.3	More Details
CVE-2026-132049	A flaw was found in lldpd. The service fails to verify that Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) location records received from network neighbors match their declared lengths. An adjacent attacker on the local network can exploit this flaw by sending crafted packets, causing client applications using the liblldpctl library to read past buffer boundaries when parsing location data. This issue can lead to an information disclosure of adjacent process memory.	4.3	More Details
CVE-2026-132247	In JetBrains TeamCity before 2026.1.3, 2025.11.7 improper authorization allowed viewing build process details via agent stacktraces	4.3	More Details
CVE-2026-130707	NocoDB from 0.301.4 before 2026.10.1 writes the nc_token session cookie from script in the Enterprise useGlobal composable without the Secure attribute. Network-positioned attackers can observe or induce plain-HTTP requests to https instances to capture the session token and impersonate the victim until expiry.	4.2	More Details
CVE-2026-108317	Tanium addressed an information disclosure vulnerability in Tanium Server.	4.2	More Details
CVE-2026-131234	In Splunk Enterprise Security versions below 8.7.0, an unauthenticated user could create and deliver an Entity Analysis dashboard link to a Splunk Enterprise Security user. If the Splunk Enterprise Security user opens the link, Splunk Enterprise Security could run attacker-controlled Search Processing Language (SPL) commands using the permissions of the Splunk Enterprise Security user. The commands could expose or modify some data available to the Splunk Enterprise Security user. The vulnerability is possible because the Entity Analysis dashboard inserts entity values from a link that an unauthenticated user creates into SPL searches without safely escaping the values. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see Entity analysis dashboard (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/exposure-analytics-and-ueba/entity-analysis-dashboard) in the Splunk documentation.	4.2	More Details
CVE-2026-133192	In JetBrains Junie GitHub Action before 1.10.0 untrusted fork code could be processed in privileged workflow runs	4.2	More Details
CVE-2026-	In JetBrains TeamCity before 2026.2 poisoning of OAuth discovery metadata via the X-Forwarded-Host header was possible	4.2	More Details

132926			
CVE-2026-131204	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the mc_investigation_read capability could overwrite Mission Control preference records for other users and remove saved filters and analyst queue settings from those records. The vulnerability is possible because the Mission Control user preferences Representational State Transfer (REST) API does not bind an update to the preference record that passes the authorization check. The attack requires the user to obtain a preference record identifier that the product does not normally expose to that user. The user should not be able to exploit the vulnerability at will. For more information see Overview of Mission Control in Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/user-guide/8.7/mission-control/overview-of-mission-control-in-splunk-enterprise-security) in the Splunk documentation.	4.2	More Details
CVE-2026-76533	In Splunk Enterprise versions below 10.4.5, 10.2.9, 10.0.12, and 9.4.17 on Ubuntu and AlmaLinux, an unauthenticated user who obtains a valid Splunk Web session cookie before the affected user logs out could continue using the cookie after logout to access data and perform actions available to that user. The vulnerability is possible because the logout flow does not invalidate the associated session token on affected deployments. The attacker must obtain the valid session cookie before the affected user logs out of Splunk Web. The issue was reproduced in Splunk Enterprise installations on Ubuntu and AlmaLinux.	4.2	More Details
CVE-2026-132252	A vulnerability was found in October CMS up to 4.4.6. This affects the function validateExternalImageHost of the file modules/system/classes/resizeimages/ValidatesExternalImages.php of the component Image Resizer. Performing a manipulation of the argument url results in server-side request forgery. The attack may be initiated remotely. The complexity of an attack is rather high. The exploitability is reported as difficult. The exploit has been made public and could be used. Upgrading to version 4.4.7 mitigates this issue. The patch is named f6372867006c03141944e7b57396a726eb02acf0. You should upgrade the affected component.	4.1	More Details
CVE-2026-130768	Obot from 0.14.0 before 0.27.0 contains a server-side request forgery vulnerability that allows Power Users to bypass outbound protections by supplying unspecified addresses, IPv6 transition addresses, or registry icon URLs. Attackers can add remote MCP servers pointing to 0.0.0.0, ::, or NAT64/6to4/Teredo addresses embedding blocked IPv4 addresses to probe internal services for reachability.	4.1	More Details
CVE-2026-120406	The Contact Form 7 WordPress plugin before 6.2.5 does not check that the post a form submission says it was sent from is publicly viewable before resolving post-related mail-tags, allowing unauthenticated users to obtain the title, URL, author name and author email of private, draft and trashed posts where such tags are used in an email or message returned to the submitter.	3.7	More Details
CVE-2026-129937	The MongoDB C# Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details
CVE-2026-133263	In JetBrains YouTrack before 2026.2.19840 missing authorization allowed guests to add issues to sprints	3.7	More Details
CVE-2026-130658	ImageMagick before 7.1.2-35 contains an information disclosure vulnerability in the RLA decoder that returns images containing uninitialized data. Attackers can supply a crafted RLA image to cause the decoded output to include uninitialized memory, potentially disclosing heap information.	3.7	More Details
CVE-2026-129929	The MongoDB Node.js Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details
CVE-2026-130665	ImageMagick before 7.1.2-35 contains a use after free vulnerability in the EXR decoder that is triggered when a specific memory allocation fails. Attackers can supply crafted EXR images processed under memory pressure to trigger the use after free, causing crashes or memory corruption.	3.7	More Details
CVE-2026-129923	The MongoDB Rust Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details
CVE-2026-	The MongoDB Ruby Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result	3.7	More

129925	in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.		Details
CVE-2026-131219	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in network traffic data. Another user who holds the "pci_user" role could later click the corresponding row in the Traffic Detail panel of the Network Traffic Activity report. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who clicked the row. The commands could access data and perform actions available to that user. The vulnerability is possible because the Network Traffic Activity report does not treat network traffic values as literal search values when it builds the Traffic Detail drilldown search. For more information see Network Traffic Activity (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-reports/network-traffic-activity) and Configure users and roles (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-the-app/configure-users-and-roles) in the Splunk documentation.	3.7	More Details
CVE-2026-131224	In Splunk Enterprise Security versions below 8.7.0, a user that holds a role with the run_collect capability could store attacker-controlled values in data that populates the Endpoint Changes report. Another user who holds the "pci_analyst" role could later click a corresponding table row. Splunk Enterprise Security could then run attacker-controlled Search Processing Language (SPL) commands using the permissions of the user who clicked the row. The commands could access or modify data available to that user. The vulnerability is possible because the Endpoint Changes report does not treat dynamic report values as literal search values when it builds table-row drilldown searches. For more information see Reports (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/user-manual/8.7/using-the-splunk-app-for-pci-compliance/reports) and Configure users and roles (https://help.splunk.com/en/splunk-enterprise-security-8/splunk-app-for-pci-compliance/installation-and-configuration-manual/8.7/configure-the-app/configure-users-and-roles) in the Splunk documentation.	3.7	More Details
CVE-2026-130662	ImageMagick before 7.1.2-35 contains a use after free vulnerability in the MSL decoder that can be triggered when processing crafted MSL image files. Attackers can supply a specially crafted MSL image to cause memory corruption, potentially crashing the application or leading to further compromise.	3.7	More Details
CVE-2026-129926	The MongoDB Python Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details
CVE-2026-130656	ImageMagick before 7.1.2-35 contains an information disclosure vulnerability in the MNG decoder that allows attackers to expose uninitialized memory by supplying a crafted image. Attackers can submit malicious MNG files for processing to leak uninitialized memory contents, potentially revealing sensitive data from the process.	3.7	More Details
CVE-2026-130691	NocoDB before 2026.10.1 contains an information disclosure vulnerability that allows unauthenticated attackers to download data exports by reconstructing predictable file paths under nc/uploads/data-export. Attackers can derive paths from the table id, base, table and view titles, and export time to read full exports, including fields hidden from public shared views.	3.7	More Details
CVE-2026-130671	Parse Server before 8.6.119 and 9.x before 9.10.8-alpha.3 contains a missing authorization vulnerability that allows LiveQuery clients without the master key to subscribe to installation and internal classes. Attackers can receive installation IDs and device tokens of other devices when these classes are in liveQuery.classNames, then modify those installations to redirect push notifications.	3.7	More Details
CVE-2026-126765	The glob functions in the GNU C Library versions 2.44 and earlier may free the same heap block twice when a memory allocation fails while glob prepends the directory name to the names it matched. The names already rewritten are freed, but their pointers stay in the result vector, and the error path frees them again with globfree. The allocation must fail after at least one name has been rewritten, so the issue can only be reached when the process is already close to its memory limit, for example under an address space limit (RLIMIT_AS) or on a 32-bit system. The attacker needs control to make the result large enough to make the program or service run out of memory, either by controlling the glob pattern and having knowledge of the filesystem layout or through other means. This issue affects all architectures and ABIs supported by the GNU C Library.	3.7	More Details
CVE-2026-129935	The MongoDB Go Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details

CVE-2026-133057	In JetBrains TeamCity before 2026.2.1, 2026.1.5 the TFS helper server was exposed on all network interfaces	3.7	More Details
CVE-2026-129933	The MongoDB Java Driver does not neutralize NUL characters in the user name and password values it places into the message it builds for SASL PLAIN authentication. A party who can choose the login values that an application passes to the driver can alter where the fields of that message begin and end. This may result in a receiving system reading a login name and secret that differ from the values the application supplied. Authenticating as a different account still requires that account's password. Only deployments that use SASL PLAIN authentication are affected.	3.7	More Details
CVE-2026-133190	MISP, a threat intelligence sharing platform, contains a vulnerability in its cryptographic key management for protected (PGP-signed) events. Two related defects allow a revoked PGP key to remain effective for signing and validating protected events. First, the validateProtectedEvent() method in the CryptographicKey model iterates over the event's stored cryptographic keys to verify a PGP signature but does not check the revoked flag on each key. A key explicitly marked as revoked is still used to validate signatures, meaning an entity whose key has been revoked can continue to sign protected events and have them accepted by the MISP instance. Second, the captureCryptographicKeyUpdate() method, responsible for synchronizing the key list when a sync user edits a protected event, contained a query bug (find('first') instead of find('all')) that caused only the first stored key to be considered. As a result, revocations of additional keys were never persisted, keys absent from the pushed list were not removed, and entries lacking an explicit fingerprint were mismatched and incorrectly deleted. Additionally, the revocation logic did not enforce the invariant that a revocation, once stored, is never lifted by subsequent pushes. A third minor issue: a signed protected event submitted without any cryptographic keys triggered an unhandled 500 error rather than a proper rejection response. Preconditions: the attacker must be an authenticated sync user with permission to add or edit protected events on the MISP instance. The impact is that revoked keys retain signing authority, key revocations propagated by sync partners are silently dropped, and the integrity of the protected-event trust chain is compromised. In a multi-instance MISP deployment, events signed by revoked keys can propagate to other instances, extending the integrity impact across the sync network.	3.7	More Details
CVE-2026-133477	A weakness has been identified in sofastack sofa-bolt up to 1.6.12. This impacts the function initSSLContext of the file src/main/java/com/alipay/remoting/connection/AbstractConnectionFactory.java of the component TLS Stack. Executing a manipulation can lead to certificate with host mismatch. The attack may be performed from remote. This attack is characterized by high complexity. The exploitability is said to be difficult. The vendor was contacted early about this disclosure but did not respond in any way.	3.7	More Details
CVE-2026-130714	NocoDB before 2026.10.1 contains a user enumeration vulnerability that allows unauthenticated attackers to determine registered email addresses by observing differing sign-up endpoint responses. Attackers can submit addresses to signup in users.service.ts and compare error messages to identify existing accounts and pending Enterprise invitations, even under invite-only signup.	3.7	More Details
CVE-2026-133181	In JetBrains ReSharper before 2026.3 server host names were not verified during certificate validation	3.7	More Details
CVE-2026-130713	NocoDB before 2026.10.1 contains a user enumeration vulnerability that allows unauthenticated attackers to identify pending invitees through a distinct sign-in error message. Attackers can submit sign-in requests where validateUser in auth.service.ts skips password comparison for saltless invited accounts, confirming the email address has an account.	3.7	More Details
CVE-2026-132993	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the notification rules editor was possible	3.7	More Details
CVE-2026-133180	In JetBrains ReSharper before 2026.3 the What's New feature did not verify server certificates	3.7	More Details
CVE-2026-133435	A flaw was found in Keycloak. When configured with a database that uses case-insensitive collation, Keycloak resolves realm URLs regardless of letter casing. A remote attacker can exploit this behavior by sending requests with alternate letter casing to bypass upstream reverse proxy access controls that rely on case-sensitive path matching. This issue can expose restricted authentication endpoints intended only for internal access, though Keycloak's internal authentication mechanisms remain intact.	3.7	More Details
CVE-2026-133179	In JetBrains Rider before 2026.3 code execution was possible via MSBuild configuration files in parent directories	3.6	More Details
CVE-2026-	PHPGurukul Beauty Parlour Management System V1.1 uses a non-cryptographic function to generate security-sensitive appointment and invoice identifiers. The identifiers are generated within a restricted numeric range and are subsequently used in record lookup operations. This implementation exposes the	3.5	More Details

128658	application to identifier prediction or guessing attacks under conditions where the identifiers are relied upon as a security boundary.		
CVE-2026-132995	In JetBrains TeamCity before 2026.2 stored XSS on the build step settings page was possible	3.5	More Details
CVE-2026-133056	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS in implicit agent requirements was possible	3.5	More Details
CVE-2026-131240	In Splunk SOAR versions below 8.8.0, an authenticated user could cause Splunk SOAR to store excessive data in the Ingestion Status history through the Representational State Transfer (REST) API. When another user with permission to view system settings retrieves the Ingestion Status history, Splunk SOAR could use additional resources and return a larger response, which could degrade the availability of administrative monitoring. The vulnerability is possible because Splunk SOAR does not limit the amount of request data that it records for failed container creation before it stores the data in the Ingestion Status history. For more information see REST Containers (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/6.4.0/container-endpoints/rest-containers), REST Roles and Permissions (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/6.4.0/role-management-endpoints/rest-roles-and-permissions), and View ingested container statistics using Ingestion Status (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.4.1/manage-your-splunk-soar-on-premises-system-health-and-performance/view-ingested-container-statistics-using-ingestion-status) in the Splunk documentation.	3.5	More Details
CVE-2026-130654	ImageMagick before 7.1.2-35 contains a data tampering vulnerability in the Distributed Pixel Cache that allows adjacent attackers to modify pixel data in transit between client and server. Attackers who can observe or modify the network path between a configured DPC client and server can tamper with pixel data to alter the resulting image.	3.5	More Details
CVE-2026-133267	In JetBrains YouTrack before 2026.2.19840 project members could detach inline comment threads of other users on articles	3.5	More Details
CVE-2026-110337	The Strava datasource plugin's resource handler POST /strava-api forwards a client-supplied endpoint string, and its query parameters, to the Strava API (https://www.strava.com/api/v3/<endpoint>) without validation. When the datasource is not in OAuth pass-through mode, the request is made with the access token stored on the datasource, not the caller's identity. Grafana authorises plugin resource calls only on the datasources:query permission, which Viewers have by default. Any such caller could therefore request any Strava API path the connected account's OAuth grant allows, for example athlete/clubs or routes/{id}, and read the response. The plugin's frontend (stravaApi.ts) only needs a fixed set of paths: athlete, athlete/activities, activities/{id}, activities/{id}/streams and segments/{id}. Requests to Strava are GET only, so the issue is read-only. Affected: plugin versions 1.2.0 through 1.7.2. The fix restricts endpoint to those five patterns and is pending release. The check runs before the access token is resolved, so a rejected request never causes a token refresh or any call to Strava. Paths with .. segments are rejected as well.	3.5	More Details
CVE-2026-133055	In JetBrains TeamCity before 2026.2.1, 2026.1.5 stored XSS on the Failure Conditions page was possible	3.5	More Details
CVE-2026-133413	A vulnerability was detected in smartalock wireguard-lwip up to cfba91666590e86bd67842c71cedf9a423be823a. This affects the function wireguard_generate_unique_index of the file src/wireguard.c of the component Session Index Allocator. The manipulation of the argument receiver_index results in allocation of resources. The attack is only possible with local access. The project was informed of the problem early through an issue report but has not responded yet.	3.4	More Details
CVE-2026-130660	ImageMagick before 7.1.2-35 contains a heap buffer over-read vulnerability in the MVG decoder that allows attackers to read out-of-bounds memory by supplying a crafted file. Attackers can provide a malicious MVG file for processing to trigger the over-read, causing application crashes or exposure of heap memory contents.	3.3	More Details
CVE-2026-130670	ImageMagick before 7.1.2-35 contains a heap buffer over-read vulnerability that allows attackers to read beyond allocated memory when BGRO image data is imported. Attackers can supply crafted BGRO data for processing to crash the application or leak adjacent heap memory contents.	3.3	More Details
CVE-2026-130669	ImageMagick before 7.1.2-35 contains a null pointer dereference vulnerability that allows attackers to crash the process by using the -subimage option on a single-image list. Attackers can supply input producing an image list with only one image while -subimage is applied, causing a denial of service.	3.3	More Details
CVE-2026-	A flaw was found in the Ansible language server shipped with the vscode-ansible extension. The regular expression used to extract documentation fragments from Ansible module files is susceptible to quadratic-time scanning. A module file containing a crafted documentation string causes the language server to consume excessive CPU while indexing it, freezing the language service in the developer's editor. An attacker	3.3	More Details

127096	who can get a developer to open or index a workspace containing an attacker-supplied collection can cause a denial of service of the editor language features. This issue has no confidentiality or integrity impact.		
CVE-2026-127095	A flaw was found in ansible-builder. The regular expression used to strip comments from a collection's requirements.txt during introspection is susceptible to super-quadratic backtracking. A requirements file containing a line with a long run of internal whitespace and no comment character causes the comment-stripping routine to consume CPU time proportional to the square of the line length. An attacker who can get a victim to build an execution environment that includes an attacker-supplied collection can cause the ansible-builder build or introspect process to hang, resulting in a denial of service. This issue has no confidentiality or integrity impact.	3.3	More Details
CVE-2026-130653	ImageMagick before 7.1.2-35 contains a policy bypass vulnerability that allows attackers to read from disallowed paths due to a missing policy check when decoding MPC files. Attackers can supply crafted MPC files that reference restricted locations to access files the configured security policy should block.	3.3	More Details
CVE-2026-130652	ImageMagick before 7.1.2-35 contains a heap buffer over-read vulnerability that allows attackers to read beyond allocated heap memory by triggering the -copy operation. Attackers can supply crafted images or arguments processed with -copy to crash the application or disclose adjacent heap memory contents.	3.3	More Details
CVE-2026-132560	A flaw has been found in MemTensor memmy-agent up to 1.1.1. This impacts an unknown function of the file /api/v1/memory/logs of the component Memory HTTP Service. This manipulation causes information disclosure. The attack is restricted to local execution. The exploit has been published and may be used. Patch name: 421ce6cb33861255d34c3887d39aced009b3ce9b. It is suggested to install a patch to address this issue.	3.3	More Details
CVE-2026-130651	ImageMagick before 7.1.2-35 contains a stack overflow vulnerability in the MSVG decoder that allows attackers to crash the application by supplying a crafted image. Attackers can submit a malicious image for processing to exhaust the stack and terminate the application, causing denial of service.	3.3	More Details
CVE-2026-133406	A vulnerability was determined in aoiflux mutant 2.3.0/2.4.0/2.5.0. This vulnerability affects unknown code of the file builtin/hive_builtins.go of the component Forensic Parser. This manipulation of the argument Count causes uncontrolled memory allocation. It is possible to launch the attack on the local host. Upgrading to version 2.6.0 is able to resolve this issue. Patch name: c4b1a7dec6c603115e2a05d1e8970a142f0b9774. Upgrading the affected component is recommended.	3.3	More Details
CVE-2026-133217	A weakness has been identified in OSGeo shapelib up to 1.6.3. This vulnerability affects the function DBFCloneEmpty of the file dbfopen.c. This manipulation causes null pointer dereference. It is possible to launch the attack on the local host. Patch name: f79050a900d2a05529501b2b714b8f89067db071. It is recommended to apply a patch to fix this issue.	3.3	More Details
CVE-2026-132125	Weblate 5.0 through 2026.11.2 contains a privilege escalation vulnerability where deleting a component list leaves stale Group.projects rows, widening team scope to whole projects. A user with componentlist.edit who deletes a list grants 'From component list' team members their full role, including project.edit, across possibly private projects.	3.3	More Details
CVE-2026-82091	The Canva Desktop App before 1.128.0 stores encrypted application state on disk. The encryption key is hardcoded as a plaintext string in the main.js file and is identical across all Canva installations on Windows and macOS.	3.2	More Details
CVE-2026-132251	In JetBrains TeamCity before 2026.1.3 privilege escalation was possible by issuing new tokens from a restricted access token	3.1	More Details
CVE-2026-133233	Quasar before 2.35.3 contains a DOM clobbering vulnerability that allows attackers to break native wrapper detection by injecting HTML elements named Capacitor, cordova, or StatusBar. Attackers can place markup such as a div with id Capacitor into server-rendered or pre-rendered HTML so app.use(Quasar) throws and the application never mounts, or openURL() and AddressbarColor.set() throw.	3.1	More Details
CVE-2026-18643	Mattermost versions 11.7.x <= 11.7.11 fail to enforce the team view permission across all filter combinations of the user listing endpoint which allows an authenticated user to enumerate the membership of a team they are not permitted to view via a crafted request to that endpoint.. Mattermost Advisory ID: MMSA-2026-00755	3.1	More Details
CVE-2026-133238	SABnzbd before 5.2.0 contains a path traversal vulnerability in the Sorter that allows attackers to write files outside the Complete folder via crafted propername, episodename or job name (%dn) values. Attackers can supply a malicious NZB with '...' sequences that collapse to '..' after single-pass replacement, moving downloaded files into arbitrary directories writable by SABnzbd.	3.1	More Details
CVE-2026-132992	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in the agent info popup was possible	3.1	More Details
	OpnForm before 2.8.2 contains a server-side request forgery vulnerability in PublicWebhookUrl::isPublicIp		

CVE-2026-133240	that fails to reject CGNAT, NAT64 and 6to4 address ranges. Authenticated users can configure webhook, form import or PDF image URLs resolving to these ranges to make the server send HTTPS requests to internal hosts.	3.1	More Details
CVE-2026-132991	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 reflected XSS on the Slack connection test page was possible	3.1	More Details
CVE-2026-123662	The ProfilePress WordPress plugin before 4.18.1 does not atomically consume the source subscription when processing a membership plan change, allowing a customer to submit concurrent plan-change checkouts against a single paid subscription and obtain several active memberships for the price of one.	3.1	More Details
CVE-2026-111518	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain an Improper Restriction of Rendered UI Layers or Frames vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Denial of Service and Launch of Phishing Attacks.	3.1	More Details
CVE-2026-132250	In JetBrains TeamCity before 2026.1.3 privilege escalation was possible via Basic authentication with a restricted access token	3.1	More Details
CVE-2026-133112	modified eCommerce Shopsoftware through 3.3.0 contains a cross-site request forgery vulnerability in gv_send.php because the action=process step never validates a CSRF token via check_secure_form(). Attackers can host an auto-submitting form that makes a logged-in customer send their gift voucher balance to an attacker-controlled email address for redemption.	3.1	More Details
CVE-2026-132990	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 stored XSS in Change Log revision tooltips was possible	3.1	More Details
CVE-2026-133252	In JetBrains Datalore before 2026.4.1 open redirect after OpenID login was possible via the origin parameter	3.1	More Details
CVE-2026-16050	Mattermost versions 11.7.x <= 11.7.11 fail to sanitize channel mentions when rendering permalinked posts, which allows an authenticated user to discover the display names of private channels they do not have access to via a permalink to a post that mentions the channel.. Mattermost Advisory ID: MMSA-2026-00721	3.1	More Details
CVE-2026-108321	Tanium addressed an improper access controls vulnerability in Tanium Server.	3.1	More Details
CVE-2026-18075	Mattermost versions 11.11.x <= 11.11.0, 11.7.x <= 11.7.11, 11.10.x <= 11.10.2 fail to invalidate cached WebSocket session state when an OAuth app is deleted, which allows a revoked OAuth integration to keep receiving real-time events, including private channel posts, via a WebSocket connection that was already open when the app was deleted.. Mattermost Advisory ID: MMSA-2026-00745	3.1	More Details
CVE-2026-133264	In JetBrains YouTrack before 2026.2.19840 private issue summaries and descriptions could be read via agile board queries for archived projects	3.1	More Details
CVE-2026-133265	In JetBrains YouTrack before 2026.2.19840 the attachment upload size limit could be bypassed via the REST API	3.1	More Details
CVE-2026-131289	In Splunk SOAR versions below 8.8.0, a user who holds a role that can edit and run playbooks in Splunk SOAR could cause Splunk SOAR to run non-read-only actions from a Safe Mode playbook, resulting in unauthorized state changes. The vulnerability is possible because Splunk SOAR does not verify that the playbook run context in an action request matches the context that the server provided.	3.1	More Details
CVE-2026-133266	In JetBrains YouTrack before 2026.2.19766 helpdesk email attachments were sent to CC recipients without a visibility check	3.1	More Details
CVE-2026-104683	The Envira Gallery WordPress plugin before 1.16.4 does not apply its own gallery-permission policy until an administrator first opens its permissions screen, instead copying gallery-management capabilities onto each role from that role's core post capabilities after activation, which lets Contributor accounts create, edit and delete galleries that the policy is meant to restrict to higher roles.	2.7	More Details
CVE-2026-131274	In Splunk SOAR versions below 8.8.0, a user with the system_settings permission could change event custom field definitions without the changes appearing in the audit trail, reducing forensic visibility and change accountability. The issue occurs because Splunk SOAR does not audit custom field definition changes.	2.7	More Details

CVE-2026-133236	Modoboa before 2.13.0 contains a CRLF injection vulnerability that allows Resellers or users with domain add or change permission to inject rspamd DKIM selector map entries by supplying an unvalidated dkim_key_selector through the relay domain API. Attackers can embed a carriage return in the selector to override another tenant's domain selector, causing rspamd to sign its outgoing mail with a nonexistent selector and fail DKIM verification.	2.7	More Details
CVE-2026-133268	In JetBrains YouTrack before 2026.2.19728 missing authorization allowed creating sub-articles without the Create Article permission	2.7	More Details
CVE-2026-133108	In JetBrains TeamCity before 2026.2.2 the current request details debug endpoint exposed authentication headers and session cookies	2.7	More Details
CVE-2026-132996	In JetBrains TeamCity before 2026.2 mute and unmute operations via XML-RPC could be attributed to other users	2.7	More Details
CVE-2026-130876	Tanium addressed an uncontrolled recursion vulnerability.	2.7	More Details
CVE-2026-130693	NocoDB from 0.253.0 before 2026.09.1 contains an incorrect authorization vulnerability in baseUserDelete that allows base Creators to remove higher-privileged Owners due to a broken role power comparison. Attackers holding the Creator role can remove other Owners of a base, since getProjectRolePower returns -1 for the target and never refuses the request.	2.7	More Details
CVE-2026-132988	In JetBrains TeamCity before 2026.2, 2026.1.4, 2025.11.8 the Cloud Profiles page disclosed cloud profiles of projects the user could not access	2.7	More Details
CVE-2026-132994	In JetBrains TeamCity before 2026.2 the REST API exposed details of users not visible in the UI	2.7	More Details
CVE-2026-130885	Tanium addressed an uncontrolled recursion vulnerability.	2.7	More Details
CVE-2026-131658	Tanium addressed an uncontrolled recursion vulnerability in Integrity Monitor.	2.5	More Details
CVE-2026-130664	ImageMagick before 7.1.2-35 contains a heap buffer overflow vulnerability that allows attackers to write beyond allocated heap memory by supplying crafted images. Attackers can provide a crafted set of images processed by the magick stream operation to corrupt heap memory and crash the application.	2.5	More Details
CVE-2026-130668	ImageMagick before 7.1.2-35 contains a double free vulnerability that allows attackers to trigger an invalid memory free by using the DPC compare metric. Attackers can supply images for comparison with the DPC metric to corrupt memory and crash the application, causing denial of service.	2.5	More Details
CVE-2026-130899	Tanium addressed a time-of-check time-of-use race condition vulnerability.	2.5	More Details
CVE-2026-133109	In JetBrains TeamCity before 2026.2.2, 2026.1.6 hTML injection in pipeline notification emails was possible via user and job names	2.4	More Details
CVE-2026-133482	A vulnerability has been found in hrshadhin CloudSchool up to 3.0.1. This affects the function EventController@store of the file app/Http/Controllers/Backend/EventController.php of the component Admin Content Management. Such manipulation of the argument data leads to cross site scripting. The attack can be launched 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.	2.4	More Details
CVE-2026-133107	In JetBrains TeamCity before 2026.2.2, 2026.1.6 hTML injection in pipeline notification emails was possible via commit messages	2.4	More Details
CVE-2026-18476	Mattermost versions 11.7.x <= 11.7.11 fail to scope an incoming webhook's direct message target to the webhook's own team, which allows a team administrator to deliver a direct message to a user on a team they have no authority over via a webhook whose posting identity belongs to both teams.. Mattermost Advisory ID: MMSA-2026-00751	2.2	More Details

CVE-2026-108319	Tanium addressed a time-of-check time-of-use race condition vulnerability in Tanium Server.	2.0	More Details
CVE-2026-111519	Dell Alienware Command Center (AWCC), versions prior to 6.14.105.0, contain an Improper Authorization in Handler for Custom URL Scheme vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Client-Side Request Forgery and Information Disclosure.	2.0	More Details
CVE-2026-99636	In the Linux kernel, the following vulnerability has been resolved: net: reject a negative TCP header offset in qdisc_pkt_len_segs_init() qdisc_pkt_len_segs_init() stores skb_transport_offset(), or the inner offset, in the unsigned hdr_len. A header in front of skb->data gives a negative offset. For an offset a little below zero, pskb_may_pull() gets a small length and succeeds. Then a 64-bit kernel reads the TCP header far after the skb and can crash. Further below zero pskb_may_pull() fails. Commit 28a57fb2c5df ("net: iptunnel: fix stale transport header during tunnel decapsulation") fixed one source of such offsets. Other code that removes bytes from the front of a packet can leave the same stale header. Return SKB_DROP_REASON_SKB_BAD_GSO for a negative offset in the TCP branch. SKB_GSO_UDP_L4 and the other GSO types read no header and still get a wrong qdisc_pkt_len(). On a 64-bit kernel only packets that cause the read after the skb today get a new result. Then __dev_queue_xmit() drops them and returns -EINVAL, but sch_handle_ingress() ignores the result. Their content can be valid. Thus the code that leaves the header stale needs its own fix.	N/A	More Details
CVE-2026-99628	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Avoid closing unowned DDC pins [WHY] dal_ddc_open() cleans up pins it acquires on later failures. If the data-pin open fails, this attempt acquired nothing. If opening the clock pin or configuring fails, dal_ddc_open() closes the pins it opened. An extra dal_ddc_close() is unnecessary and unsafe: GPIO_RESULT_ALREADY_OPENED can mean AUX owns the existing handles, and closing them here would disrupt AUX. [HOW] Remove the extra dal_ddc_close() from the failed-open path. (cherry picked from commit 175ae7772b68757b4932f34ff862923cab7abe21)	N/A	More Details
CVE-2026-99623	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Clamp the ATIF backlight curve to its 99 points amdgpu_atif_query_backlight_caps() stores the firmware's number_of_points as data_points but copies only the 99 points that luminance_data[] holds. number_of_points is a u8, so a firmware table that claims more makes the custom curve code in amdgpu_dm index past the array, and dm_query_extended_brightness_caps() copy past the end of the data_points[] buffer it fills Clamp the count to ATIF_QBTC_MAX_DATA_POINTS and warn, since a table that claims more is a firmware bug (cherry picked from commit af8a1c2b7d3de4960e10f94e8ba2fc0ccd168072)	N/A	More Details
CVE-2026-99594	In the Linux kernel, the following vulnerability has been resolved: arm64: kexec_file: Fix elf_headers memory leak in retry loop A retry allocates new elf headers and overwrites image->elf_headers, leaking the previous buffer. Free it before assigning. [catalin.marinas@arm.com: drop the redundant NULL check before vfree()]	N/A	More Details
CVE-2026-131800	A stored cross-site scripting vulnerability was identified in GitHub Enterprise Server that allowed an organization member with permission to manage custom-property definitions and edit property values to add an attacker-controlled SSH authentication key to the personal account of another member through HTML stored in a single-select custom-property option. Exploitation required another definition manager in the same organization to attempt to remove the poisoned option while it was in use and save the definition with a valid sudo session. Repository access through the added key depended on the permissions of the victim and independent authorization controls. This vulnerability affected supported GitHub Enterprise Server prior to 3.23 and was fixed in versions 3.19.16, 3.20.12, 3.21.10, and 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-99595	In the Linux kernel, the following vulnerability has been resolved: sched/fair: Clear throttle state when leaving the fair class Task-based CFS bandwidth throttling removes a task's fair sched_entity while the scheduler core still considers the task queued. p->throttled carries the fair-specific state across dequeue/enqueue operations. DEQUEUE_CLASS removes the task from cfs_rq->throttled_limbo_list but retains the bit. If the last finite CFS quota is disabled while the task runs in another class, it is absent from the unthrottle walk and the bit survives. Returning to fair while cfs_bandwidth_used() is false then performs a real fair enqueue with p->throttled still set. After bandwidth is re-enabled, a class-change dequeue treats the stale bit as limbo state and skips __dequeue_task(). p->se.on_rq therefore remains set, and put_prev_task_fair() places the entity in the EEVDF tree while the task moves to RT. On exit, the fair picker can select the TASK_DEAD task. __schedule() then returns to do_task_dead(), which hits BUG(). Keep p->throttled across migrations and group moves, but clear it on DEQUEUE_CLASS because fair-specific state must not outlive fair-class membership. A later return to fair uses the normal enqueue path. That path still respects a throttled cfs_rq. A KASAN, LOCKDEP and DEBUG_LIST v7.3 kernel reproduced the do_task_dead() BUG in 3/3 root and 3/3 delegated uid-1000 runs with CapPrm=CapEff=0 and RLIMIT_RTPRIO=1. The fix passed both matrices plus a 3/3 class-switch control with bandwidth continuously enabled, without a scheduler warning, KASAN report, lockdep splat or crash. An AI coding assistant analyzed the state mismatch and drafted the initial reproducer, fix candidate and this changelog. The source paths, trigger, controls and results were rechecked against the cited revisions.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: nvme: wait for SRCU readers before splicing the namespace list Readers walk ctrl->namespaces under ctrl->srcu, and some of them sleep there, e.g.		

CVE-2026-99596	nvme_wait_freeze() waiting for I/O to drain. nvme_remove_namespaces() moves the list onto its on-stack ns_list with list_splice_init_rcu(), which only waits for RCU readers before relinking the entries. A reader still on the last namespace then follows its next pointer to ns_list instead of ctrl->namespaces and treats the stack as a namespace. Deleting a controller while a Format NVM passthrough waits for a delayed read on the last namespace gives: BUG: KASAN: stack-out-of-bounds in nvme_wait_freeze+0x102/0x120 Read of size 8 at addr ff110001d2f3fc38 by task nvme/4413 Call Trace: nvme_wait_freeze nvme_passthru_start nvme_submit_user_cmd nvme_user_cmd nvme_dev_ioctl BUG: unable to handle page fault for address: ffffffff8af04850 Use list_splice_init_srcu() instead, which waits for the SRCU readers. This also takes the RCU grace period out of every controller deletion. Deleting 64 tcp controllers with one namespace each, median of 3 runs on each of two boots per kernel: before after nvme disconnect-all (ms) 3491 / 3509 2996 / 3165 nvme_remove_namespaces() avg (ms) 39.0 / 40.5 31.9 / 34.7	N/A	More Details
CVE-2026-99597	In the Linux kernel, the following vulnerability has been resolved: srcu: Add list_splice_init_srcu() list_splice_init_rcu() makes a list look empty to new readers, waits for the current ones with the given sync function, and only then links the entries into the new list. The sync function takes no argument, so it can wait for RCU readers but not for the readers of an srcu_struct. Passing synchronize_rcu() for a list read under SRCU still builds, but leaves SRCU readers on the entries while the last one's next pointer is moved onto the new list head. nvme_remove_namespaces() does this today. Add list_splice_init_srcu(), which waits for the readers of the given srcu_struct instead.	N/A	More Details
CVE-2026-99598	In the Linux kernel, the following vulnerability has been resolved: MIPS: sgi-ip32: Leave MACE PCI control alone in config accesses mace_pci_read_config() clears the master abort interrupt enable in the MACE PCI control register around each config read. A register read followed by two register writes can hang CRIME 1.1 under DMA load. Ack master aborts in the error register instead, and follow every write with a read so back-to-back accesses cannot form that sequence.	N/A	More Details
CVE-2026-99635	In the Linux kernel, the following vulnerability has been resolved: vhost/vsock: size receive SKBs from declared payload vhost_vsock_alloc_skb() sizes the skb from iov_length(), the total guest descriptor length. virtio_transport_inc_rx_pkt() instead adds hdr->len to rx_bytes and buf_used, and budgets each queued skb as SKB_TRUESIZE(0). The guest controls the descriptor length and hdr->len independently. A guest can therefore provide a 64 KiB descriptor while declaring a zero payload. The descriptor-sized skb remains queued, rx_bytes and buf_used remain unchanged, and only the fixed queue-entry budget limits retention. This is not a lifetime leak: the skb is freed when dequeued or at socket teardown. Repeating this across connections can nevertheless exhaust host memory. On Linux v7.2, 455 such descriptors retained 30,255,680 bytes for a 262,144-byte receive buffer; a full-payload control retained 265,984 bytes. Vhost previously read hdr->len and allocated pkt->buf from that value. The conversion from virtio_vsock_pkt to sk_buff instead allocated receive skbs from the descriptor length, introducing the allocation-accounting mismatch. Copy the header into a stack object and validate hdr->len before allocating the skb. Size the allocation as sizeof(hdr) + payload_len so descriptor capacity no longer determines queued memory. Check payload_len against len - sizeof(hdr) to avoid addition overflow on 32-bit hosts.	N/A	More Details
CVE-2026-99631	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: check amdgpu_vm_ready() in dma-buf move notify amdgpu_dma_buf_move_notify() touches the page tables without first calling amdgpu_vm_ready(). When a process is SIGKILLed, its VM entities are stopped. Teardown of an imported dma-buf then tries to clear PTEs by submitting to one of those stopped entities. That eventually hangs the task. Commit aa5fc4362fac ("drm/amdgpu: fix task hang from failed job submission during process kill") added the stopped entity check in amdgpu_vm_ready() and used it in amdgpu_cs_vm_handling(). That did not cover a call in amdgpu_dma_buf_move_notify(), what was already mentioned in: https://lists.freedesktop.org/archives/amd-gfx/2025-October/133076.html Tested by SIGKILLing a GPU workload holding imported dma-bufs. Without the patch "Trying to push to a killed entity" errors are logged, then a task hangs. With the patch there are no hung tasks. This was tested on a 6.18 kernel. (cherry picked from commit bee5c9d7a34cb1b510ad3b4c18c30f9d38b0ab7d)	N/A	More Details
CVE-2026-99633	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: keep the CS user fence BO in GTT The CS ioctl hands the GPU the GART address of the user fence BO, but since the BO keeps its own preferred and allowed domains, it no longer makes sure the BO is in GTT. A user fence in a VRAM BO, which only needs to be a page in size, gets its address from amdgpu_bo_gpu_offset(), which warns about non-contiguous VRAM. On navi31, a CS with a user fence in a 4K VRAM BO, or in one created with AMDGPU_GEM_CREATE_VM_ALWAYS_VALID: WARNING: amdgpu_object.c:1531 at amdgpu_bo_gpu_offset+0x86/0xb0 [amdgpu], CPU#11: cs_gem_checks/18393 RIP: 0010:amdgpu_bo_gpu_offset+0x86/0xb0 [amdgpu] Call Trace: <TASK> amdgpu_cs_parser_bos.isra.0+0x4eb/0x850 [amdgpu] amdgpu_cs_ioctl+0x97e/0x2510 [amdgpu] drm_ioctl_kernel+0xac/0x110 [drm] drm_ioctl+0x2c6/0x5b0 [drm] __x64_sys_ioctl+0xa0/0xf0 do_syscall_64+0xa3/0x4d0 entry_SYSCALL_64_after_hwframe+0x76/0x7e and BOs which can't be in GTT at all, like GWS or OA ones of a page in size, pass the checks as well. Only take BOs which can't be anywhere else than in GTT as user fence, which the CS then validates there along with its other BOs. While at it, detect userptrs with amdgpu_ttm_tt_is_userptr(); amdgpu_ttm_tt_get_usermm() returns NULL once the task which created the userptr is gone. v2: only take BOs which can only be in GTT instead of special casing the user fence BO in the validation, use amdgpu_ttm_tt_is_userptr() (Christian) (cherry picked from commit f16989def9561d04d0f2e77216f5319112d29bf1)	N/A	More Details

CVE-2026-99584	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to avoid folio->mapping dereference in is_{meta, node}_folio() Otherwise, it may cause NULL pointer dereference in race condition: Thread A Thread B - f2fs_setattr - f2fs_truncate - f2fs_truncate_blocks - f2fs_do_truncate_blocks - f2fs_truncate_inode_blocks - truncate_dnode - truncate_node - invalidate_mapping_pages - folio->mapping = NULL - folio_lock - is_node_folio - F2FS_F_SB(folio): dereference on folio->mapping	N/A	More Details
CVE-2026-99575	In the Linux kernel, the following vulnerability has been resolved: f2fs: return -EIO for a lost inode block after cp_error After cp_error, dirty node caches are dropped without being written. If such an inode was never written, its NAT entry is still NEW_ADDR, and reading its inode block returns -ENOENT. f2fs_get_dnode_of_data() passes that -ENOENT up, and callers take it as a hole. Since the inode block is missing, the next offset can only advance one block, so fiemap walks block by block up to maxbytes, about 2^30 blocks, while holding the inode lock. It took hours in a power-off recovery test after shutdown and blocked unlink and umount on the file. Return -EIO instead of -ENOENT when the inode block itself is missing after cp_error or shutdown. A missing direct or indirect node is still reported as -ENOENT.	N/A	More Details
CVE-2026-99576	In the Linux kernel, the following vulnerability has been resolved: f2fs: zone: fix to avoid bio split on sequential zone In the block layer, bio_split_to_limits() splits a bio if it exceeds any underlying queue limits: 1. queue_max_sectors: Splits bios larger than the transfer limit. This limit is set by the device driver/DMA controller (max_hw_sectors) and device-mapper targets. When flushing dirty pages (governed by VFS writeback_chunk_size() and f2fs nr_pages_to_skip()), f2fs merges contiguous pages into bios up to 512KB (1024 sectors) or more without checking queue_max_sectors, which will be split if the device or DM target enforces a smaller limit. 2. queue_max_segments: Splits bios exceeding scatter-gather segment limits. f2fs allocates bios with up to 256 pages, exceeding the 128-segment limit of DM targets (dm-default-key/dm-crypt) if pages are discontiguous. 3. chunk_sectors: Splits bios crossing chunk/zone boundaries. In f2fs, section_size strictly matches zone_size, and block allocation never spans across section boundaries, so bios never cross chunk_sectors. Crucially, when bio_split_to_limits() splits a bio, bio_submit_split() submits the tail half via submit_bio_noacct() before returning the head half to the caller. On a sequential zoned device, this causes the tail bio to arrive at the zone write plug before the head: dm-3: zone 7: prepare_bio wp mismatch: bio sector 30720 offset 2048 != wp_offset 248, len 248 The write plug rejects the out-of-order tail with BLK_STS_IOERR, causing f2fs_stop_checkpoint() and remounting the filesystem read-only: 2026-09-30T22:09:23.725678+08:00 localhost kernel: CPU: 9 UID: 0 PID: 243 Comm: kworker/9:1H Not tainted 6.12.81-gb3ef518fb26d #5 2026-09-30T22:09:23.725750+08:00 localhost kernel: Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.17.0-debian-1.17.0-1 04/01/2014 2026-09-30T22:09:23.725753+08:00 localhost kernel: Workqueue: dm-3_zwplugs blk_zone_wplug_bio_work 2026-09-30T22:09:23.725756+08:00 localhost kernel: Call Trace: 2026-09-30T22:09:23.725758+08:00 localhost kernel: <TASK> 2026-09-30T22:09:23.725758+08:00 localhost kernel: dump_stack_lvl+0x4d/0x70 2026-09-30T22:09:23.725759+08:00 localhost kernel: f2fs_stop_checkpoint+0x1df/0x5d0 2026-09-30T22:09:23.725759+08:00 localhost kernel: ? blk_zone_wplug_bio_work+0x4eb/0x740 2026-09-30T22:09:23.725760+08:00 localhost kernel: f2fs_cache_write_end_io+0x545/0x6e0 2026-09-30T22:09:23.725761+08:00 localhost kernel: blk_zone_wplug_bio_work+0x4eb/0x740 2026-09-30T22:09:23.725762+08:00 localhost kernel: process_one_work+0x5c6/0xde0 2026-09-30T22:09:23.725764+08:00 localhost kernel: ? assign_work+0x124/0x4e0 2026-09-30T22:09:23.725765+08:00 localhost kernel: worker_thread+0x403/0xb70 2026-09-30T22:09:23.725765+08:00 localhost kernel: ? __kthread_parkme+0x82/0x140 2026-09-30T22:09:23.725766+08:00 localhost kernel: ? __pfx_worker_thread+0x10/0x10 2026-09-30T22:09:23.725766+08:00 localhost kernel: kthread+0x235/0x2e0 2026-09-30T22:09:23.725767+08:00 localhost kernel: ? recalc_sigpending+0x128/0x1d0 2026-09-30T22:09:23.725767+08:00 localhost kernel: ? __pfx_kthread+0x10/0x10 2026-09-30T22:09:23.725769+08:00 localhost kernel: ret_from_fork+0x2f/0x70 2026-09-30T22:09:23.725769+08:00 localhost kernel: ? __pfx_kthread+0x10/0x10 2026-09-30T22:09:23.725770+08:00 localhost kernel: ret_from_fork_asm+0x1a/0x30 2026-09-30T22:09:23.725771+08:00 localhost kernel: </TASK> 2026-09-30T22:09:23.725771+08:00 localhost kernel: F2FS-fs (dm-0): Stopped filesystem due to reason: 3 To prevent write bios destined for sequential zones from being split: 1. In __bio_alloc(), bound npages to bdev_max_segments(bdev) for writes to sequential zone areas. 2. In page_is_mergeable(), check queue_max_segments() and queue_max_sectors() for sequential zone areas, so f2fs will not merge beyond the queue limits.	N/A	More Details
CVE-2026-99577	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: usb: bound what the driver feeds the after-DTIM queue Frames routed to the high queue are transmitted right after DTIM beacons only, inside the ATIM window, while they wait in the shared TX page pool. The driver puts no limit on how many it hands over, so as long as one station dozes, any sustained broadcast or multicast traffic outruns the drain and empties the pool: measured on an RTL8822BU AP with a single client in power save and ~40 frames/s of mDNS chatter, the free page count at 0x240 goes from 1803 to 16 in about 100 seconds and stays there for as long as the traffic lasts. From that point every other transmit queues behind the backlog, authentication responses arrive too late for anyone to join, and the reserved page download fails ("error beacon valid"). The AP keeps beaconing and only a reboot recovers. Feed the high queue through a small budget that refills below the measured drain rate (about 3 frames per DTIM, ~15/s at dtim_period 2 on the default 2 TU ATIM window); whatever exceeds the budget leaves on its access category queue right away. The high queue backlog is now bounded by the burst size under any load, so the page pool cannot run dry, at the price that a dozing station may miss part of a broadcast storm.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: usb: route bmc frames via the		

CVE-2026-99578	high queue only for DTIM delivery In AP mode every broadcast and multicast data frame is routed to TX_DESC_QSEL_HIGH, the after-DTIM queue, whether or not anybody is asleep. The firmware drains that queue at beacon pace, a dozen or so frames per second measured on RTL8822BU, while one associated client's mDNS/SSDP chatter alone exceeds that. The excess accumulates inside the chip until the shared TX page pool is exhausted (measured: 14 of 1803 pages left). From that point every host-sourced frame queues behind the backlog: authentication responses reach the air seconds after the client has given up, so no station can associate anymore, and the beacon reserved-page download fails the BCN_VALID poll ("error beacon valid") because it needs pages from the same pool. The AP keeps beaconing, so the failure looks like a silent RX stall and only a reboot recovers. mac80211 already decides when after-DTIM delivery is needed: it sets IEEE80211_TX_CTL_SEND_AFTER_DTIM on bmc frames only while at least one station is actually dozing. Honor that instead of routing unconditionally: flagged frames keep going through the high queue with the MORE_DATA and HGQMD handling introduced by commit 076f786a0ae1 ("wifi: rtw88: Fix AP mode incorrect DTIM behavior"), everything else leaves at line rate through the AC queues. This partially reverts the usb.c hunk of that commit, whose unconditional routing is what lets the backlog build up. On a bench AP (USB2, 20 MHz, WPA2, hostapd, a Windows client driven through disconnect/reconnect cycles): reconnects fail 0/5 on RTL8822BU and 0/3 on RTL8821CU before this change, and pass 10/10 and 5/5 with it, with the page pool staying healthy and no beacon errors logged.	N/A	More Details
CVE-2026-99579	In the Linux kernel, the following vulnerability has been resolved: mm/shmem: honour F_SEAL_EXEC when an ACL is set The owner of a memfd sealed with F_SEAL_EXEC can still change its execute bits by setting a POSIX access ACL. With CONFIG_TMPFS_POSIX_ACL, this lets an unprivileged owner make a non-executable memfd executable, bypassing vm.memfd_noexec=2. simple_set_acl() updates i_mode directly, bypassing the seal check in shmem_setattr(). Add shmem_set_acl() to check the resulting mode. Share the F_SEAL_EXEC check with shmem_setattr() using memfd_check_seals_mode() in include/linux/memfd.h. Fix generated by an LLM and sent by Chris Mason in private. [david: use common helper; simplify description]	N/A	More Details
CVE-2026-99581	In the Linux kernel, the following vulnerability has been resolved: platform/chrome: cros_usbpd_notify: Only use drvdata when parent is GOOG0004 cros_usbpd_notify_probe_acpi() takes the EC device pointer from dev_get_drvdata(dev->parent) and assumes it is a struct cros_ec_device. That only holds when GOOG0003 is a child of GOOG0004. On older devices without that hierarchy, such as Google Nami, GOOG0003 and GOOG0004 are siblings under the ACPI EC (PNP0C09): PNP0C09:00 (acpi-ec) |-- GOOG0003:00 (cros-usbpd-notify-acpi) `-- GOOG0004:00 (cros_ec_lpcs) Before commit db65a06d10b3 ("ACPI: EC: Convert the driver to a platform one") the parent platform device had no driver data, so the driver warned and continued without an EC pointer, as intended for these devices. Since that commit the ACPI EC driver stores its struct acpi_ec there, so the driver treats a struct acpi_ec as a struct cros_ec_device and calls cros_ec_cmd() on it for every USB-C PD host event. Plugging or unplugging a charger then intermittently oopses, followed by soft lockups and a hung system: BUG: unable to handle page fault for address: ffffffff9cf03220 #PF: supervisor write access in kernel mode Hardware name: Google Nami/Nami, BIOS 09/19/2019 Workqueue: kacpi_notify acpi_os_execute_deferred RIP: 0010:native_queued_spin_lock_slowpath+0x29d/0x330 Call Trace: _raw_spin_lock_irqsave+0x59/0x80 __mutex_lock.constprop.0+0x129/0x930 cros_ec_cmd_xfer+0x2a/0xf0 [cros_ec_proto] cros_ec_cmd_xfer_status+0x1a/0x90 [cros_ec_proto] cros_ec_cmd+0xb7/0x140 [cros_ec_proto] cros_usbpd_get_event_and_notify+0x42/0xc0 [cros_usbpd_notify] acpi_ev_notify_dispatch+0x4e/0x70 acpi_os_execute_deferred+0x1a/0x30 Only read the parent's driver data when the parent's ACPI node is GOOG0004, and otherwise continue without an EC pointer. While at it, drop the reference taken by fwnode_get_parent().	N/A	More Details
CVE-2026-99585	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Use kvm_s2_fault_vma_info in gmem_abort() gmem_abort() maps at s2fd->fault_ipa, which HPFAR_EL2 holds at 4K granularity whatever the page size. The generic page-table code aligns it, but pkvm_pgtable_stage2_map() looks up existing mappings over [addr, addr + size), so on a pKVM host with pages larger than 4K a guest_memfd-backed guest that faults past the first 4K of a page can find its neighbour's mapping, get -EAGAIN and take the same fault forever. The memory fault exit reports s2fd->fault_ipa too, which for a nested guest's fault is the nested IPA rather than the canonical one kvm_gmem_get_pfn() failed on. Take the addresses from kvm_s2_fault_vma_info instead, as user_mem_abort() does, and report the canonical gfn in the memory fault exit. kvm_s2_fault_get_vma_info() itself isn't called: a guest_memfd memslot's userspace_addr doesn't need to be backed by a VMA.	N/A	More Details
CVE-2026-99632	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: use the queue's process info for the WPTR BO eviction fence amdgpu_amdkfd_map_gtt_bo_to_gart() removes the eviction fence of bo->vm_bo->vm->process_info from the WPTR BO of a new queue. bo->vm_bo is the VM which opened the BO last, not necessarily the KFD process' VM, and a render node VM has no process info. On navi31, an unprivileged process which exports its WPTR BO with AMDKFD_IOC_EXPORT_DMABUF and imports it on a second render node fd oopses in AMDKFD_IOC_CREATE_QUEUE: BUG: kernel NULL pointer dereference, address: 0000000000000060 RIP: 0010:amdgpu_amdkfd_map_gtt_bo_to_gart+0x86/0x1a0 [amdgpu] init_user_queue.part.0+0x180/0x250 [amdgpu] pqm_create_queue+0x32d/0x730 [amdgpu] kfd_ioctl_create_queue+0x583/0x9f0 [amdgpu] Pass the process info of the queue instead. With the fix, the queue is created and destroyed normally. (cherry picked from commit 6b4157498a176bd1c6083f6bb5aaf8a257a2027e)	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ata: libata-scsi: reject data-in commands		

CVE-2026-99587	with a data-out buffer SG_IO can submit a data-in SCSI command with DMA_TO_DEVICE. The scatterlist can then contain pages pinned without write access. ata_scsi_rbuf_fill() writes simulated replies to that scatterlist; in the reported case, the copy tried to write the shared zero page and faulted in the kernel. The same mismatch can affect SG_IO commands translated by libata. ZBC IN and GET PHYSICAL ELEMENT STATUS completion handlers can modify the receive buffer even if the drive aborts the command. For ATAPI, the successful INQUIRY completion also rewrites the buffer in atapi_fixup_inquiry(). Reject these mismatched commands before the buffer is written or the command is issued. Report ABORTED COMMAND / DATA IN BUFFER ERROR because the CDB itself can be valid; the supplied transfer direction is wrong.	N/A	More Details
CVE-2026-99588	In the Linux kernel, the following vulnerability has been resolved: batman-adv: tt: only update flags for TT events In the first implementation of the translation table events, it was only possible to announce: * ADD: new locally detected client * DEL: disappeared local client If an ADD event was queued and then a DEL event was queued in the same timeslot, it is valid to destroy both events because nothing actually changed for remode nodes But with commit 3c4f7ab60cc2 ("batman-adv: improve the TT component to support runtime flag changes") and subsequent changes, it could also happen that there are extra ADD events: * ADD: _NEW_ locally detected client * ADD: already existing client with new flags * DEL: disappeared local client The first and the last type of event could in theory still be removed together from the (to be announced) events queue. But if the "update flags" ADD is dropped together with the DEL, it would still leave global translation table of a remote node in the incorrect state - believing that the client didn't disappear. If a detected client was previously announced, a DEL event must only overwrite the "update flags" ADD but the DEL itself must still continue to exist in the (to be anounced) events queue. For now, adjustment of the announced flags is used for everything instead of dropping already queued up events completely. Similarly, a DEL followed by an ADD caused both entries to be deleted. But the new ADD could announce flags changes which were not yet present for the client before it was marked for DELetion. The new ADD must therefore only overwrite the DEL - but still be kept.	N/A	More Details
CVE-2026-99593	In the Linux kernel, the following vulnerability has been resolved: arm64: sdei: Preserve the interrupted context's exception registers An SDEI event is dispatched by firmware regardless of PSTATE.DAIF, so it can interrupt the kernel where the exception registers are live: - exception entry, before kernel_entry has saved ELR_EL1/SPSR_EL1 and before the C handler has read ESR_EL1/FAR_EL1; - exception return, after kernel_exit has loaded ELR_EL1/SPSR_EL1 and before the ERET. An IRQ or pseudo-NMI cannot arrive there with PSTATE.I set. Firmware saves only x0-x17, PC and PSTATE around an SDEI event, so a synchronous exception taken by the handler overwrites the interrupted context's ELR, SPSR, ESR and FAR. do_sdei_event() only notices that ELR changed, and warns. The interrupted context then carries on with the handler's state. A task interrupted on its way into the kernel can be sent SIGILL for the handler's kernel data abort. A task interrupted on its way back to user space ERETs to EL1, at the handler's fixup address, with user registers loaded. With SDEI delivering NMI backtraces, this happens in normal operation. The handler runs show_regs(), which reads through possibly-NULL pointers with the _nofault() accessors. print_worker_info() faults this way for a kworker that is not running a work item, and a single sysrq-l warned once per such kworker. Save ELR, SPSR, ESR and FAR in __sdei_handler() before irqentry_nmi_enter(), which can take an exception too, and restore them after irqentry_nmi_exit(). Once they are restored, a nested exception is harmless if its handler is safe in NMI context. That holds for EL1 data aborts on translation, permission and access-flag faults, which never take the mmap lock in NMI context, and for BRK. Warn once, with the syndrome, about any other exception class.	N/A	More Details
CVE-2026-99629	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: clear RB_OVERFLOW bit when enabling interrupts for navi10_ih Port this change to navi10_ih.c: commit afbf7955ff01 ("drm/amdgpu: clear RB_OVERFLOW bit when enabling interrupts") Setting IH_RB_WPTR to 0 does not clear RB_OVERFLOW if RB_ENABLE is not set. Pulse WPTR_OVERFLOW_CLEAR after RB_ENABLE (and with WPTR_OVERFLOW_ENABLE) so overflow is cleared when interrupts are enabled. (cherry picked from commit 713fb4ab04f4d2af6120f8eb607c5673b2759b38)	N/A	More Details
CVE-2026-126965	ATEN eco DC chartDefine SQL Injection Information Disclosure Vulnerability. This vulnerability allows remote attackers to disclose sensitive information on affected installations of ATEN eco DC. Authentication is required to exploit this vulnerability. The specific flaw exists within the handling of color references when creating charts. The issue results from the lack of proper validation of a user-supplied string before using it to construct SQL queries. An attacker can leverage this vulnerability to disclose stored credentials, leading to further compromise. Was ZDI-CAN-29124.	N/A	More Details
CVE-2026-99714	In the Linux kernel, the following vulnerability has been resolved: nvme: Do not drop the cdev reference twice when cdev_device_add() fails Since commit 9952f3882ecb ("nvme: fix cdev lifetime"), nvme_add_ns_cdev() and nvme_add_ns_head_cdev() take a reference to the namespace (or namespace head) before calling nvme_cdev_add(), to be dropped by the char device's release, nvme_cdev_rel(), and drop it themselves when nvme_cdev_add() fails. But when cdev_device_add() fails, nvme_cdev_add() has already put the device, and nvme_cdev_rel() has already dropped that reference: the caller drops it a second time. With one allocation failure injected into cdev_device_add() for a namespace on an NVMe over loop controller, deleting the controller gives, with native multipath: BUG: KASAN: slab-use-after-free in nvme_put_ns_head+0x21/0x140 Write of size 4 at addr ff1100000a4ea078 by task init/1 nvme_put_ns_head nvme_put_ns nvme_remove_namespaces nvme_do_delete_ctrl nvme_delete_ctrl_sync nvme_sysfs_delete Freed by task 1: kfree nvme_ns_remove nvme_remove_namespaces refcount_t: underflow; use-after-free. Take the reference in nvme_cdev_add() instead, right before cdev_device_add(), and let the callers neither take nor drop it. Add nvme_is_ns_cdev() for the check that nvme_cdev_rel() and nvme_cdev_add() now share.	N/A	More Details

CVE-2026-99599	In the Linux kernel, the following vulnerability has been resolved: udf: don't use eloc after the end of the unallocated space table When udf_next_aext() returns 0, its eloc, elen and etype outputs hold no defined value. Two walks of the unallocated space table still use eloc after the walk has reached the end of the table. udf_table_prealloc_blocks() looks for a free extent starting at first_block and decides whether it found one by comparing eloc with first_block after the walk. If the table's list ends in a continuation descriptor that leads to an empty allocation extent, eloc is then the block of that allocation extent, which is in use, and if that is first_block, the block is handed out as preallocation. Once udf_discard_prealloc() no longer frees the wrong block, this is what still makes fsx in generic/091 and generic/263 read back bad data on a table, in about one run in four. Only take the extent if the walk stopped at it. When udf_table_free_blocks() cannot merge the freed blocks into an existing extent, it appends a new one and sets only its block number, so the partition reference comes from whatever eloc held after the walk, and with long allocation descriptors it is written to disk. For a table without extents, eloc was never set. Take the partition reference from the freed blocks instead. Short allocation descriptors do not store the partition reference, which is why this half remains a latent long-descriptor fix rather than a reproduced failure.	N/A	More Details
CVE-2026-99637	In the Linux kernel, the following vulnerability has been resolved: ovl: inherit FMODE_NONOTIFY when opening directory realfiles When fanotify hands a listener an fd for a permission event on an overlay directory, ovl_dir_open_realfile() opens the real directory via dentry_open(), not backing_file_open(), so the lower open could still raise FS_OPEN_PERM and deadlock a listener that also marks the backing filesystem. Export dentry_open_nonotify() and call it from ovl_path_open_realfile() when the overlay user file must not generate events. Ordinary ovl_path_open() callers are unchanged.	N/A	More Details
CVE-2026-99600	In the Linux kernel, the following vulnerability has been resolved: udf: fix preallocation discarding before an empty allocation extent udf_discard_prealloc() finds the last extent of a file by peeking at the next one with udf_next_aext() and stopping when that returns 0. The peek reads into eloc and elen, and the code then uses them as the last extent. But when udf_next_aext() returns 0 they hold no defined value: if the extent list ends in a continuation descriptor that leads to an empty allocation extent, they describe that continuation, i.e. the block of the empty allocation extent, one block long. The kernel creates such lists itself: when udf_delete_aext() removes the only descriptor of the last allocation extent in a list, it leaves that allocation extent empty and the continuation to it in place. If the last extent is a preallocation, udf_discard_prealloc() deletes it with udf_delete_aext(), which finds the following allocation extent empty, so it also removes the continuation and frees that allocation extent's block. udf_discard_prealloc() then frees eloc and elen: that block a second time, instead of the preallocated blocks. On a space bitmap the preallocated blocks leak and the free block count drifts. On an unallocated space table the second free adds a second free extent for the same block, which is later handed out twice: fsx as run by generic/091 and generic/263 ends up with two parts of its test file in one block and reads back bad data. Peek into separate variables, so that eloc and elen keep the last extent that udf_next_aext() returned with 1.	N/A	More Details
CVE-2026-99622	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: drop the dma-buf reference when a KFD import fails import_obj_create() takes a reference on the dma-buf for the new kgd_mem, but the error path after validating the BO doesn't drop it. Since the validation reserves the BO interruptibly, an import failing with -ERESTARTSYS leaks the exported BO with its memory, and keeps the module in use. Importing a 64 MiB VRAM BO in a loop while getting SIGUSR1 every 20 us and another thread changes the BO's mappings: INFO: 171821 imports, VRAM used before 65 MiB, after 129 MiB modprobe: FATAL: Module amdgpu is in use. Drop the reference, and free the sync object and the mutex of the kgd_mem as well. With that VRAM usage returns to 65 MiB after the same test. (cherry picked from commit b4d1eca2729840148ff2ea957724ce764ca5a94e)	N/A	More Details
CVE-2026-99614	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/nvkm: clear use.enabled when subdev init fails On Kepler boards without the video decode falcon firmware installed (msvld/mspdec/msppp), every other VA-API decode attempt binds to an engine with no ucode loaded. The submitted work never completes, which leads to a context switch timeout, a channel recovery that cannot reset the engine, and eviction waits on fences that never signal. The process ends up stuck in D state, or killed by SIGBUS when it touches a buffer evicted after the wait times out: msvld: unable to load firmware data msvld: init failed, -19 fifo: SCHED_ERROR 0a [CTXSW_TIMEOUT] vo[5444]: channel 14 killed! WARNING: drivers/gpu/drm/nouveau/nvkm/engine/fifo/runl.c:117 at nvkm_runl_rc+0x2a4/0x3b0 [nouveau] [TTM] Buffer eviction failed WARNING: drivers/gpu/drm/nouveau/nouveau_bo.c:1092 at nouveau_bo_move_ntfy.part.0+0x9b/0x100 [nouveau] Tracing nvkm_subdev_ref() and nvkm_subdev_init_() shows the pattern: a reference taken with enabled=0 runs nvkm_falcon_init(), which returns -ENODEV; the next reference sees enabled=1 and succeeds without running it. Clear use.enabled when init() fails, so that every later reference retries init and fails cleanly. The other readers of use.enabled are fini() handlers that use it to skip teardown, which they now also do after a failed init(). On a GK107 (Quadro K1100M) without video firmware, playing H.264 with mpv --hwdec=auto: without this patch the GPU channel is killed within a minute and mpv hangs; with it, VA-API decoder creation fails, mpv falls back to software decoding, and no faults are logged.	N/A	More Details
CVE-2026-126964	ATEN eco DC fwDeletes.json Directory Traversal Arbitrary File Deletion Vulnerability. This vulnerability allows remote attackers to delete arbitrary files on affected installations of ATEN eco DC. Authentication is required to exploit this vulnerability. The specific flaw exists within the handling of JSON request data supplied to the fwDeletes.json endpoint. The issue results from the lack of proper validation of a user-supplied path prior to	N/A	More Details

	using it in file operations. An attacker can leverage this vulnerability to delete files in the context of SYSTEM. Was ZDI-CAN-28621.		
CVE-2026-99706	In the Linux kernel, the following vulnerability has been resolved: mm: fix address wraparound when scanning large folio PTEs numa_rebuild_large_mapping() and do_swap_page() can be tricked into generating an address wraparound, crashing the kernel. All that's required is to map a tail page of a large folio close enough to the start/end of the address space to e.g., make unsigned long addr_start = addr - (nr << PAGE_SHIFT); to wrap around. Let's fix it by providing a better helper, folio_scan_range(), that takes care of staying within VMA and page table boundaries while handling possible wraparounds. There are some nasty details, like not being able to use pmd_addr_end() in this context. Add some helpful comments for these cases.	N/A	More Details
CVE-2026-99615	In the Linux kernel, the following vulnerability has been resolved: l2tp: only return sessions of the given tunnel in l2tp_session_get() Commit 5f77c18ea556 ("l2tp: use IDR for all session lookups") made pppol2tp_connect() and the netlink session commands find sessions with l2tp_session_get(), passing the socket of the tunnel they name. L2TPv3 session IDs are unique per network namespace, except that sessions carried over UDP are only unique per UDP socket, the socket of their tunnel. However, l2tp_v3_session_get() only checks the socket when another session has the same session ID. Otherwise it returns the session, whichever tunnel it belongs to. As of now, a session command that names the wrong tunnel acts on a session of another tunnel, and a PPPoL2TP socket for a second L2TPv3 tunnel fails with EEXIST, because it finds session 0 of the first tunnel. The issue can be reproduced with: ip l2tp add tunnel tunnel_id 3 peer_tunnel_id 3 encap udp \ local 127.0.0.1 remote 127.0.0.1 udp_sport 5003 udp_dport 6003 ip l2tp add tunnel tunnel_id 4 peer_tunnel_id 4 encap udp \ local 127.0.0.1 remote 127.0.0.1 udp_sport 5004 udp_dport 6004 ip l2tp add session tunnel_id 3 session_id 7 peer_session_id 7 ip l2tp del session tunnel_id 4 session_id 7 ip l2tp show session Before that commit, the delete fails with "No such device", and session 7 of tunnel 3 is still listed. After that commit, the delete succeeds, and session 7 of tunnel 3 is gone. This patch will make l2tp_session_get() return NULL when the L2TPv3 session it finds does not belong to the given socket. The receive path calls l2tp_v3_session_get() directly and is not changed.	N/A	More Details
CVE-2026-99573	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: vgic: Stop the target vcpu when disabling LPIs Disabling LPIs is pretty nasty, as it directly messes with the AP list of the affected CPU, which could be running... This has the potential to lead to really bad behaviours, and we should not be doing that. Take example on the way the Active state is handled and simply pause the the vcpu so that we are sure it is in a quiescent state, and the state be safely manipulated. Nobody has any expectation of performance for this anyway.	N/A	More Details
CVE-2026-99705	In the Linux kernel, the following vulnerability has been resolved: mm: memory: fix truncation removing mapped folios Ayush reported an issue that shmem/tmpfs folios are still mapped when truncating them on production: " BUG: Bad page cache in process ... pfn:... page dumped because: still mapped when deleted " I reproduced the issue with Ayush's reproducer [1] with the latest mm-unstable branch, and I got the following crash info. From the dump message, we can see that truncate_inode_folio() is really trying to remove mapped folios, which is incorrect. " [189.171977] page: refcount:3 mapcount:1 mapping:00000000fe168ee1 index:0x3880 pfn:0x191647 [189.171995] memcg:ffff0000cc073d40 [189.171997] aops:shmem_aops ino:3c01 dentry name(?):"memfd:runsc-memory" [189.172006] flags: 0x17fffef0002022d(locked|referenced|uptodate|lru| workingset|swapbacked|node=0|zone=2|lastcpupid=0x3ffff) ..... [189.172017] page dumped because: VM_BUG_ON_FOLIO(folio_mapped(folio)) [189.172026] ----------[cut here]---------- [189.172027] kernel BUG at mm/filemap.c:155! [189.172057] Internal error: Oops - BUG: 00000000f2000800 [#1] SMP ..... [189.177721] CPU: 3 UID: 0 PID: 5521 Comm: repro Kdump: loaded Tainted: 7.3.0-rc4+ #277 PREEMPT(lazy) ..... [189.183986] Call trace: [189.184124] filemap_unaccount_folio+0xf0/0x1e8 (P) [189.184391] __filemap_remove_folio+0x34/0x160 [189.184633] filemap_remove_folio+0x4c/0xb0 [189.184859] truncate_inode_folio+0x34/0x58 [189.185087] shmem_undo_range+0x220/0x658 [189.185308] shmem_fallocate+0x2f0/0x470 [189.185527] vfs_fallocate+0x128/0x328 ..... [189.187078] el0t_64_sync+0x184/0x188 " After analysis, I believe the race exists between truncation and MADV_DONTNEED, and shmem's ->map_pages() merely makes the issue easier to reproduce. Since MADV_DONTNEED synchronously releases the pagetable page before calling tlb_flush_rmaps(), this could cause another thread's truncation to skip zap_pte_range() (due to pmd is none) but still observe the folio's mapcount as non-zero. A possible race scenario is as follows: CPU 0 CPU 1 shmem_fallocate madvise_dontneed_single_vma unmap_mapping_range .... (filemap_map_pages() remap folios) zap_pte_range truncate_inode_folio zap_empty_pte_table(pmd clear) unmap_mapping_folio ..... zap_pmd_range (saw pmd none, skip) filemap_remove_folio BUG_ON(folio_mapped) tlb_flush_rmaps(remove rmaps) To fix this issue, we should move the pmd clear operation (via zap_empty_pte_table()) to after tlb_flush_rmaps() and use sync_with_folio_pmd_zap() to sync with concurrent zapping, so that when we find a pmd_none() while unmapping a folio without holding the PTL, the folio's mappings are guaranteed to have been removed via tlb_flush_rmaps(). Meanwhile move the second pmd_none() check to where it belongs: if splitting zapped the PMD.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: mm/page-writeback: avoid 32-bit overflow in dirty threshold shares The ratios used by __wb_calc_thresh() are scaled to parts per million, but the threshold and ratio operands are unsigned long. On 32-bit kernels, their multiplication can wrap before the result is assigned to u64. For example, with a threshold of 25600 pages and the default maximum ratio of	N/A	More Details

99704	1000000, wb_max_thresh becomes 4125 pages instead of 25600. The shared threshold and minimum share calculations can overflow in the same way, producing incorrect per-writeback dirty limits. Use a u64 threshold for the three ratio multiplications and div_u64() to retain the existing division order and rounding.		
CVE-2026-99703	In the Linux kernel, the following vulnerability has been resolved: mm/hugetlb: fix mremap deadlock when PMD sharing is possible move_hugetlb_page_tables() holds the mapping's i_mmap_rwsem for write while moving hugetlb PTEs. This prevents races with truncation and protects huge_pmd_unshare(), but the move loop must not call helpers that can acquire the same rwsem for read. huge_pte_alloc() can do that when the destination can use PMD sharing: it may call huge_pmd_share(), which takes i_mmap_rwsem for read while searching for a shareable PMD table. Calling it while the move path already holds i_mmap_rwsem for write can deadlock. Avoid that path for PMD-sized hugetlb mappings that can use PMD sharing. Instead of allocating or sharing the destination page table while moving, zap the source range and let the destination fault path populate or share the PMD lazily. Reuse the hugetlb unmap logic for the source-side zap by splitting the core loop out of __unmap_hugepage_range(). The wrapper keeps the existing tlb_start_vma(), tlb_end_vma() and huge_pmd_unshare_flush() behavior, while move_hugetlb_page_tables() can reuse the unmap loop with its existing mmu_gather.	N/A	More Details
CVE-2026-99617	In the Linux kernel, the following vulnerability has been resolved: arm64: mm: validate /dev/mem mmap ranges without overflow An arm64 syzkaller reproducer creates a FUSE character device with major:minor 1:1, so opening it uses the /dev/mem file operations. It maps one page from offset 0xfffffffffffff000 at user address zero. The next openat() with a NULL pathname triggers a synchronous external abort (SEA, ESR 0x96000010) in strncpy_from_user() via do_getname(), panicking the kernel. valid_mmap_phys_addr_range() checks the exclusive end by adding the mapping size to the starting address. For this offset, the sum wraps to zero and passes the PHYS_MASK test, allowing an out-of-range PFN to reach remap_pfn_range(). The inclusive-end overflow check in mmap_mem_prepare() does not catch it because the end is ULONG_MAX. Reject starting PFNs outside PHYS_MASK and compare the inclusive length against the remaining physical range without overflowing. This also allows mappings that end exactly at PHYS_MASK, which the old check incorrectly rejects. [catalin.marinas@arm.com: changed the check to look more like __access_ok()]	N/A	More Details
CVE-2026-99702	In the Linux kernel, the following vulnerability has been resolved: mm/vma: write lock source VMA before duplicating it in __split_vma() __split_vma() duplicates the source VMA before taking its VMA write lock, and until the VMA is write-locked, a concurrent thread might set VMA_MAYBE_GUARD_BIT, resulting in the duplication missing the flag. If this happens, then fork may fail to duplicate guard markers and, if a file-backed mapping, THP collapse may also discard the markers. Resolve this by taking the VMA write lock before vm_area_dup() and, since all callers of vm_area_dup() now hold the VMA write lock, assert that this is the case there, and add a comment to explain why it's needed. [ljs@kernel.org: add VMA write assert in vm_area_dup(), commit message touchup]	N/A	More Details
CVE-2026-99701	In the Linux kernel, the following vulnerability has been resolved: mm/ksm: fix userfaultfd protection during merge and unmerge KSM loses userfaultfd protection when replace_page() builds a new PTE from vma->vm_page_prot. Dropping its uffd bit lets a subsequent write go through ordinary COW instead of userfaultfd write protection handling. For RWP, a merged page also becomes readable without faulting. Keeping the uffd bit and restoring PAGE_NONE fixes merging, but we also have to handle unmerging. An existing KSM page can be protected with RWP before MADV_UNMERGEABLE, so unmerging already needs fixing independently of preserving RWP during merging. Unsharing replaces a shared page with a private one and should leave userfaultfd protection in place, as do_wp_page() already does for write protection. Currently, an unshare fault reaches do_uffd_rwp() before it can proceed to do_wp_page(). - With synchronous RWP, handle_userfault() returns VM_FAULT_SIGBUS. break_ksm() then reports success with a KSM page still merged. - With asynchronous RWP, protection is cleared before unsharing. Preserve the uffd bit when merging, and restore PAGE_NONE for RWP. Let FAULT_FLAG_UNSHARE skip RWP fault handling and keep PAGE_NONE on each private copy as well. Unsharing can then complete with RWP still armed for a subsequent access.	N/A	More Details
CVE-2026-99619	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: don't drop the gang dependency fence twice drm_sched_job_add_dependency() drops the fence reference when it fails, but amdgpu_cs_submit() drops it again, so the scheduled fence of a gang member loses a reference it still needs. A GFX+COMPUTE gang waiting for another context, with the xarray node allocation for the leader's dependency failed through failslab: WARNING: drivers/dma-buf/dma-fence.c:589 at dma_fence_release+0x34d/0x730, CPU#1: amd_uapi_regres/2993 drm_sched_job_cleanup+0x2c2/0x340 [gpu_sched] amdgpu_job_free+0xcb/0x3e0 [amdgpu] amdgpu_cs_parser_fini+0x435/0x710 [amdgpu] amdgpu_cs_ioctl+0xcd2/0x7a60 [amdgpu] Oops: general protection fault, probably for non-canonical address 0xdffffc0000000001: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f] RIP: 0010:dma_fence_signal_timestamp_locked+0x1b8/0x610 (cherry picked from commit 5aa070a83c8378810b906cc8bfc650ce69822a7c)	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm/hugetlb_vmemmap: map the shared tail page writable page_ref_add_unless_zero() reads the refcount and then updates it with a cmpxchg. A lockless get, such as a page cache lookup or a PFN scanner, can read a non-zero refcount from a page that is then freed and reused as a hugetlb tail page. If HVO remaps its struct page to the shared tail page in between, the cmpxchg writes to read-only vmemmap: CPU0 (lockless get) CPU1 c = atomic_read(&page->_refcount) page freed page allocated as a hugetlb tail HVO remaps the struct page RO try_cmpxchg(&page->_refcount, &c, c + 1) -> write fault on read-only vmemmap The refcount is zero by then and the compare		

CVE-2026-99700	fails, but a compare-and-swap instruction needs write access whether or not it succeeds. The kernel oopses when hugetlb pages are allocated at runtime with HVO enabled while a lockless get races with it. Commit bd225530a4c7 ("mm/hugetlb_vmemmap: fix race with speculative PFN walkers") closed the window with RCU on both sides. The RCU was removed together with fake heads, on the wrong belief that it was only needed for them. Map the shared tail page writable instead of restoring the RCU. Every struct page on it is a tail with a zero refcount, so a late cmpxchg fails and leaves the page unchanged. RCU would put rcu_read_lock() back around every get and synchronize_rcu() back into hugetlb allocation and freeing. The cost is the protection against stray writes to tail struct pages. The shared page is shared by all hugetlb folios of the same order in the zone, so such a write now corrupts all of them instead of oopsing. Bootmem folios optimized when the memmap is set up keep their tails read-only. Nothing does a lockless get on them before hugetlb takes them over.	N/A	More Details
CVE-2026-99698	In the Linux kernel, the following vulnerability has been resolved: mm/cma: accept memory of CMA areas reserved at a fixed address CMA areas allocated from memblock have their memory accepted by memblock_alloc_range_nid(), but areas reserved at a fixed address with memblock_reserve() are not. If a fixed CMA area is requested by the cma=<size>@<base> command line parameter, it may be in regular RAM and on TDX and SEV-SNP guests it must be accepted. Accept the memory of fixed CMA areas as well, so that all CMA areas are accepted regardless of how they were reserved.	N/A	More Details
CVE-2026-99699	In the Linux kernel, the following vulnerability has been resolved: mm/cma: accept memory of multi-range CMA areas Patch series "mm/cma: accept memory of CMA areas reserved with memblock_reserve()", v1. Pages of allocated CMA areas are released to the buddy allocator with __free_pages(), which does not check for unaccepted memory. This usually works because memblock accepts the memory it allocates. However, multi-range CMA areas, used by hugetlb_cma, and fixed CMA areas requested with cma=<size>@<base> are reserved with memblock_reserve() and their memory is never accepted. On TDX and SEV-SNP guests the first access to such a page is fatal. To ease the backporting, each case is addressed by a separate patch: - patch 1 accepts the memory of multi-range CMA areas - patch 2 does the same for the fixed CMA areas The second patch won't apply cleanly to 6.6 and 6.12, but I believe the stable team LLMs will take care of that :) This patch (of 2): On systems with unaccepted memory, memblock allocations accept the allocated range, but when cma_declare_contiguous_multi() uses open coded allocation of multiple ranges, it reserves them with memblock_reserve() and they are never accepted. These pages are released to the buddy allocator by cma_activate_area() with __free_pages() that does not check for unaccepted memory, so the first access to such a page in a TDX or SEV-SNP guest is fatal. Accept the memory of each range when it is reserved.	N/A	More Details
CVE-2026-99620	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: don't walk the CS chunks when allocating them failed amdgpu_cs_pass1() sets p->nchunks before allocating p->chunks, and the allocation failure path skips resetting it. amdgpu_cs_parser_fini() then frees the kdata of nchunks chunks from a NULL array. On navi48 with the chunk array allocation failed through failslab: Oops: general protection fault, probably for non-canonical address 0xdfffffc0000000001: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f] RIP: 0010:amdgpu_cs_parser_fini+0x35d/0x710 [amdgpu] amdgpu_cs_ioctl+0xcae/0x7a20 [amdgpu] drm_ioctl_kernel+0x180/0x320 [drm] Set nchunks once the array exists. (cherry picked from commit 0c828445757b86aaa3ed2638832d76f450022bdf)	N/A	More Details
CVE-2026-99621	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: unlock the VM and BO when attaching a KFD BO fails kfd_mem_attach() locks the page directory and the BO with a drm_exec transaction on its stack, but leaves through the unwind label without drm_exec_fini() when locking fails, e.g. when MAP_MEMORY_TO_GPU gets a signal while the page directory or the BO is busy. Reservations already acquired then remain locked with their ww_acquire_ctx pointing into the expired stack frame. With one thread doing GEM_VA map/unmap on the KFD VM and another one mapping KFD BOs while getting SIGUSR1 every 20 us: BUG: KASAN: stack-out-of-bounds in __ww_mutex_lock.constprop.0+0x34f5/0x3da0 ww_mutex_lock_interruptible+0x2f/0x150 drm_exec_lock_obj+0x107/0x2d0 [drm_exec] kfd_mem_attach+0x4dd/0x1b60 [amdgpu] and the process ends up in D state. Finish the transaction before unwinding. (cherry picked from commit 7b885fa27e4b9fa5cbdc64e66fc35872aee296f5)	N/A	More Details
CVE-2026-99624	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Report MES ring wait timeouts. Return -ETIMEDOUT when the fence wait before submission expires instead of reporting success without submitting the packet. (cherry picked from commit 34341594a87b4a3f28be41e3e70a4787f02b3d9e)	N/A	More Details
CVE-2026-99708	In the Linux kernel, the following vulnerability has been resolved: ipv6: reject wrapped temporary address preferred lifetimes ipv6_create_tempaddr() raises a temporary address's preferred lifetime to regen_advance + age + 1 when the configured lifetime is too short. The result is stored in a u32, while regen_advance and age are unsigned long. If the sum wraps or is truncated, the stored lifetime can still be less than or equal to regen_advance + age. addrconf_verify_rtnl() then sees the new address as due for regeneration and creates another one. Temporary addresses are not limited by max_addresses since commit 4b08a8f1bd8c ("ipv6: remove max_addresses check from ipv6_create_tempaddr"), so this can cause unbounded address growth. With dad_transmits set to -1, a baseline kernel reached 2859 loopback addresses in two seconds. After storing the adjusted lifetime, verify that it is still greater than regen_advance + age. Reject the new address if the check fails. On an x86-64 KASAN kernel, the normal configuration created three IPv6 addresses. The negative dad_transmits, positive multiplication overflow, and large retransmission-time cases each left only the loopback and public addresses. An i386 defconfig and an x86-64 allmodconfig both built addrconf.o with	N/A	More Details

	W=1.		
CVE-2026-99716	In the Linux kernel, the following vulnerability has been resolved: hwmon: (coretemp) Set nr_cores only after core_data allocation succeeds In init_temp_data(), pdata->nr_cores is populated from topology_num_cores_per_package() before allocating pdata->core_data. If that allocation fails, init_temp_data() returns NULL with pdata->core_data still NULL and pdata->nr_cores set to a positive value. coretemp_add_core() logs the error without failing coretemp_cpu_online(), so the CPU still comes online with pdata->hwmon_dev registered. When the CPU is later taken offline (or the module is unloaded), coretemp_cpu_offline() calls get_temp_data(pd, cpu), which loops up to pdata->nr_cores and dereferences the NULL pdata->core_data pointer. Store the core count in a local variable and assign pdata->nr_cores only after pdata->core_data is allocated, preserving the invariant that pdata->nr_cores is zero whenever pdata->core_data is NULL.	N/A	More Details
CVE-2026-99607	In the Linux kernel, the following vulnerability has been resolved: ext2: preserve IOMAP_F_DIRTY for extending direct writes ext2_iomap_begin() sets IOMAP_F_DIRTY for writes that may extend i_size, but clears the flag after ext2_get_blocks(). When the write uses an existing mapped block, iomap can then use FUA and skip generic_write_sync(), even though ext2_dio_write_end_io() updates the inode size. An O_DSYNC write can therefore return without persisting the new size. For example, on a FUA-capable device, a 512-byte O_DIRECT|O_DSYNC write at offset 512 into a fsynced one-byte file returns successfully and extends the in-memory size to 1024, while an immediate read of the on-disk inode still shows size 1. An explicit fsync persists the new size. Move setting of IOMAP_F_DIRTY after the initialization.	N/A	More Details
CVE-2026-126980	Philips Hue Bridge HomeKit Observable Discrepancy Information Disclosure Vulnerability. This vulnerability allows network-adjacent attackers to disclose sensitive information on affected installations of Philips Hue Bridge. Authentication is not required to exploit this vulnerability. The specific flaw exists within the HomeKit Authentication service, which listens on TCP port 8080 by default. When processing crafted pairing requests, the service discloses internal information via an observable discrepancy between possible responses. An attacker can leverage this in conjunction with other vulnerabilities to execute arbitrary code in the context of the device. Was ZDI-CAN-28470.	N/A	More Details
CVE-2026-99604	In the Linux kernel, the following vulnerability has been resolved: drm/etnaviv: Flush the tile status cache through both the PE and BLT Since commit f232d9ec029c ("drm/etnaviv: fix TS cache flushing on GPUs with BLT engine"), jobs on GPUs with a BLT engine use a BLT command to flush the tile status cache instead of TS_FLUSH_CACHE, which is broken on most of these GPUs. GC7000 r6008/r6009 and GC8200 r6401 are the other way around: TS_FLUSH_CACHE works, but the BLT command does not flush the cache. The vendor driver handles these GPUs using TS_FLUSH_CACHE, but the etnaviv kernel driver has no way to identify them. As a result, dirty tile status cache lines survive the end of a job. They are written back by the next cache flush, using the page table of the job issuing it. If that job belongs to another process, the write goes to the wrong address space. On GC7000 r6009 this causes MMU faults and GPU recoveries blamed on an innocent process. I see around 1100 per run of a multisample depth/stencil test running alongside two other GPU clients. Emit both flushes at the end of every job on GPUs with a BLT engine. Whichever does not work on a given GPU is a no-op. Since dirty lines are only written back once, this costs one additional state load per job.	N/A	More Details
CVE-2026-99713	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/mmu: Conditionally resched when retrying gPTE walk on A/D update failure Conditionally reschedule when retrying a guest PTE walk because the CMPXCHG to update A/D bits failed, e.g. because a different task/vCPU modified the relevant gPTEs. On kernels without involuntary preemption of some kind, a misbehaving guest can "capture" a pCPU effectively indefinitely and trigger soft lockups by continuously modifying its gPTEs such that the result of the walk is remains valid.	N/A	More Details
CVE-2026-99712	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Conditionally resched when restarting REP string emulation When restarting emulation, e.g. of a REP-prefixed string instruction, conditionally reschedule as emulating guest writes can be deceptively slow and expensive, e.g. when shadow paging is in use and each write can trigger write tracking. In a perfect world, restarting emulation (or not) would be less arbitrary, e.g. would honor not only host events (NEED_RESCHED, signals, etc.) but also guest events (e.g. by checking kvm_vcpu_has_events()). Unfortunately, as pointed out by Sashiko, pio_in_emulated()'s read-ahead logic is baked into the restart logic, as KVM relies on restarting emulation to drain the cache and avoid duplicate reads from the I/O port. I.e. re-entering the guest on kvm_vcpu_has_events() or __xfer_to_guest_mode_work_pending() would introduce architecturally-incorrect, guest-visible changes.	N/A	More Details
CVE-2026-99711	In the Linux kernel, the following vulnerability has been resolved: KVM: nVMX: Process each host MSR entry at most once when unwinding failed VM-Entry Add a missing "break" in the inner for-loop of the open coded mess that unwinds MSRs back to the pre-VM-Enter state when VMLAUNCH/VMRESUME VM-Fails due to a consistency check that is caught by hardware but missed by KVM. I.e. move on to the next MSR in the VM-Entry load list after performing a write using a value found in the VM-Exit load list. Processing *every* matching entry in the VM-Exit load list makes the processing quadratic, which is problematic if the load lists are full of MSRs that are expensive to emulate. E.g. writing APIC_BASE triggers kvm_recalculate_apic_map() and a whole pile of unnecessary work. Functionally, arbitrarily restoring the value of the first matching entry is "fine", because none of this is architecturally correct: the load lists are never supposed to have been used in the first place. In practice, it's likely a moot point as no well-behaved VMM will have multiple, diverging	N/A	More Details

	entries in the load lists.		
CVE-2026-99605	In the Linux kernel, the following vulnerability has been resolved: drm/etnaviv: Fix NULL dereference on unbind without PPU flop reset flop_reset_data_ppu is only allocated on GPUs that need the PPU flop reset. However, etnaviv_unbind(..) frees it unconditionally, and etnaviv_cmdbuf_free(..) dereferences cmdbuf->suballoc before checking whether cmdbuf is NULL. As a result, unloading the driver on GPUs that do not need the PPU flop reset causes a NULL pointer dereference at address 0.	N/A	More Details
CVE-2026-99606	In the Linux kernel, the following vulnerability has been resolved: net/smc: don't spin in smc_rx_recvmsg() while a pipe holds spliced data While bytes spliced off the RMB are still in a pipe, recvmsg() waits for the pipe to release them before it copies the data behind them. The path to that wait skips the timeout and signal checks, and smc_rx_wait() returns at once on RCV_SHUTDOWN. With a signal pending, after SHUT_RD, or once the peer closes or shuts down writing, recvmsg() spins on a CPU until the pipe is drained, and SIGKILL cannot stop it. MSG_DONTWAIT and an expired SO_RCVTIMEO wait there too instead of returning -EAGAIN. Check the timeout and signals before that wait, and end it when the pipe releases the data or the connection is killed, not on RCV_SHUTDOWN. A zero-length read returns 0 without waiting for the pipe; a non-blocking one returns -EAGAIN while poll() still reports the socket readable (splice() can still take those bytes). The wait sets SOCKWQ_ASYNC_WAITDATA like smc_rx_wait(), so a byte arriving behind the pipe raises no SIGIO on a reader already asleep for it, and a full barrier orders its wait queue entry before its read of splice_pending, so no wake is lost.	N/A	More Details
CVE-2026-99608	In the Linux kernel, the following vulnerability has been resolved: ALSA: usx2y: Guard the playback copy usx2y_urb_play_prepare() sums the packet lengths of a whole URB, copies that many frames with a single wrap, and sizes each playback packet from the capture packet's actual_length. Three things are unchecked: that the URB fits in the ring buffer, that a packet fits in the playback endpoint's maxpacksize, from which tmpbuf is allocated, and that actual_length is a whole number of frames. The ISO descriptors take the raw actual_length while transfer_buffer_length and the copies take count * stride, so a packet that is not a multiple of the stride makes the descriptors add up to more than the buffer, by up to stride - 1 bytes per packet. Return -EPIPE, as the existing packet-size check does, when the URB is larger than the buffer or a packet is larger than maxpacksize, and size each descriptor from whole frames, so the descriptors add up to transfer_buffer_length and each stays within maxpacksize. The buffer case is prevented by the previous patch and this is the guard for it; the other two depend only on the device. The first case is reachable on the unpatched driver with ordinary packets: on 7.3-rc7, a USB/IP device presenting the US-122 interface layout and sending 48-frame packets, with a 100-frame playback buffer (two 50-frame periods, inside the stock constraints), KASAN reports: BUG: KASAN: out-of-bounds in i_usx2y_urb_complete+0x79b/0x1ac0 Read of size 18446744073709551280 at addr ffff88800e6b02e0 by task vhci_rx/80 With both patches, the same device and a 200-frame playback buffer is refused at hw_params with EINVAL, and a 400-frame buffer plays for 30 s with the position advancing inside the buffer, no KASAN report and neither guard message. The maxpacksize check is not exercised: the emulated device reports the same wMaxPacketSize on both endpoints. The descriptor case is visible with capture packets of 195 bytes at stride 4 (48 frames and 3 bytes): without this change the four descriptors add up to 780 bytes for 768 bytes of data, and USB/IP's completion check refuses the URB ("total length of iso packets 780 not equal to actual length of buffer 768", -EPROTO), after which the driver stops the streams. With this patch every descriptor is 192 bytes, the descriptors add up to the data in every URB for 30 s, the position advances inside the buffer and KASAN reports nothing. The 48-frame playback run above is unchanged.	N/A	More Details
CVE-2026-99613	In the Linux kernel, the following vulnerability has been resolved: ALSA: cs46xx: Handle SCB creation failures before dereferencing With CONFIG_PM_SLEEP enabled, cs46xx_dsp_create_scb() can return NULL if kmemdup() fails to allocate the SCB data saved for resume. cs46xx_dsp_scb_and_task_init() and _dsp_create_generic_scb() dereference the returned pointer without checking it, causing a NULL pointer dereference on allocation failure. Check both return values before accessing the SCB descriptors. Return -ENOMEM when creating the null SCB fails and propagate NULL from _dsp_create_generic_scb(). Found by Linux Verification Center (linuxtesting.org) with SVACE static analysis tool.	N/A	More Details
CVE-2026-126979	Philips Hue Bridge HomeKit Accessory Protocol chacha20_decrypt Integer Underflow Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of Philips Hue Bridge. Authentication is not required to exploit this vulnerability. The specific flaw exists within the HomeKit Accessory Protocol service, which listens on TCP port 8080 by default. When parsing pairing requests, the process does not properly validate user-supplied data, which can result in an integer underflow before writing to memory. An attacker can leverage this vulnerability to execute code in the context of the device. Was ZDI-CAN-28332.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/mmu: Cap number of parent PTEs per SP at 1000 (the arbitrary rmap limit) When potentially reusing a shadow page, zap the shadow page if its number of parent PTEs exceeds KVM's (mostly) arbitrary rmap recycling threshold. Letting the guest create a (mostly) unbounded number of parent PTEs allows the guest to coerce KVM into creating lists containing millions of PTEs, effectively bounded only by the maximum number of shadow pages. Because PTE lists are simple singly linked lists, removal is O(n). Even worse, if a list is constructed just so, teardown can be O(n^2), e.g. if SPTEs are built such that zapping child SPTEs always zaps the last entry in the list. To cap things off, all paths that walk PTE lists do so while holding mmu_lock, without conditionally rescheduling. Using a doubly		

CVE-2026-99710	linked list would not address the issue, as removal would remain O(n) since KVM still wouldn't know where in the list to remove from. Removal could be made O(1) by storing a per-SPTE pointer, but that would largely defeat the purpose of reusing SPs, and would penalize well-behaved guests, especially for nested TDP where aliasing page tables is uncommon. So, cap the maximum number of parent PTEs at the rmap limit. In practice, well-behaved guests will likely never run afoul of the limit. E.g. when running Linux as L1 with TDP, the number of parent PTEs for a single SP tops out at under 200 (though undoubtedly obscure workloads can push it higher). Deliberately use a "greater than or equal" as opposed to rmap's strictly "greater than", as the rmap code checks the threshold immediately *after* inserting into the list, i.e. the limit is 1000 for rmaps, KVM just very briefly blows past it. Apply the check to both indirect and direct SPs, because while a fully direct MMU naturally limits the number of parent PTEs to '1', direct SPs in indirect MMUs have no such constraint. E.g. if the guest maps a huge PTE into many locations in its page tables, and KVM is backing the hugepage with a direct SP, i.e. with a series of smaller pages, then KVM could end up with a direct SP with many parents. Note, shadow_mmu_try_split_huge_pages() always restarts the rmap walk or bails entirely if it gets to the point where it tries to grab a SP, i.e. there is *no* pre-existing bug that would be tripped by zapping a SP and thus rmaps when getting a SP.	N/A	More Details
CVE-2026-99610	In the Linux kernel, the following vulnerability has been resolved: ALSA: caiaq: free the EP4 URB when input setup fails The three controllers that use the EP4 input URB allocate it before checking the endpoint type and before registering the input device. On failure they jump to exit_free_idev, which clears cdev->input_dev but leaves cdev->ep4_in_urb set. The card teardown then calls snd_usb_caiaq_input_free(), which returns early because input_dev is NULL, so the URB is never freed. Release the URB before dropping the input device on the error path.	N/A	More Details
CVE-2026-99626	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix NULL deref in dcn401_update_clocks_update_dtb_dto [Why] resource_get_otg_master_for_stream() can return NULL if no OTG master pipe is found for a given stream. The current code dereferences this pointer without a NULL check. This may lead to a possible NULL pointer dereference. [How] Add an explicit NULL check and skip the stream if the OTG master is not found. This is similar to dcn60_update_clocks_update_dtb_dto(). Found by Linux Verification Center (linuxtesting.org) with SVACE. (cherry picked from commit ffde3d74c1faeeb645b66d02d02fdae3e073941b)	N/A	More Details
CVE-2026-126978	Philips Hue Bridge hap_pair_setup_handler Heap-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of Philips Hue Bridge. Authentication is not required to exploit this vulnerability. The specific flaw exists within the HomeKit Accessory Protocol service, which listens on TCP port 8080 by default. When parsing pairing requests, the process does not properly validate the length of user-supplied data prior to copying it to a fixed-length, heap-based buffer. An attacker can leverage this vulnerability to execute code in the context of the device. Was ZDI-CAN-28434.	N/A	More Details
CVE-2026-99612	In the Linux kernel, the following vulnerability has been resolved: ALSA: hda/tas2781: Only accept CRC-less V1 data if all entries are valid Commit 408249089c48 ("ALSA: hda/tas2781: Accept V1 calibration data without CRC") only checked the first device entry. On machines where the BIOS populates fewer entries than p->ndev, the remaining all-zero entries are applied and tas2781_cali_preproc() divides by zero on playback open. Accept CRC-less V1 data only if every expected entry is non-zero and self-consistent (R0 * 1/R0 close to 1.0), otherwise reject it as before.	N/A	More Details
CVE-2026-99709	In the Linux kernel, the following vulnerability has been resolved: inet: frags: limit the truesize of a reassembly queue The truesize of fragments in a reassembly queue is only bounded by the per netns ipfrag_high_thresh sysctl, which has no upper limit. inet_frag_reasm_finish() sums the truesize of all fragments in an int. If this sum reaches 2^32, it wraps, and the final sk_wmem_alloc adjustment of the socket owning the reassembled skb can cancel its own reference, leading to a use-after-free of this socket. This can be triggered from an unprivileged user/network namespace, raising ipfrag_high_thresh and using padded TUN packets, which keep a large truesize after being trimmed by ip_rcv_core(). Track the sum of the truesize of fragments in each queue, and refuse a fragment that would make it bigger than INT_MAX / 2. Callers kill the queue on this error. This bound is far above the needs of legitimate traffic.	N/A	More Details
CVE-2026-99616	In the Linux kernel, the following vulnerability has been resolved: ovl: look up the index by the decoded origin A file opened by file handle on an overlay with nfs_export=on can show the stale lower data of a file that was copied up to the upper layer. While the file opened this way stays in the inode cache, a lookup of the copied up file by path fails with ESTALE for all users. The open can also succeed for a file that was removed, where it has to fail with ESTALE. This happens with a handle that has one byte changed. The handle comes from an NFS client, or from a process that calls open_by_handle_at() with CAP_DAC_READ_SEARCH in the initial user namespace. When a file in a lower layer is about to be modified, overlayfs first copies it to the upper layer (copy up). With nfs_export=on, each copy up also creates an entry in the index directory, a directory that overlayfs keeps for itself on the upper filesystem. The handle of an overlay file that comes from a lower layer wraps the handle of its lower inode. Overlayfs calls this lower inode the origin. The overlay handle stays valid after a copy up. To decode such a handle, ovl_lower_fh_to_d() first decodes the wrapped handle to the origin. If no overlay inode is cached for the origin, it looks in the index directory: - An index entry there leads to the upper file. - A whiteout, the marker that overlayfs uses for a removed file, means that the overlay file was removed. The name of an index entry is the handle that copy up encoded for the origin, as a hex string. ovl_lower_fh_to_d() builds the name for this lookup from the bytes that the caller passed in. But handles with	N/A	More Details

	different bytes can decode to the same inode: - Overlayfs ignores some flag bits of the handle. - A lower filesystem can accept a different fid type, extra bytes, or generation 0 for any generation. Such a handle decodes to the origin but misses the index entry or the whiteout. So the decode returns an overlay inode that has only the lower file. While that inode is cached, a path lookup of a copied up file fails with ESTALE, since the upper file that the lookup finds does not match the cached inode, which has none. Reproduced on tmpfs and on ext4 with open_by_handle_at(), after a copy up and a mount cycle. A handle in which one flag bit, the fid type, the length or the generation was changed read the old lower data, and path lookups returned ESTALE. Fix this by re-encoding the decoded origin and building the index name from that handle.		
CVE-2026-99681	In the Linux kernel, the following vulnerability has been resolved: blk-mq: fix out-of-bounds read in blk_mq_free_rqs A KASAN slab-out-of-bounds read can occur in blk_mq_free_rqs() when there is a mismatch between the number of hardware queues allocated for the IO scheduler (et->nr_hw_queues) and the number of hardware queues in the block tag set (set->nr_hw_queues). This mismatch can happen if blk_mq_update_nr_hw_queues() fails halfway through (e.g., due to memory pressure during blk_mq_prealloc_tag_set_tags()). In this error path, the elevator is restored with a larger number of hardware queues than the block tag set actually has. When the elevator is later freed, blk_mq_free_sched_tags() iterates up to the new et->nr_hw_queues and calls blk_mq_free_rqs(). In blk_mq_free_rqs(), it attempts to access set->tags[hctx_idx] to get the driver tags. Because set->nr_hw_queues was not updated due to the earlier failure, set->tags still has the old (smaller) size, leading to an out-of-bounds read. BUG: KASAN: slab-out-of-bounds in blk_mq_free_rqs+0xde/0x680 Read of size 8 at addr ffff88818d67fc28 by task syz-executor117/5839 Call Trace: blk_mq_free_rqs+0xde/0x680 blk_mq_free_map_and_rqs+0x40/0xf0 blk_mq_free_sched_tags blk_mq_free_sched_res+0xeb/0x280 elevator_change_done+0x1d2/0x5c0 elevator_change+0x34f/0x480 elv_iosched_store+0x504/0x630 queue_attr_store+0x207/0x2b0 To fix this, explicitly check if hctx_idx < set->nr_hw_queues before accessing set->tags[hctx_idx]. If hctx_idx >= set->nr_hw_queues, the hardware queue doesn't exist in the tag set, meaning there are no driver tags to clear mappings from. In this case, safely set drv_tags to NULL. The subsequent call to blk_mq_clear_rq_mapping() already handles a NULL drv_tags pointer and will safely return.	N/A	More Details
CVE-2026-99572	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to avoid potential reference count leak in do_garbage_collect() In do_garbage_collect(), we have reserved reference counts on summary cache of all segments, once gc_thread encounter freezing, we need to release all reference counts.	N/A	More Details
CVE-2026-18883	ServiceNow has addressed an indirect prompt injection security issue that was identified in the Email Response AI Agent of Now Assist for CSM. We are not currently aware of malicious exploitation against ServiceNow instances. This remediation was deployed centrally to the Email Response AI Agent (Now Assist for CSM) and does not require a platform version upgrade or hot fix; no version-specific backports apply.	N/A	More Details
CVE-2026-99689	In the Linux kernel, the following vulnerability has been resolved: xfs: fix temporary cursor leak in btree shift corruption paths In xfs_btree_lshift and xfs_btree_rshift, when XFS_IS_CORRUPT detects an invalid firstrec/lastrec result after xfs_btree_dup_cursor has successfully allocated a temporary cursor, the error path jumps to error0 which returns without freeing tcur. This leaks the cursor memory, the xfs_group reference, and the per-level buffer references held by the duplicated cursor. Jump to error1 instead, which calls xfs_btree_del_cursor(tcur, XFS_BTREE_ERROR) to properly release all resources.	N/A	More Details
CVE-2026-126574	Missing authorization vulnerability in ASUS Control Center Datacenter and Control Center Enterprise allows remote authenticated users with a low-privileged Viewer role to change the password of an Administrator account, thereby obtaining administrative privileges, via a crafted password update request. Refer to the ' Security Update for ASUS Control Center ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-119469	Incorrect default permissions in the ASUS DisplayWidget Center allow a local, low-privileged user to achieve arbitrary code execution in the security context of another user via a crafted executable placed in that user-writable directory, which the application launches when the other user runs it and uses a particular feature.Refer to the ' Security Update for ASUS Display Widget Center ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-116643	An SSH certificate authentication bypass vulnerability in GitHub Enterprise Server allowed an authenticated organization member to read and modify another user's private repositories within the trusted certificate authority's scope. Exploitation required an instance using built-in authentication and a valid certificate issued to the attacker by a trusted SSH certificate authority. This vulnerability affected supported GitHub Enterprise Server releases prior to 3.23 and was fixed in versions 3.19.16, 3.20.12, 3.21.10, and 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-99670	In the Linux kernel, the following vulnerability has been resolved: can: dev: can_fill_info(): don't wake up controller if interface is down If the interface is down, the CAN controller might be powered down and its clocks gated. Don't call do_get_berr_counter() to read the bus error counters if the interface is down, and adjust can_get_size() to match. This removes one trigger for a synchronous external abort seen on i.MX8QXP when "ip link show" is run on a flexcan interface that was never brought up. The underlying cause is platform-specific: on i.MX8QXP the CAN clocks are gated through the SCU LPCG and the subsystem shares the IMX_SC_R_CAN_0 power domain, and there is a window after pm_runtime_resume_and_get() returns where the clock has not reached the module yet. This patch avoids the netlink query path while the interface is down, but does not root-cause the LPCG window itself. As a result, IFLA_CAN_BERR_COUNTER is no longer	N/A	More Details

	reported for down interfaces. Userspace must treat the attribute as optional. The issue was triggered on an i.MX8QXP board by running "ip link show" without ever bringing the interface up: Internal error: synchronous external abort: 0000000096000210 pc : flexcan_read_le+0x0/0x18 lr : flexcan_get_berr_counter+0x4c/0x8c Call trace: flexcan_read_le+0x0/0x18 flexcan_get_berr_counter+0x4c/0x8c can_fill_info+0x1f8/0x434 rtnl_fill_ifinfo+0x8fc/0xbb4 rtnl_dump_ifinfo+0x364/0x448 This is the same class of problem fixed for m_can in commit 91a55c72a821 ("can: m_can: m_can_get_berr_counter(): don't wake up controller if interface is down").		
CVE-2026-99718	In the Linux kernel, the following vulnerability has been resolved: hwmon: (peci/dimmtemp) Fix devres leak and UAF on concurrent unbind When a peci_dimmtemp device is unbound concurrently with its delayed initialization work (detect_work), newly allocated devm resources and the hwmon device are leaked, leading to a use-after-free when reading the lingering hwmon sysfs attributes. In peci_dimmtemp_probe(), detect_work is registered for cancellation via devm_delayed_work_autocancel(). During driver detach, devres_release_all() first moves all existing entries from dev->devres_head onto a local todo list before invoking their release callbacks in reverse order. If create_dimm_temp_info_delayed() is already running when devm_delayed_work_drop() calls cancel_delayed_work_sync(), create_dimm_temp_info() can allocate dimmtemp_label strings via devm_kzalloc()/devm_kasprintf() and register the hwmon device via devm_hwmon_device_register_with_info(), appending new devres entries to the now-empty dev->devres_head. Once the work finishes, devres_release_all() frees priv and priv->name from the local todo list without re-evaluating dev->devres_head, leaving the late devres allocations and registered hwmon device behind with a dangling priv pointer. Cancel detect_work synchronously in a driver .remove() callback instead of relying on devm_delayed_work_autocancel() so that the work is fully drained before devres_release_all() detaches dev->devres_head, allowing any devm resources registered by the delayed work to be released in proper reverse order. Also only schedule detect_work when check_populated_dimms() returns -EAGAIN so that probe failure paths do not leave delayed work pending.	N/A	More Details
CVE-2026-126972	KiCad Plugin Directory Traversal Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of KiCad. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the handling of plugin packages. The issue results from the lack of proper validation of a user-supplied path prior to using it in file operations. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current user. Was ZDI-CAN-33085.	N/A	More Details
CVE-2026-116637	A site administrator could modify files on the appliance host with root privileges and disclose private repository data due to improper path validation in GitHub Enterprise Server. Exploitation required a crafted pre-receive hook environment. This vulnerability affected supported GitHub Enterprise Server versions prior to 3.23 and was fixed in versions 3.19.16, 3.20.12, 3.21.10, and 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-99723	In the Linux kernel, the following vulnerability has been resolved: vhost-vdpa: fix pinned page leak in vhost_vdpa_pa_map() error path When pin_user_pages() or vhost_vdpa_map() fails, vhost_vdpa_pa_map() is supposed to unpin the pages pinned after the last successfully mapped chunk, i.e. the range [map_pfn, last_pfn]. The unpin is guarded by nchunks, but nchunks is reset to 0 every time a chunk is mapped from inside the current page list, so a failure after any such boundary leaves that pending range pinned forever. With physically scattered memory nearly every page list contains a boundary, so the leak happens on practically every failed IOTLB update. The pending range is non-empty exactly when a page has been pinned, which last_pfn already tells us, so drop nchunks and test last_pfn.	N/A	More Details
CVE-2026-114205	Incorrect authorization in ASUS Control Center Express allows a remote authenticated user with a low-privileged role to create an administrator account, resulting in administrative control of the management console and of the clients it manages. Refer to the ' Security Update for ASUS Control Center Express Agent ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-114204	Use of Hard-coded Cryptographic Key in ASUS Control Center Express allows a remote user to bypass authentication by forging a session token for an account with a recent valid login, resulting in administrative control of the management console and of the clients it manages. Refer to the ' Security Update for ASUS Control Center Express Agent ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-112115	An untrusted search path in ASUS Armoury Crate allows a local, low-privileged user to achieve privilege escalation by executing a crafted library at a high integrity level, via user-controlled environment settings inherited by an elevated process. Refer to the ' Security Update for Armoury Crate App ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-126973	KDE Gwenview PNG File Parsing Integer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of KDE Gwenview. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the parsing of PNG files. The issue results from the lack of proper validation of user-supplied data, which can result in an integer overflow before allocating a buffer. An attacker can leverage this vulnerability to execute code in the context of the current process. Was ZDI-CAN-29373.	N/A	More Details
	Hinemos Agent Improper Certificate Validation Remote Code Execution Vulnerability. This vulnerability allows		

CVE-2026-126963	network-adjacent attackers to execute arbitrary code on affected installations of Hinemos Agent. Authentication is not required to exploit this vulnerability. The specific flaw exists within the Hinemos Agent service. The issue results from the lack of proper validation of the certificate presented by a server. An attacker can leverage this vulnerability to execute code in the context of the service account. Was ZDI-CAN-28837.	N/A	More Details
CVE-2026-114202	Use of Hard-coded Cryptographic Key in ASUS Control Center Express allows a remote user to gain administrative privileges via a forged authentication token. Refer to the ' Security Update for ASUS Control Center Express Agent ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-99666	In the Linux kernel, the following vulnerability has been resolved: mm/khugepaged: fix PMD table lock leak in try_collapse_pte_mapped_thp() For a private mapping with userfaultfd, try_collapse_pte_mapped_thp() takes the PMD table lock before it maps the page table for the second time. If that (pte_offset_map_rw_nolock) fails, because the page table vanished in the meantime due to concurrent page table reclaim, it leaves through the abort label, where the PMD table lock is unlocked if it differs from PTE table lock. In case the PMD table lock and the PTE table lock match, which is the case without CONFIG_SPLIT_PTE_PTLOCKS after pte_offset_map_rw_nolock() failed, we miss to unlock the PMD table. To fix it, forget the PTE table lock when we unmap+unlock the PTE table. Fix generated by an LLM and sent by Josef Bacik in private. [david: move the ptl = NULL handling; rewrite description, drop wrong comments]	N/A	More Details
CVE-2026-112117	A Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') in ASUS System Control Interface , allows a local, low-privileged user to gain full control over arbitrary existing files and folders outside the allowed directories via a crafted path containing dot-segment sequences that satisfies a folder prefix check, which can lead to unauthorized read and modification of protected files and may result in local privilege escalation. Refer to the ' Security Update for ASUS System Control Interface ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-99665	In the Linux kernel, the following vulnerability has been resolved: mm/khugepaged: preserve dirty bits when collapsing PTE-mapped THPs try_collapse_pte_mapped_thp() clears PTEs without transferring their dirty bits to the folio. A shmem read fault can map a clean folio writable, so a later store only dirties the PTE. Losing that dirty bit allows reclaim to discard modified data. The assumption that regular files are read-only is also obsolete since commit 1b6444331bebb ("mm/khugepaged: enable clean pagecache folio collapse for writable files"). In any case, transferring PTE dirty bits to the folio when removing mappings (and while the folio is still mapped) is the right thing to do. Use get_and_clear_ptes() to collect the dirty bits and mark the folio dirty before unmapping it. Update the comment. This was reproduced with shmem using MADV_COLLAPSE followed by MADV_PAGEOUT. Fix generated by an LLM and sent by Chris Mason in private. [david: adjust comment; simplify description; cover file-thp as well]	N/A	More Details
CVE-2026-99664	In the Linux kernel, the following vulnerability has been resolved: mm/gup: recheck the PUD under PTL in follow_pud_mask() follow_pud_mask() reads the PUD without a lock. If it is a leaf, it takes the lock and calls follow_huge_pud(), which reads the entry again, but only tests whether it is present. A leaf can be zapped concurrently and a page fault can fill the empty entry with a page table. This currently only applies when device-dax makes use of PUD entries, and the fault after e.g., MADV_DONTNEED would find THPs to be disabled. That requires a sequence like 1) Fault in the PUD 2) Disable THP 3) Have GUP race against concurrent MADV_DONTNEED + refault Which is rather unlikely to happen in practice, which is why we never received a report about such a problem. Before commit 1b1676180246 ("mm/gup: handle huge pud for follow_pud_mask()"), we performed that check in follow_devmap_pud() by testing for pud_devmap(). Let's fix it by just structuring follow_pud_mask() similar to follow_pmd_mask(), minus protnone stuff that does not apply and minus FOLL_SPLIT_PMD handling. That makes it much easier to spot any such inconsistencies. Perform the pud_bad() check early in follow_pmd_mask() and just forward the PUD value. We can drop the pud_present() check in follow_huge_pud(), just like we have in follow_huge_pmd(). Even though we are performing more changes than strictly necessary, the fix should be easy to backport.	N/A	More Details
CVE-2026-99663	In the Linux kernel, the following vulnerability has been resolved: mm/gup: fix clearing folios in memfd_pin_folios() memfd_pin_folios() pins folios obtained directly from the page cache without checking whether they are uptodate. However, shmem fallocate() can leave newly allocated folios uncleared and not uptodate. Callers such as udmabuf can then expose uninitialized memory to userspace. Clear non-uptodate folios and mark them uptodate under the folio lock before returning them to the caller. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify description and comment]	N/A	More Details
CVE-2026-126974	llama.cpp GGUF File Parsing Integer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of llama.cpp. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the parsing of GGUF files. The issue results from the lack of proper validation of user-supplied data, which can result in an integer overflow before writing to memory. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current process. Was ZDI-CAN-34553.	N/A	More Details
CVE-2026-99661	In the Linux kernel, the following vulnerability has been resolved: mm/gup: check for overflow when converting nr_pages to bytes in GUP-fast The overflow check in gup_fast_fallback() only covers the addition to start. The preceding nr_pages << PAGE_SHIFT can already have truncated the length on 32bit. Users such as RapidIO's rio_dma_transfer() can derive such page counts from a 64-bit userspace length. Let's just	N/A	More Details

	check_shl_overflow() to reject any such overflow with -EOVERFLOW early.		
CVE-2026-99671	In the Linux kernel, the following vulnerability has been resolved: can: raw, isotp: copy socket options outside of lock_sock() The CAN_RAW and CAN_ISOTP specific socket options are copied from and to user space while holding lock_sock(). A copy that faults or stalls, e.g. on a userfaultfd or FUSE backed mapping, holds lock_sock() for its whole duration. raw_notify() and isotp_notify() take lock_sock() of a socket that is bound to the vanishing device while holding RTNL. As the CAN_RAW options can also be changed on a bound socket, a stalled copy can block RTNL on an unregistration or namespace change of the bound device. Before "can: protect the receive lists with a per-device lock", raw_setsockopt() even held RTNL itself during the copy. CAN_ISOTP only copies its options for unbound sockets, which isotp_notify() skips without taking lock_sock(). Keep the user space copies out of the locked sections: Check the option name and length and copy the option values, including the CAN_RAW_FILTER filters, to kernel memory before lock_sock(). raw_setsockopt_locked() and isotp_setsockopt_locked() get the copied values directly and drop their now redundant length checks and copies. As a side effect, a failing copy can not leave partially overwritten options behind anymore. In raw_getsockopt() duplicate the filters under lock_sock() and copy them to user space after release_sock(). As the option name and length are now checked before lock_sock(), the order of the checks in isotp_setsockopt() and accordingly the order of the returned error values changes. This is in line with POSIX, which does not define an order in which multiple error conditions are detected and reported. The generic SOL_SOCKET options in sk_setsockopt() are not affected by this patch.	N/A	More Details
CVE-2026-117634	Pleasanter contains an incorrect authorization vulnerability in Publishing feature. If this vulnerability is exploited, a remote unauthenticated attacker who can access a site published using the product may obtain, by sending a specially crafted request, information on records of unpublished sites (e.g., titles), as well as the names of users, groups, and organizations registered in the same tenant. Note that the vulnerability only affects configurations where both of the following conditions are met: a site is published using the Publishing feature, and the published site contains link items (items linked to other sites, users, groups, or organizations).	N/A	More Details
CVE-2026-99722	In the Linux kernel, the following vulnerability has been resolved: vringh: bound top-level re-entry into indirect tables In __vringh_iov(), the F_INDIRECT branch descends into the indirect table and continues before the descriptor accounting runs. A top-level indirect descriptor is therefore never charged to count. If such a descriptor sets NEXT to point back at itself, returning from the indirect table resumes at the same top-level descriptor, which is again not counted, so the walk never makes forward progress. Because count stays flat and indirect_count is reset to 0 on every return to the top-level table, neither bound in the loop check trips. A guest can spin the vringh worker at 100% CPU inside __vringh_iov(), an uninterruptible host DoS. Move the descriptor accounting above the indirect switch so that a top-level indirect descriptor is charged one top-level traversal step before the walk descends into its table, bringing this re-entry under the existing vrh->vring.num bound. Each top-level descriptor is still charged at most one step and indirect_count still bounds a single table, so legitimate chains (including the multiple-indirect case in tools/virtio/vringh_test.c) stay within vring.num, while a cyclic indirect descriptor is now rejected with -ELOOP.	N/A	More Details
CVE-2026-99721	In the Linux kernel, the following vulnerability has been resolved: hwmon: (aquacomputer_d5next) reject short status reports The HID raw_event callback runs before HID core validates the received report length. aqc_raw_event() reads device-specific status fields at fixed offsets, including multi-byte values, without first ensuring that the received buffer contains every field. A truncated report can therefore cause out-of-bounds reads and update hwmon state with unrelated data. The required length differs between supported devices and cannot be derived from the HID report descriptor. Define the minimum length needed by each asynchronous status format, select it while configuring the device, and reject reports that do not contain the last field accessed by that format. Keep the published length zero until probe configuration is complete, which also excludes legacy devices whose status is retrieved through feature reports.	N/A	More Details
CVE-2026-99682	In the Linux kernel, the following vulnerability has been resolved: virtio_blk: use synchronous quiesce in virtblk_freeze_priv() blk_mq_run_work_fn() can call blk_mq_sched_dispatch_requests() through blk_mq_run_dispatch_ops(). The dispatcher checks QUEUE_FLAG_QUIESCED while the flag is clear. It can then be preempted before it gets a request. At that time, it has no queue usage reference, so blk_mq_freeze_queue() does not wait for it. virtblk_freeze_priv() then calls blk_mq_quiesce_queue_nowait() and blk_mq_unfreeze_queue(). A new request can enter the queue. The old dispatcher can run again, use its old flag check, get the new request, and call virtio_queue_rq(). At the same time, virtblk_freeze_priv() can reset the device, delete the virtqueues, and free vblk->vqs. The old dispatcher can then use a deleted virtqueue or a NULL vblk->vqs, and the kernel can crash. Fix this by using blk_mq_quiesce_queue(), which waits for dispatch code that was already running. Keep the existing freeze and unfreeze order. The queue remains quiesced until virtblk_restore_priv() calls blk_mq_unquiesce_queue().	N/A	More Details
CVE-2026-131820	A server-side request forgery vulnerability in GitHub Enterprise Server allowed an authenticated user with repository write access to send requests to internal services through Slack webhook validity checks, potentially leading to remote code execution on the appliance. Exploitation required Secret Protection with secret scanning and validity checks for partner patterns enabled. This vulnerability affected all versions prior to 3.23 and was fixed in versions 3.22.5, 3.21.10, and 3.20.12. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: hwmon: (lm87) Fix drvdata type confusion		

CVE-2026-99719	in vrm_show() and vrm_store() When an LM87 or ADM1024 device is probed with VID support enabled, reading from or writing to the vrm sysfs attribute triggers a driver data type confusion in vrm_show() and vrm_store() in drivers/hwmon/lm87.c. In lm87_probe(), the driver stores the struct lm87_data pointer on the I2C client via i2c_set_clientdata() and passes the struct i2c_client pointer as the drvdata argument to devm_hwmon_device_register_with_groups(), which attaches the attribute groups to the registered hwmon device with client as its driver_data. Unlike the other sysfs callbacks that retrieve client via dev_get_drvdata() and then obtain data via i2c_get_clientdata(), vrm_show() and vrm_store() assign dev_get_drvdata(dev) directly to struct lm87_data *. Dereferencing data->vrm therefore accesses the struct i2c_client instance at offsetof(struct lm87_data, vrm), causing vrm_show() to leak an unrelated byte from struct i2c_client to userspace and vrm_store() to corrupt a byte within struct i2c_client while failing to update the actual VRM value. Fix the problem by assigning the return value from dev_get_drvdata(dev) to struct i2c_client *client and using i2c_get_clientdata() to get the actual client data.	N/A	More Details
CVE-2026-19399	Improper Authentication and Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') in ASUS Router firmware allow an unauthenticated nearby user to perform file operations or read files in an unintended directory via crafted HTTP requests during an active user session. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-88711	Improper neutralization of special elements used in an OS command in certain ASUS Router firmware models allows an remote attacker to to execute arbitrary commands with root privileges by inducing a logged-in administrator to visit a malicious website. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-99683	In the Linux kernel, the following vulnerability has been resolved: vduse: fix VA/PA confusion in the coherent page fault path vduse_domain_get_coherent_page() assumes that iotlb entries above bounce_size were installed by vduse_domain_alloc_coherent(), so it treats map->addr as a physical address. But vduse registers with use_va = true, and when the device is bound to vhost_vdpa, vhost_vdpa_va_map() stores a userspace virtual address in map->addr, which vduse_domain_set_map() then copies into the same iotlb. A process holding the domain file descriptor across a rebind can therefore fault in an arbitrary physical page with read/write access, ignoring the entry permissions. Additionally, since the domain file mmap lacks VM_IO, the domain file itself can be mapped and passed to vhost_vdpa_va_map(), producing self-referential entries that bypass the coherent path check. To fix, only translate entries whose backing file is the domain file itself in vduse_domain_get_coherent_page(), reject such entries in vduse_domain_set_map() since only external mappings should be imported, and set VM_IO on the domain file mmap so vhost_vdpa_va_map() refuses to map it in the first place.	N/A	More Details
CVE-2026-99677	In the Linux kernel, the following vulnerability has been resolved: xfs: clamp the pNFS layout range to the maximum file size xfs_fs_map_blocks() lets two bad layout ranges through. An offset of 2^63 or more is negative and passes "offset > limit". A length above the limit, such as NFS4_UINT64_MAX, wraps "offset > limit - length" and stays unclamped. With NFS4_UINT64_MAX, a read layout is built from an uninitialized mapping. Reject a negative offset and one at the limit, and clamp the length. This needs "xfs: cap the block count of xfs_iomap_write_direct() at one extent".	N/A	More Details
CVE-2026-99684	In the Linux kernel, the following vulnerability has been resolved: hwrng: virtio - prevent use-after-free after a failed restore remove_common() frees the driver data during freeze, but does not clear vdev->priv, which points to freed memory. If probe_common() then fails during restore, a later remove or freeze can dereference the dangling pointer. Clear vdev->priv on the probe_common() error cleanup path and before freeing the driver data in remove_common(). Also reset the device and return early from remove_common() if vdev->priv is already NULL for a safe teardown after a failed restore.	N/A	More Details
CVE-2026-126969	SmarterTools SmarterMail Insufficient Entropy Authentication Bypass Vulnerability. This vulnerability allows remote attackers to bypass authentication on affected installations of SmarterTools SmarterMail. Authentication is not required to exploit this vulnerability. The specific flaw exists within the generation of Application Passwords. The issue results from using insufficient entropy in the generation of credentials. An attacker can leverage this vulnerability to bypass authentication on the system. Was ZDI-CAN-33568.	N/A	More Details
CVE-2026-126968	SmarterTools SmarterMail Missing Authorization Information Disclosure Vulnerability. This vulnerability allows remote attackers to disclose sensitive information on affected installations of SmarterTools SmarterMail. Authentication is required to exploit this vulnerability. The specific flaw exists within the handling of email attachments. The issue results from the lack of authorization prior to allowing access to functionality. An attacker can leverage this vulnerability to disclose sensitive information in the context of a target user. Was ZDI-CAN-33567.	N/A	More Details
CVE-2026-99675	In the Linux kernel, the following vulnerability has been resolved: nvmet: reject keep-alive timeouts that would overflow nvmet_set_feat_kato() and nvmet_alloc_ctrl() convert the 32-bit millisecond KATO from the host to seconds with DIV_ROUND_UP(kato, 1000). The addition in DIV_ROUND_UP() is done in 32 bits, so any KATO above 4294966296 ms wraps around and rounds to 0 seconds. A Set Features command with KATO 0xffffffff then succeeds but disables the keep-alive timer: Get Features returns 0 and Keep Alive fails with Keep Alive Timeout Invalid. A Connect command with such a KATO creates a controller without a keep-alive timer. The converted value has more consumers that can overflow. Get Features multiplies it by 1000, and the timer paths multiply it by HZ in 32 bits and arm a delay that must not exceed MAX_JIFFY_OFFSET. Reject	N/A	More Details

	a KATO whose rounded value does not fit all of these. Return Keep Alive Timeout Invalid from both Set Features and Connect. Validate Set Features before stopping the current timer so a failed update leaves its timeout unchanged. Round in 64 bits. A KATO of 0 keeps its current meaning.		
CVE-2026-99687	In the Linux kernel, the following vulnerability has been resolved: vhost-scsi: zero misaligned I/O bounce pages The misaligned I/O fallback copies the full bounce scatterlist to the request buffer during completion. If a data-in command transfers less than requested, part of a newly allocated bounce page can remain unwritten. Allocate bounce pages with __GFP_ZERO so all bytes copied during completion have a defined value. This path is already the slow fallback for misaligned I/O, so the added initialization cost is acceptable. Found during an LLM-assisted source audit and validated with a QEMU/KVM vhost-scsi target using misaligned data-in commands.	N/A	More Details
CVE-2026-88760	Missing Authentication for Critical Function in certain ASUS Router firmware allows unauthenticated nearby users to obtain router information and resources (e.g., multimedia files) via HTTP requests. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-126967	SmarterTools SmarterMail File Edit API Directory Traversal Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of SmarterTools SmarterMail. Authentication is required to exploit this vulnerability. The specific flaw exists within the FileStorageController class. The issue results from the lack of proper validation of a user-supplied path prior to using it in file operations. An attacker can leverage this vulnerability to execute code in the context of SYSTEM. Was ZDI-CAN-30742.	N/A	More Details
CVE-2026-93458	Insertion of sensitive information into sent data in certain ASUS Router firmware allows a remote attacker to obtain router’s UID or login authorization tokens for the victim's linked account via a crafted web page, which requires a logged-in router administrator to actively complete the third-party authorization flow. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-93459	Use of externally-controlled format string in certain ASUS Router firmware allows a nearby administrator to read or write memory, potentially resulting in a denial-of-service, via crafted format specifiers sent to the logging/summary function. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-126970	KiCad Altium SCHLIB File Parsing Out-Of-Bounds Write Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of KiCad. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the parsing of SCHLIB files. The issue results from the lack of proper validation of user-supplied data, which can result in a write past the end of an allocated object. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current process. Was ZDI-CAN-33213.	N/A	More Details
CVE-2026-131647	A CRLF injection vulnerability was identified in GitHub Enterprise Server that allowed a user with write access to a repository to bypass branch protections and update branches or tags in that repository through specially crafted input when resolving pull request conflicts. This vulnerability affected all versions of GitHub Enterprise Server prior to 3.23 and was fixed in versions 3.19.16, 3.20.12, 3.21.10, and 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-99673	In the Linux kernel, the following vulnerability has been resolved: can: replace ARPHRD_CAN type checks with can_get_ml_priv() Commit 4e096a18867a ("net: introduce CAN specific pointer in the struct net_device") introduced an explicit way to assign the midlayer private pointer (dev->ml_priv) to named users like ML_PRIV_CAN. A valid pointer returned by can_get_ml_priv() therefore identifies a real CAN device. The referenced commit only converted the j1939 and proc code to this check, while the other parts of the CAN subsystem still rely on dev->type == ARPHRD_CAN. The TUNSETLINK ioctl allows a down TUN/TAP device to set its hardware type to ARPHRD_CAN while dev->ml_priv remains NULL. Such a repurposed netdevice passes the ARPHRD_CAN checks. Commit 8aa59e355949 ("can: af_can: fix NULL pointer dereference in can_rx_register()") already made can_rx_register() reject such devices. But can_rx_unregister() still only checks dev->type and dereferences the NULL ml_priv when it walks the receive lists of the device. CAN_ISOTP reaches this path: isotp_bind() ignores the return values of can_rx_register() and binds the socket anyway. When the socket is released or the device is unregistered, isotp_release() and isotp_notify() remove the filters that have never been registered and trigger the NULL pointer dereference in can_rx_unregister(). This only requires owning a TUN/TAP device or CAP_NET_ADMIN in an unprivileged user namespace. Convert the remaining ARPHRD_CAN checks in af_can, bcm, gw, isotp and raw to can_get_ml_priv().	N/A	More Details
CVE-2026-99720	In the Linux kernel, the following vulnerability has been resolved: hwmon: chipcap2: fix IRQ teardown ordering The threaded IRQs are currently requested before the hwmon device is registered. Since devres releases resources in reverse order, the hwmon device is unregistered before the IRQs are freed during driver teardown. An IRQ may therefore still run after the hwmon device has been released and access the stale data->hwmon pointer. Separate IRQ discovery from IRQ registration so that the IRQ numbers and alarm visibility flags are initialized before registering the hwmon device while the IRQs are released before the hwmon device during devres cleanup.	N/A	More Details
	KiCad Project Archive File Handling Directory Traversal Remote Code Execution Vulnerability. This		

CVE-2026-126971	vulnerability allows remote attackers to execute arbitrary code on affected installations of KiCad. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the handling of project archive files. The issue results from the lack of proper validation of a user-supplied path prior to using it in file operations. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current user. Was ZDI-CAN-33212.	N/A	More Details
CVE-2026-99717	In the Linux kernel, the following vulnerability has been resolved: hwmon: (pmbus/q54sj108a2) Add missing NULL pointer check When CONFIG_DEBUG_FS is enabled and per-client PMBus debugfs directory creation fails during pmbus_do_probe(), q54sj108a2_probe() passes the NULL pointer returned by pmbus_get_debugfs_dir() to debugfs_create_dir() without validation, creating its debugfs directory at the debugfs root instead of under the client directory. If pmbus_init_debugfs() fails to create the client directory, data->debugfs is left as NULL or an ERR_PTR, which pmbus_get_debugfs_dir() converts to NULL. Passing a NULL parent to debugfs_create_dir() places q54sj108a2_dir in the root of the debugfs filesystem and populates it with attribute files whose private data pointers reference psu->debugfs_entries. Because this root-level directory is outside the client debugfs hierarchy, PMBus core cleanup does not remove it on device unbind, leaking the debugfs entries and causing a use-after-free of the devm-allocated struct q54sj108a2_data if the orphaned files are accessed via q54sj108a2_debugfs_read() or q54sj108a2_debugfs_write() after device removal. Fix the problem by adding the missing NULL pointer check.	N/A	More Details
CVE-2026-99660	In the Linux kernel, the following vulnerability has been resolved: mm/userfaultfd: write-lock the VMA before waking waiters when unregistering userfaultfd_unregister() wakes pending missing faults before clearing the VMA's userfaultfd registration. With per-VMA fault locking, holding mmap_lock for write does not prevent another fault from queuing after the wakeup but before userfaultfd_clear_vma() write-locks the VMA. That waiter can then remain asleep indefinitely. Call vma_start_write() before waking waiters. Existing fault handlers queue themselves before dropping the VMA read lock, and new readers are excluded until mmap_write_unlock().	N/A	More Details
CVE-2026-99571	In the Linux kernel, the following vulnerability has been resolved: nvme-tcp: Validate the queue state under the queue lock in the address attribute Reading the "address" sysfs attribute calls kernel_getsockname() on the admin queue socket, which can race with the queue teardown during connection error recovery: a reader can observe NVME_TCP_Q_LIVE set, then block on the queue lock while nvme_tcp_free_queue() releases the socket and clears queue->sock without holding the lock. Once the reader acquires the lock it dereferences the NULL socket pointer. Commit 76d54bf20cdc ("nvme-tcp: don't access released socket during error recovery") had the queue liveness check inside the lock. Commit 782373ba2766 ("nvme: tcp: avoid race between queue_lock lock and destroy") moved it back out because the lock itself is destroyed in nvme_tcp_free_queue(), so a reader arriving late trips mutex_destroy(). Check the queue liveness inside the queue lock, and clear queue->sock under the same lock. The teardown clears NVME_TCP_Q_LIVE under the lock well before the socket is released, so a reader either sees the bit cleared or a socket that cannot be released while it holds the lock. To also close the lock lifetime race, the queue locks are now owned by the controller: they are initialized with the queue array and only destroyed when the controller is freed, which cannot happen while a sysfs attribute read is still in progress. The queue count is computed by the NVME_TCP_OPT_NR_QUEUES() macro from the original connect options for both initialization and destruction, as reconnects may shrink ctrl->queue_count.	N/A	More Details
CVE-2026-99639	In the Linux kernel, the following vulnerability has been resolved: ovl: fix whiteout creation in casefolded directories unlink and rmdir fail with EINVAL on an overlay with casefolded (case-insensitive) layers, if the name exists in a lower layer. Casefolded layers need CONFIG_UNICODE. Root can mount such an overlay, and so can an unprivileged user in a user namespace, with tmpfs layers. Overlayfs does not change a lower layer. To remove a name that exists there, it hides the name with a marker file, a whiteout. It creates the whiteout in the work directory, a scratch directory on the upper filesystem, and renames it over the name in the upper directory: - ovl_whiteout() creates the whiteout with mknod, or as a hard link to an earlier whiteout. - ovl_cleanup_and_whiteout() then prepares the rename with start_renaming_two_dentries(), which returns -EINVAL if the source or the target dentry is unhashed, that is, not in the dentry hash table. With casefolded directories, both dentries can be unhashed: - The target. Lookup in a casefolded directory leaves a negative dentry unhashed, so the rename target is unhashed if the name does not exist in the upper directory. - The source, which is the new whiteout. tmpfs hashes a dentry when mknod or link makes it positive, but ext4 does not, so on ext4 the source is unhashed as well. The cleanup after the failure also refuses an unhashed dentry, so on ext4 the whiteout is left behind in the work directory. Reproduced with rm and rmdir on casefolded tmpfs and ext4 layers, also as uid 1000 in a user namespace. Fix this by creating the whiteout with ovl_create_temp(), which repeats the lookup if the new dentry is unhashed. Prepare the rename with start_renaming_dentry(), which looks up the target name with the directories locked. The caller looked the name up without a lock, so return -ESTALE if the name now refers to a different inode. [amir: pass the name used for lookup instead of taking a name snapshot]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm/vma: call ->may_split() before a partial merge A driver can use the ->may_split() hook to forbid a VMA boundary at certain addresses. A device-dax device is a character device that gives raw access to a range of memory through mmap(). device-dax refuses any address that is not a multiple of the device's alignment through vm_ops->may_split. vma_merge_existing_range() can move boundaries, but does not call vm_ops->may_split. If a range at the edge of a VMA gets the attributes of the adjacent VMA, the function moves only that range over to the neighbour: one VMA shrinks and the other grows. This is a partial merge, and vma_modify() tries it before any		

CVE-2026-99648	split. Take a device-dax VMA that lies next to another VMA of the same file. The user calls mprotect() on the one page of the first VMA that touches the second VMA, and asks for the protection that the second VMA already has. The result is a partial merge. It puts the boundary inside a 1 GiB mapping, which is a single leaf entry at the PUD level. As part of the merge, the kernel looks for a huge PMD to split at the new boundary. That code does not expect a PUD leaf: vma_merge_existing_range() commit_merge() vma_adjust_trans_huge() split_huge_pmd_address() mm_find_pmd() treats the PUD leaf as a pointer to a PMD table, so the "table" is the device memory that the leaf maps __split_huge_pmd() reads a PMD that the user wrote to the device memory folio_put() on the folio that this PMD names Bad. While the code must at some point be taught to split PUD entries as well, the issue here is that we move VMA boundaries such that they are incompatible with the driver. Fix this by calling ->may_split() on the address that a partial merge would move the boundary to, and by abandoning the merge if the driver refuses. vma_modify() then tries the split, which fails with -EINVAL. Factor out existing split handling into check_split() and reuse it. Fix generated by an LLM and sent by Chris Mason in private. [david: improve description; cleanup the code]	N/A	More Details
CVE-2026-22062	StorageGRID (formerly StorageGRID Webscale) versions 12.1.0 and higher are susceptible to an information exposure vulnerability which allows a remote unauthenticated attacker to acquire low-sensitivity credentials via a Man-in-the-middle attack when certain non-default features are enabled.	N/A	More Details
CVE-2026-99647	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: split huge PUDs when adjusting VMA boundaries PFN mappings such as VFIO PCI BAR mappings can use PUD leaves while allowing VMA splits at page granularity. However, vma_adjust_trans_huge() only handles PMDs. After faulting a writable PUD-sized VFIO BAR mapping, changing its first page to read-only with mprotect() makes the PMD walk treat the PUD leaf as a PMD table. Applying MADV_DONTFORK to the first PMD-sized chunk leaves the PUD spanning two VMAs: the PMD-aligned boundary skips splitting. Walk the page tables once at each adjusted boundary and split a huge PUD or PMD leaf as needed. Handle PUD leaves even when the boundary is PMD-aligned.	N/A	More Details
CVE-2026-99646	In the Linux kernel, the following vulnerability has been resolved: mm/vma: don't leak the charge for stack growth acct_stack_growth() charges stack expansion even without VM_ACCOUNT. This can happen for read-only or MAP_NORESERVE mappings created with MAP_GROWSDOWN. Unmapping these VMAs does not release the charge, leaking Committed_AS. Only check and charge growth for VM_ACCOUNT VMAs, matching mremap(). Unaccounted growth skips the overcommit check while remaining subject to the other stack-growth limits.	N/A	More Details
CVE-2026-99644	In the Linux kernel, the following vulnerability has been resolved: mm/memory-tiers: unregister empty tiers outside memory_tier_lock When offlining a node's last memory empties its tier, clear_node_memory_tier() calls device_unregister() while holding memory_tier_lock. A concurrent sysfs read can hold a kernfs active reference and block in nodelist_show() on that mutex. device_unregister() then waits in kernfs_drain() for the read to finish and release its active reference, while the reader waits for the offlining task to release the mutex. Return the node's former tier from clear_node_memory_tier(), making the helper more similar to set_node_memory_tier(), which also returns a memory_tier pointer. Let the caller unlink an empty tier under memory_tier_lock before rebuilding demotion targets, then unregister its device after dropping the lock.	N/A	More Details
CVE-2026-99643	In the Linux kernel, the following vulnerability has been resolved: khugepaged: fix readahead lock inversion in collapse_file() collapse_file() can call page_cache_sync_readahead() while still holding source folios locked. If truncate holds mapping->invalidate_lock exclusively and waits for one of those folios, readahead waits for invalidate_lock and neither side can make progress. We need to release those folio locks before starting readahead. On a page cache miss, roll back first, populate the collapse range, and retry once. Doing readahead after rollback avoids an extra pass when all folios are already present. If another miss occurs on retry, give up rather than keep retrying while truncate or reclaim removes folios. Keep page_cache_sync_readahead() and wait for read I/O with filemap_lock_folio(), releasing each folio before moving to the next. The existing uptodate check during collapse still handles read errors. Shmem uses shmem_get_folio() for swapin and fallocated folios without taking mapping->invalidate_lock, so it is unaffected.	N/A	More Details
CVE-2026-99642	In the Linux kernel, the following vulnerability has been resolved: mm/gup: check folio->mapping under folio lock in memfd_pin_folios() memfd_pin_folios() pins folios from the page cache without taking the folio lock. khugepaged or MADV_COLLAPSE can therefore replace a folio after checking its refcount while memfd_pin_folios() still pins the old folio. The caller then receives a folio which is no longer part of the memfd. For udmabuf and iommufd, this can cause data corruption: DMA accesses the detached folio rather than the live memfd contents. DMA reads can return stale data, while DMA writes update the detached copy and are lost from the live memfd. The detached folio remains pinned, so these accesses do not constitute a use-after-free or target another application's memory. For memfd live update, the stale contents can reappear in the memfd after handover. Lock the folio and recheck folio->mapping before pinning it, like other page cache users do. If the folio was removed, retry the folio lookup from the index that was removed.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ovl: fail the mount if the casefold flag has changed An unprivileged user in a user namespace can repeatedly trigger a kernel warning with their own overlay mounts. They do so by changing the casefold flag of the topmost layer's empty root, the top directory of that layer, after the layers are set and before the mount is created. The casefold flag of a directory makes name lookups in that directory case-insensitive. All layers of an overlay must agree on casefolding.		

CVE-2026-99641	ovl_mount_dir_check() reads the casefold flag of a layer root when fsconfig() sets that layer, and saves the result in ofs->casefold. The mount is created later, by another fsconfig() call with FSCONFIG_CMD_CREATE. The casefold flag of an empty directory can be changed. The casefold flag of the overlay root must match ofs->casefold. ovl_get_root() copies the flag from the topmost layer root, as it is when the mount is created. On a mismatch, ovl_get_root() only does a WARN_ON(), and the mount continues. The check is not a WARN_ON_ONCE(), so each such mount warns again. An unprivileged user needs CONFIG_UNICODE, and a tmpfs mounted with the casefold option in their namespace. Reproduced by changing the casefold flag of the empty upper directory after the layers were set and before FSCONFIG_CMD_CREATE, with tmpfs and ext4 layers, and as uid 1000 in a user namespace. Fix this by replacing the WARN_ON() with a check that fails the mount with -EINVAL.	N/A	More Details
CVE-2026-102430	An uncontrolled recursion vulnerability exists in the LINE iOS app prior to version 26.19.0. The Flex Message processing component does not adequately validate the structure of externally supplied message content. As a result, an attacker who is able to send a crafted message to a victim could cause the application to crash when the message is received or displayed. A server-side mitigation has been deployed that also protects existing iOS clients that have not been updated to version 26.19.0.	N/A	More Details
CVE-2026-99640	In the Linux kernel, the following vulnerability has been resolved: ovl: do not let chattr change the casefold flag chattr +F or -F on an overlay directory can make the directory, or the whole mount, unusable. Lookups of the directory, and lookups in it, then fail. From code inspection, the caller only needs to own the directory. The upper filesystem must support casefolding: ext4 or f2fs with the casefold feature, or tmpfs with the casefold option. chattr +F and -F set and clear FS_CASEFOLD_FL, the casefold flag of a directory. The flag enables casefolding (case-insensitive name lookup) for the directory. All directories of an overlay must match the casefold state that was recorded at mount time, and ovl_lookup_single() enforces this: - A lookup in a directory that does not match fails with -ESTALE. - A lookup of the directory itself fails with -EREMOTE. ovl_fileattr_set() copies the directory up, which creates it in the upper layer if it only existed in a lower layer. It then passes FS_CASEFOLD_FL on to the upper filesystem. The upper filesystem accepts the change if the upper directory is empty, regardless of what the merged directory contains. The merged directory is the one that the user sees, with the names from the upper and the lower directories. After the change the upper directory no longer matches the mount. If the directory is the root of the overlay, the whole mount is affected. The flag can also change without chattr +F or -F. chattr reads all flags of a directory, changes one and writes all of them back. For a directory that only exists in a lower layer, overlayfs reads the flags from the lower filesystem. xfs with the ascii-ci feature reports FS_CASEFOLD_FL on every inode, but never sets S_CASEFOLD, the inode flag that overlayfs compares with the mount. So its directories count as not casefolded. With a lower layer on such an xfs and an upper layer on ext4 with the casefold feature, chattr +A on a lower-only directory therefore sets FS_CASEFOLD_FL on the new, empty upper directory. Reproduced as root on tmpfs and on ext4 with chattr +F or -F on an overlay directory whose upper directory was empty. Afterwards lookups below it failed with ESTALE, and lookups of the directory itself failed with EREMOTE. The chattr +A case was seen once, on a kernel with an earlier form of this patch that still passed the caller's flag on: lookups of the directory failed with EREMOTE. Fix this by returning -EOPNOTSUPP if FS_IOC_SETFLAGS, the only call that can change the flag, asks for a casefold flag on a directory that does not agree with the casefold state of the mount. This also refuses the chattr +A case above. [amir: check the flag against the casefold state of the mount instead of reading back the flags]	N/A	More Details
CVE-2026-99638	In the Linux kernel, the following vulnerability has been resolved: net/sched: clamp skb->queue_mapping before indexing the tx queue This issue was caught by Sashiko (nipa). A producer-side path cap is already merged (ea4d4b5dddb5); this patch adds the matching consumer-side caps. The issue: The tx queue is picked and capped in netdev_core_pick_tx() before the __dev_xmit_skb() qdisc enqueue path. A tc BPF program running inside an egress qdisc (not clsact) can then change __sk_buff->queue_mapping; that store is never re-checked. bpf_convert_ctx_access() bounds it only by NO_QUEUE_MAPPING, not by the device's queue count. Two consumers read the stale value: - On dequeue, qdisc_restart() calls netdev_get_tx_queue(), which indexes dev->_tx[]. netdev_get_tx_queue() guards it with just a DEBUG_NET_WARN_ON_ONCE, so sch_direct_xmit() takes the out-of-bounds txq->_xmit_lock and writes txq->xmit_lock_owner past the end of the allocation on non-lltx devices such as ifb. Drivers that index their rings from skb_get_queue_mapping() (ex: ifb) go out of range the same way. - On enqueue, taprio_enqueue() indexes q->qdiscs[] with the same value; the array is sized dev->num_tx_queues. A duplicating netem on a child qdisc clones the skb and re-enqueues the clone at the taprio root (qdisc_root_bh()) through a direct rootq->enqueue() call, so the value reaches taprio_enqueue() before any dequeue cap runs. The Fix: Cap the mapping at the consumers and write the capped value back, so the qdisc's _tx[] lookup, the taprio q->qdiscs[] lookup and any later driver read stay in range. On the dequeue side the cap folds into the dequeue path instead of a separate walk over the returned list: dequeue_skb() caps the head, and each bulk helper caps the element it appends. For a qdisc that is not TCQ_F_ONETXQUEUE, netdev_cap_txqueue() logs a ratelimited notice and selects queue 0. A TCQ_F_ONETXQUEUE qdisc is bound to a single tx queue: dequeue_skb() uses q->dev_queue for the frozen check, the BQL limit and the trace point, while qdisc_restart() transmits on the skb's queue. Capping an out-of-range mapping to queue 0 would therefore send a per-queue child's skb through a different queue than the one it was dequeued for. On mq, mqprio, taprio and htb offload children, a stopped queue 0 could then park the list in q->gso_skb, and the later wake would schedule queue 0's own child instead of this qdisc. Pin a TCQ_F_ONETXQUEUE mapping to the qdisc's own queue to keep the documented invariant. The pin covers every mapping, not only an out-of-range one: an in-range mapping that a filter sets on one of those per-queue children is rewritten to the child's own queue, and this path does not call netdev_cap_txqueue(), so it logs no notice. On a single-queue device qdisc_create() sets TCQ_F_ONETXQUEUE, so the one-queue trigger	N/A	More Details

	described below also takes this branch and is corrected without a notice. That also keeps a bulk list on one queue by construction: every element of a TCQ_F_ONETXQUEUE qdisc, including a follower changed by a filter, is pinned to q->dev_queue before it is appended. The fast bulk helper therefore needs no head comparison and no skb_bad_txq requeue; the slow helper, used only by non-ONETXQUEUE qdiscs, still compares and requeues a mismatched follower, where a wake on any queue reschedules q. On the enqueue side taprio_enqueue() caps the mapping against real_num_tx_queues and writes it back before selecting q->qdiscs[queue], so a duplicating child cannot steer the clone out of the array. Conditions to recreate the bug: The dequeue path: load a SCHED_CLS or SCHED_ACT program that stores __sk_buff->queue_mapping (a tc BPF action is enough), attach it to a transmit qdisc (for example a matchall filter on a prio root), and send a packet. With a 1-queue device a store of 1 indexes one past dev->_tx[]. Under KASAN with panic_on_warn=1 the unfixed kernel faults with "BUG: KASAN: slab-out- ---truncated---		
CVE-2026-99726	In the Linux kernel, the following vulnerability has been resolved: drm/v3d: Widen cache_clean_lock over the whole L2TCACTL sequence v3d_clean_caches() and v3d_flush_l2t() both write the single L2TCACTL register and poll its status bits. The mutex cache_clean_lock exists to serialize them, but v3d_clean_caches() only took the lock around its final FLM_CLEAN write. These functions run concurrently: v3d_flush_l2t() is issued from the BIN/RENDER/CSD invalidate path while v3d_clean_caches() runs from the CACHE_CLEAN queue, and each queue's scheduler uses its own ordered workqueue, so their run_job callbacks execute in parallel. Because clean locked only its final write, a concurrent flush can write L2TCACTL during clean's unlocked phase. Both use non read-modify-write writes to the one register, so whichever lands last wins: clean's TMUWCF write can land on the flush's in-flight L2TFLS invalidate, triggering the GFXH-1897 hazard of writing L2TCACTL while a flush is pending. Hold cache_clean_lock across the entire L2TCACTL access sequence so it is fully mutually exclusive with v3d_flush_l2t(), which already takes the lock around its own write.	N/A	More Details
CVE-2026-113347	Multiple products produced by Ruijie Networks Co., Ltd. contain an OS command injection vulnerability that may lead to arbitrary OS commands being executed by an attacker with the administrative privilege.	N/A	More Details
CVE-2026-15223	An improper export control defect in ChatEventDispatchActivity (com.vivo.agent.business.chatmode.activity.ChatEventDispatchActivity) of Vivo Assistant allows a local malicious application to execute voice commands without user interaction via a crafted Intent, which may lead to unauthorized phone calls, SMS sending, access to private data such as location and calendar entries, and modification of system settings	N/A	More Details
CVE-2026-130493	A hard-coded cryptographic key vulnerability exists in affected Ethernet swiches. The use of a shared cryptographic key with insufficient key strength could allow an unauthenticated remote attacker, under certain conditions, to impersonate affected devices and potentially compromise the confidentiality and integrity of communications.	N/A	More Details
CVE-2026-130494	An improper firmware authenticity verification vulnerability exists in the affected Ethernet Switches products. An authenticated remote attacker with administrative privileges could exploit the lack of cryptographic signature verification to install unauthorized firmware, potentially leading to arbitrary code execution on the affected device.	N/A	More Details
CVE-2026-99566	In the Linux kernel, the following vulnerability has been resolved: assoc_array: keep gc's cleanup out of the live tree on allocation failure assoc_array_gc() builds a copy of the tree. Where the copy reaches a shortcut, it allocates the copy of the shortcut with kmalloc_flex(), fills it with memcpy(), links it into the new tree, and only then allocates the copy of the node after it. The memcpy() leaves the new shortcut's next_node pointing at the old tree's node, so if that second allocation fails, the cleanup, assoc_array_destroy_subtree() on the new tree, follows the new shortcut into the old tree, and the BUG_ON() on the old node's back pointer fires. Allocate the node first, and the shortcut once the node exists, freeing the node again if the shortcut cannot be allocated. A copied shortcut is then only linked into the new tree when the copy of its node is there to follow it. Nothing else changes: with no allocation failing, the trees gc builds are the same as before. Found by a userspace test that makes each allocation of each operation fail in turn.	N/A	More Details
CVE-2026-99568	In the Linux kernel, the following vulnerability has been resolved: iommufd/iova_bitmap: Ensure that iova_bitmap_set() handles the full IOVA Generally userspace is going to step the dirty bitmap through the IOVA space. However if the bitmap ends in the middle of a large PTE this is problematic. The first step will clear the dirty bit and report dirty to the end of its bitmap while the second step will see the clear'd dirty and report clean for the start of the bitmap. If it happens this ends up losing dirties. Further this function is missing bounds checking so if this does happen ther is a good chance it will write memory beyond what the userspace was expecting. Return an error code if this illegal condition happens.	N/A	More Details
CVE-2026-99569	In the Linux kernel, the following vulnerability has been resolved: f2fs: release a removed xattr node in roll-forward recovery When xattrs shrink enough to fit in the inline xattr area, write_all_xattrs() truncates the xattr node and clears i_xattr_nid. If the file is then fsynced without a checkpoint, only the inode block is written, with i_xattr_nid of 0. Roll-forward recovery starts from the checkpointed inode, whose i_xattr_nid still points to the old xattr node. f2fs_recover_inline_xattr() copies only the inline xattr area, and f2fs_update_inode() writes the old i_xattr_nid back. So the inode keeps the stale xattr node: - The node block is never released, and fsck reports "nonzero bytes in xattr space after end of list". - If the inline xattr list ends exactly at the end of the inline area, its terminator is read from the stale node, and the removed xattrs come	N/A	More Details

	back after recovery. When the fsynced inode has no xattr node but the checkpointed one has, release the old xattr node as truncate_node() does: check its block address, invalidate the block and drop its node cache, which may still be dirty.		
CVE-2026-99696	In the Linux kernel, the following vulnerability has been resolved: ocfs2: find the xattr entry again after defragmenting a bucket ocfs2_xattr_set_entry_bucket() hands loc.xl_entry to ocfs2_xa_set() twice: once for the bucket as it stands, and, if that reports -ENOSPC, once more after ocfs2_defrag_xattr_bucket() has packed the name/value pairs down. Defragmenting ends by sorting the entry array on xe_name_hash. Buckets are kept sorted, so the sort only permutes entries that share a hash, but lib/sort.c is an unstable heapsort and it does permute them. loc.xl_entry is an address inside xh_entries[], so afterwards it denotes whichever entry landed in that slot rather than the one the search found. The second ocfs2_xa_set() then works on the wrong entry. Nothing on that path compares the entry's name with xi->xi_name, so what happens depends on how the two entries compare in size: ocfs2_xa_reuse_entry() trips its BUG_ON() when the padded name sizes differ, and otherwise the new value is written into the neighbouring attribute's value slot while the attribute being set is left untouched. setxattr() reports success either way. Remember the entry's hash before defragmenting and look the entry up again afterwards, the way ocfs2_xattr_bucket_find() does. An entry that cannot be found again means the bucket no longer holds what the search found, so fail with -EFSCORRUPTED rather than writing through the stale pointer.	N/A	More Details
CVE-2026-125411	Pleasanter contains an incorrect authorization vulnerability in the access control of Items. If this vulnerability is exploited, a remote authenticated attacker who can view records in the affected product may obtain the values of Items for which read access is restricted by Column access control. Note that the vulnerability affects configurations where read access conditions are set in Column access control settings.	N/A	More Details
CVE-2026-99570	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to invalidate curseg if change_curseg() fails Our consistency stress test (fsstress + fio + atomic write workers + "f2fs_io shutdown 1") triggers below kernel BUG: F2FS-fs (vdd): Stopped filesystem due to reason: 0 ------------[cut here]------------ kernel BUG at fs/f2fs/segment.c:3989! Oops: invalid opcode: 0000 [#1] SMP KASAN PTI Workqueue: writeback wb_workfn (flush-253:48) RIP: 0010:f2fs_allocate_data_block+0x257a/0x38e0 Call Trace: <TASK> do_write_block+0x155/0xc10 f2fs_outplace_write_data+0x192/0x340 f2fs_do_write_data_page+0xc57/0x1880 f2fs_write_single_data_page+0xda9/0x15a0 f2fs_write_cache_pages+0x973/0x1f20 __f2fs_write_data_pages+0x6c4/0x1890 do_writepages+0x21e/0x560 __writeback_single_inode+0x119/0x11b0 writeback_sb_inodes+0x67d/0x17a0 wb_writeback+0x207/0x820 wb_workfn+0x128/0xc30 process_one_work+0x925/0x1600 worker_thread+0x60a/0xff0 kthread+0x35f/0x460 ret_from_fork+0x474/0x7e0 ret_from_fork_asm+0x1a/0x30 </TASK> The BUG is f2fs_bug_on(sbi, curseg->next_blkoff >= BLKS_PER_SEG(sbi)) in f2fs_allocate_data_block(), next_blkoff was 512. The root cause is a race in between atomic write commit and shutdown as below: Thread A (commit atomic write) Thread B (shutdown) Thread C (writeback) - f2fs_ioc_commit_atomic_write - f2fs_commit_atomic_write - __f2fs_commit_atomic_write - __replace_atomic_write_block - f2fs_replace_block - __f2fs_replace_block - f2fs_ioc_shutdown - f2fs_do_shutdown : F2FS_GOING_DOWN_METASYNC - f2fs_sync_fs - f2fs_write_single_data_page : pass f2fs_cp_error() check - f2fs_stop_checkpoint : set CP_ERROR_FLAG - change_curseg - reset_curseg : curseg->segno = new_segno : curseg->alloc_type = SSR : curseg->next_blkoff = __next_free_blkoff() which is BLKS_PER_SEG if there is no free slot in new_segno - f2fs_get_sum_cache : return -EIO due to f2fs_cp_error() : return error w/o restoring curseg : goto out_unlock, skip restoring old segno, blkoff and alloc_type - f2fs_do_write_data_page - f2fs_outplace_write_data - do_write_block - f2fs_allocate_data_block : curseg->segno is valid : f2fs_bug_on(next_blkoff >= BLKS_PER_SEG) Since commit 86f33603f8c5 ("f2fs: handle errors of f2fs_get_meta_page_nofail"), change_curseg() can fail after the curseg has already been switched to the new segment, then the curseg is left in an inconsistent state: it points to a segment which may have no free slot, its summary block is zeroed, and the caller can not restore it. Subsequent block allocation on the curseg will trigger the BUG. Moreover, the writeback thread dies with curseg_lock, curseg_mutex and sentry_lock held, so the commit, shutdown and other I/O threads hang in D state. Since change_curseg() only fails once CP_ERROR_FLAG is set, no checkpoint will be committed afterward, let's invalidate the curseg via reset_curseg_fields() on failure, as new_curseg() does for -ENOSPC, then f2fs_allocate_data_block() will fail with -ENOSPC on the curseg instead of triggering BUG, and do_write_block() can handle the error since CP_ERROR_FLA ---truncated---	N/A	More Details
CVE-2026-61827	netty-incubator-codec-ohttp implements Oblivious HTTP (OHTTP) gateway and client functionality using Netty. Prior to 0.0.23.Final, BinaryHttpParser does not enforce limits on encoded variable-length fields controlled by a remote peer. A remote peer can cause the decoder to buffer data indefinitely until memory is exhausted, causing denial of service. This issue is fixed in version 0.0.23.Final.	N/A	More Details
CVE-2026-121649	Use of cache containing sensitive information vulnerability in Apache Tomcat's ExpiresFilter caused by the filter removing cache-control header fields set by the authentication layer. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.27, from 10.1.0-M1 through 10.1.61, from 9.0.0.M1 through 9.0.123. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.1 through 7.0.109. Older unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.29, 10.1.62 or 9.0.124, which fix the issue. Note: Tomcat 11.0.28 was not released	N/A	More Details
CVE-2026-	Pleasanter contains an issue with improper limitation of a pathname to a restricted directory ('Path Traversal') in the API. If this vulnerability is exploited, an arbitrary file on the server where the product is	N/A	More

121811	running may be obtained by a remote authenticated attacker.		Details
CVE-2026-99649	In the Linux kernel, the following vulnerability has been resolved: mm/ksm: check large swap-cache folios before mapping them KSM folios are always small, but their shared swap slots can be read back as part of a large folio in the swap cache. ksm_might_need_to_copy() skips large folios, so another process can map a page with an incompatible anon_vma or page index. Reverse mapping then cannot find the new mapping, preventing migration or reclaim. Apply the same compatibility checks to large folios as to small ones. Pass the page corresponding to the swap entry and copy only that page when necessary. In do_swap_page()'s copy path, decrement only that slot's swap count, because the source swap-cache folio can now be large. Fix generated by an LLM and sent by Chris Mason in private.	N/A	More Details
CVE-2026-130163	Collabora Online is a collaborative online office suite based on LibreOffice technology. Prior to 24.04.19.3, 25.04.13.3, and 26.04.3.3, the RouteToken query parameter is reflected into the unauthenticated admin-console page without HTML encoding. The admin-console page is served without authentication. A victim must load a crafted URL, and token exposure requires the victim to be an administrator already signed in to the console; the reflected value can run script in the Collabora Online origin and expose the console authentication token. This issue is fixed in versions 24.04.19.3, 25.04.13.3, and 26.04.3.3.	N/A	More Details
CVE-2026-126966	SmarterTools SmarterMail Missing Authorization Privilege Escalation Vulnerability. This vulnerability allows remote attackers to escalate privileges on affected installations of SmarterTools SmarterMail. Authentication is required to exploit this vulnerability. The specific flaw exists within the svcDomainAdmin.asmx endpoint. The issue results from the lack of authorization prior to allowing access to functionality. An attacker can leverage this vulnerability to escalate privileges to resources normally protected from the user. Was ZDI-CAN-33566.	N/A	More Details
CVE-2026-110307	An incorrect authorization vulnerability was identified in GitHub Enterprise Server that allowed an authenticated user whose numeric user ID matched an enterprise's numeric ID to manage that enterprise's GitHub Apps and generate private keys without enterprise membership. GitHub App ownership checks compared numeric IDs without checking the owner's account type. Exploitation required control of a session for the colliding user account and an enterprise-owned GitHub App. A generated private key could be used to authenticate as the App and access repositories where it was installed, within its granted permissions. On standard fresh installations, the documented collision involved the non-loginable built-in ghost account; the reported GHES reproduction used site administrator impersonation. This vulnerability affected supported GitHub Enterprise Server releases prior to 3.23 and was fixed in versions 3.19.16, 3.20.12, 3.21.10, and 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-99658	In the Linux kernel, the following vulnerability has been resolved: mm/mmu_notifier: finish non-blocking interval invalidation on error When an ordinary MMU notifier rejects a non-blocking invalidation, the interval invalidation started beforehand is left unfinished: the caller skips mmu_notifier_invalidate_range_end(). Subsequent interval notifier operations can then wait forever. The only in-tree trigger is nouveau SVM, which uses both notifier types and rejects non-blocking invalidations. Reaping via the OOM reaper or process_mrelease() can therefore leave the GPU fault worker stuck, holding a reference that prevents the dying process's address space and memory from being freed. End the interval invalidation when an ordinary notifier fails a non-blocking invalidation, like mn_itree_invalidate() already does when an interval notifier fails. Make that handling clearer in the documentation of mmu_notifier_invalidate_range_start_nonblock(). Fix generated by an LLM and sent by Josef Bacik in private. [david: simplify description, update function doc]	N/A	More Details
CVE-2026-99727	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: properly handle refcounted PUD mappings on fork copy_huge_pud() copies device-DAX entries without taking a folio reference or updating mapcounts and the child's RSS. For example, MADV_GUARD_INSTALL sets VM_MAYBE_GUARD on a device-DAX VMA, forcing copy_page_range() to traverse the whole VMA on fork. Any remaining huge PUDs are then copied without accounting for the new mappings. Unmapping them in the child decrements accounting that was never added. Use vm_normal_page_pud() to identify refcounted mappings. Take a folio reference and add the child's PUD rmap and RSS accounting before installing the copied entry. Pass src and dst VMAs consistently instead of src and dst MM. Move the TODO.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2/cluster: unlink socket containers before releasing resources A socket container's node and socket must remain alive while the container is on sock_containers. The debugfs formatter reads sc_node and sc_sock under o2net_debug_lock without taking their references, but sc_kref_release() releases both before taking that lock to unlink the container. With CONFIG_DEBUG_FS enabled, a synchronous sender can retain its last socket-container reference after o2net_disconnect_node() has detached the connection and flushed o2net_wq. The caller runs outside that workqueue. After the sender removes its configfs dependency, node removal can proceed without waiting for the sender's final sc_put(): Sender (final sc_put) Configfs rmdir Debugfs reader ---------- ---------- ---------- o2nm_undepend_item() Drop node item refs o2nm_node_put() [node can be freed] sc_seq_show() Lock debug lock Read sc_node->nd_name Unlock debug lock sc_node = NULL o2net_debug_del_sc() The earlier socket release has a similar window. The container is still selectable while its socket has been released and sc_sock has not yet been cleared: Sender (final sc_put) Debugfs reader ---------- ---------- sock_release(sc->sc_sock) [socket can be freed] sc_seq_show() Lock debug lock Read sc_sock->sk fields Unlock debug lock sc_sock = NULL ... o2net_debug_del_sc() Move o2net_debug_del_sc() before releasing		

CVE-2026-99691	either resource. Unlinking under the existing debug lock waits for an active formatter to finish and prevents later readers from selecting the container. This covers both socket fields and node fields, including nd_num in the stats formatter. KASAN report as below: BUG: KASAN: slab-use-after-free in string+0x42b/0x470 Read of size 1 at addr ffff88810f56fc90 by task python3/506 CPU: 0 UID: 0 PID: 506 Comm: python3 Not tainted 7.3.0-rc4-next-20260921-pmb-ocfs2-functional-v1+ #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 Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 kasan_report+0xe0/0x110 string+0x42b/0x470 vsnprintf+0x330/0x1110 seq_printf+0x1a8/0x260 sc_seq_show+0x604/0x9f0 seq_read_iter+0x480/0x1110 seq_read+0x362/0x5a0 full_proxy_read+0x11f/0x180 vfs_read+0x1a0/0xc20 ksys_read+0x111/0x200 do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f (User register dump omitted.) </TASK> Allocated by task 497: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0xaa/0xb0 __kmalloc_cache_noprof+0x28d/0x610 o2nm_node_group_make_item+0xea/0x1a0 configfs_mkdir+0x762/0xca0 vfs_mkdir+0x399/0x7e0 filename_mkdirat+0x129/0x450 __x64_sys_mkdir+0x70/0xa0 do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f Freed by task 505: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 kasan_save_free_info+0x3b/0x60 __kasan_slab_free+0x5f/0x80 kfree+0x308/0x580 config_item_cleanup+0x14b/0x230 config_item_put+0x79/0x90 sc_put+0x454/0x6e0 o2net_send_message_vec+0x63d/0x1560 o2net_send_message+0xa6/0xe0 validation_send_write+0x17f/0x1d0 full_proxy_write+0x11f/0x180 vfs_write+0x25a/0x1010 ksys_write+ ---truncated---	N/A	More Details
CVE-2026-99692	In the Linux kernel, the following vulnerability has been resolved: ocfs2: stop recovery before shutting down local allocators Recovery completion needs the local allocators until it has finished reclaiming a failed node's allocation window and truncate log. During unmount, quota recovery is stopped first, but that leaves recovery_state at OCFS2_REC_QUOTA_DISABLED. New non-quota recovery can still start. ocfs2_truncate_log_shutdown() drains existing completion work and releases osb_tl_bh before ocfs2_recovery_exit() closes that publication path: Unmount Recovery ------- -------- stop quota recovery; disable quotas truncate_log_shutdown() flush existing completion work recover another failed node queue completion flush local log; release buffer complete_truncate_log_recovery() use released osb_tl_bh recovery_exit(): disable and drain An extra buffer reference alone would still let completion append new records after the final truncate-log flush. Split stopping and draining recovery from freeing its state. Stop all recovery producers and drain completion work before shutting down either local allocator. Keep recovery-state release at its existing later point, so its lifetime relative to other teardown users stays unchanged. The second stop at recovery_exit() is harmless once recovery is disabled. KASAN report as below: BUG: KASAN: slab-use-after-free in ocfs2_truncate_log_needs_flush+0x23d/0x360 Read of size 8 at addr ffff8881093711e8 by task kworker/u16:3/53 CPU: 0 UID: 0 PID: 53 Comm: kworker/u16:3 Not tainted 7.3.0-rc4-next-20260921-pmb-ocfs2-functional-v1+ #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: ocfs2_wq ocfs2_complete_recovery Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 ? ocfs2_truncate_log_needs_flush+0x23d/0x360 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x20e/0x420 ? ocfs2_truncate_log_needs_flush+0x23d/0x360 kasan_report+0xe0/0x110 ? ocfs2_truncate_log_needs_flush+0x23d/0x360 ocfs2_truncate_log_needs_flush+0x23d/0x360 ? __pfx_ocfs2_truncate_log_needs_flush+0x10/0x10 ? __pfx_down_write+0x10/0x10 ocfs2_complete_truncate_log_recovery+0x315/0x780 ? trace_hardirqs_on+0x18/0x160 ? __pfx_ocfs2_complete_truncate_log_recovery+0x10/0x10 ? ocfs2_complete_recovery+0x666/0x11b0 ? srso_alias_return_thunk+0x5/0xfbef5 ? kfree+0x308/0x580 ocfs2_complete_recovery+0x676/0x11b0 ? __pfx_ocfs2_complete_recovery+0x10/0x10 ? reacquire_held_locks+0xdd/0x200 ? process_one_work+0x7c8/0x1b40 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_acquire+0x190/0x300 ? __pfx_autoremove_wake_function+0x10/0x10 ? process_one_work+0x834/0x1b40 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 ? srso_alias_return_thunk+0x5/0xfbef5 process_one_work+0x9a8/0x1b40 ? __pfx_process_one_work+0x10/0x10 ? lock_acquire+0x190/0x300 ? lock_is_held_type+0x8f/0x100 ? srso_alias_return_thunk+0x5/0xfbef5 worker_thread+0x65c/0xe40 ? srso_alias_return_thunk+0x5/0xfbef5 ? __kthread_parkme+0x177/0x220 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_worker_thread+0x10/0x10 kthread+0x351/0x460 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x659/0x940 ? __pfx_ret_from_fork+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __switch_to+0x74f/0xf70 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1a/0x30 </TASK> Allocated by task 538: kasan_save_stack+0x33/0x60 kasan ---truncated---	N/A	More Details
CVE-2026-112740	A command injection vulnerability was identified in GitHub Enterprise Server that allowed an authenticated site administrator to execute arbitrary commands as the host git user on peer Git server nodes outside the pre-receive hook sandbox by supplying a crafted pre-receive environment download URL. Exploitation required a clustered deployment using custom pre-receive hook environments and at least one peer Git server node. This vulnerability affected all versions of GitHub Enterprise Server prior to 3.23 and was fixed in versions 3.22.5, 3.21.10, 3.20.12, and 3.19.16. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2/dlm: serialize migration MLE cleanup on failure A migration MLE has a base reference for the mastering answer and an in-use reference held by dlm_migrate_lockres(). The finishing assertion consumes the base reference, but migration failure drops it again without checking whether the assertion already consumed it. The target queues its migration worker		

CVE-2026-99693	before it acknowledges the ALL_DONE packet, so its finishing assertion can be handled while the sender still waits for the status. If the connection fails before the status arrives, dlm_send_one_lockres() returns an error for a migration that the target has already completed: Old-master sender Finishing assertion ---------- ---------- publish MLE; take in-use reference send ALL_DONE; wait for status take lookup reference unlink MLE put lookup and base references send returns a connection error put base reference again; free MLE put in-use reference; access freed MLE dlm_assert_master_handler() holds dlm->spinlock from the MLE lookup to the final puts, and it takes dlm->master_lock around the unlink and the puts. Decide under the same locks: if the MLE is already unhashed, the assertion consumed the base reference, so drop only the sender's in-use reference. Otherwise unlink it before consuming the base reference so a later assertion no longer finds it. An assertion that got there first has also made the target the owner and cleared DLM_LOCK_RES_MIGRATING. Treating the attempt as failed at that point would leave the lockres owned by the target while this node still holds the other nodes' locks and refmap bits, so it could neither be purged nor migrated again. Finish our side as a completed migration does instead: drop those locks under res->spinlock if the target still owns the lockres, and recalculate its usage. Use the same cleanup for the pre-send failure, the send failure and the target-death failure paths. The sender's in-use reference keeps the MLE alive until its final put outside the nested locks. Successful migration retains its existing in-use-only put. A target that has died cannot assert after the sender has seen its death. The o2net node-down callback runs before the DLM one; it disconnects the node and flushes the workqueue that runs the message handlers, so an assertion already received has been handled by the time the DLM removes the node from the domain map. A node that leaves the domain cleanly sends its exit message only after its assertions have been answered. Two windows of the hand-over itself are left as they are: the handler switches the owner and the sender drops the other nodes' locks moments later, and an assertion from a live target that arrives after the sender has given up still meets the owner check in the handler. Waiting for the target before giving up would only stretch the first window into the time in which recovery of the target can observe it. KASAN report, only the relevant frames are kept: BUG: KASAN: slab-use-after-free in dlm_put_mle_inuse+0xce/0xe0 Read of size 8 at addr ffff8881190e3c20 by task umount/593 Call Trace: <TASK> dlm_put_mle_inuse+0xce/0xe0 dlm_migrate_lockres+0x16af/0x20d0 dlm_empty_lockres+0x60a/0xbd0 dlm_unregister_domain+0x53b/0x11d0 o2cb_cluster_disconnect+0xa5/0x110 ocfs2_cluster_disconnect+0x73/0xc0 ocfs2_dlm_shutdown+0x311/0x440 ocfs2_dismount_volume+0x696/0x940 generic_shutdown_super+0x15f/0x490 kill_block_super+0x40/0xa0 deactivate_locked_super+0xae/0x120 deactivate_super+0xa2/0xc0 cleanup_mnt+0x217/0x450 task_work_run+0x144/0x230 </TASK> Allocated by task 593: kmem_cache_alloc_noprof+0x20b/0x5c0 dlm_migrate_lockres+0x20e/0x20d0 dlm_empty_lockres+0x60a/0xbd0 dlm_unregister_domain+0x53b/0x11d0 o2cb_cluster_disconnect+0xa5/ ---truncated---	N/A	More Details
CVE-2026-99694	In the Linux kernel, the following vulnerability has been resolved: ocfs2: do not index before the entry array when dividing an empty bucket ocfs2_divide_xattr_bucket() appends an empty bucket carrying a sentinel hash when the source bucket has no divide position, taking that hash from xh_entries[start - 1]. ocfs2_xattr_find_divide_pos() returns count when every entry shares one hash, and its loop does not run at all for count < 2, so an empty bucket comes back with start == count == 0 and the read is xh_entries[-1]. That is outside the __counted_by_le(xh_count) bounds of xh_entries[] and lands on the header, so the value picked up is xh_count | xh_free_start << 16. One above it is written into the new bucket's xh_entries[0] and committed, which leaves a bucket whose first hash does not follow the bucket before it. ocfs2_xattr_bucket_find() binary searches on xh_entries[0], so the broken key hides the range the sentinel was meant to cover: getxattr() answers -ENODATA for an attribute that is stored, and an insert can go to the wrong bucket and leave two copies of a name. The cross cluster divide hands the same value back as a tree record's e_cpos, which then claims a hash range it does not cover. Read xh_entries[0] instead. The branch is only reached when every entry shares one hash, so for a non-empty bucket it holds the same value, and for an empty one it is the slot that carries the bucket key. An empty bucket reaches the divide from ocfs2_check_xattr_bucket_collision(), which reports no collision for one, through ocfs2_add_new_xattr_bucket() and ocfs2_extend_xattr_bucket(). It does not take a failed set to get there: a divide that finds no position is itself what leaves an empty sentinel bucket behind, and a later extend can pick it as the target.	N/A	More Details
CVE-2026-99695	In the Linux kernel, the following vulnerability has been resolved: ocfs2: find the xattr entry again after converting a block to a bucket ocfs2_xattr_create_index_block() turns a non-indexed xattr block into a tree by copying it into a bucket, and the copy sorts the entries: a block holds them in insertion order, a bucket by name hash. ocfs2_xattr_update_xattr_search() nevertheless relocates the caller's search by the index the entry had in the block, so unless the two orders agreed it points at an unrelated entry, and nothing between there and the write compares that entry's name with the one being set. Setting an xattr whose new value no longer fits the block arrives here with xs->not_found clear. The block gave up because ocfs2_xa_block_can_reuse() wants an exact namevalue size match, while ocfs2_xa_bucket_can_reuse() only wants an entry at least as large as the new one, so the bucket accepts a pair the block rejected and the update lands on whichever entry the sort left at that index. Equal padded name lengths get that other attribute's value replaced in place and its name left alone, so setxattr returns success while the attribute asked for keeps its old value. A longer one trips ocfs2_xa_reuse_entry()'s BUG_ON, and a shorter one is refused by the bucket as well, so ocfs2_xa_prepare_entry() wipes the entry and writes the new name over it, leaving xe_name_hash describing a name that is no longer there. Find the entry again by the name being set, matching the type as well. The copy carries every entry and every name over unchanged, and (type, name) pairs are unique within a block, so the match is exact. The type matters because the handlers strip their	N/A	More Details

	prefix before calling ocfs2_xattr_set(), so user.foo and trusted.foo share a stored name, a stored length and a hash, and both can sit in the same block. The name has to come from xi rather than from xs->here: a failed set can remove the entry xs->here names, and only the "growing" caller sets ctxt->set_abort, which is what stops the conversion. Scan every entry rather than narrowing on the hash the way ocfs2_find_xe_in_bucket() does, since an existing entry keeps the hash it was stored under, which a kernel that sign-extended the name bytes made different from the one computed now. An empty result means the entry did not survive the failed set and this has become an insert; leaving xs->here alone would keep it pointing into the block, which is rewritten as a tree root before the caller reads it again.		
CVE-2026-111801	An authorization bypass through user-controlled key vulnerability was identified in GitHub Enterprise Server that allowed authenticated users to read review-thread data from private repositories they could not access. An attacker required an authenticated account, a readable repository and pull request to use as the request context, and target review-thread identifiers. The thread-preview endpoint resolved supplied identifiers without restricting them to the authorized pull request, exposing private file paths and, when positioning and diff data were available, source-code excerpts, commit identifiers, and line metadata without victim interaction. This vulnerability affected supported GitHub Enterprise Server releases in the 3.21, 3.22 series and was fixed in versions 3.21.10, 3.22.5. This vulnerability was reported via the GitHub Bug Bounty program.	N/A	More Details
CVE-2026-126975	llama.cpp GGUF File Parsing Heap-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of llama.cpp. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the parsing of GGUF files. The issue results from the lack of proper validation of the length of user-supplied data prior to copying it to a heap-based buffer. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current process. Was ZDI-CAN-34558.	N/A	More Details
CVE-2026-99656	In the Linux kernel, the following vulnerability has been resolved: mm/madvise: fix data loss from MADV_FREE on a KSM page MADV_FREE can clear the dirty flag of a singly mapped KSM folio. The folio remains in the stable tree and can gain new mappings without being dirtied again. Once marked lazyfree, reclaim can discard data that those mappings still need. KSM always marks a folio dirty when converting it to a KSM folio in try_to_merge_one_page(). Marking a folio lazyfree clears PG_swapbacked but leaves PG_dirty unchanged. Preserving PG_dirty makes reclaim restore PG_swapbacked instead of discarding the folio. This also protects folios that were already lazyfree when KSM converted them. Skip KSM folios before clearing their dirty flag, checking under the folio lock to serialize with KSM conversion.	N/A	More Details
CVE-2026-99655	In the Linux kernel, the following vulnerability has been resolved: mm/memory: reject remote COW faults on PFN mappings When handling a COW fault without a normal source page (PFN mappings), __wp_page_copy_user() copies from the faulting user address. For remote faults, this accesses the wrong address space, resulting in wrong data getting copied. For example, user_events updates an enable bit in the registering process's memory. If that memory needs to be faulted in, user_event_enabler_fault_fixup() calls fixup_user_fault() from a work item. Placing the enable bit in a private PFN mapping can trigger this remote COW path. Reject these faults with VM_FAULT_SIGBUS when the target mm differs from current->mm: remote faults. Ordinary (non-PFN) mappings are not affected. The user-address copy dates back to commit 6aab341e0a28 ("mm: re-architect the VM_UNPAGED logic"); likely, actual ways to get remote faults in there on COW PFN mappings appeared afterwards. But the fix is simple and straight forward. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify comment and description]	N/A	More Details
CVE-2026-126976	llama.cpp MTMD Multimodal GGUF File Parsing Integer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of llama.cpp. User interaction is required to exploit this vulnerability in that the target must open a malicious file. The specific flaw exists within the parsing of audio multimodal projector GGUF files. The issue results from the lack of proper validation of user-supplied data, which can result in an integer overflow before allocating a buffer. An attacker can leverage this vulnerability to execute arbitrary code in the context of the current process. Was ZDI-CAN-30980.	N/A	More Details
CVE-2026-99653	In the Linux kernel, the following vulnerability has been resolved: mm: fix unkillable fault loop on MAP_DROPPABLE memory handle_mm_fault() suppresses VM_FAULT_OOM for VM_DROPPABLE mappings. For user-mode faults, pending signals are processed before retrying the access. Kernel-mode faults on user buffers instead return directly to the faulting instruction. With persistent allocation failures (esp. memcg limits), a syscall, such as write(), can therefore loop while not dropping locks. Only suppress VM_FAULT_OOM when FAULT_FLAG_USER is set, allowing kernel-mode faults and GUP to fail and signal allocation errors. Fix generated by an LLM and sent by Chris Mason in private. [david: simplify description]	N/A	More Details
CVE-2026-92182	Netgate pfSense CE XMLRPC Incorrect Authorization Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of Netgate pfSense. Authentication is required to exploit this vulnerability. The specific flaw exists within the XMLRPC endpoint. The issue results from incorrect authorization prior to allowing access to functionality. An attacker can leverage this vulnerability to execute code within the context of the root. Was ZDI-CAN-32353.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm: fix deadlock of mremap() against UFFDIO_MOVE With per-VMA locks, UFFDIO_MOVE can run concurrently with mremap() on different VMAs in		

CVE-2026-99652	the same mm. Those VMAs can share page tables, so both operations can acquire the same pair of page-table locks. UFFDIO_MOVE orders the locks by address, while mremap() takes the source lock first. These conflicting orders can deadlock; mremap() holding mmap_lock for writing no longer helps. Move double_pt_lock() and double_pt_unlock() to mm/internal.h and use them for all page-table moves in mremap(). Preserve the source entry checks after acquiring both locks. Remove the comments claiming that exclusive mmap_lock makes lock ordering irrelevant. In move_ptes(), make sure to just recheck source and destination PMD after taking the locks and drop the overly-long comment. Fix generated by an LLM and sent by Chris Mason in private. [david: use new helpers consistently; simplify description and comments; move_ptes() fixup]	N/A	More Details
CVE-2026-99651	In the Linux kernel, the following vulnerability has been resolved: mm/ksm: reset the unstable trees of all nodes Each full scan resets ksm_nr_node_ids unstable tree roots. Enabling merge_across_nodes resets this count to 1, even if per-node roots were previously allocated. So when toggling "merge_across_nodes" 0 -> 1 -> 0, a scan that starts while the setting is 1 and continues after switching back to 0 can encounter stale unstable tree roots, resulting in UAF. Only root is allowed to change "merge_across_nodes", and the documented sequence is: "merge_across_nodes setting can be changed only when there are no ksm shared pages in the system: set run 2 to unmerge pages first, then to 1 after changing merge_across_nodes, to remerge according to the new setting." Let's fix it by tracking in ksm_nr_node_ids the actual number of allocated tree roots. remove_all_stable_nodes() and ksm_check_stable_tree() will now scan more tree roots (skipping the empty ones), which is harmless.	N/A	More Details
CVE-2026-126977	Philips Hue Bridge HomeKit Pair-Setup Default Password Authentication Bypass Vulnerability. This vulnerability allows network-adjacent attackers to bypass authentication on affected installations of Philips Hue Bridge. Authentication is not required to exploit this vulnerability. The specific flaw exists within the HomeKit Authentication service, which listens on TCP port 8080 by default. Crafted data in a pairing request can cause the device to allow authentication using a blank password. An attacker can leverage this vulnerability to bypass authentication on the system. Was ZDI-CAN-28323.	N/A	More Details
CVE-2026-125276	Pleasanter contains an issue with insertion of sensitive information into sent data in the API. If the affected product receives a specially crafted API request sent by a remote authenticated attacker, the authentication credentials of other users, including administrators, such as API keys, may be obtained. As a result, the attacker may perform arbitrary operations on the affected product with the privileges associated with the obtained credentials.	N/A	More Details
CVE-2026-99381	In the Linux kernel, the following vulnerability has been resolved: smb: client: reject short or out-of-range FIND response parameters CIFSFindFirst() and CIFSFindNext() read the SearchHandle, SearchCount, EndofSearch and LastNameOffset of a TRANS2_FIND_FIRST2 or TRANS2_FIND_NEXT2 response at ParameterOffset. validate_t2() bounds the offset at 1024 and ParameterCount below 512, and nothing compares them with the number of bytes SendReceive() copied into the buffer or with the size of the response parameters. smb_init() places the request and the response in the same cifs_request allocation and cifs_buf_get() clears only the start of it, so a short response with a ParameterOffset past its end has the search handle, the entry count, the end-of-search flag and the last-entry offset read out of the request that was just sent, or out of whatever the allocation held before, and a ParameterCount smaller than the parameters has them read past the end of the parameter block. The reads stay inside the allocation, but what they return is not part of the response, and CIFSFindFirst() then hands back a search handle the server never sent. That handle is sent back to the server in the next FIND_NEXT2 or FIND_CLOSE2, so two bytes of the buffer leak at an offset the server picks. SMB1 is not negotiated by default; reaching this code requires an explicit vers=1.0 mount. Reject a response whose parameter area is smaller than the response parameters or extends past the received length, returning smb_EIO2() with a new smb_eio_traces value for each of the two responses, like the other response checks this file gained since v6.19. Backports to trees before v6.19 need smb_EIO2() replaced with -EINVAL and the trace.h hunk dropped. A conforming server is not affected. The parameter block of a transaction response is part of the SMB, so ParameterOffset + ParameterCount cannot exceed the smbCalcSize() bytes that SendReceive() copies, and the response parameters are 10 bytes for FIND_FIRST2 and 8 bytes for FIND_NEXT2, which is what T2_FFIRST_RSP_PARMS and T2_FNEXT_RSP_PARMS hold. This patch depends on "smb: client: reject FIND data areas that run past the received response", whose declarations it extends.	N/A	More Details
CVE-2026-125288	ServiceNow has addressed a code injection security issue that was identified in the ServiceNow AI Platform. This issue could enable an authenticated user with template editing privileges to execute code in the context of another user's session. ServiceNow deployed an update to hosted instances, and provided the update to our partners and self-hosted customers. We are not currently aware of malicious exploitation against ServiceNow instances. We recommend customers promptly apply appropriate updates or upgrade to a patched release if they have not already done so.	N/A	More Details
CVE-2026-87868	MRemoteNG uses single MD5 hash with no salt and no iterations for key-derivation. An attacker who obtains the encrypted data can brute-force a weak or moderate passphrase offline at extreme speed and then decrypt every stored connection password. When the default passphrase is in use, no cracking is needed at all; the key is derived from a publicly known value. Vendor did not provide the range of vulnerable versions. Only version 1.76.20 was tested and confirmed vulnerable. Other versions may also be vulnerable.	N/A	More Details
CVE-			

2025-59910	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-133175	MISP contains an authorization bypass in its server-to-server proposal synchronization mechanism. Two related flaws were identified: 1. The events/pushProposals endpoint located the target event by UUID without enforcing any access-control check. A sync user with perm_sync could push shadow-attribute proposals to any event on the local instance, including events with restricted distribution (e.g., distribution level 0, visible only to the owning organisation). The stored proposal also carried the sender-supplied event_uuid and event_org_id rather than the local event's actual identity, corrupting the data model. 2. When accepting an edit proposal (old_id != 0), the code resolved the target live attribute by UUID alone (Attribute.uuid = shadow.uuid) without scoping the lookup to the proposal's own event. A malicious sync user could therefore craft a proposal whose uuid matched an attribute in a different event; upon acceptance, the wrong attribute would be modified or deleted. The same unscoped lookup existed in the ShadowAttribute::acceptProposal path. Preconditions: an authenticated user with sync permission (perm_sync) on a MISP instance that participates in server-to-server synchronisation. The attacker must know or enumerate the UUID of a target event or attribute. Security impact: unauthorised read access to restricted events, unauthorised insertion of proposals onto events the user cannot see, and unauthorised modification or deletion of attributes belonging to other events. This compromises the confidentiality and integrity of threat-intelligence data across organisational boundaries within a MISP deployment.	N/A	More Details
CVE-2026-133177	MISP contains an authorization flaw in the sharing group capture logic. When a sharing group is re-captured with a newer modified timestamp and an updated roster, the application unconditionally overwrote the type, nationality, sector, and contacts fields of every organisation listed in the roster, including the host organisation. This behaviour was reachable by any user holding the perm_sync role, and additionally by any external server pulled from or any feed fetched, because those operations execute with the privileges of the site administrator who initiated them. As a result, a low-privileged sync user or a malicious/compromised external feed or pull source could silently corrupt organisation records on the receiving MISP instance without any site-admin or internal-sync authorization. The vulnerability does not disclose data or affect availability; the impact is limited to integrity of organisation metadata (type, nationality, sector, contacts). The fix restricts organisation detail updates to internal sync operations (pulls over links flagged as internal, or senders holding perm_sync_internal) and site administrators; all other sync, pull, and feed operations may only match existing organisations by UUID or create new ones, without modifying their stored details.	N/A	More Details
CVE-2026-133178	flowintel allows users with editor-level access to author case and task fields, including aggregated tags. When these records are exported as CSV, the field values are written verbatim into the output file. Because no sanitisation is applied to user-supplied text before CSV serialisation, an attacker can store values that begin with spreadsheet formula triggers (such as =, +, -, @, or their fullwidth Unicode equivalents, or leading control/whitespace characters). When a victim opens the exported CSV in a spreadsheet application (e.g. Microsoft Excel, LibreOffice Calc), the stored text is interpreted and executed as a live formula rather than displayed as literal text. This can lead to unintended data manipulation within the spreadsheet, potential data exfiltration via formula functions, or execution of downstream commands depending on the spreadsheet application and its configuration. The vulnerability requires the victim to actively open the CSV file in a spreadsheet application; the flowintel application itself is not directly compromised. Affected: <4.1.0	N/A	More Details
CVE-2026-133183	MISP contains an authorization flaw in the handling of remote pull and feed operations. When a non-internal pull or a feed processes incoming data and encounters an existing sharing group, the application did not verify that the remote server's organization was actually a member of that sharing group before allowing the group's metadata to be updated. The captureRemote mechanism previously stored only a boolean flag indicating whether the link was internal, and the authorization check in captureSGExisting relied solely on that flag. For non-internal pulls and feeds, the check was insufficient: a remote server whose organization was not a member of the target sharing group could still cause the group record to be modified (for example, altering the local flag, the modified timestamp, or other group properties). A feed, which has no associated server organization at all, could likewise trigger an update to an existing sharing group's metadata. The security impact is an integrity violation: an authenticated remote entity with sync or feed privileges can alter sharing group attributes that govern how data is distributed among organizations, potentially changing the local flag or other access-relevant properties of groups it does not belong to. The vulnerability requires an authenticated user with the ability to initiate a pull or feed (a sync user or feed administrator) and a target sharing group that exists on the local instance. No unauthenticated access or user interaction is required.	N/A	More Details
CVE-2026-133184	MISP contains an authorization and input-validation flaw in its event push mechanism. When a user with push permissions initiates a push of events to a remote MISP instance, the system fails to enforce two critical checks: (1) it does not verify that an event is published before including it in the push, allowing unpublished (i.e., not yet publicly shared) events to be transmitted to the remote server; and (2) it does not verify that the pushing user has read access to the event, causing events the user cannot fetch to be sent as empty event shells, thereby disclosing the existence and metadata of events the user is not authorized to view. Additionally, the remote server's reply selecting which offered events to accept is used without validating that the selected identifiers actually correspond to events in the offered set, potentially allowing the remote to reference events outside the intended scope. The vulnerability requires an authenticated user with push permissions on the source MISP instance and a configured push target. The security impact is unauthorized disclosure of confidential threat-intelligence data (unpublished events and events restricted by user	N/A	More Details

	permissions) to a remote MISP instance that should not receive them.		
CVE-2026-133188	MISP contains a vulnerability in the cryptographic signature verification of protected events. When a MISP instance signs a protected event for synchronization, it strips all whitespace from the JSON body before computing the PGP signature (preg_replace("/\s+/", "", $data)). The receiving instance performs the same whitespace-stripping before verification. Because the signature covers only the whitespace-free representation, a network relay positioned between two MISP instances can insert whitespace characters into JSON object keys (effectively renaming fields) or into identifier values such as UUIDs, timestamps, and distribution fields without invalidating the signature. The receiving instance accepts the modified event as authentic. The vulnerability affects the protected event synchronization path: event creation (EventsController::add), event editing (EventsController::edit), event viewing (EventsController::view1), and server-to-server pull synchronization (ServerSyncTool, Server model). An attacker with a network position to intercept and modify traffic between MISP instances can alter the semantic content of protected events while the PGP signature still validates, undermining the integrity guarantees that protected events are designed to provide. The fix introduces two complementary mechanisms. First, a post-verification scan rejects any received body that contains whitespace in JSON keys (except under Galaxy cluster 'meta' keys) or in identifier-like values (uuid, id, timestamp, distribution, date). Second, a new exact-byte signature scheme (x-pgp-signature-v2) signs the body as transmitted, without whitespace normalization. MISP 3.0 and later peers exchange this exact-byte signature; older peers retain the legacy whitespace-stripped signature. A signature_version field is embedded in protected events so that a 3.0 receiver refuses a body carrying that marker if only the legacy signature is present. Both mechanisms together close the whitespace-injection vector while preserving backward compatibility with pre-3.0 instances.	N/A	More Details
CVE-2025-59909	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-15353	SOAP information disclosure through malformed array dimensions. The SOAP extension reads an uninitialized array position when decoding an encoded array whose arraySize attribute is empty, whitespace-only or non-numeric, or when such dimensions come from a WSDL definition. Four bytes of unrelated heap data can then be used as an integer key in the decoded array. An unauthenticated remote client can trigger this with a crafted request to a SoapServer endpoint, and a malicious SOAP server with a response to SoapClient. Heap memory is disclosed only if the application returns or otherwise exposes the decoded array. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-87867	MRemoteNG establishes encrypted but unauthenticated channel with Microsoft SQL Server backend. Client accepts any certificate the server presents, including a self-signed certificate supplied by an attacker. An attacker positioned between the client and the SQL server (for example via ARP/DNS spoofing) can transparently pintercept the connection (man-in-the-middle) and read all traffic, including the encrypted connection-password table read and the database login. Vendor did not provide the range of vulnerable versions. Only version 1.76.20 was tested and confirmed vulnerable. Other versions may also be vulnerable.	N/A	More Details
CVE-2026-99321	In the Linux kernel, the following vulnerability has been resolved: rust_binder: wake epoll on self-work queued to poll thread When a thread pushes work to its own todo list, InnerThread::push_work() returns PushWorkRes::OkNotifyPoll if the thread uses epoll, but four call sites in thread.rs discard this return value. As a result, a non-blocking BC_* write that enqueues work onto the calling thread's todo list without also reading will not wake epoll. Check the return value and call self.process.notify_poll(false) when PushWorkRes::OkNotifyPoll is returned. This is similar to commit 97830f3c3088 ("binder: signal epoll threads of self-work"), but in C Binder pending thread work also prevented subsequent process work from waking epoll, causing a permanent hang, which is not the case in Rust Binder. This means that in practice, it merely delays notifications rather than completely failing to send them.	N/A	More Details
CVE-2026-99333	In the Linux kernel, the following vulnerability has been resolved: HID: amd_sfh: Use report type from request node Use report type from request node while sending feature and input reports to HID core. The cli_data report type array is never assigned and feature reports are sent as input reports.	N/A	More Details
CVE-2025-59908	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-99323	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix invalid pointer dereference when get_inode_acl() fails When get_inode_acl() fails, it returns an ERR_PTR(), but ksmbd_acls_fattr() stored that pointer in smb_fattr without checking for errors. Later callers then passed the error pointer to security descriptor helpers or posix_acl_release(), resulting in an invalid pointer dereference. This was reproduced on an XFS-backed share when xfs_get_acl() returned ERR_PTR(-ENOMEM). Querying the SMB ACL from a CIFS client triggered the crash: BUG: unable to handle page fault for address: fffffffffffffff8 RIP: smb_acl_sec_desc_scratch_len+0x58/0x90 [ksmbd] Call Trace: smb2_query_info+0x40d/0x17c0 [ksmbd] handle_ksmbd_work+0x1ca/0x770 [ksmbd] Fix this by checking the return values from get_inode_acl() in ksmbd_acls_fattr(). Treat -EOPNOTSUPP as indicating that no ACL is available, and propagate all other errors to the callers. Update smb2_open() and smb2_get_info_sec() to handle these errors.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix out-of-bounds read when		

CVE-2026-99324	padding the final compound response Compound response handling pads the last response of a compound to an 8-byte boundary by growing the final iov and the RFC1002 length. smb2_allocate_rsp_buf() allocates exactly response_sz. A command can fill the response buffer to within 7 bytes of its end, for example a QUERY_INFO whose output is bounded only by the space left. The padding can then run past the allocation, and ksmbd_tcp_writev() reads those bytes when it sends the response. BUG: KASAN: slab-out-of-bounds in _copy_from_iter+0x744/0x10e0 Call Trace: _copy_from_iter+0x744/0x10e0 tcp_sendmsg_locked+0x2290/0x3f80 ksmbd_tcp_writev+0x10c7/0x1ab0 [ksmbd] Allocate room for the padding when the request is a compound, as smb2_read() and smb2_read_pipe() do for their payloads. response_sz is unchanged, so commands still fill only up to it, and kvzalloc() zeroes the padding.	N/A	More Details
CVE-2026-99325	In the Linux kernel, the following vulnerability has been resolved: HID: hid-lenovo-go-s: Add suspend function and gate access under bool. Currently a device can suspend while delayed work is pending or executing. Cancel delayed work if it was pending and restart it if it hasn't run yet. Gate access to show/store functions behind a bool to prevent access prior to cfg_setup() being completed.	N/A	More Details
CVE-2026-99326	In the Linux kernel, the following vulnerability has been resolved: HID: hid-lenovo-go-s: Prevent deadlock if device removed during setup hid_gos_cfg_remove() calls a mutex before canceling delayed work. That delayed work function calls mcu_property_out(), which holds the same mutex. Prevent deadlock my removing the mutex prior to stopping the work. Additionally, since cancel_delayed_work has the potential to re-arm, switch to disable_delayed_work.	N/A	More Details
CVE-2026-99327	In the Linux kernel, the following vulnerability has been resolved: HID: roccat: kovaplus: reject short button reports The raw event callback runs before HID core validates and zero-pads the report. It passes the buffer to helpers which inspect data[0] and read a complete struct kovaplus_mouse_report_button for button reports. A truncated report can therefore cause an out-of-bounds read. Require enough data for the button report before calling either helper. Commit 47669bec44fe ("HID: asus: refactor the two workqueues and init sequence") added the same kind of raw-event length validation to hid-asus.	N/A	More Details
CVE-2026-99328	In the Linux kernel, the following vulnerability has been resolved: HID: picoLCD: check for fb_deferred_io_init() error fb_deferred_io_init() allocates deferred I/O state, populating info->fbdefio_state, or leaving it NULL if an error occurs. Currently picolcd_init_framebuffer() ignores its return value. Therefore if an error arises in fb_deferred_io_init() (for instance, due to an allocation failure) info->fbdefio_state is left NULL. When the file is subsequently opened, fb_open() will dereference a NULL pointer (calling fb_deferred_io_open()). Fix this by checking for the error. Also correct cleanup ordering - defio is initialised after the sysfs file, so cleanup defio first.	N/A	More Details
CVE-2026-99329	In the Linux kernel, the following vulnerability has been resolved: HID: mcp2221: always initialize and cancel init_work The writable gpio_mode_enforce parameter can change between probe and remove. If it changes from false to true after probe, remove skips cancellation of init_work. Pending or running work can then access devm-managed driver state after teardown. Changing it from true to false can make remove cancel an uninitialized delayed work item. Initialize init_work whenever IIO is reachable and cancel it unconditionally on remove. Keep the parameter check only for scheduling the initialization. The false-to-true cancellation gap was previously reported by Sashiko. We also encountered this issue during our static-analysis research.	N/A	More Details
CVE-2026-99330	In the Linux kernel, the following vulnerability has been resolved: HID: mcp2200: validate READ_ALL response length The HID core invokes raw_event callbacks before validating the report length. mcp2200_raw_event() casts a READ_ALL response to struct mcp_read_all_resp and reads fields through offset ten without checking that the response contains the structure. A one-byte READ_ALL response reproduced the issue under KASAN: BUG: KASAN: slab-out-of-bounds in mcp2200_raw_event+0x24b/0x3a0 Read of size 1 by task hidtrigger/89 Call Trace: mcp2200_raw_event+0x24b/0x3a0 kasan_report+0x139/0x170 Reject an incomplete READ_ALL response with a protocol error before reading its fields. The common completion path then wakes the command waiter with the error instead of leaving it to time out. Commit 47669bec44fe ("HID: asus: refactor the two workqueues and init sequence") added the same kind of raw_event length validation to hid-asus.	N/A	More Details
CVE-2026-99331	In the Linux kernel, the following vulnerability has been resolved: HID: cp2112: validate response report lengths The HID core invokes raw_event callbacks before validating the report length. cp2112_raw_event() reads every field of a transfer status response without checking size. For a data response, it reads the length at offset two and copies that many bytes from offset three, again without ensuring that the input report contains them. Require a complete transfer status structure. For data responses, first require the header and then ensure the clamped payload length fits in the received report before copying it. Commit 47669bec44fe ("HID: asus: refactor the two workqueues and init sequence") added the same kind of raw_event length validation to hid-asus.	N/A	More Details
CVE-2025-59911	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-132039	In Eclipse Ditto versions before 3.9.9, Thing events delivered via WebSocket, Server-Sent Events and connections disclose field values which the policy does not allow the subscriber to read. A READ revoke for one subject of the subscriber's authorization context is not applied when another of its subjects is granted READ. For subscribers with partial READ access, subscription RQL filters, the SSE `_context`, `thing-json` placeholders and extra fields are evaluated on the unfiltered event. This allows an authenticated subscriber	N/A	More Details

	with at least partial READ access to a Thing to obtain the values of fields it may not read when events for that Thing are emitted.		
CVE-2026-129916	In Eclipse Ditto versions 3.9.0 to 3.9.8, the resolved policy view (`GET /api/2/policies/{policyId}?policy-view=resolved`) does not restrict merged policy content to what the caller may read in the policy it originates from. Imported and namespace-root entries are returned in full as soon as the caller may read any part of them in the source policy, entries of transitive imports and of the imports of namespace-root policies are only checked against the intermediate policy, and content copied into an entry by `references`, also from another entry of the same policy, is not checked against the referenced entry. An authenticated caller with READ on a policy can therefore read subjects, resources and granted/revoked permissions of policy entries that are hidden from them, including by explicit revokes, when reading the originating policy directly.	N/A	More Details
CVE-2026-129702	In Eclipse Ditto versions before 3.8.15 and 3.9.0 to 3.9.8, the policy enforcers do not let a revoke win over a grant in two kinds of permission checks. Checks of several permissions at once compare the highest grant weight with the highest revoke weight across all requested permissions instead of per permission, so the "check permissions" API (`POST /api/2/checkPermissions`, available since 3.7.0) can report a set of permissions as granted although one of them is revoked. Partial permission checks, which authorize retrieving Things and Policies and determine the receivers of events, count a grant on a sub-resource although the same permission is revoked on that sub-resource for another subject of the same authorization context (or, with the memory-optimized policy evaluator, for the same subject). A requester can thereby learn about the existence of Things and Policies and receive metadata without any readable content, and an application relying on the check permissions result may grant access the policy denies. The content of retrieved Things and Policies is still filtered correctly, and modifications are not affected.	N/A	More Details
CVE-2026-115560	Session serializer mismatch can lead to remote code execution. Session upload progress, when enabled, can store attacker-controlled data in a session key using the php_binary or php_serialize session serializer. If the same session is later read with the php serializer, because session.serialize_handler is changed at runtime or differs between applications sharing a session store, bytes inside the key can be passed to the unserializer. A remote attacker able to submit a multipart upload can inject objects, which may lead to code execution when the reading application contains a suitable gadget chain, or to session corruption or denial of service. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-120753	In Eclipse Ditto versions 1.1.0 to 3.9.8, when pre-authentication is enabled, the gateway service accepts the pre-authenticated subject not only from the `x-ditto-pre-authenticated` HTTP header set by a trusted reverse proxy, but also from an undocumented `x-ditto-pre-authenticated` URL query parameter that a client can set. A reverse proxy that only overwrites the header therefore does not stop clients from supplying their own subject. With the default configuration of the Ditto Helm chart (`global.jwtOnly: false`, which implicitly enables pre-authentication), the bundled nginx skips HTTP Basic authentication for any request with an `Authorization: Bearer` header and forwards it without the pre-authentication header. An unauthenticated remote attacker can send an invalid Bearer token together with the query parameter to authenticate as any chosen policy subject. Pre-authentication succeeds before the JWT is validated, so the attacker can read and modify Things and Policies with that subject's permissions. As a mitigation, disable pre-authentication in the gateway (`ENABLE_PRE_AUTHENTICATION=false`) or, when using the Helm chart, configure `global.jwtOnly: true`.	N/A	More Details
CVE-2025-59912	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-126086	In Eclipse Dataspace Components (EDC) – Connector (`org.eclipse.edc:decentralized-claims-*`, `org.eclipse.edc:jwt-verifiable-credentials`, `org.eclipse.edc:ldp-verifiable-credentials`) there's a credential verification bypass in Decentralized Claims Protocol (DCP) presentation handling: a participant can present credentials appearing to be issued by any trusted issuer without holding them.	N/A	More Details
CVE-2026-126714	Partially overwritten session files can lead to remote code execution. The files session save handler, which is the default, writes updated session data into the existing session file without first clearing it. If a write is interrupted, the file can contain the beginning of the new session followed by the end of the old one, and bytes of an ordinary string can later be unserialized as objects. An attacker who can influence the session data, for example through session upload progress, and interrupt a write at a suitable point may inject objects, which can lead to code execution when the application contains a suitable gadget chain. Session data can also be corrupted. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-127960	Denial of service through folded HTTP response headers. The HTTP stream wrapper joins each continuation line of a folded response header to the accumulated value by allocating a new string and copying the whole value, so a response with many short continuation lines takes disproportionate CPU time to parse. An attacker who controls an HTTP or HTTPS response fetched through file_get_contents(), fopen() or other stream functions, by making the application request a chosen URL or by influencing a server it already contacts, can keep PHP workers busy. Repeated or concurrent requests can cause a denial of service. This issue affects PHP versions from 8.1.32 through 8.1.34, from 8.2.28 before 8.2.36, from 8.3.18 before 8.3.37, from 8.4.5 before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details

CVE-2026-128064	Mixed-case HTTPS and FTPS URLs can expose data in plaintext. The HTTP and FTP stream wrappers recognise URL schemes case-insensitively, but enable TLS only when the scheme contains a lowercase "s" at a fixed position. A URL such as HTTPS:// or FTPS:// passed to file_get_contents(), fopen() or other stream functions, or received as an HTTP redirect target, is therefore handled over unencrypted HTTP or FTP. An attacker able to observe the network traffic can read data the application expected TLS to protect, including credentials, request bodies and transferred files. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-128090	PHAR OpenSSL signature bypass via algorithm downgrade. The Phar extension selects the signature verification method from the signature type stored in the archive, and only the OpenSSL types consult the separate .pubkey file holding the trusted public key. An attacker who can modify an OpenSSL-signed Phar archive, but not its public key file, can switch the signature type to a keyless hash such as SHA-256 with a recomputed digest, so the altered archive passes verification even with phar.require_hash enabled. An application that loads the archive may then execute attacker-controlled PHP code. Native, tar-based and zip-based phars are affected. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129182	Use-after-free via repopulation of last_fatal_error_backtrace global after shutdown. A fatal error raised while modified INI settings are restored during request shutdown can repopulate the stored fatal error backtrace after the executor has been cleaned up. The request memory is then freed, and a diagnostic emitted before executor initialization in a later request handled by the same PHP-FPM or mod_php worker can access the stale backtrace. The shutdown error requires per-path configuration of max_memory_limit, or of zlib.output_compression together with output_handler (which .user.ini can also set). The early diagnostic requires Apache per-directory configuration or, for PHP-FPM, a PHP_VALUE or PHP_ADMIN_VALUE parameter sent by the front end. This can cause memory corruption and worker crashes, and may allow remote code execution. This issue affects PHP versions 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129289	PHP-FPM request smuggling through an unread HTTP request body. PHP does not discard the unread part of an HTTP request body at request shutdown when the application has opened php://input without reading it to the end. When the front end reuses backend connections, which is not the default in nginx or Apache (for example nginx with fastcgi_keep_conn on, or Apache mod_proxy_fcgi with enablereuse=on), PHP-FPM can parse the leftover bytes as a new FastCGI request. An unauthenticated remote attacker can use a single crafted multipart/form-data request to forge a backend request with chosen parameters such as SCRIPT_FILENAME, REMOTE_ADDR and PHP_ADMIN_VALUE, which can lead to execution of attacker-controlled PHP code. With LSAPI, subsequent requests on the connection can be disrupted. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129291	PHP-FPM epoll use-after-free can lead to local privilege escalation. When catch_workers_output is enabled and the epoll event backend is used, the PHP-FPM master can close a worker's output pipe in fpm_child_close() without removing the associated event, and the delayed cleanup can free the worker record while its epoll registration is still active. A later event is then dispatched through a pointer into freed memory. A local attacker who can execute PHP code in an FPM worker can crash the master or execute code with its privileges, commonly root. catch_workers_output is disabled by default but enabled in the official PHP-FPM Docker images. This issue affects PHP versions from 8.1.20 through 8.1.34, from 8.2.7 before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129292	PHP-FPM status endpoint XSS via unescaped user and script fields. The PHP-FPM full status output includes the request method, authentication user and script filename of each worker without escaping them for the HTML, XML or JSON format. When the status endpoint is enabled, an unauthenticated remote attacker can store markup in a worker's scoreboard entry with an ordinary request, for example through the username of an HTTP Basic authentication header that the web server forwards to PHP-FPM, or through the script filename where the attacker can influence it. When a user views the full status page as HTML, the injected JavaScript runs in the origin of that page. The attacker does not need access to the status endpoint. Malformed XML or JSON output can also disrupt monitoring tools. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129438	Incorrect RFC 3986 URI serialization for oversized input. The Uri\Rfc3986\Uri class does not correctly handle the output-size limit and errors of the uriparser library when converting a URI back to a string. For a URI whose string representation reaches approximately 2 GiB, toString(), toRawString() and other serializing operations can return a truncated or invalid string instead of an error. An application that accepts attacker-controlled URI data of that size, has enough memory to process it and trusts the result may make incorrect validation or routing decisions, and the large allocations can exhaust memory. parse_url() and Uri\WhatWg\Url are not affected. This issue affects PHP versions 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129474	Use-after-free in magic method backtraces can lead to remote code execution. When collecting the arguments of a method call handled by __call() or __callStatic() for a backtrace, PHP can read stale values that may refer to freed memory. If a fatal error such as exceeding memory_limit occurs while an application enters such a call, with fatal_error_backtraces enabled (the default) and zend.exception_ignore_args disabled, a remote attacker who controls enough request data to influence memory allocation may cause memory corruption, crash the PHP worker or, with knowledge of memory addresses, execute arbitrary code. php.ini-production enables zend.exception_ignore_args, which prevents this. This issue affects PHP versions	N/A	More Details

	8.5.* before 8.5.13.		
CVE-2026-129481	Use-after-free after interrupted session shutdown. The session extension records the filename of the executing script when a session starts. If a fatal error occurs in an application callback while the session variables are destroyed at shutdown, such as stream_close() of a user-defined stream wrapper, the pointer is retained after the request memory is freed, and a later request handled by the same worker can free the stale string. An attacker who can trigger such a failing cleanup may crash the worker or corrupt memory; code execution has not been demonstrated. This requires the filename to be allocated in request memory, for example when the script is not cached by OPcache. This issue affects PHP versions 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129496	Timezone cache use-after-free can lead to remote code execution. In non-thread-safe builds, a diagnostic written to a file-based error_log during module startup, for example for an invalid date.timezone value or a session configuration warning, causes the timezone cache to be allocated in startup memory that is freed while a reference to it is kept. A diagnostic raised early in the first request of a worker, such as the max_input_vars warning, then accesses the freed cache. On such a deployment an unauthenticated remote attacker can crash PHP-FPM, Apache mod_php or CGI workers with a crafted request. Arbitrary code execution may be possible and has been demonstrated with known memory addresses. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129501	Double free during DOM XML serialization. The XML serializer of Dom\XMLDocument and Dom\HTMLDocument can share ownership of a generated namespace prefix between the namespace-prefix maps of an element and a descendant, so the prefix is used after being freed and then freed a second time. The document does not need to be malformed. An application that parses attacker-controlled XML or HTML with these classes and serializes it with saveXml() can be crashed, and repeated requests can cause denial of service. Code execution has not been demonstrated. This issue affects PHP versions 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129506	SOAP use-after-free when processing DTD entity references. The SOAP extension frees the document type definition (DTD) of a parsed XML document while attribute values can still refer to entities it owns, and can then read freed memory. Removing the DTD also bypasses the SOAP server's rejection of documents that contain one. An unauthenticated remote attacker who can send a crafted request to a SOAP endpoint, or who controls a WSDL or schema document loaded by the application, can crash the PHP process and cause denial of service. Memory disclosure and code execution have not been confirmed. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129538	PHP-CGI out-of-bounds read and write during PATH_INFO correction. The PHP CGI binary (php-cgi), in CGI and FastCGI mode with cgi.fix_pathinfo enabled (the default), can perform an out-of-bounds read and write when the path removed from SCRIPT_FILENAME or PATH_TRANSLATED is longer than the supplied PATH_INFO, or when PATH_INFO is empty. This requires a web server mapping that passes such inconsistent values, or a FastCGI listener reachable by untrusted clients. An unauthenticated remote attacker can crash the process, may obtain process memory if the application exposes PATH_INFO or PATH_TRANSLATED, and may corrupt memory. PHP-FPM and the Apache module are not affected. This issue affects PHP versions before 8.2.36, 8.3.* before 8.3.37, 8.4.* before 8.4.28 and 8.5.* before 8.5.13.	N/A	More Details
CVE-2026-129699	In Eclipse Ditto versions 2.0.0 to 3.8.14 and 3.9.0 to 3.9.8, the authorization check for JSON merge patch commands on Things (HTTP `PATCH` / Ditto Protocol `merge`) does not check permissions for fields deleted or set to an empty object by the patch: fields with a `null` value or an empty object value (`{}`) are skipped by the check, and Ditto's regex-based key deletion with the documented `{{ ~<regex>~ }}` syntax is not recognized as a deletion. An authenticated user with WRITE permission on any part of a Thing can therefore delete other attributes, features and the definition of that Thing, or overwrite or create fields as empty objects, including parts the policy revokes or never grants WRITE on. -->	N/A	More Details
CVE-2026-129700	In Eclipse Ditto versions 2.4.0 to 3.9.8, Thing fields a requester is not allowed to read can be disclosed through the `condition` and `live-channel-condition` headers, which are evaluated against the full Thing: the READ check of `live-channel-condition` is skipped when a `condition` header is also present on Thing query commands (3.0.0 to 3.9.8); both headers are evaluated without any authorization on WoT Thing Description requests (`Accept: application/td+json`) when WoT TD permission filtering is not active (by default in 3.0.0 to 3.8.14, in 2.4.x with WoT integration enabled, and in 3.9.x with filtering disabled); the `condition` of live messages and live events is evaluated without a READ check, and of claim messages without any permission (3.1.0 to 3.9.8); and the twin response of a Thing query with a non-matching `live-channel-condition` is returned without filtering it by the policy (3.0.0 to 3.9.8). This allows an authenticated user to read Thing fields they are not allowed to read, in some cases without any permission on the Thing.	N/A	More Details
CVE-2026-99332	In the Linux kernel, the following vulnerability has been resolved: HID: core: Fix out-of-bounds left shift by 32 in s32ton() Syzbot reported an out-of-bounds shift (32-bit integer shifted left by 32) in the s32ton() routine in the HID core. This routine converts a signed 32-bit integer to a signed n-bit integer, and it caps the value of n at 32. But when n == 32 there's no need to do anything; it can simply return the input value. This will avoid calculating (1 << n) when n is 32, which was the cause of the invalid shift.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: KVM: PPC: Book3S HV: Add		

CVE-2026-99334	preempt_disable() around virtual-mode HPTE bit-lock users kvmppc_hv_find_lock_hpte() requires virtual-mode callers to run with preemption disabled, because it can return with HPTE_V_HVLOCK still held until the caller later unlocks the HPTE. Existing virtual-mode callers in book3s_64_mmu_hv.c already follow that rule. Two virtual-mode call paths do not: 1. kvmppc_handle_exit_hv() calls kvmppc_hpte_hv_fault() for hash-mode data-side and instruction-side faults after guest exit, with preemption already enabled. kvmppc_hpte_hv_fault() calls kvmppc_hv_find_lock_hpte() directly. 2. kvmppc_pseries_do_hcall() executes virtual-mode handlers for HPT hcalls via kvmppc_pseries_do_hpt_hcall() with preemption enabled. The handlers for H_ENTER, H_REMOVE, H_READ, H_CLEAR_MOD, H_CLEAR_REF, H_PROTECT, and H_BULK_REMOVE spin on try_lock_hpte() or lock_rmap(), and H_ENTER also reaches kvmppc_do_h_enter(), which uses arch_spin_lock() on kvm->mmu_lock. That raw lock choice is intentional because kvmppc_do_h_enter() is also used by real-mode callers, so the correct fix is to establish the proper preemption context at the virtual-mode caller boundary. If a vCPU thread is preempted while holding HPTE_V_HVLOCK, any other thread on the same CPU trying to acquire the same bit-lock can spin indefinitely. The same applies to lock_rmap() users and to the H_ENTER path while holding the raw mmu_lock. Fix this by adding preempt_disable()/preempt_enable() pairs around the two kvmppc_hpte_hv_fault() call sites in kvmppc_handle_exit_hv(), and around the kvmppc_pseries_do_hpt_hcall() invocation in kvmppc_pseries_do_hcall(). Targeting only the HPT hcall helper avoids wrapping non-HPT hcalls (such as H_RTAS, H_CONFER, H_REGISTER_VPA) that perform page faults, blocking memory allocations, or context switches.	N/A	More Details
CVE-2026-99403	In the Linux kernel, the following vulnerability has been resolved: xfs: mark overfull dabtree node blocks as corrupt LOLLM observes that we need to range-check the dabtree node record count before using it to index an array. Fix that.	N/A	More Details
CVE-2026-19707	A bug in the Airlock Gateway (8.5.2 and earlier, and 8.6.0) allows attacker-controlled HTML in a backend response to exploit the Smart Form Protection feature, possibly leading to a Denial of Service (DoS) or Remote Code Execution (RCE) attack against the Airlock Gateway. The issue has been fixed in Airlock Gateway version 8.5.3 and 8.6.1. For an attack to be executable, the Airlock Gateway needs to have 'Smart Form Protection' (or the corresponding expert settings) enabled on a mapping whose backend returns a response containing certain multi-byte HTML character references. The buffer overflow then happens during the parsing of that HTML response. Such a response can come either from a malicious or compromised backend, or from an otherwise trusted backend that returns attacker-supplied HTML, e.g. an application that stores and serves plain user-provided HTML files. Successfully exploiting the vulnerability can enable an attacker to perform a DoS attack or, in the worst case, an RCE attack against the Airlock Gateway in front of the affected backend. Compromising the functionality of the Airlock Gateway could in this way also affect the availability and protection of other well-intentioned backends deployed behind the same Gateway.	N/A	More Details
CVE-2026-99368	In the Linux kernel, the following vulnerability has been resolved: NFSD: Do not send CB_RECALL_ANY to NFSv4.0 clients deleg_reaper() sends CB_RECALL_ANY to every ACTIVE client holding delegations, but CB_RECALL_ANY is an NFSv4.1 operation. An NFSv4.0 client's callback service accepts only CB_GETATTR and CB_RECALL, so it replies OP_ILLEGAL. The decoder maps the unexpected opnum to -EIO, and nfsd4_cb_done() marks the client's callback channel down. Nothing brings the channel back. nfsd4_run_cb_work() sets NFSD4_CB_UP only for a minor version above zero, and the only nfsd4_probe_callback() call site an NFSv4.0 client reaches is nfsd4_setclientid_confirm(). One visit from the reaper therefore leaves the channel marked down until the client re-establishes its clientid. RENEW then returns NFS4ERR_CB_PATH_DOWN for as long as the client holds delegations. nfsd4_cb_channel_good() stops returning true, so the client is granted no further delegations. Skip clients at minor version zero.	N/A	More Details
CVE-2026-99369	In the Linux kernel, the following vulnerability has been resolved: NFSD: Destroy a recalled delegation the client does not hold A client that answers CB_RECALL with NFS4ERR_BADHANDLE or NFS4ERR_BAD_STATEID has no record of the delegation, so the FREE_STATEID that clears it from cl_revoked never arrives. Every later SEQUENCE reply carries SEQ4_STATUS_RECALLABLE_STATE_REVOKED, and the client loops issuing TEST_STATEID. Destroy such a delegation when it is reaped rather than revoking it onto cl_revoked. RFC 8881 Section 20.2.4 completes the recall at the reply when its status is neither NFS4_OK nor NFS4ERR_DELAY, so a rejected recall leaves nothing to revoke. An administrative revoke keeps that path, since NFS4ERR_ADMIN_REVOKED reports it. A destroyed stateid returns NFS4ERR_BAD_STATEID instead of NFS4ERR_DELEG_REVOKED. A client that rejects the recall but still holds the delegation gets no notice that its state was revoked. CB_RECALL can outrun the reply that granted the delegation, so honor a rejection only once the client has seen that grant. Per RFC 8881 Section 2.10.6.3, retirement of the slot that carried the grant is that proof; retry until then, and revoke when the retries lapse.	N/A	More Details
CVE-2026-99371	In the Linux kernel, the following vulnerability has been resolved: nfsd: clear XDR padding in GETXATTR replies nfsd4_vbuf_to_stream() uses a raw memcpy() for the first fragment of an attribute value. xdr_reserve_space() aligns the reservation but does not initialize the extra bytes. If the complete value fits in that fragment, the loop that clears padding for the final fragment is never entered. A remote client reading an attribute whose length is not a multiple of four can therefore receive one to three stale bytes from the response page. Encode the first fragment with xdr_encode_opaque_fixed(), matching the final-fragment path and ensuring that any XDR padding is zeroed.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: NFSD: cap the number of listeners accepted in listener_set nfsd_nl_listener_set_doit() matches each requested listener against the existing set. The nested loop that does this is O(N * M), where N is the requested count and M is the existing count. The loop		

CVE-2026-99372	runs under sv_lock with bottom halves disabled. A userland request with a very large listener list can therefore spin in atomic context for a long time. Reject a request that carries more than NFSD_NL_LISTENER_MAX (1024) entries. The check goes in nfsd_nl_validate_listeners(), before the code takes any lock. The limit is far above any realistic configuration. This patch does not cap M. Only the message size bounded N; real sockets bound M. A listener_set result set is the requested set, so that path also holds M at the cap, but __write_ports_addxprt() adds two listeners per call and removes none, so repeated calls can push M past it. The worst case under sv_lock is therefore 1024 * M, plus 1024 nla_parse_nested() calls. Both interfaces require CAP_NET_ADMIN.	N/A	More Details
CVE-2026-99373	In the Linux kernel, the following vulnerability has been resolved: NFS/localio: detect a short read or write before the iterator has moved nfs_local_call_read() and nfs_local_call_write() issue a misaligned DIO as up to three segments and must stop at the first short one, or the next segment lands at the wrong file offset. They test for that by comparing the bytes transferred against iov_iter_count() of the segment, read after read_iter() or write_iter() has advanced the iterator, when the count is the residual rather than the request: before the call: count == expected after the call: count == expected - status The test is therefore status < expected - status, and it fires only when less than half of the segment was transferred. A short I/O that covers between 50% and 99% of a segment slips through: the loop moves on with ki_pos advanced by the short amount, the next segment is issued at the wrong offset, and the header's byte count is over-reported to the caller. A write that slips through also never sets NFS_CONTEXT_WRITE_SYNC, which is what makes the writes that follow a short one synchronous. Take the segment's count before issuing it and compare against that. This is the same defect that commit 250ec14932d5 ("nfsd: fix partial-write detection in nfsd_direct_write") fixed in NFSD's version of this loop.	N/A	More Details
CVE-2026-99374	In the Linux kernel, the following vulnerability has been resolved: NFSv4/flexfiles: Use the full 64-bit stripe_unit ff_layout_alloc_lseg() decodes stripe_unit as the 64-bit value the protocol defines and struct nfs4_ff_layout_segment stores it as one, but both consumers narrow it back to 32 bits: nfs4_ff_layout_calc_dss_id() divides with do_div(), which casts the divisor, and ff_layout_pg_test() copies it into a u32 first. A stripe_unit that does not fit is silently truncated, so the client stripes on a unit the server did not specify -- or divides by zero, if the low 32 bits happen to be clear. Divide by the full value, with div64_u64() and div64_u64_rem().	N/A	More Details
CVE-2026-99375	In the Linux kernel, the following vulnerability has been resolved: NFSv4/flexfiles: fail I/O to the data server once a lock is lost When a client loses a byte-range lock, because its lease expired and it may not recover the lock (nfs.recover_lost_locks=0) or because reclaim failed, it marks the lock state NFS_LOCK_LOST and fails further I/O by that lock owner. Otherwise an application that believes it holds a lock over a range is never told that the lock is gone, and can overwrite data that another client has since written under its own lock on that range. The NFSv4 READ/WRITE and file layout I/O paths apply this in rpc_call_prepare, where nfs4_set_rw_stateid() returns -EIO for a lost lock. The flexfiles data server I/O path does not: the NFSv3 data server path never checked, and the check on the NFSv4 path was removed when the data server stateid began to be taken from the layout. Fail read and write RPCs to the data server in rpc_call_prepare when the lock is lost. The data server path has no use for the stateid that nfs4_set_rw_stateid() selects, and selecting it looks up a delegation or open stateid on every I/O. Instead add nfs4_lock_context_lost(), which only tests NFS_LOCK_LOST, and check it alongside the open state. Its lookup, __nfs4_lookup_lock_state(), is split out of __nfs4_find_lock_state() and takes no reference: a lock state is only unlinked from state->lock_states under state->state_lock, so one found while holding that lock cannot be freed before it is released. Use it in nfs4_copy_lock_stateid() as well, which only reads the lock state under that lock. That drops a reference get and put from stateid selection for all NFSv4 I/O without changing the stateid selected. Terminating the RPC in rpc_call_prepare is not enough on its own. The done callback reads the -EIO as a data server connection error, forgets the device ID, and resets the I/O to pNFS (or, for an NFSv4 data server, possibly to the MDS). After a reset to pNFS a new LAYOUTGET is granted and rpc_call_prepare fails the I/O again, without end. The existing NFS_CONTEXT_BAD check does not loop, since the open state is then invalid and pnfs_update_layout() fails the reset I/O, but it still forgets the device ID. So when an RPC that was never sent fails and the open context is bad, the open state is invalid, or the lock is lost, complete it with -EIO without error handling. An RPC that was sent and failed is handled as before, so a real data server error is still reported to the MDS.	N/A	More Details
CVE-2026-99376	In the Linux kernel, the following vulnerability has been resolved: NFSv4/pnfs: Free the netid when draining a data-server address list nfs4_pnfs_ds_addr_free() releases all three of an address entry's allocations, but the four sites that drain a decoded list before it reaches a data server open-code the loop and free only da_remotestr and the entry, leaking da_netid. Both layout drivers drain that list when nfs4_pnfs_ds_add() returned a data server already in the cache, and again on their error exits. Lift the drain into nfs4_pnfs_ds_addr_list_free() so an address list has one owner. A striping mount reaches the cached case on nearly every GETDEVICEINFO, since many deviceIDs name the same data servers.	N/A	More Details
CVE-2026-99377	In the Linux kernel, the following vulnerability has been resolved: NFSv4/pnfs: bound the CB_NOTIFY_DEVICEID array count before allocating decode_devicenotify_args() hands the server's notification count straight to kmalloc_objs() as the array length, without checking it against the message that carried it. Bound it the way nfs4xdr.c bounds an attribute length, against xdr_stream_remaining().	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: pnfs/blocklayout: Complete a device upcall when the pipe is closed Once blkmapd has read a whole upcall, rpc_pipe_read() unlinks the message from		

CVE-2026-99378	every pipe list, so the purges in rpc_pipe_release() and rpc_close_pipes() no longer reach it. bl_upcall_ops supplies no .release_pipe callback, and bl_pipe_destroy_msg() returns without completing the upcall because a downcall is expected to follow. A daemon that exits between the read and the write therefore strands its waiter: bl_resolve_deviceid() sleeps in wait_for_completion() and never returns, and it holds nn->bl_mutex across that wait, so every later device resolution in the net namespace blocks behind it. Add a .release_pipe that completes an upcall blkmapd has consumed and left unanswered. Completing an upcall retires it, so a message the framework has already purged is no longer in flight and is not completed a second time here.	N/A	More Details
CVE-2026-99382	In the Linux kernel, the following vulnerability has been resolved: smb: client: reject FIND data areas that run past the received response CIFSFindFirst() and CIFSFindNext() take the directory entries of a TRANS2_FIND_FIRST2 or TRANS2_FIND_NEXT2 response from DataOffset for DataCount bytes. validate_t2() bounds DataOffset at 1024 and ParameterCount + DataCount by the byte count, and nothing compares DataOffset + DataCount with the number of bytes SendReceive() copied into the buffer. smb_init() places the request and the response in the same cifs_request allocation and cifs_buf_get() clears only the start of it, so a short response whose data area extends past its end has srch_entries_start point at bytes of the request that was just sent, or at whatever the allocation held before, and those bytes are parsed as directory entries. cifs_readdir() and find_cifs_entry() walk the entries with nxt_dir_entry(), which checks every entry but the first of a response against the end of the SMB; cifs_query_path_info() and cifs_backup_query_path_info() read the first entry with no check at all. Those reads stay inside the allocation, but what they return is not part of the response. The last entry, which LastNameOffset locates relative to the data area, is bounded against the received response by an earlier patch of this series, as is every entry that cifs_fill_dirent() parses; the walk itself and the two query functions are not. SMB1 is not negotiated by default; reaching this code requires an explicit vers=1.0 mount. Reject a response whose data area extends past the received length, with the -EINVAL that validate_t2() returns for the other malformed transaction responses. cifs_query_path_info() and cifs_backup_query_path_info() still read a whole entry from a data area that may be smaller than one; that is not changed here. A conforming server is not affected. The data block of a transaction response is part of the SMB, so DataOffset + DataCount cannot exceed the smbCalcSize() bytes that SendReceive() copies; coalesce_t2() appends each secondary response to the data area and adds its length to both DataCount and the byte count, which keeps that relation.	N/A	More Details
CVE-2026-99335	In the Linux kernel, the following vulnerability has been resolved: KVM: guest_memfd: Establish memslot<=>guest_memfd bindings *after* memslot is ready Wait to bind a memslot to a guest_memfd instance until *after* the memslot is fully prepared, as creating the binding in guest_memfd will effectively expose the memslot to readers. As pointed out by Sashiko, binding the memslot before it's ready to be exposed to the rest of the world can break various memslot assumption and rules. E.g. x86 could observe a NULL rmap pointer if a PUNCH_HOLE hit the guest_memfd after the binding was created, but before KVM made it through kvm_prepare_memory_region().	N/A	More Details
CVE-2026-99384	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix OOB read of the resume name sent in TRANS2_FIND_NEXT2 cifs_fill_dirent() reads the fixed part of a directory entry and the name length it holds, or for SMB_FIND_FILE_UNIX scans the name for its terminating NUL, without knowing where the response ends. cifs_save_resume_key() calls it for the entry at the LastNameOffset the server chose, after every FindNext and after the FindFirst of a search rewind, and records the name it returns as the resume name of the search without bounding it; the next CIFSFindNext() copies up to PATH_MAX - 1 bytes from that name into its request and sends them to the server. cifs_filldir() bounds the name against the end of the response since commit f8cf09a53a0d ("smb: client: bound dirent name against end of SMB response in cifs_filldir"), but only after cifs_fill_dirent() has read the entry, and nxt_dir_entry() checks the entries it walks only against sizeof(FILE_DIRECTORY_INFO), for every level but SMB_FIND_FILE_INFO_STANDARD, and never sees the first entry of a response. Bounding where last_entry starts, as the preceding patch does, does not bound the read. A last entry that starts within the final bytes of the response still has its fixed part read past the end of the received data, and the name length it declares, a __le32 for every level but SMB_FIND_FILE_INFO_STANDARD and SMB_FIND_FILE_UNIX, is recorded without being checked against anything. CIFSFindNext() rejects only a length of PATH_MAX or more, then copies the name, up to 4095 bytes, into its request and sends it, and as the report below shows that copy can run past the end of the cifs_request allocation that holds the response. For SMB_FIND_FILE_UNIX the name is in addition scanned for its terminating NUL over up to PATH_MAX + 1 UTF-16 units, or PATH_MAX bytes, before any check. The FindNext that carries the name is only sent by the SMB1 dialect, which is not negotiated by default; reaching it requires an explicit vers=1.0 mount. With KASAN enabled, a response whose last entry declares a name longer than the bytes that follow it gives: == BUG: KASAN: slab-out-of-bounds in CIFSFindNext+0x8ca/0x14c0 Read of size 4080 at addr ffff88807c823f98 by task ls/83 CPU: 0 UID: 0 PID: 83 Comm: ls Not tainted 7.3.0-rc4-00457-gf14572c203d5 #69 PREEMPT(lazy) Call Trace: <TASK> kasan_report+0xdf/0x1a0 kasan_check_range+0x10f/0x1e0 __asan_memcpy+0x23/0x60 CIFSFindNext+0x8ca/0x14c0 cifs_readdir+0xf51/0x29c0 iterate_dir+0x1c0/0x570 __x64_sys_getdents64+0x133/0x270 do_syscall_64+0x109/0x5d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 83 on cpu 1 at 13.497119s: cifs_buf_get+0x36/0x90 smb_init+0x4f/0x110 CIFSFindNext+0xf7/0x14c0 cifs_readdir+0xf51/0x29c0 The buggy address belongs to the object at ffff88807c820000 which belongs to the cache cifs_request of size 16588 The buggy address is located 16280 bytes inside of allocated 16588-byte region [ffff88807c820000, ffff88807c8240cc)	N/A	More Details

	========== The read is the memcpy() that copies the recorded name into the next request, so those bytes go to the server. Pass the end of the response to cifs_fill_dirent() and have it reject an entry whose fixed part does not fit or whose name extends past the end, before either is used, so that both callers are covered; cifs_save_resume_key() computes the end the same way cifs_readdir() does. The name bound compares the bytes left between the name and the end with the name length rather than adding the length to the name pointer, since the length is a 32-bit value taken from the entry and the sum could wrap on 32-bit systems. cifs_fill_dirent() is also used on the SMB2 readdir path; there smb2_parse_query_directory() has already checked every entry against the end of the response, so a vali ---truncated---		
CVE-2026-99387	In the Linux kernel, the following vulnerability has been resolved: smb: client: bound POSIX lock data area by response length CIFSSMBPosixLock() receives a TRANS2 response whose server-provided DataOffset and DataCount are not required by validate_t2() to describe an interval inside the response. A server can therefore make the client read a cifs_posix_lock beyond the response buffer and return those kernel bytes through F_GETLK. Use smbCalcSize() to obtain the full response length after any multipart TRANS2 reassembly, then reject data areas that extend past it before dereferencing them.	N/A	More Details
CVE-2026-99390	In the Linux kernel, the following vulnerability has been resolved: xfs: fix correction of unwritten extents in quota file data fork LOLLM noticed a couple of problems with the code that fixes unwritten mappings in the quota file's data fork: first, we should zero the ondisk contents as part of converting them to "unwritten" just in case we crash shortly after the conversion. Second, and more importantly, it noticed that we fail to resynchronize the iext tree cursor after modifying the iext tree. If the tree shape changes, we could be in a world of hurt. Third, we should force a quotacheck if we reset an ondisk dquot. Fix all issues.	N/A	More Details
CVE-2026-99391	In the Linux kernel, the following vulnerability has been resolved: xfs: don't try to scrub self-referential dirent in metapath checker LOLLM observes that the _ilock_both helper will loop forever trying to double-lock the same inode if mpath->dp == sc->ip. Callers are not supposed to call us with parent == child because that makes no sense. Trigger an assertion, mark the scan incomplete, and return IO error so that this fails noisily.	N/A	More Details
CVE-2026-99392	In the Linux kernel, the following vulnerability has been resolved: xfs: don't subtract excess usage twice when computing agf_btreeblks LOLLM noticed that the free space btree repair code calculates agf_btreeblks from the number of blocks we wrote to rebuild the free space btrees (minus the roots). That's correct. However, it observes that we subtract the "excess" reservation, which is not correct. Fix this accounting problem.	N/A	More Details
CVE-2026-99393	In the Linux kernel, the following vulnerability has been resolved: xfs: don't overrun incore inode when cleaning attr forks LOLLM noticed that xrep_dinode_check_afork doesn't check for a di_forkoff value that is so large that the attr fork would begin after the end of the inode space. This causes afork_{ptr,size} to be nonsense. What we really ought to do here is zap the attr fork.	N/A	More Details
CVE-2026-99394	In the Linux kernel, the following vulnerability has been resolved: xfs: complain loudly when dirtree checker finds broken links LOLLM observed that the directory tree walk code bails out if it finds broken parent pointers along the way. However, it flattens the return code to zero, which means that we can end up in the repair code having not completed the scan. We should never proceed with a repair in this state, so let's mark the scrub incomplete and exit.	N/A	More Details
CVE-2026-99396	In the Linux kernel, the following vulnerability has been resolved: xfs: bound da-node entry count against the correct geometry xfs_da3_node_verify() bounds the node entry count against the larger of the directory and attribute geometries because, as a buffer verifier, it cannot tell whether the block belongs to the directory or the attribute tree. When the directory block size exceeds the fs block size (e.g. mkfs.xfs -n size=64k -b size=4k), an attribute node buffer is a single fs block that holds only m_attr_geo->node_ents entries, yet a crafted attr node may claim a count up to m_dir_geo->node_ents and still pass the verifier. xfs_da3_node_lookup_int() then indexes btree[] up to that count during its binary search -- an out-of-bounds read via getxattr/listxattr on a mounted crafted image. The buffer verifier is the wrong place to tighten this: it has no fork context, and the transaction-less read path used by getxattr does not run xfs_da3_node_set_type() either. xfs_da3_node_lookup_int(), on the other hand, always runs on that path and holds args->geo, the geometry of the fork actually being searched. Bound the entry count against args->geo->node_ents there, before walking the entries.	N/A	More Details
CVE-2026-99397	In the Linux kernel, the following vulnerability has been resolved: xfs: always report inode repair setup failures LOLLM noticed that xchk_setup_inode silently discards errors being returned by xrep_setup_inode. Fix that, and fix the stub macro so that it actually "returns" a value.	N/A	More Details
CVE-2026-99398	In the Linux kernel, the following vulnerability has been resolved: xfs: always forget cached ACLs if the attr fork swap succeeds LOLLM complains that xrep_xattr_rebuild_tree omits the xfs_forget_acl call if the attr fork swap succeeds (and hence the acls might have changed) but reaping the old attr fork fails. Fix it so that we always purge the cached acls if the swap succeeds, even if subsequent steps do not. This avoids some cache coherency problems with access controls.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: NFSD: Don't complete a cld upcall the daemon has not read cld_pipe_downcall() matches a reply to its upcall on the xid alone. A write that arrives before the daemon has read that upcall completes __cld_pipe_upcall() while its struct rpc_pipe_msg is still		

CVE-2026-99367	queued on the pipe. The waiter returns, its stack frame goes away, and the pipe keeps a list_head that points into it. The next rpc_queue_upcall() walks that pointer: BUG: KASAN: vmalloc-out-of-bounds in __list_add_valid_or_report Read of size 8 at addr ffffc90007027950 by task syz.0.96/13066 __list_add_valid_or_report+0x6a/0x130 rpc_queue_upcall cld_pipe_upcall nfsd4_cld_grace_start nfsd4_cld_tracking_init The pipe is mode 0600 in rpc_pipefs, so the trigger is a broken or hostile nfsdcld rather than an unprivileged task. Record whether the daemon has consumed the whole message and complete only when it has. cld_pipe_destroy_msg() runs on every path by which the pipe releases a message, so it can maintain that flag under cn_lock. nfsdcld sends its -EINPROGRESS downcalls only after reading the upcall, so the new check does not break the GraceStart record stream.	N/A	More Details
CVE-2026-99365	In the Linux kernel, the following vulnerability has been resolved: nfsd: Fix id-to-name cache entry leak in idtoname_parse() idtoname_lookup() returns the id-to-name cache entry with a new reference. The name is parsed between the lookup and the update call, and when the parsed name is malformed (negative length or longer than IDMAP_NAMESZ), idtoname_parse() jumps to the error label without dropping the reference returned by idtoname_lookup(), leaking the cache entry. Release the reference with cache_put() before taking the error exit.	N/A	More Details
CVE-2026-99364	In the Linux kernel, the following vulnerability has been resolved: NFSD: Fix out-of-bounds read in the rpc_status dump nfsd_nl_rpc_status_get_dumpit() loads args->opcnt and args->ops separately, then walks ops[] before rechecking rq_status_counter. A COMPOUND that completes between the two loads runs nfsd4_release_compoundargs(), which zeroes opcnt and points ops back at the eight-entry inline array. A dump that already sampled an opcnt of 200 clamps it to the sixteen slots in rq_opnum, then indexes iops[0..15]. iops is the last member of struct nfsd4_compoundargs and rq_argp is allocated at exactly that size, so the walk runs off the end of the allocation. The trailing recheck discards the sampled data, but the read has already happened. NFSD_CMD_RPC_STATUS_GET carries no GENL_ADMIN_PERM, so an unprivileged local user can repeat the dump against a busy server until it lands in the window. Sample opcnt and ops into locals, then finish the counter recheck before dereferencing ops. An unchanged counter means both came from the same COMPOUND, where opcnt cannot exceed what ops holds.	N/A	More Details
CVE-2026-99363	In the Linux kernel, the following vulnerability has been resolved: NFSD: Map flex file layout IDs through the request's user namespace nfsd4_ff_proc_layoutget() and nfsd4_ff_encode_layoutget() translate the file's owner and group with init_user_ns, but every other identity nfsd places on the wire goes through nfsd_user_namespace(). When the transport carries a credential from a non-initial user namespace, the flex file layout reports host-global IDs. The client copies those IDs into the AUTH_SYS credential it presents to the data server, and svcauth_unix_accept() resolves that credential in the transport's namespace, so data server I/O runs under an identity unrelated to the file's owner. Switching to the request's namespace introduces a second hazard. from_kuid() returns (uid_t)-1 when the target namespace has no mapping for the owner, and the IOMODE_READ arm adds one to that result to derive an identity for which the data server denies writes. The addition would wrap to zero, handing the client uid 0 instead of an identity distinct from the owner. Translate both IDs in nfsd4_ff_proc_layoutget(), which has the svc_rqst, and carry the wire values in struct pnfs_ff_layout. from_kuid_munged() substitutes overflowuid for an unmapped owner and thus never returns (uid_t)-1, so the increment cannot wrap to zero.	N/A	More Details
CVE-2026-99337	In the Linux kernel, the following vulnerability has been resolved: KVM: TDX: Enable Bus Lock VM exit Enable Bus Lock VM exit functionality for TDX guests. Bus Lock VM exit is an existing feature supported by KVM. Userspace can enable Bus Lock VM exit through KVM_BUS_LOCK_DETECTION_EXIT when it's reported as supported. However, KVM reports the support of this CAP just based on the hardware capability but doesn't differentiate between VMX and TDX. This leads to the issue that userspace can enable this cap for TDX guests without getting an error, but the feature is not actually enabled because KVM doesn't call the TDX module API to program the relevant TD VMCS fields. Enable Bus Lock VM exit for TDX guests by programming the BUS_LOCK_DETECTION control in the TD VMCS and by adding the exit handler. Note, there is no enumeration bit for this feature by TDX module because all TDX modules support it and allow to set the TD VMCS as long as the hardware supports the feature.	N/A	More Details
CVE-2026-99338	In the Linux kernel, the following vulnerability has been resolved: KVM: TDX: Enable Notify VM exit Enable Notify VM exit functionality for TDX guests. Notify VM exit is an existing feature supported by KVM. Userspace can enable Notify VM exit through KVM_CAP_X86_NOTIFY_VMEXIT when it's reported as supported. However, KVM reports the support of this CAP just based on the hardware capability but doesn't differentiate between VMX and TDX. This leads to the issue that userspace can enable this cap for TDX guests without getting an error, but the feature is not actually enabled because KVM doesn't call the TDX module API to program the relevant TD VMCS fields. Enable Notify VM exit for TDX guests by: - Invoking TDX module API calls to set NOTIFY_VM_EXITING and Notify Window in TD VMCS. It's done in tdx_vcpu_init() where other TD VMCS bits are set. Since TDX vCPU cannot be reset, it only needs to be configured once when initializing the TDX vCPU. - Adding corresponding exit handler for TDX Notify VM Exit. Notify VM exit can happen when executing the IRET instruction. If the IRET unblocks the NMI blocking state, bit 12 of the exit qualification is set. In this case, the VMM needs to restore the "blocked by NMI" state when it decides to re-enter the guest. For TDX, KVM cannot manage the GUEST_INTERRUPTIBILITY_INFO and it's TDX module's responsibility to handle it. Extract the common part without NMI blocking handling into a helper in common.h so that it can be shared between VMX and TDX. Note, KVM uses "pre-production" terminology for the feature formally called Notify VM-Exit. All public versions of the SDM refer to the feature as Instruction Timeout. This will be remedied in the near	N/A	More Details

	future, for now, use KVM's terminology for consistency. Note, #2, there is no enumeration bit for Notify VM exit by TDX module because all TDX modules support it, and allow to set the corresponding TD VMCS fields as long as the hardware supports the feature.		
CVE-2026-99339	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Zero reserved synic flags when delivering message to avoid stack leak Ensure all reserved flags are zeroed when setting the msg_pending flag as part of SynIC message delivery, as setting only msg_pending can leak seven bits of kernel stack data to the guest. E.g. gcc typically uses a bitwise-OR to set the msg_pending flag, without initializing the on-stack variable. Precisely zero the flags, e.g. as opposed to zeroing the entire structure, to make it somewhat more obvious that the flags *need* to be zeroed. E.g. it would be quite easy to misread the initial read of guest memory and think that zeroing the entire structure is completely superfluous.	N/A	More Details
CVE-2025-59907	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-99341	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Request event processing for exception VM-Exits Request event processing whenever kvm_queue_exception_vmexit() queues an exception VM-Exit. Most callers reach the helper through kvm_multiple_exception(), which already sets KVM_REQ_EVENT, but nested asynchronous page faults call it directly. Without another event request, vcpu_enter_guest() skips kvm_check_and_inject_events() and re-enters L2. A subsequent VM-Exit can clear the pending exception while reconstructing vectoring state, losing the synthetic #PF VM-Exit. The APF reason remains PAGE_NOT_PRESENT and a later regular #PF can consume the stale value. Setting KVM_REQ_EVENT in the common helper ensures that L1 receives the queued VM-Exit before KVM re-enters L2.	N/A	More Details
CVE-2026-99342	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/xen: Enforce 4-byte alignment of vcpu_info registration Xen's map_guest_area() enforces that vcpu_info is aligned to sizeof(xen_ulong_t). KVM has no such check, allowing a guest to register vcpu_info at an arbitrary byte alignment. Enforce unconditional 4-byte alignment regardless of the current shinfo mode. This is sufficient because subsequent commits ensure that all locked atomic operations on vcpu_info fields use at most 32-bit accesses. Return -ENXIO on failure, matching Xen's map_guest_area() behaviour for unaligned requests. Originally observed in review of an unrelated patch: https://lore.kernel.org/all/20260604193554.1BA311F00893@smtp.kernel.org/	N/A	More Details
CVE-2026-99344	In the Linux kernel, the following vulnerability has been resolved: KVM: nSVM: Ignore EFER.LMA if EFER.LME=0 when preparing L2 state Force EFER.LMA=0 if EFER.LME=0 when preparing L2 state for VMRUN, i.e. mimic real hardware's behavior of ignoring EFER.LMA if EFER.LME=0. VMRUN unfortunately allows the nonsensical combination, i.e. doesn't fail, but KVM itself has an invariant EFER.LMA can be set et if and only if EFER.LME is set. Breaking that invariant can lead to a variety of issue, particularly in MMU code that keys off EFER.LMA when determining whether to emulate/virtualization 4/5-level paging versus PAE paging.	N/A	More Details
CVE-2025-59906	Rejected reason: This CVE identifier has been rejected or withdrawn by its CVE numbering authority. New CVEs are now being reserved.	N/A	More Details
CVE-2026-99347	In the Linux kernel, the following vulnerability has been resolved: ALSA: x86: hdmi: fix chmap array missing terminator entry had_build_channel_allocation_map() allocates a single struct snd_pcm_chmap_elem via kzalloc() and stores it as intelhaddata->chmap->chmap. struct snd_pcm_chmap.chmap is expected to be a zero-terminated array (see snd_pcm_std_chmaps[]), but had_register_chmap_ctls() never overrides the TLV callback installed by snd_pcm_add_chmap_ctls(), so the generic pcm_chmap_ctl_tlv() from ALSA core is left in place. pcm_chmap_ctl_tlv() is reachable by any local user via SNDRV_CTL_IOCTL_TLV_READ on the "Playback Channel Map" control and walks the array with "for (map = info->chmap; map->channels; map++)". Since the driver's single element is never terminated, this walks past the allocation. The same pointer is also freed on hot-unplug and rebuilt on hot-plug under intelhaddata->mutex, while the TLV callback runs under card->controls_rwsem without that mutex, so the walk can also race a concurrent free of the array. Fix this by allocating a second, zeroed element so the array is properly terminated, and by dropping the inherited TLV callback entirely at registration time, since the driver never implements one of its own.	N/A	More Details
CVE-2026-99348	In the Linux kernel, the following vulnerability has been resolved: afs: Fix afs to abort the rxrpc call on send error Fix afs_send_empty_reply() and afs_send_simple_reply() to always try to abort the rxrpc call rather than just aborting on -ENOMEM and otherwise abandoning it. If the call is already complete due to network failure or a received abort, this will do nothing. Also make afs_make_call() always abort on send error; again, it does nothing if the rxrpc call is already dead.	N/A	More Details
CVE-2026-99349	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix lack of short-send handling in rxrpc_kernel_send_data() Fix rxrpc_kernel_send_data() to loop around if it detects a short send. David Laight suggested doing it here rather than wrapping all the calls in loops. Further, remove the len argument and use the iterator count instead and return 0 on success, not the amount copied. Note this is also a prerequisite for changing the way rxrpc_send_data() works to return a short send rather than an error if some data was buffered.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix update of call->tx_pending without holding lock Currently, rxrpc_send_data() updates call->tx_pending just before it returns - but it won't		

CVE-2026-99350	be holding the call->user_mutex when it does this if a wait was interrupted by a signal. This would allow a parallel sendmsg() to race. Further, both the callers of rxrpc_send_data() call it with the lock held, and then it returns an indication through the parameter list to say whether it has dropped the lock or not - after which the callers both just drop the lock if it's still held. Fix this by: (1) Moving the release of call->user_mutex down into rxrpc_send_data() and get rid of the indicator parameter. This makes it easier to see where the lock is held. (2) After waiting, if the attempt to reacquire the mutex is interrupted, just return directly there rather than going to out_unlock (3) Restricting the txb variable to inside the buffering loop and leaving ->tx_pending set until we've queued the buffer. Note that there's a slight change in behaviour in that wait_for_space failure now doesn't check for completion because it doesn't hold the call user_mutex. The caller, however, should re-issue the send and pick up any error at a second attempt.	N/A	More Details
CVE-2026-99351	In the Linux kernel, the following vulnerability has been resolved: sunrpc: treat empty auth.unix.gid replies as negative entries When rpc.mountd cannot resolve a uid (getpwuid() or getgrouplist() failure, e.g. while winbind or sssd is briefly unreachable), it answers the auth.unix.gid upcall with zero groups. unix_gid_parse() installs that as a valid positive entry, and svcauth_unix_set_client() then replaces the credential's group list with the empty one on every request, RPCSEC_GSS included via svcauth_gss_set_client(). One failed lookup strips that uid of all supplementary groups on every export for up to mountd's configured TTL (30 minutes by default), long after the NSS backend has recovered. mountd cannot send an empty list for a successful lookup, since getgrouplist(3) always includes at least the user's primary group, so a zero-group reply can only mean the lookup failed. Record it as a negative entry: unix_gid_find() then returns -ENOENT and svcauth_unix_set_client() keeps the groups the RPC credential already carries. This is the fallback that commit 3fc605a2aa38 ("[PATCH] knfsd: allow the server to provide a gid list when using AUTH_UNIX authentication") promised when no answer is available, and the same state try_to_negate_entry() already creates when no listener holds the channel open.	N/A	More Details
CVE-2026-99352	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: trace an accepted transport before publishing it svc_handle_xprt() publishes a newly accepted transport before tracing it. svc_add_new_temp_xprt() places the transport on sv_tempsocks and clears XPT_BUSY, so another nfsd thread can dequeue it at once. If the peer has already closed the connection, that thread runs svc_delete_xprt(), which drops the list reference, the only reference the new transport holds. trace_svc_xprt_accept() then reads the transport's addresses and network namespace from freed memory. Before publication the accepting thread owns the transport exclusively: XPT_BUSY is still set from svc_xprt_init() and the transport is on no list. Emit the tracepoint there. [cel: trace before publication instead of taking an extra reference]	N/A	More Details
CVE-2026-99353	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Reject a socket that already has an svc_sock attached Writing the same socket descriptor to /proc/fs/nfsd/portlist twice attaches a second svc_sock to one socket. svc_setup_socket() saves the socket's callbacks before installing its own, so the second attach records svc_write_space() as the old write_space callback. svc_udp_init() invokes that callback by way of svc_sock_setbufsize(), and svc_write_space() then calls itself until the kernel stack is exhausted: BUG: TASK stack guard page was hit at ffffc900037d7ff8 svc_write_space+0x90/0x2b0 net/sunrpc/svcsock.c:429 svc_write_space+0xe6/0x2b0 net/sunrpc/svcsock.c:430 ... 700 more ... svc_sock_setbufsize+0x18d/0x220 net/sunrpc/svcsock.c:386 svc_udp_init net/sunrpc/svcsock.c:854 [inline] svc_setup_socket+0xb2f/0x1090 net/sunrpc/svcsock.c:1498 svc_addsock+0x2fd/0x760 net/sunrpc/svcsock.c:1547 __write_ports_addfd fs/nfsd/nfsctl.c:742 [inline] write_ports+0xa5b/0xcc0 fs/nfsd/nfsctl.c:861 nfsctl_transaction_write+0x106/0x1a0 fs/nfsd/nfsctl.c:112 svc_data_ready() and svc_tcp_state_change() chain through their saved callbacks the same way, so a TCP descriptor added twice recurses on the next incoming segment instead. Reaching any of this takes a writer on portlist, and the nfsd filesystem sets no FS_USERNS_MOUNT, so the reproducer needs CAP_SYS_ADMIN in the initial user namespace. Reject a socket that already carries sk_user_data. svc_setup_socket() overwrites that field unconditionally, so a socket some other consumer has claimed is one NFSD would corrupt whether or not the callbacks recurse.	N/A	More Details
CVE-2026-99354	In the Linux kernel, the following vulnerability has been resolved: sunrpc: pin gss module across auth_domain RCU free A gss auth_domain stores a raw pointer to svcauthops_gss in dom->flavour, and its release queues svcauth_gss_domain_release_rcu() via call_rcu(). Both live in auth_rpcgss.ko text, and nothing pins the module while a domain referencing it is reachable. rmmod auth_rpcgss can therefore unmap that text while an NFSD export still holds a reference to the domain. The export cache is one such holder. When the export is finally released, the domain reference drops from module text: svc_export_put() auth_domain_put(exp->ex_client) auth_domain_release() dom->flavour->domain_release(dom) /* into module text */ svcauth_gss_domain_release() call_rcu(&dom->rcu_head, svcauth_gss_domain_release_rcu) Once auth_rpcgss is unloaded, the domain_release call lands in unmapped memory. The rcu_barrier() in exit_rpcsec_gss() does not help: it waits only for callbacks already queued, and this one is not queued until the export goes away. NFSD's NFSv4 objects import symbols from auth_rpcgss.ko, so only a server built without CONFIG_NFSD_V4 can unload auth_rpcgss while exports still name a gss domain. Take a module reference in svcauth_gss_register_pseudoflavor() before publishing the domain and drop it at the end of svcauth_gss_domain_release_rcu(). The rcu_barrier() then holds the module image until this callback has returned. [cel: rewrote description]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Copy the deferred RPC Call from the head buffer svc_defer() snapshots the RPC Call for replay once a pending cache upcall resolves. It copies the Call starting (rq_arg.len - head[0].iov_len) bytes before head[0].iov_base, a back-up that once covered a		

CVE-2026-99357	transport header. No transport has had one since commit 983084b2672c ("SUNRPC: Remove svc_rqst::rq_xprt_hlen"), so the two lengths are meant to be equal. An RPCSEC_GSS Call breaks that assumption. Unwrapping the integrity service truncates rq_arg.len by the length of the trailing checksum, and unwrapping the privacy service leaves rq_arg.len short by an amount the client sets with trailing bytes after the wrap token. Neither adjusts head[0].iov_len. Both lengths are size_t, so the subtraction underflows and the copy starts past head[0].iov_base. Every deferred integrity Call replays a shifted snapshot. A deferred replay skips GSS verification, so a shifted snapshot runs as an authenticated Call. With the privacy service the offset is client-chosen. A forged Call placed in that window, naming another user's GSS context handle, runs under that identity. Copy the Call from head[0].iov_base. Refuse to defer a Call whose length runs past the head buffer, so the copy cannot read out of bounds.	N/A	More Details
CVE-2026-99358	In the Linux kernel, the following vulnerability has been resolved: nfsd: zero NFSv4 COMPOUND tag padding nfs4svc_encode_compoundres() copies the tag bytes into reserved XDR space and skips directly to the aligned end of the field. xdr_reserve_space() rounds the reservation up but does not initialize the padding bytes. A remote client can choose a tag length that is not a multiple of four, causing one to three stale bytes from the response page to be returned in the COMPOUND reply. Use xdr_encode_opaque_fixed() to copy the tag and clear its XDR padding.	N/A	More Details
CVE-2026-99359	In the Linux kernel, the following vulnerability has been resolved: NFSD: validate transport name in listener_set before serv creation nfsd_nl_listener_set_doit() holds nfsd_mutex for the whole listener teardown and rebuild. The code checks NFSD_A_SOCK_TRANSPORT_NAME for presence only, and not for content. An arbitrary name therefore reaches svc_xprt_create_from_sa(). There, a name that matches no registered class calls request_module("svc%s", name). That call is a TASK_KILLABLE usermode helper upcall, and it runs under nfsd_mutex. Check the name against the classes that NFSD can create: tcp, udp and rdma. The check goes in nfsd_nl_validate_listeners(), which runs before the code takes nfsd_mutex. The rejection names the offending attribute through extack, since -EPROTONOSUPPORT on its own does not say which entry carried the bad name. This narrows the upcall. It does not remove it. NFSD accepts "rdma" without a condition, so on a kernel that does not build svcrdma the name still reaches request_module("svcrdma") under nfsd_mutex. That is necessary for the modular case, where the autoload is legitimate.	N/A	More Details
CVE-2026-99360	In the Linux kernel, the following vulnerability has been resolved: NFSD: Set nn->cld_net before registering the cld pipe rpc_mkpipe_dentry() instantiates the pipe's dentry before it returns, so the cld pipe can be opened as soon as nfsd4_cld_register_net() has run. __nfsd4_init_cld_pipe() does not store the new cld_net in nn->cld_net until after that call. cld_pipe_downcall() reaches the cld_net only through that pointer and takes cn_lock without checking it, so a write() that lands in the window oopses. Store nn->cld_net once @cn is fully initialized and before the pipe is registered, and clear it again when registration fails.	N/A	More Details
CVE-2026-99361	In the Linux kernel, the following vulnerability has been resolved: nfsd: propagate SETXATTR value decode errors nfsd4_decode_setxattr() stores the status returned while assembling a non-empty attribute value, but then unconditionally returns nfs_ok. An allocation failure can therefore leave a NULL value buffer paired with a non-zero length and allow compound processing to continue. Return the decoder status so that the operation stops on the original error.	N/A	More Details
CVE-2026-133165	MISP contains an authorization and scope flaw in the ObjectReference model's captureReference method. When a user saves an event (via REST API, sync, or feeds), any object reference included in the payload is looked up by its UUID across the entire MISP instance rather than being scoped to the event being saved. An authenticated user with permission to add or edit events can therefore include the UUID of an object reference belonging to a different event—including events in other organizations that the user cannot view—and cause the system to overwrite that foreign reference's relationship_type, comment, deleted flag, and timestamp with attacker-supplied values. The affected reference is neither deleted nor recreated in the attacker's event; its fields are silently rewritten in place. This constitutes a cross-tenant integrity violation: data integrity of object references in events the attacker has no read or write access to is compromised. The vulnerability is reachable through any event save path (REST add/edit, inter-instance sync, feed ingestion) and requires only a valid authenticated session with event modification privileges and knowledge or discovery of a target reference UUID.	N/A	More Details
CVE-2026-133163	MISP contains an authorization bypass affecting the event protection mechanism. In MISP, a protected event is designed to be modifiable only through the events/edit endpoint with a valid cryptographic signature. However, the protection check was enforced solely on that single route. All other write endpoints that operate on an event—attributes (add, edit, delete), objects (delete), events (populate, publish, unprotect, upload_sample, upload_analysis_file), tags (attach, detach), event reports (edit, delete), galaxy clusters (attach, detach), and cryptographic keys (add, delete)—lacked the equivalent check. As a result, an authenticated user holding the sync permission (perm_sync) who is neither a site administrator nor an internal sync user (perm_sync_internal), and who belongs to the same organization as the event's creator (orgc_id), could modify, delete, or add data to a protected event without providing a valid signature. This included the ability to unprotect the event entirely, publish it to other instances, and add cryptographic signing keys that would subsequently permit signed edits, effectively nullifying the protection mechanism. The vulnerability requires an authenticated sync user account within the event's creator organization. No user interaction or special network conditions are required; standard API calls to the affected endpoints are sufficient. The security impact is a loss of integrity over protected threat-intelligence events: an attacker can inject, alter, or remove attributes, objects, tags, reports, and galaxy cluster associations, and can remove the	N/A	More Details

	protection flag or add signing keys to enable persistent unauthorized modification.		
CVE-2026-133157	flowintel contains a CSV injection vulnerability in its audit log CSV export functionality. Attacker-influenced text recorded in audit logs could be interpreted as spreadsheet formulas when an administrator opened the exported CSV file in a spreadsheet application. The vulnerable behavior allowed an attacker who could influence audit log content (for example by generating network traffic that is captured and logged) to embed formula-like strings (such as =CMD() or =HYPERLINK()) into audit entries. When the resulting CSV export was opened in a spreadsheet application, these strings could be executed as formulas, potentially leading to data exfiltration, data manipulation, or execution of further malicious actions within the spreadsheet context. Preconditions: - An attacker must be able to influence content that is recorded in flowintel audit logs. - An administrator or user must export the audit logs to CSV and open the file in a spreadsheet application that interprets formulas. Security impact: - Potential information disclosure from the spreadsheet or linked resources. - Potential integrity impact on data within the spreadsheet application. - No direct compromise of the flowintel application itself. Affected version: <4.1.0	N/A	More Details
CVE-2026-61584	Chamilo LMS is an open-source learning management system. Prior to 1.11.40, main/inc/ajax/record_audio_rtc.ajax.php concatenates the user-controlled audio_dir parameter with the course document root and passes the result to recursive directory creation without path normalization. An authenticated course user can select directory paths outside the intended course storage tree wherever the web-server account has permission to create directories. The resulting filesystem changes can interfere with application layout or support a separate file-write or code-execution chain, although this endpoint alone is not shown to write attacker-selected file content outside the course tree. This issue is fixed in version 1.11.40.	N/A	More Details
CVE-2026-132216	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesEventos' parameter in'/Clientes/ClientesEventos/'.	N/A	More Details
CVE-2026-132217	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesEventos' parameter in'/Clientes/ClientesEventos/'.	N/A	More Details
CVE-2026-132218	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesEventos' parameter in'/Clientes/ClientesEventos/'.	N/A	More Details
CVE-2026-132219	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'DocumentosOpcionesImprimir' parameter in'/Gestion/DocumentosOpcionesImprimir/'.	N/A	More Details
CVE-2026-132220	Stored Cross-Site Scripting (XSS) vulnerability in winCRM. Exploiting this vulnerability could allow a remote user to send a specially crafted URL to the victim and steal their session data by allowing JavaScript code to be injected in 'Tarea' parameter in '/Tareas/TareasCrear?ddv=cal'.	N/A	More Details
CVE-2026-59159	Linux Mint Update Manager manages software updates on Linux Mint systems. From 5.4.8 until 5.5.4.3, 5.5.9.1, 5.8.8.1, 5.9.8.1, 6.0.9.1, and 6.1.5.1, mint-refresh-cache passes a user-controlled --parent-window-id value directly to the privileged Synaptic process without validating it as an integer. A local user permitted to execute mint-refresh-cache through sudo can inject additional Synaptic command-line options and achieve arbitrary code execution as root. This issue is fixed in versions 5.5.4.3, 5.5.9.1, 5.8.8.1, 5.9.8.1, 6.0.9.1, and 6.1.5.1.	N/A	More Details
CVE-2026-128661	PHPGurukul Employee Leaves Management System v2.1 contains an insufficient verification vulnerability in the password recovery functionality in forgot-password.php. The application allows an unauthenticated user to access the password-change workflow by providing a valid employee ID and corresponding email address without requiring an additional out-of-band verification mechanism.	N/A	More Details
CVE-2026-61578	Chamilo LMS is an open-source learning management system. Prior to 1.11.40 on the 1.11 release line and in 2.0.0, plugin/cleandeletedfiles/src/ajax.php and public/plugin/CleanDeletedFiles/src/ajax.php pass the request path parameter and list parameter directly to unlink without canonicalization, filename validation, or confinement to the plugin's scanned directories. An authenticated administrator can delete arbitrary files writable by the web server; a cross-site request forgery against an administrator session can reach the same operation. Deleting application source, uploaded content, cache, or configuration files can cause data loss and denial of service. This issue is fixed in version 1.11.40 on the 1.11 release line. No fixed 2.0 release is established as of this review.	N/A	More Details
CVE-2026-61582	Chamilo LMS is an open-source learning management system. Prior to 1.11.40, main/inc/ajax/user_manager.ajax.php exposes the update_users action without an authorization check and stores attacker-controlled firstname, lastname, and email values without HTML sanitization. An unauthenticated attacker can modify another user's profile fields, and the stored content executes when user_portal.php, the administrative user listing, user information, social profile, course participant list, or message-composer recipient header renders the modified account. Targeting an administrator account can expose administrator-session data or permit actions in that session. This issue is fixed in version 1.11.40.	N/A	More Details

CVE-2026-61583	Chamilo LMS is an open-source learning management system. Prior to 1.11.40, main/dropbox/dropbox_functions.inc.php accepts a literal .htaccess filename in the course dropbox and stores it in a directory served by Apache, where the file can alter handler behavior under the required AllowOverride All configuration. Any authenticated course participant can upload .htaccess and then upload a PHP payload disguised as a .gif, resulting in code execution as the web-server user when the dropbox tool is enabled. This can expose or modify the Chamilo installation, course content, and database credentials. This issue is fixed in version 1.11.40.	N/A	More Details
CVE-2026-75521	Miren Runtime provides the runtime component for the Miren application platform. Prior to 0.11.1, the QUIC/HTTP3 RPC server on the default port 8443 uses tls.RequestClientCert and does not verify that presented certificates chain to the configured cluster certificate authority. The RPC authenticators derive a caller identity from PeerCertificates[0].Subject.CommonName without requiring a non-empty VerifiedChains value, and certificate identities bypass all RBAC checks. A network attacker who can reach the listener can present a self-signed certificate with an issuer distinguished name matching the public cluster CA subject and gain cluster superuser access, allowing disclosure, modification, or disruption of applications, sandboxes, disks, credentials, and other managed resources. This issue is fixed in version 0.11.1.	N/A	More Details
CVE-2026-126962	open62541 Subscription Use-After-Free Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of open62541. Authentication is required to exploit this vulnerability. The specific flaw exists within the handling of subscriptions. The issue results from the lack of validating the existence of an object prior to performing operations on the object. An attacker can leverage this vulnerability to execute code in the context of the service account. Was ZDI-CAN-30231.	N/A	More Details
CVE-2026-128667	A Cross-Site Request Forgery (CSRF) vulnerability in the admin panel of PHPGurukul News Portal Project V4.1 allows remote attackers to hijack the authentication of administrators and perform unauthorized state-changing actions, such as adding, editing, or deleting posts and sub-admins, or changing passwords, via crafted HTTP requests. This occurs because the application fails to implement anti-CSRF tokens for admin actions and permits destructive operations via HTTP GET requests.	N/A	More Details
CVE-2026-76090	Tapir is a Scala library for describing HTTP APIs and exposing, consuming, and documenting them using open standards. Prior to 1.13.27, tapir-openapi-codegen-core generates Scala source from untrusted OpenAPI documents by interpolating object property names, query, header, and path parameter names, and operation tags without context-appropriate validation or escaping. Crafted names or values can break out of generated identifier or string-literal contexts and inject arbitrary Scala that executes when generated code is compiled or endpoints are assembled. This is a code-generation and build-time issue and does not make a running tapir server directly exploitable. This issue is fixed in version 1.13.27.	N/A	More Details
CVE-2026-122119	Rejected reason: ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: CVE-2026-122227. Reason: This candidate is a reservation duplicate of CVE-2026-122227 Notes: All CVE users should reference CVE-2026-122227 instead of this candidate. All references and descriptions in this candidate have been removed to prevent accidental usage.	N/A	More Details
CVE-2026-130089	A caller who can send requests to the temporal-proxy gRPC listener can obtain the static API key of an upstream whose address is templated from the request namespace, and then use that key directly against the upstream Temporal service. temporal-proxy does not validate the request namespace before rendering it into a templated upstream address, so when an upstream's hostPort is templated from the namespace, as in the published Temporal Cloud configuration example, a crafted namespace changes the host the proxy dials. If that upstream verifies TLS against the system trust store and sets no literal server name, a caller-controlled host with a publicly trusted certificate passes verification and receives the key. The built-in static-token and JWKS authenticators do not restrict namespaces, so any caller they admit can do this; an external authenticator that admits only listed namespaces prevents it. temporal-proxy versions 0.1.0 through 1.2.0 are affected. The issue is fixed in 1.0.2, 1.1.2, and 1.2.1, and releases before 1.0.0 receive no fix. Upgrade to the fixed release for the deployed minor release line, and rotate any upstream API key that untrusted callers could reach.	N/A	More Details
CVE-2026-130090	When two templated upstreams in temporal-proxy render the same address, a request routed to one upstream can be sent under the other upstream's credential, namespace translation, and payload encryption keys. temporal-proxy pools connections for templated upstreams by rendered address and TLS server name only, although each upstream's outbound credential, namespace translation, and payload encryption keys, where configured, are attached to its connection. A request routed to one upstream can therefore be sent over the connection the other created, and the upstream Temporal service authorizes it as the other upstream's principal. Configuration validation rejects only identical hostPort values, so templates that differ in text but render the same address are accepted. This mixes the two upstreams' ordinary traffic, and a caller that an external authenticator restricts to one upstream's namespaces can use it to act under the other upstream's credential. Deployments with at most one templated upstream are not affected. temporal-proxy versions 0.1.0 through 1.1.1 are affected. The issue is fixed in 1.0.2, 1.1.2, and 1.2.0, and releases before 1.0.0 receive no fix. Upgrade to the fixed release for the deployed minor release line.	N/A	More Details
	A caller with write permission in a namespace can terminate temporal-proxy through memory exhaustion by storing highly compressed payloads in Workflow history, and the stored history causes the same failure on		

CVE-2026-130091	every later read. temporal-proxy 1.2.0 decompresses every payload marked as compressed in the responses it forwards, whether or not compression is configured. Each payload is limited to 4 MiB after decompression, but the total for a response is not, and the decoded payloads are held in memory together. The caller can store many small, highly compressed payloads through the proxy or directly at the upstream Temporal service. When that history is read through the proxy, for example by a Worker processing a task, the proxy can run out of memory and be terminated. The codec server's /decode endpoint is affected in the same way when it is enabled. Only temporal-proxy 1.2.0 is affected; earlier versions do not decompress payloads. The issue is fixed in 1.2.1. Upgrade to 1.2.1 or later.	N/A	More Details
CVE-2026-130092	A caller who can reach the temporal-proxy gRPC listener, with or without a valid credential, can grow the proxy's memory until the process is terminated by sending requests with many distinct gRPC method names. temporal-proxy labels its request metrics with the method name exactly as the caller sent it, and records every request, including those that authentication or routing reject. Each distinct method name adds metric series that are never removed. The series grow whether or not the metrics endpoint is scraped, and metric labels configured from request metadata multiply the growth. temporal-proxy versions 0.1.0-alpha.20260716 through 1.2.0 are affected. The issue is fixed in 1.0.2, 1.1.2, and 1.2.1, and releases before 1.0.0 receive no fix. Upgrade to the fixed release for the deployed minor release line.	N/A	More Details
CVE-2026-132144	MISP contains an authorization flaw in the user login-email change logic. When an external authentication backend is configured (LDAP, Shibboleth, ApacheSecureAuth, or CustomAuth), the login email address serves as the key by which the directory service resolves a user's identity, organization, and role on each login. The __canChangeLogin() method in UsersController only checked the MISP.disable_user_login_change setting and site-admin status; it did not account for whether an external authentication backend was active. As a result, any authenticated non-admin user (including org admins) could change their own login email to an address that matches a higher-privilege account in the external directory. On the user's next login, the external backend would resolve that email to the target directory identity, granting the attacker the organization and role associated with that identity. This constitutes an account-identity capture and privilege escalation: the attacker effectively assumes the identity of any directory account whose email they can guess or discover. The vulnerability requires an authenticated session with a valid API key or session token and a MISP instance configured with at least one external authentication backend. The impact is limited to the scope of the captured directory identity's privileges within MISP.	N/A	More Details
CVE-2026-132559	When payload encryption is enabled, failure messages and stack traces that temporal-proxy is configured to keep from the upstream Temporal service can reach it in plaintext, where anyone who can read the Workflow's history at the upstream can read them. When the proxy returns a failure to a Worker, it decrypts the failure's encoded attributes and copies the message and stack trace back onto the failure, keeping the attributes. A Worker that sends that failure back through the proxy, as SDKs do when an Activity failure becomes the cause of a Workflow failure, sends the restored text with it, and the proxy forwards the text unchanged because the failure already has encoded attributes. The upstream then stores the text in plaintext beside its encrypted copy. This affects deployments that enable failure encryption in the proxy, and deployments that leave it off while the SDK encodes failure attributes itself. Payloads are not affected. temporal-proxy versions 0.9.0 through 1.2.0 are affected. The issue is fixed in 1.0.2, 1.1.2, and 1.2.1, and releases before 1.0.0 receive no fix. Upgrade to the fixed release for the deployed minor release line.	N/A	More Details
CVE-2026-132636	MISP contains an authorization bypass in its event search and filter functionality. When a user performed searches through events/restSearch (value, value prefix, searchall, type, uuid, attribute tags, object filters), objects/restSearch, or the event index tag, galaxy, and quick filters, the underlying SQL subqueries matched against attributes and objects without enforcing the per-attribute distribution, sharing-group, or soft-deleted visibility rules that apply when fetching individual elements. A user who could view an event (for example, a community-distributed event) could therefore use search to determine whether organisation-only, sharing-group-restricted, or soft-deleted attributes and objects existed within that event. Because the value filter supports prefix matching (LIKE), an attacker could enumerate the full value of a hidden attribute one character at a time. Tags attached exclusively to hidden attributes also leaked through the tag, galaxy, and quick-filter paths on the event index. The vulnerability affects the EventsController quick-filter and index-filter logic, the Event model wildcard/value/uuid/tag filter methods, the MispAttribute tag filter, the MispObject buildFilterConditions and restSearch methods, and the ReportFromEvent object filter. Site administrators and users belonging to the event's own organisation retain full visibility, and sync users continue to match soft-deleted elements as intended. The issue was reported by ENISA.	N/A	More Details
CVE-2026-132215	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesEventos' parameter in'/Clientes/ClientesEventos/'.	N/A	More Details
CVE-2026-132214	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesEventos' parameter in'/Clientes/ClientesEventos/'.	N/A	More Details
CVE-2026-132213	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'TareasCrear' parameter in'/Tareas/TareasCrear/'.	N/A	More Details

CVE-2026-132212	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesClasificaciones' parameter in'/Clientes/ClientesClasificaciones/'.	N/A	More Details
CVE-2026-109857	CAST Imaging is vulnerable to authentication bypass in the TenantFilter class. The TenantFilter processes the X-ISSUER-LOCATION HTTP header and dynamically registers an OpenID Connect issuer based on its value. An unauthenticated remote attacker can register a malicious identity provider, create users matching existing administrative usernames, and obtain valid JWTs which will be accepted by the application. This results in complete authentication bypass and privilege escalation to administrative access. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109858	CAST Imaging has multiple path traversal vulnerabilities in various endpoints. The application does not properly validate the path parameters, allowing an authenticated attacker to use "../" sequences to escape the intended directory and read arbitrary files. Critically, by reading "/proc/self/environ", the attacker can obtain environment variables containing sensitive credentials such as Keycloak and PostgreSQL passwords, leading to full server compromise and administrative access. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109859	CAST Imaging has a missing authentication vulnerability in the "/api/auth/keycloak/users/{username}" endpoint. An unauthenticated remote attacker can modify the email address of an arbitrary Keycloak user and access sensitive account information, including the API authentication token hash. This can lead to account takeover, user impersonation and user enumeration. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109860	CAST Imaging is vulnerable to Path Traversal vulnerability in "/api/server-folders" and "/api/server-folders/download" endpoints. These file-management endpoints accept user-supplied filesystem paths and allow selected system commands to be executed at the specified location, but they do not enforce sufficient authorization boundaries or restrict access to a predefined set of resources. An authenticated attacker can enumerate server directories via the "ls" command and retrieve arbitrary files outside of the intended application workspace, such as "/etc/passwd", exposing sensitive data. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109861	CAST Imaging contains a hardcoded authentication token in the "KeyCloakController" class of auth-service.jar. An attacker who gains access to the source code, application package, logs, or configuration files can reuse this token to call auth-service endpoints without authentication. Using the token with the "/api/auth/keycloak/users" endpoint, the attacker can enumerate users across Keycloak realms and retrieve user attributes containing API token hashes, enabling reconnaissance and further exploitation. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109862	CAST Imaging does not properly check whether a user is authorized to modify another user's notification settings in the "/api/admin/user/update-subscriber" endpoint. An authenticated, low-privileged attacker can enable or disable email notifications on another user's behalf, which can suppress important communications such as security alerts and disrupt business processes. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109863	CAST Imaging allows authenticated, low-privileged users to access parts of the administrator panel. When a user navigates to any URL path the application doesn't recognize, a wildcard route automatically redirects them to the administrative view at "/admin/applications". This allows an attacker to perform certain administrator-restricted actions, including triggering app-to-app link generation, configuring regexes, and changing searchable object properties, resulting in unauthorized modification of application configuration. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-109864	CAST Imaging has an access control flaw in the "/api/jobs/<uuid>/logs/download" endpoint. An authenticated low-privilege user can download log archives for any job whose UUID they know, even if they are not authorized to access it. The vulnerability has only been confirmed in version 3.6.2 but may also affect other versions. The vendor did not provide the affected version range, but assured that the vulnerability is fixed "in newer versions".	N/A	More Details
CVE-2026-127142	LMS (Lan Management System) is vulnerable to Remote Code Execution via unrestricted upload of dangerous file types in the service ticket feature of the client panel. The application does not verify the type or name of files attached to service tickets. An authenticated attacker can upload an .htaccess file to bypass the access restrictions on the storage directory, together with a PHP script. Visiting the uploaded PHP file's URL leads to arbitrary code execution with the privileges of the web server. This issue was fixed in commit 89dd3f7.	N/A	More Details
	LMS (Lan Management System) is vulnerable to privilege escalation in the lmsd daemon. The daemon loads		

CVE-2026-127143	the libraries listed in the LMS database without checking that they come from a trusted location. An attacker with access to the web server process or with write access to the LMS database can register a malicious library and trigger a configuration reload, causing lmsd to load and execute it. Because the example systemd unit does not set a separate user, lmsd runs as root, allowing the attacker to execute arbitrary code with root privileges. This issue was fixed in commit 89dd3f7. The lmsd component was removed. For existing installations that use lmsd, the recommended solution is to disable and remove this component.	N/A	More Details
CVE-2026-129961	LMS (Lan Management System) is vulnerable to Remote Code Execution in its e-mail-based Helpdesk ticket intake. The mail parser does not reject Apache .htaccess configuration files or PHP-executable files when saving ticket attachments. A remote unauthenticated attacker can e-mail a Helpdesk ticket with such an attachment; in the default configuration, anyone can request the uploaded file's URL which executes attacker-supplied code with the privileges of the web server process. No operator interaction or LMS client session is required. The issue was fixed in commit d5059b1.	N/A	More Details
CVE-2026-122042	WarehousePG (WHPG) 6.27.1 through 6.27.7 and 7.2.1 through 7.6.1 are affected by an out-of-bounds write vulnerability (CWE-787) in gpfdist, the file server used to load and unload external tables. When gpfdist attaches a request to a session, it sizes two internal arrays from the client-supplied X-GP-SEGMENT-COUNT header and writes to them at the index given by the X-GP-SEGMENT-ID header, without checking that the index is non-negative and smaller than the count (CWE-129). gpfdist does not authenticate callers unless gpfdists is used with client certificate verification, so an attacker who can reach the gpfdist port can crash it with a single HTTP request, failing every external table load or unload that uses that gpfdist instance. gpfdist is normally reachable only from the cluster's internal network.	N/A	More Details
CVE-2026-15263	A server-side request forgery (SSRF) vulnerability in the Pure Storage FlashArray extension for Windows Admin Center versions could allow an authenticated Gateway User to direct the management gateway to an unintended host, potentially exposing stored FlashArray credentials and API tokens and allowing unauthorized access to the managed FlashArray.	N/A	More Details
CVE-2026-128719	An issue in PHPGurukul Employee Leaves Management System (ELMS) Version 2.1 allows a remote attacker to escalate privileges via the password recovery functionality	N/A	More Details
CVE-2026-48728	The GLPI Inventory Plugin handles network discovery, inventory, software deployment, and data collection for GLPI agents. Prior to 1.6.8, the /b/deploy, /b/collect, and /b/esx agent API endpoints do not authenticate requests before processing Deploy, Collect, and ESX task operations. An unauthenticated remote attacker can enumerate jobs and manipulate their status through these endpoints, disclosing job information and interfering with task processing. This issue is fixed in version 1.6.8.	N/A	More Details
CVE-2026-112160	In PcVue version 17.0.0 through 17.0.1 included, user account security settings could be modified. This could lead to unauthorized modification of password policies or other security-related account settings by privileged users. The update aligns user account management with the existing security-altering enforcement model and prevents unauthorized security configuration changes.	N/A	More Details
CVE-2026-130007	k8s-mcp-server manages Kubernetes clusters through the Model Context Protocol. Prior to 1.1.0, SERVER_MODE defaults to SSE, with the listener on 0.0.0.0:8080, but main.go creates the SSE server without an authentication option or context function. The GET /sse handler creates a sessionId that is only a correlation value, and the POST /message handler accepts a loadable sessionId without an authorization or Origin check before dispatching tools/call; consequently, any client able to reach the port can enumerate and invoke registered Kubernetes tools without credentials. Read-only mode is disabled by default, leaving write-capable Kubernetes tools available to run with the server's configured cluster privileges. This issue is fixed in version 1.1.0.	N/A	More Details
CVE-2026-132208	Hardcoded credentials in the executable file. These can be extracted using the “Strings.exe” tool and could allow an attacker unauthorized access to the database and SMTP server. winCRM.exe	N/A	More Details
CVE-2026-132209	Hardcoded credentials in the executable file. These can be extracted using the “Strings.exe” tool and could allow an attacker unauthorized access to the database and SMTP server. winSAT.exe	N/A	More Details
CVE-2026-132210	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'idCliente' parameter in'/Gestion/DocumentosCrear/'.	N/A	More Details
CVE-2026-132211	Insecure Direct Object Reference (IDOR) vulnerability. Exploiting this vulnerability allows a user to view the names of other users/companies on the platform by iterating over the numerical value that identifies the customer in 'ClientesCamposAdicionales' in '/Clientes/ClientesCamposAdicionales/'.	N/A	More Details
	MISP contains an authorization bypass in the correlations/eventCorrelations endpoint and the underlying correlation engines (Default and OnDemand). The CorrelationsController::eventCorrelations action did not verify that the requesting user was permitted to view the event specified by the event ID parameter. Any authenticated user could supply the ID of an event belonging to another organisation (distribution level 0,		

CVE-2026-132637	organisation-only) and receive its full correlation results, including attribute identifiers and values. Additionally, for events the user could partially view, the correlation engines filtered only the related (target) side of each correlation row and did not vet the queried event's own source attributes. As a result, correlations keyed by attributes hidden from the user (organisation-only attributes, attributes in sharing groups the user is not a member of, or attributes inside organisation-only objects) were returned with their IDs and values. The Default engine's related-event listing (used in events/view) also included events related solely through such hidden attributes. The vulnerability was reported by ENISA. Successful exploitation requires a valid authenticated session with any role and knowledge or enumeration of a target event ID. The security impact is disclosure of threat-intelligence attribute values (IP addresses, domain names, etc.) and attribute identifiers that the user is not authorised to see, as well as confirmation of the existence of otherwise invisible events.	N/A	More Details
CVE-2026-132641	MISP's EventDelegationsController delegateEvent method accepted a user-supplied sharing_group_id parameter without verifying that the requesting user was a member of the referenced sharing group. When a delegation request was submitted with distribution set to 4 (sharing-group distribution), the raw sharing_group_id value was stored directly into the EventDelegation record. The delegation view subsequently rendered the stored sharing group's name back to the requester. Because no authorization check was performed, any authenticated user with delegation permission could enumerate sharing group identifiers (by guessing sequential IDs or other values) and read the names of sharing groups to which their organisation did not belong, including hidden or restricted groups. Hidden, missing, and malformed identifiers were not distinguished in the response, but the successful case still leaked the group name. The vulnerability required an authenticated session with the perm_delegate capability and knowledge of a valid event ID to delegate. The security impact is limited to disclosure of sharing group names (a low-impact confidentiality breach) and does not grant access to the group's contents or modify any data.	N/A	More Details
CVE-2026-132642	MISP contains an access control flaw in its tag management endpoints: events/addTag (single and bulk), events/removeTag, and attributes/removeTag. These endpoints loaded their target event or attribute using an unscoped database query (recursive = -1) that did not apply the calling user's ACL visibility rules. As a result, two distinct security issues existed: (1) Information disclosure: the removeTag endpoint returned different responses depending on whether a given tag was present on a hidden event, allowing an authenticated user to enumerate which tags were applied to events they had no visibility into. (2) Unauthorized modification: users belonging to the host organisation, who are permitted to set local tags on records they can see, could add or remove local tags on events and attributes they could not see, because the canModifyTag ACL check operated on the unscoped object without first verifying the caller's visibility. The fix replaces the unscoped find() calls with ACL-aware fetch methods (fetchSimpleEvent and fetchAttributeSimple) that apply the user's visibility before the object is returned, and adds an explicit empty-object check in removeTag so that a hidden event responds identically to a non-existent one. Affected components include the EventsController (addTag, removeTag), the AttributesController (removeTag), and the ACLComponent (canModifyTag). The vulnerability requires an authenticated user with at least the tagger role (perm_tagger) or host-organisation membership.	N/A	More Details
CVE-2026-71536	Mareforma is a local verification layer for AI-assisted research. From 0.3.0 until 0.3.8, verify_bundle in mareforma/export_bundle.py validates the outer DSSE envelope and exporter-signed subject digests, but JSONLDExporter._claim_node in mareforma/exporters/jsonld.py supplies no signatureBundle for individual claims. Because the exporter signs both the predicate nodes and their digest list, the check proves only internal bundle consistency and allows claims with missing, invalid, or forged asserter signatures to be accepted as authentic. This issue is fixed in version 0.3.8.	N/A	More Details
CVE-2026-91138	The Deco BE65 v3 firmware contains a shared hard-coded password hash for the root account in /etc/shadow. Where the corresponding credential is accepted by a production authentication path, an attacker with existing unprivileged shell access may recover or otherwise obtain the plaintext password and use it to escalate privileges. Successful exploitation may grant root-level access to the affected device.	N/A	More Details
CVE-2026-91139	Deco BE65 firmware contains a cryptographic weakness in the local web-management authentication flow due to the use of insufficient RSA key sizes. The /login?form=auth endpoint returns a 512-bit RSA public key and sequence value. Successful exploitation may compromise the cryptographic protections applied to local web-management authentication traffic.	N/A	More Details
CVE-2026-132798	Improper Limitation of a Pathname to a Restricted Directory in the Como log viewer in Temenos Application Framework Java (TAFJ) R18.SP76.0 allows remote authenticated attackers to read arbitrary files via path traversal	N/A	More Details
CVE-2026-117678	Inefficient Algorithmic Complexity vulnerability in benoitc idna allows an attacker who controls a hostname passed to the library to exhaust CPU via an over-long xn-- label. idna:ulabel/1 in src/idna.erl hands every ASCII label that starts with xn-- to punycode:decode/1 without applying check_label_length/1, the 63-octet DNS label limit. punycode:decode/1 in src/punycode.erl builds its output with list appends in insert/4, so its cost grows super-linearly with the label length. A label of a few kilobytes keeps a scheduler busy for seconds to minutes before the label is rejected as invalid punycode. Every decode entry point reaches this path: decode/1, decode/2, to_unicode/1 and from_ascii/1. The encode path reaches it as well, because alabel/1 calls ulabel/1 on an ASCII xn-- label before its own length check, so encode/1, encode/2 and to_ascii/1 are affected for such input. This issue affects idna: from 0.5.0 before 8.0.0.	N/A	More Details

CVE-2026-110461	The Comelit 1456B Multi-User Gateway, in firmware versions up to and including 2.10.0, is configured with static default passwords for multiple privileged roles. None of these passwords are required to be changed during initial configuration, and they remain at their default values unless the administrator changes them manually. An attacker with network access to the device can authenticate using these default credentials and gain administrative control.	N/A	More Details
CVE-2026-110462	The Comelit 1456B Multi-User Gateway, in firmware versions up to and including 2.10.0, allows a user authenticated to the web interface with the Installer role to export a system configuration backup. The resulting backup file contains the configured passwords for all privileged roles	N/A	More Details
CVE-2026-19609	Cross-site scripting in the embedded co-browsing feature in Unblu inc. Unblu Spark 7.0 through 7.57.9 and 8.0 before 8.40.2, on on-premises and cloud deployments, allows an unauthenticated remote attacker in the visitor role to execute arbitrary JavaScript in the browser of an agent taking part in the attacker's co-browsing session, via manipulated synchronized page attributes of the co-browsed page. The JavaScript runs in the context and under the identity of the agent's session in the Agent Desk or the Single Conversation Desk. Exploitation requires an agent to join an embedded co-browsing session with the attacker. Mobile co-apping isn't affected.	N/A	More Details
CVE-2026-132043	In Eclipse Yasson versions 3.0 to 3.0.5, a race condition (CWE-362) exists in the handling of @JsonbNumberFormat. Yasson creates one java.text.NumberFormat instance per annotated property and shares it across all threads that use the same Jsonb instance. NumberFormat is not thread-safe, and the deserializer called parse() on it without synchronization. When multiple threads deserialize concurrently, a thread can receive a numeric value that belongs to another thread's input, receive an incorrect value, or fail with an exception. A remote attacker who can send concurrent requests to an application that deserializes untrusted JSON this way may cause incorrect data to be processed or cause another user's value to be returned. Exploitation requires that the application use @JsonbNumberFormat with a shared Jsonb instance and that the attacker win a timing race. This issue is fixed in version 3.0.6 by synchronizing access to the shared NumberFormat instances.	N/A	More Details
CVE-2026-92613	In Eclipse iceoryx-rs, versions 0.0.11 to 0.0.13 and 0.1.0 as well as on the latest main branch. `SampleReceiver::<T>::take()` and `Sample::<T>::deref()` are safe functions that hand the caller a `&T` pointing at a chunk written by another iceoryx process, but the subscriber path neither requires `T: ShmSend` nor verifies that the received chunk is at least `size_of::<T>()` bytes / `align_of::<T>()`-aligned. A peer process that publishes an undersized chunk, or any chunk at all when the subscriber chose a pointer-bearing `T`, drives a `#![forbid(unsafe_code)]` subscriber into out-of-bounds reads of adjacent shared-memory chunks or dereference of a publisher-controlled pointer. Resolution: It will not be fixed since this is a known limitation in iceoryx as it is the users responsibility to specify the same data type on sender and receiver side. [iceoryx2](https://github.com/eclipse-iceoryx/iceoryx2) is the successor of iceoryx-rs where the API is designed to prevent this by default. It is strongly recommended to migrate to iceoryx2 which is actively maintained and receives security support. The iceoryx-rs package has officially reached End-of-life and will not receive any features, bugfixes or security fixes.	N/A	More Details
CVE-2026-112216	Vulnerability involving the exposure of sensitive information in the URLs of the UOC Virtual Campus's document upload feature. After completing a document upload, the application includes identifiers related to the active session—such as `apjsessionid`—and other internal parameters associated with requests and users in the HTTP GET request parameters. This data may be stored in the browser history and in server or proxy logs, or transmitted via the HTTP Referer header, depending on the configuration and browsing flow. Its inclusion in the URL could allow for the unauthorized disclosure of information to third parties with access to those logs or recipients.	N/A	More Details
CVE-2026-129517	A cross-site scripting vulnerability exists in Win Bird JYUGYO SHIEN browser extensions for Google Chrome and Microsoft Edge, which allows an attacker to execute arbitrary script in the user's web browser.	N/A	More Details
CVE-2026-132989	Rejected reason: This candidate is a duplicate of CVE-2026-132923. All CVE users should reference CVE-2026-132923 instead of this candidate.	N/A	More Details
CVE-2026-133001	FlowIntel contains an authorization flaw in its duplicate alert ingestion logic. When a duplicate alert is received and its content is updated, the ownership and visibility scope fields (owner user and owner organization) were not refreshed to reflect the connector's current scope. As a result, an alert that was originally created under a broader visibility scope retained that broader scope even after the connector's scope was narrowed or changed. This means that users or organizations that should no longer have visibility into a given alert could continue to access it, because the stale ownership fields were never updated during the duplicate-update path. The issue affected all duplicate updates, including those involving alerts in terminal review states and those recovered from uniqueness conflicts. Security impact: - Unauthorized disclosure of security alert content to users or organizations outside the intended visibility scope. - Stale access persists indefinitely until the alert is independently re-scoped or deleted. Affected version: <4.1.0.	N/A	More Details
CVE-	The Guardian 2 alarm unit lacks a mechanism to authenticate control commands sent over 433.92 MHz, which allows an attacker to replay the static commands to impersonate activity. The attacker can also		More

2026-82948	disarm or arm the system at will if a Nexsmart keypad (product link) is in use, thus compromising the functionality of the alarm system completely.	N/A	Details
CVE-2026-82949	The static codes used for information transmission throughout the 433.92 MHz system of the Guardian 2 system allow an attacker in proximity to monitor the movements of users inside the premises where the sensors are set up. This is further aided by the fact that codes transmitted reveal sensor type, although sensor location cannot be inferred from this traffic alone.	N/A	More Details
CVE-2026-82950	The device only supports Wi-Fi protocols vulnerable to deauthentication attacks (WEP/WPA/WPA2). Should any event which would typically be logged to the user’s mobile device (such as arming or disarming the alarm or any tamper events) be incurred whilst the device is temporarily disconnected from its Wi-Fi network (such as via a deauthentication attack), those alerts will not be sent even after connection is regained.	N/A	More Details
CVE-2026-132129	Passing a different column type while using deprecated range predicates, allow OOB read of any memory address. Users should update to 1.17.2 or 1.18.2	N/A	More Details
CVE-2026-115342	Argument injection in OpenVPN version 2.4.0 through 2.6.23 and 2.7_alpha1 through 2.7.8 on Windows allows local authenticated users to bypass interactive service option validation and add arbitrary configuration options via a crafted command line.	N/A	More Details
CVE-2026-115684	In specific versions of Octopus Tentacle, Tentacle fails to unset registration credentials as environment variables, resulting in deployment scripts potentially having the ability to read those credentials.	N/A	More Details
CVE-2026-87850	nVision sends user's token via query parameters of GET requests. An attacker who obtains valid token can gain access to the victim's account and perform actions on their behalf. This issue was fixed in version 18.	N/A	More Details
CVE-2026-132069	MISP (Malware Information Sharing Platform) contains an authorization bypass in the SharingGroup model's captureSGExisting method. The authorization check used a logical AND condition that effectively exempted users with the sync permission (perm_sync) from the membership verification. Specifically, the original condition 'if (!checkIfAuthorised(...) && !perm_sync) return false' meant that any user holding the sync role would pass the check regardless of whether they were actually a member of the target sharing group. A sync user who knows or can discover the UUID of a non-local sharing group could exploit this flaw through the /sharing_groups/add endpoint or by embedding a SharingGroup object in pushed sync data. The attacker could rewrite the sharing group record to add their own organisation as a member or insert an all_orgs server entry for their MISP instance, thereby gaining read access to all events contained in that group. This exposes threat intelligence data that was intended to be restricted to the group's legitimate members. The vulnerability requires an authenticated account with the sync role and knowledge of a valid sharing group UUID. No user interaction is required. The impact is unauthorized disclosure of shared threat intelligence events and limited integrity modification of the sharing group record.	N/A	More Details
CVE-2026-77283	obsidian-sync-mcp is an MCP server that gives AI agents access to Obsidian vaults. Prior to 0.5.7, deployments running without MCP_AUTH_TOKEN install no serverOptions.authenticate check for the HTTP /mcp endpoint, so requests are dispatched without validating the Host or Origin headers. A malicious website visited by the operator can use DNS rebinding to reach loopback, LAN, VPN, or container-published addresses, or directly issue a cross-origin request to loopback, and invoke the MCP tool surface to read, create, overwrite, delete, or move vault notes. Deployments using MCP_AUTH_TOKEN are not affected because browsers do not attach the bearer token. This issue is fixed in version 0.5.7.	N/A	More Details
CVE-2026-71534	Mareforma is a local verification layer for AI-assisted research. From 0.3.0 until 0.3.8, restore logic in mareforma/db/restore.py derives transparency_logged from the unsigned rekor.uuid field inside signature_bundle without requiring a verified claim-bound rekor_inclusions entry, so an attacker who modifies claims.toml can falsely make a claim eligible for REPLICATED convergence. The same logic resets signed claims in non-Rekor graphs from transparency_logged 1 to 0 after restore, permanently preventing convergence because no Rekor refresh is configured. This issue is fixed in version 0.3.8.	N/A	More Details
CVE-2026-132643	MISP contains an information disclosure vulnerability in the Sighting model when the Sightings_anonymise setting is enabled. The listSightings and sightings/restSearch endpoints accepted a caller-supplied org_id filter parameter and applied it directly to the SQL query without verifying whether the requesting user was authorized to filter by that organisation. Because the anonymization mechanism masks the reporter's organisation in the returned records, a non-site-admin user could bypass the anonymization by specifying another organisation's ID in the filter. The query would then return only that organisation's sightings, effectively naming the reporter and defeating the purpose of the anonymization feature. A negated filter (NOT IN) similarly revealed which sightings belonged to the target organisation by exclusion. The vulnerability requires an authenticated user with sighting read permissions (perm_sighting) who is not a site administrator, and the Sightings_anonymise plugin setting must be enabled. The security impact is the disclosure of which organisation reported a given sighting, undermining the privacy guarantees that the anonymization feature is designed to provide. This issue was reported by ENISA.	N/A	More Details

CVE-2026-71533	Mareforma is a local verification layer for AI-assisted research. From 0.3.7 until 0.3.8, search_claims in mareforma/db/core.py does not call _row_verified_on_read and therefore does not provide the same high-trust verification projection as query_claims. A forged or tampered REPLICATED or ESTABLISHED row that query_claims excludes can be returned by search_claims at its persisted trust level, and the missing verified marker prevents consumers from detecting that signature verification was skipped. The include_unverified option does not authorize this high-trust bypass. This issue is fixed in version 0.3.8.	N/A	More Details
CVE-2026-55169	Open OnDemand is an open-source high-performance computing portal. Prior to 4.1.6 and 4.2.3, the Open OnDemand dashboard includes noVNC passwords in query parameters generated by apps/dashboard/app/helpers/batch_connect/sessions_helper.rb and apps/dashboard/app/views/batch_connect/sessions/connections/_novnc.html.erb. The server records those query strings in per-user PUN access logs, and insufficiently protected log directories can allow another local user to recover and replay a launch URL. Replaying the credential grants access to the running remote desktop and can result in horizontal or vertical privilege escalation according to the privileges of the desktop owner. This issue is fixed in versions 4.1.6 and 4.2.3.	N/A	More Details
CVE-2026-55171	Open OnDemand is an open-source high-performance computing portal. Prior to Open OnDemand 4.1.6 and 4.2.3, and prior to ood-ansible 4.2.1, the portal generator defaults OIDCCryptoPassphrase to SHA1(servername), and the templates/ood_portal.yml.j2 role checks a different variable from the documented passphrase setting. Administrators who set ood_oidc_crypto_passphrase can therefore believe a strong secret is configured while the deployment continues to use an attacker-derivable value. In OIDCSessionType client-cookie deployments, an unauthenticated remote attacker can derive the encryption key and forge a mod_auth_openidc_session cookie for another user, while server-cache deployments permit lower-impact tampering with protected OIDC state and nonce cookies. This issue is fixed in Open OnDemand versions 4.1.6 and 4.2.3 and ood-ansible version 4.2.1.	N/A	More Details
CVE-2026-55172	Open OnDemand is an open-source high-performance computing portal. Prior to 4.1.6 and 4.2.3, apps/dashboard/config/initializers/session_store.rb stores dashboard sessions by default in a predictable /var/tmp/<username> directory, or in another location configured by an administrator, but omits PUN user directory ownership and mode 0700 validation. A local attacker can pre-create another user's directory, read the victim's session identifier and CSRF token after login, or stage a predetermined session identifier. ActiveSupport file-store deserialization of attacker-controlled cache data may also occur in the victim's context, enabling file modification or code execution with the victim's privileges. This issue is fixed in versions 4.1.6 and 4.2.3.	N/A	More Details
CVE-2026-76100	OwnPay is a self-hosted open-source payment gateway automation platform. In 0.2.0 and earlier, modules/gateways/alipay/AlipayGateway.php fails open when an Alipay callback lacks a signature or the merchant has no configured public key because verify(array $callbackData, array $credentials) falls back to an undefined mode credential that defaults to test mode, while verifyWebhook() returns true unconditionally. An unauthenticated request to POST /webhook/{gateway} can provide a pending order reference and matching amount, pass both verification gates without a valid payment signature, and cause the order and corresponding payment-received ledger entries to be marked completed. No fixed version is available as of this review.	N/A	More Details
CVE-2026-77295	MyMangaDB is a self-hosted manga collection manager with automatic metadata, MyAnimeList import, Docker support, and collection statistics. Prior to 2.7.2, the unauthenticated /api/v1/database/export handler in backend/app/api/v1/endpoints/database.py returns a ZIP containing the live SQLite database and the cover-image library. The database includes the users table and hashed_password values, allowing a remote attacker to obtain all application data and conduct offline password attacks without a session or token. This issue is fixed in version 2.7.2.	N/A	More Details
CVE-2026-77296	MyMangaDB is a self-hosted manga collection manager with automatic metadata, MyAnimeList import, Docker support, and collection statistics. Prior to 2.7.2, the unauthenticated /api/v1/images/manga/save handler in backend/app/api/v1/endpoints/image.py accepts a client-controlled filename and joins it to IMAGE_SAVE_PATH without confining the result. Traversal sequences can escape the image directory and overwrite process-writable files, including DATABASE_PATH for the live SQLite database, while the client-supplied image content type does not validate the uploaded bytes. An attacker can replace the database to plant an administrative user and can overwrite other writable files, causing full application compromise and availability loss. This issue is fixed in version 2.7.2.	N/A	More Details
CVE-2026-77297	MyMangaDB is a self-hosted manga collection manager with automatic metadata, MyAnimeList import, Docker support, and collection statistics. Prior to 2.7.2, the unauthenticated /api/v1/database/import handler in backend/app/api/v1/endpoints/database.py accepts a ZIP archive and moves an included database file to DATABASE_PATH while also replacing IMAGE_PATH content. An attacker can supply a database containing an attacker-controlled users table and administrative account, then authenticate as that account; invalid replacement data can also destroy application data or availability. This issue is fixed in version 2.7.2.	N/A	More Details
CVE-2026-106094	Improper neutralization of special elements used in an OS command ('OS command injection') vulnerability in Westell Technologies, Inc Remote series products on Linux allows OS Command Injection. This issue affects several generations of Westell Remotes: through firmware version 9.14. It has been reproduced across multiple devices. However, it is confirmed to affect only devices configured for RAS (Remote Access Server -	N/A	More Details

	e.g., RADIUS, TACACS).		
CVE-2026-121720	OpenVPN 2.4.0 through 2.6.23 and 2.7_alpha1 through 2.7.8 on Windows allows local authenticated users to execute arbitrary code via a privileged service using user-controlled environment variables.	N/A	More Details
CVE-2026-132644	MISP contains an authorization flaw in the GalaxyCluster model's captureCluster() function, which handles the import and synchronization of Galaxy Clusters. Prior to the fix, any authenticated user holding the galaxy editor role could replace an existing locked Galaxy Cluster by supplying its UUID in an import request, regardless of which organization owned the cluster. The function unconditionally called captureOrganisationAndSG() before the save operation, overwriting the cluster's org_id and orgc_id fields with the importing user's organization, effectively transferring ownership of the cluster to the attacker's organization. Additionally, the save operation was not restricted to a specific set of fields, allowing arbitrary field modification. The vulnerable behavior required the attacker to possess a valid MISP account with the perm_galaxy_editor role and to know or discover the UUID of a target cluster belonging to another organization. No site-admin or owner-organization privilege was enforced for the update path. The result was that a low-privilege galaxy editor in one organization could modify the content (value, description, distribution, etc.) of clusters owned by other organizations and seize ownership, corrupting the integrity of shared threat-intelligence data across the MISP instance. The vulnerability was reported by ENISA and CERT.pl. The fix introduces an explicit authorization gate that restricts cluster updates to the owner organization, a sync user who can see the cluster via sharing-group rules, or a site administrator, and enforces a field allow-list that excludes org_id and orgc_id from update writes, preserving original ownership.	N/A	More Details
CVE-2026-132645	MISP contains an authorization bypass vulnerability in the Galaxy import/sync functionality (captureGalaxy in app/Model/Galaxy.php). The captureGalaxy() method matched an incoming galaxy payload to an existing galaxy by UUID and unconditionally overwrote the matched record, regardless of whether the importing user belonged to the owning organization or whether the galaxy was a system-provided default galaxy. Any authenticated user holding the perm_galaxy_editor role could therefore: (1) demote a default galaxy by setting its default flag to false, (2) reassign the galaxy's org_id and orgc_id to the attacker's organization, effectively taking ownership, and (3) modify other galaxy fields (name, description, distribution, version). Additionally, the role-level permission check (perm_site_admin or perm_galaxy_editor) was executed after the captureOrganisationAndSG() call, meaning organization and sharing-group data was already captured before the authorization decision was made. The vulnerability requires an authenticated user with the perm_galaxy_editor role and the ability to submit a galaxy import payload containing a UUID that matches an existing galaxy. The security impact is loss of data integrity across the MISP instance: default galaxy definitions are corrupted, cross-organization ownership is silently transferred, and the integrity of shared threat-intelligence galaxy metadata is compromised for all users of the instance.	N/A	More Details
CVE-2026-63411	Jane is a set of PHP libraries that generates models and API clients from JSON Schema and OpenAPI specifications. Prior to 7.12.0, the Jane open-api-3 generator emits an OpenAPI parameter description verbatim in generated Client/Endpoint method docblocks. A parameter description containing a PHPDoc closing delimiter can terminate the comment early and place attacker-controlled PHP in class-body position, including an __construct method that runs when the generated Client is instantiated. Exploitation requires generating a client from an attacker-influenced OpenAPI document and then using the generated SDK, and can execute arbitrary PHP with the consuming application's privileges. This issue is fixed in version 7.12.0.	N/A	More Details
CVE-2026-63412	Jane is a set of PHP libraries that generates models and API clients from JSON Schema and OpenAPI specifications. Prior to 7.12.0, the Jane open-api-3 generator emits an OpenAPI query parameter name verbatim as an array key in generated Client/Endpoint method docblocks. A query parameter name containing a PHPDoc closing delimiter can terminate the comment early and place attacker-controlled PHP in class-body position, including an __construct method that runs when the generated Client is instantiated. Exploitation requires generating a client from an attacker-influenced OpenAPI document and then using the generated SDK, and can execute arbitrary PHP with the consuming application's privileges. This issue is fixed in version 7.12.0.	N/A	More Details
CVE-2026-63413	Jane is a set of PHP libraries that generates models and API clients from JSON Schema and OpenAPI specifications. Prior to 7.12.0, the Jane open-api-3 generator emits an OpenAPI header parameter name verbatim as an array key in generated Client/Endpoint method docblocks. A header parameter name containing a PHPDoc closing delimiter can terminate the comment early and place attacker-controlled PHP in class-body position, including an __construct method that runs when the generated Client is instantiated. Exploitation requires generating a client from an attacker-influenced OpenAPI document and then using the generated SDK, and can execute arbitrary PHP with the consuming application's privileges. This issue is fixed in version 7.12.0.	N/A	More Details
CVE-2026-63415	Jane is a set of PHP libraries that generates models and API clients from JSON Schema and OpenAPI specifications. Prior to 7.12.0, Jane's Reference.php::doResolve passes an attacker-controlled OpenAPI $ref target to file_get_contents without restricting its scheme or resolved local path. When a victim generates code from an attacker-influenced document, an HTTP or HTTPS reference can cause requests to attacker-chosen or internal services, while an absolute or parent-relative path can read local JSON or YAML files and inline the retrieved content into generated output. Exploitation occurs during generation or CI processing and can disclose local files and internal-service responses. This issue is fixed in version 7.12.0.	N/A	More Details

CVE-2026-68522	QPixel is a Q&A-based community knowledge-sharing software. Prior to 0.13.0, the unauthenticated comments#show action in app/controllers/comments_controller.rb resolves user-controlled comment IDs with Comment.unscoped, bypassing the Comment community default scope, and performs no deletion or post-visibility authorization check. A remote unauthenticated attacker can enumerate sequential IDs and read the full content of comments belonging to other communities, soft-deleted comments, and comments under access-restricted categories. The JSON response exposes the stored comment columns because the model has no limiting as_json override. This issue is fixed in version 0.13.0.	N/A	More Details
CVE-2026-75518	DOMShell is a Chrome automation project whose MCP server exposes Unix-style tools for AI agents. Prior to 2.0.8, the DOMShell MCP server classifies the domshell_eval tool and eval command as read-tier operations that can run without --allow-write, even though eval shares the Runtime.evaluate dispatch path with the write-tier js command and has no side-effect gate. An MCP caller with tool-session access can therefore use domshell_eval to mutate DOM, window, or global state while the host approval interface, server authorization flags, denial messages, and audit records represent the operation as read-only. The bypass does not expose a capability beyond domshell_js for callers already authorized to write, but it violates the configured permission boundary for callers lacking that authorization. This issue is fixed in version 2.0.8.	N/A	More Details
CVE-2026-125353	An incorrect permission assignment vulnerability in Omada SDN Controller prior to 6.2.20 may cause a database backup generated during the uninstall workflow to remain readable by other local users on the host. Successful exploitation requires local, low-privileged access to the affected host and a retained uninstall backup. An attacker meeting these conditions could access the backup and disclose controller data without controller credentials.	N/A	More Details
CVE-2026-129315	Dokploy is a free, self-hostable Platform as a Service (PaaS). In 0.31.0 and earlier, the hosted installer and manual installation instructions set /etc/dokploy to mode 0777, allowing a local non-root user on a shared host to replace root-owned schedule or Traefik configuration directories. Dokploy subsequently consumes those paths as trusted configuration and can execute substituted scheduled-task scripts as root, resulting in local privilege escalation.	N/A	More Details
CVE-2026-129985	Dokploy is a free, self-hostable Platform as a Service (PaaS). Prior to 0.31.1, a user-controlled schedule appName is joined into the schedule directory path and interpolated without shell quoting into commands that create the directory and write the schedule script. An authenticated user who can create or update a dokploy-server schedule can inject operating-system commands that execute with the Dokploy process privileges when the schedule is saved, including root privileges on typical installations. This issue is fixed in version 0.31.1.	N/A	More Details
CVE-2026-71500	Mareforma is a local verification layer for AI-assisted research. From 0.3.0 until 0.3.8, _maybe_update_replicated_unlocked in mareforma/db/core.py does not check t_invalid in its new-claim gate, peer-selection query, or final update, so a claim marked invalid by a signed contradiction verdict can still be promoted to REPLICATED through convergence. The invalidated claim can also act as the corroborating peer that promotes an honest claim, bypassing the trust ladder's terminal refutation state. This issue is fixed in version 0.3.8.	N/A	More Details
CVE-2026-99402	In the Linux kernel, the following vulnerability has been resolved: netfilter: nft_payload: restrict checksum offsets to known values We need to prevent userspace from corrupting e.g. tcp->doff, because many locations in conntrack and conntrack helpers rely on nf_conntrack_in() having validated the packet headers. nft_payload allows to alter headers later which invalidates this assumption. The 'Fixes' commit restricts writes to safe fields, but there is another side channel: the checksum location. Restrict this too. Reported via sashiko/gemini.	N/A	More Details
CVE-2026-99404	In the Linux kernel, the following vulnerability has been resolved: xfs: check fork type against dabtree block magic number LOLLM complains that we fail to flag dabtree blocks that should never be associated with the file fork -- data forks can't point to xattr leaves, and attr forks can't point to dirent blocks. Add the missing checks.	N/A	More Details
CVE-2026-118977	Authorization bypass through a user-controlled key in the PuppetDB facts terminus used by Puppet Server allows a client holding any valid agent certificate to overwrite the facts, including trusted facts, that PuppetDB stores for a different node. In Puppet Enterprise deployments where task permissions are scoped to node groups defined by facts or trusted facts, this can allow a user to run tasks on nodes outside their permitted scope.	N/A	More Details
CVE-2026-132151	Rejected reason: Reason for rejection.	N/A	More Details
CVE-2026-123146	Rejected reason: Inline mitigations exists.	N/A	More Details
CVE-2026-127925	Rejected reason: Inline mitigations exists.	N/A	More Details

CVE-2026-128884	Rejected reason: Inline mitigations exists.	N/A	More Details
CVE-2026-131807	United Chargers Grizzl-E Smart 40A OCPP GetConfiguration Stack-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the handling of OCPP messages. When parsing GetConfiguration requests, the process does not properly validate the length of user-supplied data prior to copying it to a fixed-length stack-based buffer. An attacker can leverage this in conjunction with other vulnerabilities to execute code in the context of the device. Was ZDI-CAN-28968.	N/A	More Details
CVE-2026-131808	United Chargers Grizzl-E Smart 40A Wi-Fi Configuration Server Missing Authentication Vulnerability. This vulnerability allows network-adjacent attackers to modify Wi-Fi configuration on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the handling of Wi-Fi configuration requests. The issue results from the lack of authentication prior to allowing access to functionality. An attacker can leverage this in conjunction with other vulnerabilities to execute arbitrary code in the context of the device. Was ZDI-CAN-29068.	N/A	More Details
CVE-2026-131809	United Chargers Grizzl-E Smart 40A OCPP Backend Configuration Authentication Bypass Vulnerability. This vulnerability allows network-adjacent attackers to bypass authentication on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the handling of OCPP backend configuration requests. The issue results from the lack of authentication prior to allowing access to functionality. An attacker can leverage this in conjunction with other vulnerabilities to execute arbitrary code in the context of the device. Was ZDI-CAN-29063.	N/A	More Details
CVE-2026-131810	United Chargers Grizzl-E Smart 40A Download of Code Without Integrity Check Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the handling of OCPP messages. The issue results from the lack of proper verification of firmware integrity prior to installation. An attacker can leverage this in conjunction with other vulnerabilities to execute code in the context of the device. Was ZDI-CAN-28957.	N/A	More Details
CVE-2026-131811	United Chargers Grizzl-E Smart 40A OCPP Backend Configuration Hard-Coded Credentials Authentication Bypass Vulnerability. This vulnerability allows network-adjacent attackers to bypass authentication on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the configuration of the OCPP backend settings interface. The issue results from the use of hard-coded credentials. An attacker can leverage this vulnerability in conjunction with other vulnerabilities to execute arbitrary code in the context of the device. Was ZDI-CAN-29056.	N/A	More Details
CVE-2026-131812	United Chargers Grizzl-E Smart 40A FTP URL Path Parsing Stack-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the parsing of FTP URL paths during firmware updates. The issue results from the lack of proper validation of the length of user-supplied data prior to copying it to a fixed-length stack-based buffer. An attacker can leverage this in conjunction with other vulnerabilities to execute code in the context of the device. Was ZDI-CAN-28960.	N/A	More Details
CVE-2026-131813	United Chargers Grizzl-E Smart 40A OCPP UpdateFirmware Out-Of-Bounds Write Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of United Chargers Grizzl-E Smart 40A EV chargers. Authentication is not required to exploit this vulnerability. The specific flaw exists within the handling of OCPP UpdateFirmware requests. The issue results from the lack of proper validation of user-supplied data, which can result in a write past the end of an allocated buffer. An attacker can leverage this in conjunction with other vulnerabilities to execute code in the context of the device. Was ZDI-CAN-29487.	N/A	More Details
CVE-2026-62255	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, cram_index_load does not verify that the number of reference sequences in a CRAM .crai index does not exceed the CRAM header reference count. A malformed index can drive unbounded index-array allocation, and when the number of references in the index is within two of INT_MAX, the growth calculation can wrap, read beyond a heap buffer, and use the resulting value as a pointer for an unchecked memory write. Any region query can trigger the flaw, causing denial of service or memory corruption. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-129400	Joomla Extension - vi-solutions - Unauthenticated verification mail amplification in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The verification mail task accepted arbitrary form IDs and recipients without rate limiting, allowing unauthenticated attackers to trigger mass verification mails and database entries.	N/A	More Details
	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, fai_retrieve does not verify that the line width recorded in a FASTA .fai index is greater than or equal to the		

CVE-2026-62256	number of bases per line. When the line width is exactly one less than the base count, retrieval allocates an undersized heap buffer and writes one byte beyond its end; larger discrepancies instead cause an allocation failure that is handled. This issue can corrupt heap memory and crash the process. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-62257	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, the experimental CRAM v4 XPACK and XRLE decoders cache expanded block pointers at indexes beyond the 512-element slice->block_by_id array. An array index partially controlled by data in the input file first causes an out-of-bounds read to determine whether a block is cached; a null value then causes allocation of a block and an out-of-bounds pointer write to the same element. The invalid read may crash the process, while the write may corrupt heap state and could permit arbitrary code execution. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-62258	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, the CRAM decoder validates the 254-character query-name limit only after decoding the name into a 1024-byte stack array. An attacker-controlled CRAM record can provide an overlong query name and overflow the local stack buffer before the later validation runs; all CRAM format versions are affected. Opening such a file can crash the process, corrupt stack data, and may permit arbitrary code execution. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-62259	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, the experimental CRAM v4 functions cram_xpack_decode_char and cram_xrle_decode_char do not validate that attacker-controlled decode sizes remain within their expanded input arrays. Crafted CRAM data can therefore make either decoder read beyond the supplied input buffer, potentially disclosing process-state information or crashing the process through invalid memory access. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-62260	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, the CRAM BETA decoder functions cram_beta_decode_long, cram_beta_decode_int, and cram_beta_decode_char, and the experimental CRAM v4 XPACK functions cram_xpack_decode_long and cram_xpack_decode_int, have bounds-checking flaws caused by integer overflow after a multiplication. The resulting incorrect bounds check can allow decoding to read beyond the supplied input buffer, potentially leaking process-state information or crashing the process. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-63390	HTSlib is a library for reading and writing bioinformatics file formats. Prior to 1.21.2, 1.22.3, 1.23.2, and 1.24, the BCF reader can experience an integer overflow in a record-length check that verifies whether a record contains enough bytes for the expected number and size of stored items. A crafted BCF record can bypass this check, causing reads beyond the stored data and, in some cases, writes beyond the heap buffer when the reader sets the least significant bit of accessed bytes. Opening such a file can crash the process, corrupt data or heap structures, and may permit arbitrary code execution. This issue is fixed in versions 1.21.2, 1.22.3, 1.23.2, and 1.24.	N/A	More Details
CVE-2026-105802	Git LFS is a Git extension for versioning large files. From 0.6.0 until 3.8.1, Git LFS reads Git configuration options defined for a submodule without checking certain values for embedded line-ending control characters. A malicious repository can insert control characters such as line feed (LF) into those values to mimic options that cause Git LFS to execute an attacker-controlled command or binary while fetching or adding objects. This issue is fixed in version 3.8.1.	N/A	More Details
CVE-2026-105803	Git LFS is a Git extension for versioning large files. From 1.3.0 until 3.8.1, Git LFS does not correctly filter the limited set of options permitted in a repository's .lfsconfig file when option keys contain certain additional components. Crafted keys can mimic options that configure Git LFS to use a custom command to transfer objects, causing Git LFS to execute an attacker-selected command or binary while fetching objects. This issue is fixed in version 3.8.1.	N/A	More Details
CVE-2026-129396	Joomla Extension - vi-solutions - Arbitrary file write leading to remote code execution in CSS editor in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The CSS editor did not restrict file names, extensions and paths of created and saved files, allowing privileged backend users to write arbitrary PHP files.	N/A	More Details
CVE-2026-129397	Joomla Extension - vi-solutions - Arbitrary file deletion via path traversal in CSS editor in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The delete and rename tasks of the CSS editor lacked authorization checks, CSRF protection and path validation, allowing backend users to delete files outside the CSS directory.	N/A	More Details
CVE-2026-102770	Improper neutralization of special elements in data query logic (Cypher injection) in the fulltext and vector index query generator of the sharded property database in Neo4j Enterprise Edition 2026.07.0 before 2026.10.0 allows an authenticated user holding ACCESS, CREATE INDEX and NAME MANAGEMENT privileges to inject Cypher that is accepted for execution on a property shard, in the session of any user who later queries the index, via a crafted index name containing a single quote. The impact is bounded by the privileges of the querying user.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ksmbd: wait for negotiation before		

CVE-2026-99565	receiving another PDU The receive thread queues a request and can read the next PDU before the negotiation worker finishes. While need_neg is set, the next PDU can cause ksmbd_init_smb_server() to replace conn->ops. A pipelined SMB2 request can therefore make the worker handle an earlier, shorter SMB1 request with SMB2 operations and read beyond its buffer. An SMB1 multi-protocol negotiate may be followed by an SMB2 NEGOTIATE that changes conn->ops again. Wait for work queued while the connection is new or needs negotiation to complete before reading another PDU. This keeps the operations stable for both stages of negotiation without changing when need_neg is cleared.	N/A	More Details
CVE-2026-99562	In the Linux kernel, the following vulnerability has been resolved: ksmbd: reject FSCTL_SET_ZERO_DATA on a directory handle An authenticated client with a writable tree connection can open a directory with FILE_WRITE_DATA and send FSCTL_SET_ZERO_DATA on that handle. The handler only checks fp->daccess & FILE_WRITE_DATA_LE, but on a directory FILE_WRITE_DATA is the legitimate add-file right, so the check passes. ksmbd_vfs_zero_data() then calls kernel_write() on the directory file, which has no FMODE_WRITE, and __kernel_write_iter() hits WARN_ON_ONCE(!(file->f_mode & FMODE_WRITE)), panicking a panic_on_warn kernel. Before commit 25e414db7033 ("ksmbd: report holes in allocated range queries") the dense-file path used vfs_fallocate(), which rejects directories with -EISDIR. The kernel_write() loop that replaced it has no such check. Reject a directory handle in ksmbd_vfs_zero_data() before writing, as ksmbd_vfs_write() already does for SMB2_WRITE. WARNING: fs/read_write.c:608 at __kernel_write_iter+0x74b/0x8c0 Workqueue: ksmbd-io handle_ksmbd_work Call Trace: kernel_write+0x1ef/0x4d0 ksmbd_vfs_zero_data+0x413/0x590 smb2_ioctl+0x21fa/0x66a0 handle_ksmbd_work+0x61b/0x1ff0	N/A	More Details
CVE-2026-99561	In the Linux kernel, the following vulnerability has been resolved: fork: don't create io threads once PF_POSTCOREDUMP is set zap_process() skips every thread that already has PF_POSTCOREDUMP set. Such a thread is past synchronize_group_exit(), so a coredump can't catch it anymore. It isn't counted in core_state->threads_remaining and it isn't sent SIGKILL. do_exit() sets PF_POSTCOREDUMP in synchronize_group_exit() and calls io_uring_files_cancel() right after that. The cancellation runs task work and a create_worker_cb() that io-wq queued before the exit creates a new io-wq worker from there. If another thread started a coredump in the meantime that worker joins a thread-group which is already being dumped. zap_process() never saw it so it was never counted. But when it exits it sees signal->core_state in synchronize_group_exit() and decrements threads_remaining like any other thread. So the count reaches zero one thread early and the dumper leaves coredump_wait_inactive() while a thread it counted is still running. The task has no fatal signal pending because zap_process() deliberately didn't send it one, and it never went through get_signal() so it doesn't have PF_SIGNALED either. Refuse to create an io thread when the creator has PF_POSTCOREDUMP. That costs nothing. io_uring_files_cancel() raises IO_WQ_BIT_EXIT before it runs any task work, so a worker created from there only ever gets to exit again, and io_should_retry_thread() doesn't retry -EINTR. Clearing PF_POSTCOREDUMP for the new thread in copy_process() was the other option. It makes the worker a thread like any other, but it only helps while the dump hasn't started. A worker born after zap_process() has run is invisible to it whatever its flags say and still decrements threads_remaining on the way out. Refusing to create it covers both, and then no task is ever born with the flag, so there is nothing to clear. Moving io_uring_files_cancel() ahead of synchronize_group_exit() closes the same window. It runs the cancellation, and whatever that can block on, before the thread announces itself to the dumper. A thread that blocks there never reaches coredump_task_exit() at all, so the dumper ends up waiting for a thread that never parks. Refusing the creation leaves the cancellation where it is. But it's very ugly to run io_uring work even before we did all the generic exit work.	N/A	More Details
CVE-2026-99533	In the Linux kernel, the following vulnerability has been resolved: ksmbd: validate AppInstanceId takeover target AppInstanceId lookup currently selects the first open with the matching identifier. CREATE for an unrelated share or path can therefore force-close another session's open before authorization of the new target. Matching only the base path also conflates alternate data streams with the unnamed stream. MS-SMB2 requires the old open to match the share and path, belong to a different ClientGuid, and be closed only when the requester has GENERIC_READ access. Search all opens with the ID and match share, resolved path, stream, and the current connection's client GUID. Retain a share configuration reference in ksmbd_file and reuse its name instead of allocating a name for every open. This keeps matching safe during tree connection teardown and preserves name-based matching across configuration updates. Refresh the share reference on durable reconnect and restore the old reference if reopening fails. Serialize connection and share comparisons with durable disconnect/reconnect, and recheck the target before closing it. Keep the original ClientGuid for durable replay and lease checks. Check AppInstanceVersion before querying access, even when GENERIC_READ is absent, and before disconnected delete-on-close cleanup. Rejected CREATE must leave other durable opens untouched. Calculate maximal access on the matched open's referenced path so a rename cannot substitute another file's authorization. Require read access with the effective VFS credentials as well as the NT DACL before takeover. Access denial only skips takeover and leaves CREATE to its normal access checks. Restore filesystem credentials before returning a takeover error. Handle takeover before checking FILE_CREATE collisions, and re-resolve the path after closing the old open. This lets delete-on-close remove the old file before a new one is created. Include SYNCHRONIZE in the baseline maximal-access mask so read-only files report GENERIC_READ correctly.	N/A	More Details
CVE-2026-99534	In the Linux kernel, the following vulnerability has been resolved: ksmbd: prevent stale data from being written by SMB2 WRITE requests Initialize the compatibility byte in request buffers without zeroing the entire allocation. Keep the one-byte compatibility exception for SMB2 messages so valid zero-length WRITE requests remain compatible, and validate the data range in the WRITE handlers against the current	N/A	More Details

	compound request before accessing it. This prevents stale request-buffer data from being written to files or sent to IPC pipes.		
CVE-2026-99535	In the Linux kernel, the following vulnerability has been resolved: ksmbd: prevent share root removal through empty paths A FileLinkInformation request can supply "/" as the new link name. smb2_get_name() strips the trailing slash and passes an empty string to ksmbd_vfs_kern_path_start_removing(). ksmbd_vfs_path_lookup() resolves an empty pathname as the configured share path, without LOOKUP_BENEATH. With ReplaceIfExists set, smb2_create_link() can therefore remove the share root from its parent before link creation fails. An authenticated client with a writable share and DELETE access to a file can trigger this if the share root is empty and the mapped user has write and search permission on its parent directory. Return -EINVAL for empty pathnames when for_remove is set, before resolving the share path or acquiring removal resources. Ordinary lookups continue to accept empty pathnames for opening the share root.	N/A	More Details
CVE-2026-99537	In the Linux kernel, the following vulnerability has been resolved: ksmbd: preserve unreadable named streams on open A write-only client can open an existing named stream with FILE_OPEN_IF without permission to read the underlying file. The case-insensitive lookup finds the xattr name, but querying its value length then fails with -EACCES. smb2_set_stream_name_xattr() treats every negative result as a missing stream and replaces the existing xattr with an empty value. A failed xattr list is not proof that the stream is absent either: for example, -ENOMEM can lead to the same empty replacement. Preserve listing errors and create a stream only for -ENOENT or -ENODATA. The latter can occur if the xattr is removed between listing its name and querying its value. Return other lookup errors without changing the stream. A write-only FILE_OPEN now reports access denied instead of name not found. FILE_OVERWRITE_IF and FILE_SUPERSEDE on an unreadable stream also fail instead of appearing to replace it. Stream writes and EOF changes already need to read the old value, so such handles cannot update the stream after opening either.	N/A	More Details
CVE-2026-99538	In the Linux kernel, the following vulnerability has been resolved: ksmbd: ignore network devices of other namespaces With no interface list configured, ksmbd binds every interface that comes up. ksmbd_netdev_event() does not look at the device's namespace, and create_socket() opens the listener in the namespace of the task raising NETDEV_UP. A container with its own address on a bridge brings up its interface and finds port 445 in its namespace already taken by ksmbd, so an SMB server running inside the container fails with "Address already in use". Interfaces are also matched by name only, so a container's eth0 going down would stop the listener on the host's eth0. Only handle netdev events for devices in init_net.	N/A	More Details
CVE-2026-99539	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix session reference leak on compounded SESSION_SETUP A compound request can place SESSION_SETUP after another command. That earlier command leaves a reference to its session in work->sess. smb2_sess_setup() then overwrites work->sess with the session it looks up or creates, and the dispatcher later puts work->sess only once. The inherited reference is never dropped. That leaks the session, together with its setup connection and socket, and every repeat of the compound leaks another. unreferenced object (size 512): __kmalloc_cache_noprof+0x32c/0x3a0 ksmbd_smb2_session_create+0xf6/0x9c0 [ksmbd] smb2_sess_setup+0xda4/0x6280 [ksmbd] handle_ksmbd_work+0x88c/0x1cb0 [ksmbd] Save the inherited session on entry. After the request has been processed, release it once, but only when smb2_sess_setup() has replaced work->sess with its own session or cleared it. Do this after ksmbd_conn_unlock() so the put does not run under srv_mutex. An error path may instead leave the inherited session in work->sess. Keep it there, so the dispatcher still puts it and the error response is still signed or encrypted with it as before.	N/A	More Details
CVE-2026-99542	In the Linux kernel, the following vulnerability has been resolved: ksmbd: deny access for an empty DACL instead of granting the request smb_check_perm_dacl() returns the access mask it computes to the SMB2 CREATE handler, which uses it as the granted access for the open. For a security descriptor whose DACL is present but contains no ACEs, the function is supposed to deny access: an empty (non-NULL) DACL grants no access to anyone except the object owner's implicit READ_CONTROL and WRITE_DAC (MS-DTYP 2.4.5). The empty-DACL branch only denies when both - the DACL has no trailing bytes (pdacl_size == sizeof(struct smb_acl)), and - the request contains bits beyond READ_CONTROL / WRITE_DAC. In every other case it takes the shared "goto err_out" with rc still 0, which returns success without modifying *pdaccess, so the caller grants the full requested access. Three problems follow: - The READ_CONTROL / WRITE_DAC carve-out models the owner's implicit rights, but the check runs before is_owner is computed, so those rights are handed to *every* caller, not just the owner. For READ_CONTROL the SMB2 CREATE path does not fall back to inode_permission(), so any user can read the security descriptor of an object whose empty DACL is meant to deny them. - A request of exactly READ_CONTROL and/or WRITE_DAC is granted to any caller. WRITE_DAC lets the handle rewrite the DACL. - A DACL that declares zero ACEs but carries trailing bytes (pdacl_size > sizeof(struct smb_acl)) skips the deny entirely and is granted whatever was requested. On shares where the NT ACL is the authoritative access control -- e.g. "force user" shares, where every client maps to one POSIX identity so POSIX cannot distinguish users -- this is an access-control bypass. Compute is_owner before the empty-DACL branch and deny access from an empty DACL unless the caller is the owner requesting no more than READ_CONTROL / WRITE_DAC.	N/A	More Details
CVE-2026-99543	In the Linux kernel, the following vulnerability has been resolved: backlight: qcom-wled: Fix NULL pointer dereference in wled_remove() wled_probe() doesn't set the driver data for the platform device. As a result, platform_get_drvdata() in wled_remove() will return NULL, and the first dereference of it crashes on unbind. Set the platform device driver data in wled_probe().	N/A	More Details

CVE-2026-99545	In the Linux kernel, the following vulnerability has been resolved: driver core: take parent lock on deferred probe USB interfaces should be probed with the parent USB device lock held, but this is not honoured when probe is deferred. Add a new bus_probe_device_deferred() helper to handle deferred probe where the parent device lock is not held.	N/A	More Details
CVE-2026-99547	In the Linux kernel, the following vulnerability has been resolved: thunderbolt: Do not RPM complete unrelated subtrees on unplug When a Thunderbolt device is unplugged, that switch and every switch subsequent to it on the bus device list is `complete()`d, not just its child devices. In the case of multiple domains, it may cause, eg, an attempt to `complete()` the uninitialized domain 1 root switch RPM completion on unplug of a switch from domain 0 (though ordering is not guaranteed numerically sorted). I initially noticed this issue on a Dell XPS9720 when the system would seem to sporadically not wake after unplugging a Thunderbolt dock. I tried multiple kernel versions including v7.2.7 and v7.3-rc4 to see if that resolved the issue, but the issue persisted. Eventually I found that I could consistently generate a kernel Oops reporting a page fault at 0xffff_ffff_ffff_fff8 on unplug from one side of the laptop and went bug hunting. I became suspicious of the `bus_for_each_dev()` walk that I ultimately ended up changing. To confirm this issue, I did an eBPF trace of `bus_for_each_dev()`, `complete_rpm()`, and `complete()` while unplugging in a variety of situations and found that an unplug from domain 0 tries to `complete_rpm()` domain 1's switch, whose RPM completion is not initialized and thus is holding a null wait-list pointer, triggering the page fault. Here are the key sections of that eBPF trace. Note that HEAD denotes the completion's own wait-list HEAD: ``` RESCAN_ENTER domain=0 COMPLETE rescan_domain=0 target=0-1 target_domain=0 route=0x1 parent=0-0 unplugged=1 COMPLETION done=0 head=HEAD next=HEAD prev=HEAD BUS_WALK domain=0 start=0-1 callback=complete_rpm CALLBACK_ENTER rescan_domain=0 dev=1-0 parent=domain1 is_switch=1 COMPLETE rescan_domain=0 target=1-0 target_domain=1 route=0x0 parent=domain1 unplugged=0 COMPLETION done=0 head=HEAD next=0 prev=0 complete+5 complete_rpm+43 bus_for_each_dev+133 icm_free_unplugged_children+250 icm_rescan_work+42 ``` This patch gates the `complete()` on an additional condition - that the Thunderbolt domain matches the domain of the device that was unplugged. I only have a single system with multiple Thunderbolt domains - the Dell XPS9720. I am consistently able to reproduce the Oops on that system. Since running with this patch, I have not seen the Oops reoccur. I have also tested it on a Thinkpad X1-Gen12, and that system boots and runs well with it, but it only has a single Thunderbolt domain, so it does not exercise the multi-domain failure case.	N/A	More Details
CVE-2026-99548	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix early procfs exposure race Currently, init_cifs() calls cifs_proc_init() near the beginning of the function, before INIT_LIST_HEAD, leading to a momentarily uninitialized, but exposed, procfs. If a userspace process reads this empty procfs, iterating over a zeroed list_head leads to a NULL pointer dereference and a kernel oops/panic. For example, if a delay is added between cifs_proc_init() and INIT_LIST_HEAD, and a userspace process is polling /proc/fs/cifs/DebugData, the following oops occurs: BUG: kernel NULL pointer dereference, address: 0000000000000608 Oops: Oops: 0000 [#1] SMP PTI CPU: 0 UID: 0 PID: 3477 Comm: cat Not tainted 7.3.0-02212-g4df4d2a577db-dirty #6 PREEMPT(lazy) RIP: 0010:cifs_debug_data_proc_show+0x103/0xd80 [cifs] Call Trace: <TASK> ? seq_read_iter+0x86/0x400 ? seq_read_iter+0x86/0x400 seq_read_iter+0x183/0x400 proc_reg_read_iter+0x62/0x90 copy_splice_read+0x215/0x310 splice_file_to_pipe+0xec/0x1c0 do_splice+0x4b8/0x780 __se_sys_splice+0x18f/0x250 do_syscall_64+0x109/0x330 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Fix this by moving cifs_proc_init() down in init_cifs(), ensuring it runs only after all internal lists, locks, counters, and workqueues are fully ready, but before registering the filesystem. Update the error cleanup ladder and exit_cifs() to maintain proper LIFO order.	N/A	More Details
CVE-2026-99550	In the Linux kernel, the following vulnerability has been resolved: ocfs2: pin owning modules for open debugfs files The OCFS2 debugfs file operation tables do not set owner. The full debugfs proxy obtains real_fops with fops_get() at open and keeps it until final release. Without an owner, this does not acquire a reference to the module containing the table and its callbacks. Removing the debugfs entries does not drain retained open files. An open connected_nodes file can therefore outlive ocfs2_nodemanager and make full_proxy_release() read real_fops->release from freed module memory. Heming also reproduced the issue with fs_state and asked for the other OCFS2 debugfs operation tables to be checked. Set owner to THIS_MODULE in all eleven explicit debugfs operation tables: five in ocfs2_nodemanager, four in ocfs2_dlm, and two in ocfs2. The existing debugfs get/put pair then holds the appropriate module until final release. Leave the open, read and release callbacks unchanged. The blockcheck table already sets owner through DEFINE_DEBUGFS_ATTRIBUTE(), so it needs no change. The original controlled test used an unloadable OCFS2 module and case-local probe filters without changing the candidate open, unload or release logic. Closing a retained connected_nodes file after normal unload produced: BUG: unable to handle page fault for address: ffffbfff804cfef Oops: Oops: 0000 [#1] SMP KASAN NOPTI RIP: 0010:full_proxy_release+0xbf/0x190 Call Trace: <TASK> __fput+0x39a/0xa60 fput_close_sync+0xff/0x200 __x64_sys_close+0x8c/0xf0 do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Modules linked in: [last unloaded: ocfs2_nodemanager]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2: drain quorum work before removing the cluster The configured cluster must remain available until an in-flight quorum decision has finished using its fencing method. o2quo_make_decision() drops qs_lock before calling o2quo_fence_self(), which marks heartbeat regions for unclean stop and then reads o2nm_single_cluster->cl_fence_method. Cluster removal clears the singleton without waiting for this independent system-workqueue item. When a failed-quorum decision overlaps removal of the heartbeat regions, nodes and cluster through configfs, the worker can be		

CVE-2026-99551	preempted after marking the regions and resume after the singleton has been cleared: Quorum worker Configfs removal select fence; drop qs_lock o2hb_stop_all_regions() remove regions; join heartbeat threads remove nodes; stop o2net clear singleton; release cluster read cl_fence_method The resumed read dereferences NULL instead of applying the configured fencing method. The heartbeat thread's dependency protects the local node only until the thread exits, and draining o2net does not wait for qs_work. Add o2quo_flush_work() to drain pending or running quorum decisions before clearing the singleton in o2nm_cluster_group_drop_item(). By this point the heartbeat regions and nodes have been removed, so their work publishers are stopped. The worker can finish while the cluster and its configured fencing method remain available. It does not acquire the configfs mutex held by the removal path. Keep o2quo_exit() as the counterpart of o2quo_init() for module teardown and initialization failure, and have it call the same drain helper. Cluster removal uses only the explicitly named drain operation. KASAN report as below: Oops: general protection fault, probably for non-canonical address 0xdfffffc000000011b: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x00000000000008d8-0x00000000000008df] CPU: 0 UID: 0 PID: 72 Comm: kworker/0:2 Not tainted 7.3.0-rc4-next-20260921-pmb-ocfs2-functional-v1+ #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: events o2quo_make_decision RIP: 0010:o2quo_fence_self+0x65/0x200 Register dump omitted. Call Trace: <TASK> o2quo_make_decision+0x345/0x3a0 ? __pfx_o2quo_make_decision+0x10/0x10 ? process_one_work+0x834/0x1b40 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 ? srso_alias_return_thunk+0x5/0xfbef5 process_one_work+0x9a8/0x1b40 ? __pfx_process_one_work+0x10/0x10 ? lock_acquire+0x190/0x300 ? lock_is_held_type+0x8f/0x100 ? srso_alias_return_thunk+0x5/0xfbef5 worker_thread+0x65c/0xe40 ? srso_alias_return_thunk+0x5/0xfbef5 ? __kthread_parkme+0x177/0x220 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_worker_thread+0x10/0x10 kthread+0x351/0x460 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x659/0x940 ? __pfx_ret_from_fork+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __switch_to+0x74f/0xf70 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1a/0x30 </TASK> Modules linked in: ---[end trace 0000000000000000]---	N/A	More Details
CVE-2026-99552	In the Linux kernel, the following vulnerability has been resolved: ocfs2/dlm: drain join handlers before releasing module text The code for a registered DLM handler must remain resident until its last invocation finishes. dlm_unregister_net_handlers() only removes dlm_join_handlers from the o2net registry. o2net_handler_get() takes a kref on the wrapper, but that reference does not pin ocfs2_dlm text. o2net_process_message() calls nh_func after dropping o2net_handler_lock. With CONFIG_OCFS2_FS_O2CB=m, ocfs2_dlm can be loaded independently of ocfs2_stack_o2cb while nodemanager continues handling messages from a configured, heartbeating peer. If no dependent module or module reference prevents normal unload, the following ordering is possible: Receive worker Module unload ------------ ------------ o2net_handler_get() take wrapper kref drop handler lock <preempted> dlm_exit() unregister handlers release DLM module text call saved nh_func in released module text Use o2net_unregister_and_flush_handler_list() in the shared DLM unregister helper. It first prevents new lookups, then waits for the ordered o2net workqueue to drain, including a concurrent destruction of that queue. The acquired invocation therefore finishes before dlm_exit() returns and the loader releases the callback code. The same helper covers partial registration and dlm_init() error unwind before module memory is released. Skip the drain when the private join-handler list is empty. Early dlm_init() failures can reach this helper before any handler has been published. Partial-registration failure also reaches it a second time after the first call has already unregistered and drained the list. Neither case needs to take o2net_wq_mutex or drain unrelated work. The list is changed only by serialized module init/exit paths; a nonempty list still uses the full unregister-and-flush helper. Oops report as below: Oops: int3: 0000 [#1] SMP KASAN NOPTI CPU: 2 UID: 0 PID: 54 Comm: kworker/u16:3 Not tainted 7.3.0-rc4-next-20260921-pmb-ocfs2-functional-v1+ #1 PREEMPT(lazy) Hardware identification omitted. Workqueue: o2net o2net_rx_until_empty [ocfs2_nodemanager] RIP: 0010:0xffffffffc0423871 [Instruction byte dump omitted.] Register dump omitted. Call Trace: <TASK> ? o2net_rx_until_empty+0x19b1/0x3300 [ocfs2_nodemanager] ? __pfx_o2net_rx_until_empty+0x10/0x10 [ocfs2_nodemanager] ? lock_acquire+0x190/0x300 ? process_one_work+0x935/0x1b40 ? process_one_work+0x834/0x1b40 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 ? srso_alias_return_thunk+0x5/0xfbef5 ? process_one_work+0x9a8/0x1b40 ? __pfx_process_one_work+0x10/0x10 ? lock_acquire+0x190/0x300 ? lock_is_held_type+0x8f/0x100 ? srso_alias_return_thunk+0x5/0xfbef5 ? worker_thread+0x65c/0xe40 ? srso_alias_return_thunk+0x5/0xfbef5 ? __kthread_parkme+0x177/0x220 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_worker_thread+0x10/0x10 ? kthread+0x351/0x460 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_kthread+0x10/0x10 ? ret_from_fork+0x659/0x940 ? __pfx_ret_from_fork+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __switch_to+0x74f/0xf70 ? __pfx_kthread+0x10/0x10 ? ret_from_fork_asm+0x1a/0x30 </TASK> Modules linked in: ocfs2_nodemanager [last unloaded: ocfs2_dlm] ---[end trace 0000000000000000]---	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix race between client_info_show() and free_client() client_info_show() renders /proc/fs/nfsd/clients/<id>/info and walks clp->cl_sessions under clp->cl_lock to print each session's slot counts. free_client() tears down the same list without taking cl_lock, and is the only unlocked mutator of cl_sessions. A reader can observe a client mid-teardown because get_nfsdfs_clp() pins the nfs4_client but not its sessions: free_client() frees every session before calling nfsd_client_rmdir(), so an in-flight seq_file reader can follow a list_del()'d node whose ->next now holds LIST_POISON1 and take a general protection fault: Oops: general protection fault, probably for non-canonical		More

2026-99553	address 0xdead000000000014c CPU: 1 UID: 0 PID: 132488 Comm: cat RIP: 0010:client_info_show+0x2bf/0x3d0 RAX: dead000000000100 Call Trace: seq_read_iter+0x12a/0x4b0 seq_read+0xf1/0x130 vfs_read+0xbf/0x350 ksys_read+0x6f/0xf0 do_syscall_64+0x8b/0xcb0 entry_SYSCALL_64_after_hwframe+0x76/0x7e Kernel panic - not syncing: Fatal exception Detach cl_sessions onto a local reaplist under cl_lock, then free the sessions after dropping the lock. Removing entries from cl_sessions under cl_lock matches unhash_session(), and the detach-then-reap shape matches how __destroy_client() reaps cl_delegations. The sessions cannot be freed while cl_lock is held, since free_session() calls nfsd4_del_conns(), which re-acquires it.	N/A	Details
CVE-2026-99554	In the Linux kernel, the following vulnerability has been resolved: net: usb: pegasus: do not advance rx_skb->data on dropped frames For the 0x8513 chip, read_bulk_callback() skips the 2-byte receive header by advancing pegasus->rx_skb->data in place. This is done before the pkt_len > PEGASUS_MTU check, and when that check drops the frame the good path resubmits the same skb with the advanced rx_skb->data as the transfer buffer, still PEGASUS_MTU bytes long. Every frame with an oversized length field therefore moves the receive buffer two more bytes towards the end of the skb data area. A device that keeps sending such frames makes the host controller write received data over skb_shared_info and past the end of the allocation, and the corrupted skb_shared_info is used when the skb is eventually freed. With an emulated device sending frames filled with 0x5a on a KASAN kernel: BUG: KASAN: slab-out-of-bounds in dummy_timer+0xade/0x3f00 [dummy_hcd] Write of size 1536 at addr ffff888105821202 by task swapper/1/0 ... Oops: general protection fault, probably for non-canonical address 0xeb4b474b4b4b4b4c: 0000 [#1] PREEMPT SMP KASAN NOPTI KASAN: maybe wild-memory-access in range [0x5a5a5a5a5a5a5a60-0x5a5a5a5a5a5a5a67] RIP: 0010:skb_release_data+0x36f/0x800 Only skip the header with skb_reserve() once the frame is going to be passed up, so the receive buffer always starts at the beginning of the skb data area. Unlike the old pointer adjustment, skb_reserve() also moves skb->tail, so the two header bytes now count against the buffer size: lower the length limit for the 0x8513 accordingly, otherwise a pkt_len of PEGASUS_MTU - 1 or PEGASUS_MTU would trip skb_over_panic(). Found by code review with LLM assistance and reproduced with an emulated USB device (raw-gadget + dummy_hcd).	N/A	More Details
CVE-2026-99555	In the Linux kernel, the following vulnerability has been resolved: media: usbtv: fix null-pointer dereference on disconnect A null-pointer dereference can occur in usb_make_path() when called from usbtv_querycap() (and other ioctl handlers) during device disconnection. Oops: general protection fault, probably for non-canonical address 0xdffffc000000000a: 0000 [#1] SMP KASAN NOPTI KASAN: null-ptr-deref in range [0x0000000000000050-0x0000000000000057] RIP: 0010:usb_make_path include/linux/usb.h:985 [inline] RIP: 0010:usbtv_querycap+0x9c/0x100 drivers/media/usb/usbtv/usbtv-video.c:612 ... Call Trace: <TASK> v4l_querycap+0x236/0x470 drivers/media/v4l2-core/v4l2-ioctl.c:1106 __video_do_ioctl+0x8af/0xc70 drivers/media/v4l2-core/v4l2-ioctl.c:3133 video_usercopy+0x860/0x1430 drivers/media/v4l2-core/v4l2-ioctl.c:3475 v4l2_ioctl+0x18d/0x1e0 drivers/media/v4l2-core/v4l2-dev.c:366 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:597 [inline] __se_sys_ioctl+0xfc/0x170 fs/ioctl.c:583 do_syscall_x64 arch/x86/entry/syscall_64.c:63 [inline] do_syscall_64+0x174/0x580 arch/x86/entry/syscall_64.c:94 This happens due to a race condition between the USB device disconnect routine (usbtv_disconnect()) and V4L2 ioctls. When the device is disconnected, usbtv_disconnect() sets usbtv->udev = NULL without holding the usbtv->v4l2_lock mutex. Concurrently, an ioctl handler like usbtv_querycap() can be executing under the v4l2_lock and attempt to use usbtv->udev, leading to a crash. To fix this, move the usbtv->udev = NULL assignment into usbtv_video_free() and protect it with mutex_lock(&usbtv->v4l2_lock). This ensures that no non-queue ioctl handler can run concurrently with the pointer being nullified. If an ioctl is already running, the mutex will block the disconnect path until the ioctl finishes. If an ioctl is blocked waiting for the lock, it will acquire it after the disconnect routine releases it, but will then check video_is_registered() (which was cleared by vb2_video_unregister_device()) and safely return -ENODEV. Acquiring v4l2_lock after vb2_video_unregister_device() avoids reintroducing a historical deadlock issue, and usbtv_audio_free() remains safe as it is called before usbtv_video_free().	N/A	More Details
CVE-2026-99557	In the Linux kernel, the following vulnerability has been resolved: ceph: fail readdir when the MDS reply carries no dirfrag A ceph mount oopses in ceph_readdir() when the MDS answers a READDIR request with a reply whose extra blob is empty. parse_reply_info() skips reply blobs whose length is zero, so parse_reply_info_readdir() never runs, r_reply_info.dir_dir stays NULL, and the reply still parses as successful. ceph_mdsc_do_request() returns 0 and ceph_readdir() then reads rinfo->dir_dir->frag with no NULL check: /* extra */ ceph_decode_32_safe(&p, end, len, bad); if (len > 0) { ... err = parse_reply_info_extra(&p, p + len, req, features, s); ... } The oops on v7.3.0-rc4-00537-ga3ff15db6820, one getdents64(2) on the mount root after an ms_mode=legacy mount, trimmed: BUG: KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] Oops: general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASAN NOPTI CPU: 1 UID: 0 PID: 266 Comm: ls Tainted: G N 7.3.0-rc4-00537-ga3ff15db6820 #1 PREEMPT(lazy) RIP: 0010:ceph_readdir+0xbff/0x3ea0 Call Trace: <TASK> ? __pfx_ceph_readdir+0x10/0x10 wrap_directory_iterator+0xa2/0xe0 iterate_dir+0x1da/0x5d0 __x64_sys_getdents64+0x13b/0x2a0 do_syscall_64+0xdd/0x4a0 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> The unconditional frag read came in with 81c6aea5275ea ("ceph: handle frag mismatch between readdir request and reply"), so the bug is old. Treat a reply with no dirfrag as a protocol violation: warn once, drop the request and return -EIO, the same way ceph_readdir() already handles the two malformed-reply cases later in the function. Tested on v7.3.0-rc4-00537-ga3ff15db6820 (KASAN). Unpatched, the getdents64(2) above takes the machine down (oops=panic) on the first READDIR reply. With the patch it warns once and getdents64(2) returns -EIO, no oops, no hang. A	N/A	More Details

	conformant reply for an empty directory — dirfrag present, zero entries, FRAG_END|FRAG_COMPLETE — still lists "." and ".." and returns EOF on both kernels, so normal readdir behavior is unchanged. The oops is still present in v7.3-rc5 (72d3fcf802c45).		
CVE-2026-99558	In the Linux kernel, the following vulnerability has been resolved: blk-crypto: fix NULL deref and UAF in __blk_crypto_evict_key() Sashiko reported a possible NULL pointer dereference, triggered by the following flow: dm-inlinecrypt ctor # first use of the fallback ├─ blk_crypto_init_key() # succeeds, ctx->key.size != 0 ├─ blk_crypto_start_using_key() │ ├─ hw-wrapped key: -EOPNOTSUPP │ │ # profile never allocated (NULL) │ └─ raw key: fallback_start_using_mode() │ └─ blk_crypto_fallback_init() fails │ └─ kfree(profile) # not NULLed ├─ goto bad └─ inlinecrypt_dtr() # evicts whenever key.size is set └─ blk_crypto_evict_key() └─ blk_crypto_config_supported_natively() = false └─ blk_crypto_fallback_evict_key() └─ __blk_crypto_evict_key(blk_crypto_fallback_profile) └─ profile->num_slots # BOOM!!! This issue can be reproduced using the following commands: dd if=/dev/zero of=/tmp/backing bs=1M count=64 LOOP=$(losetup -f --show /tmp/backing) SZ=$(blockdev --getsz $LOOP) KEY=$(head -c 64 /dev/urandom | od -An -tx1 | tr -d ' \n') dmsetup create ic --table "0 $SZ inlinecrypt aes-xts-plain64 \ $KEY 0 $LOOP 0 1 keytype:hw-wrapped" And the following oops is triggered: == blk-crypto: loop0: no support for wrapped keys Oops: general protection fault, probably for non-canonical address 0xdffffc000000000b: 0000 [#1] SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000058-0x000000000000005f] CPU: 5 UID: 0 PID: 901 Comm: dmsetup 7.3.0-rc3 #1 PREEMPT(lazy) RIP: 0010:__blk_crypto_evict_key+0x2e/0x3b0 Call Trace: blk_crypto_evict_key+0x24c/0x320 inlinecrypt_dtr+0xf6/0x160 inlinecrypt_ctr+0x2d3/0xbd0 dm_table_add_target+0x661/0xf40 table_load+0x1c4/0x490 ctl_ioctl+0x32f/0x5f0 dm_ctl_ioctl+0xe/0x20 == dm-inlinecrypt is the only caller that can evict a key that was never started, as fscrypt frees such a key without evicting it. Fix this by returning early from blk_crypto_fallback_evict_key() when tfms_inited[] is not set for the key's mode. tfms_inited[] is only set after blk_crypto_fallback_init() has successfully initialized the profile, so the crash above is no longer reachable. The bug has been latent since commit c984ff1423ae ("blk-crypto: dynamically allocate fallback profile") was merged in v6.5, and only became actually triggerable when commit e7f57d2c47e2 ("dm-inlinecrypt: add target for inline block device encryption") was merged in v7.2.	N/A	More Details
CVE-2026-99559	In the Linux kernel, the following vulnerability has been resolved: ARM: 9524/1: mm: align hugetlb mappings to the huge page size mmap(MAP_HUGETLB) on 32-bit Arm can return an address which is not aligned to the huge page size. In a QEMU virt guest, unmapping the VMA then hits the alignment check in __unmap_hugepage_range(): kernel BUG at mm/hugetlb.c:5213! Internal error: Oops - BUG: 0 [#1] SMP ARM PC is at __unmap_hugepage_range+0x540/0x584 LR is at __zap_vma_range+0x13a0/0x1410 Commit 7bd3f1e1a9ae ("mm: make hugetlb mappings go through mm_get_unmapped_area_vmflags") made hugetlb mappings use the architecture's arch_get_unmapped_area(). The generic implementations handle hugetlb alignment, but the Arm implementations only account for the optional SHMLBA cache-colouring constraint. A 4 KiB-page kernel can therefore align the mapping to 4 KiB or 16 KiB instead of the required 2 MiB. Set the alignment mask from the file's hstate for hugetlb mappings in both Arm implementations. Huge-page alignment also satisfies the SHMLBA constraint.	N/A	More Details
CVE-2026-99560	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7925: fix reset panic caused by the WM2 rx irq mask mt7925_irq_map sets wm2_complete_mask, but mt7925_dma_layout has no tx done ring, so the MT_RXQ_MCU_WA queue and its NAPI are never set up. On a chip reset, mt7925e_mac_reset() calls napi_disable() on that NAPI and crashes: BUG: kernel NULL pointer dereference, address: 0000000000000a40 Workqueue: mt76 mt7925_mac_reset_work [mt7925_common] RIP: 0010:mutex_lock+0x1c/0x50 Call Trace: napi_disable+0x1d/0x50 mt7925e_mac_reset+0x327/0x420 [mt7925e] mt7925_mac_reset_work+0x9a/0x1a0 [mt7925_common] Remove the WM2 mask from mt7925_irq_map and from MT7925_INT_RX_DONE_ALL.	N/A	More Details
CVE-2026-129398	Joomla Extension - vi-solutions - Authorization bypass in saveAjaxTextarea task in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The saveAjaxTextarea task did not check edit permissions and accepted arbitrary table and column names, allowing backend users to modify the configuration of any form, including notification recipients.	N/A	More Details
CVE-2026-129401	Joomla Extension - vi-solutions - Open redirect on invalid form submission in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The return parameter of an invalid form submission was not validated, allowing redirects to external sites.	N/A	More Details
CVE-2026-99405	In the Linux kernel, the following vulnerability has been resolved: xfs: abort dirtree repair if the live update hook dies LOLLM notices that if the live update hook for the directory tree repair code hits an error, it will set dl->aborted and stops working. If the repair code ever sees that condition, it ought to bail out immediately instead of continuing with obsolete data.	N/A	More Details
CVE-2026-86094	WatchGuard Dimension records unredacted session identifiers for logged-in users in its web UI diagnostic log. A low-privileged Dimension Administrator can retrieve this log and extract a Super Administrator's session token while that administrator is logged in, enabling account takeover. This vulnerability is a secondary unmitigated attack path for CVE-2026-78174.	N/A	More Details

CVE-2026-103580	An OS command injection vulnerability in WatchGuard Dimension's package update functionality allows an authenticated Dimension Administrator to execute arbitrary operating system commands as root by submitting a specially crafted package update request.	N/A	More Details
CVE-2026-118900	An OS command injection vulnerability exists in the Android application "Imperial Palaces Guide" provided by The Imperial Household Agency. An attacker may be able to execute arbitrary OS commands with the privileges of the application.	N/A	More Details
CVE-2026-126037	A broken access control vulnerability in WatchGuard Dimension's Device Summary feature allows an authenticated, low-privileged user to retrieve a device's detailed summary data by bypassing restrictions that would otherwise hide this information.	N/A	More Details
CVE-2026-117657	A Cleartext Transmission of Sensitive Information (CWE-319) vulnerability exists in the Android application "Imperial Palaces Guide" provided by The Imperial Household Agency. A man-in-the-middle attacker may be able to tamper with data transmitted between the application and the server, resulting in the display of modified content to users.	N/A	More Details
CVE-2026-129418	Joomla Extension - phoca.cz - Stored XSS through YouTube links in Phoca Gallery extension 4.2.0-6.1.6 - The YouTube upload feature did not validate the supplied video URL and embedded the derived video ID without escaping, allowing users with frontend upload permission to inject HTML attributes into the gallery detail view.	N/A	More Details
CVE-2026-129419	Joomla Extension - akeeba.com - Account takeover through missing OAuth state validation in Akeeba SocialLogin extension 1.0.0-4.11.0 - The OAuth callback did not validate a session-bound state parameter, allowing an attacker to link their own social provider identity to the account of a logged-in user who opens a crafted callback link and subsequently log in as that user.	N/A	More Details
CVE-2026-129422	Joomla Extension - joomlatools.com - Unauthenticated blind server-side request forgery in FILEman extension < 6.1.4 - The frontend file view did not restrict requested files to their container and passed URLs to the framework URL fetch, allowing unauthenticated attackers to make the server issue requests to arbitrary addresses, including internal network resources.	N/A	More Details
CVE-2026-81430	A vulnerability in WatchGuard Dimension's system upgrade process allows an authenticated administrator to create or overwrite arbitrary files on the appliance filesystem by uploading a specially crafted upgrade archive.	N/A	More Details
CVE-2026-81434	A reflected Cross-Site Scripting (XSS) vulnerability in WatchGuard Dimension's Device Configuration History page allows an authenticated attacker with view permission on a device to execute arbitrary JavaScript in another authenticated user's browser by enticing the victim to open a specially crafted link.	N/A	More Details
CVE-2026-84322	WatchGuard Dimension returns unredacted session identifiers for logged-in users in a management API response. A low-privileged Dimension Administrator can retrieve the active session IDs of other administrators, including Super Administrators, while they are logged in, enabling account takeover.	N/A	More Details
CVE-2026-86126	A Cross-Site Request Forgery (CSRF) vulnerability in WatchGuard Dimension's user management interface allows a remote, unauthenticated attacker to create a persistent local Super Administrator account by tricking an authenticated Super Administrator into visiting a malicious web page.	N/A	More Details
CVE-2026-129402	Joomla Extension - vi-solutions - Unauthenticated SQL injection in reload task in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Input placeholders in SQL field statements were inserted without quoting in the public reload task, allowing unauthenticated SQL injection in read queries.	N/A	More Details
CVE-2026-86130	A missing authentication vulnerability in WatchGuard Dimension's log collection service allows a remote, unauthenticated attacker to execute arbitrary SQL statements and ultimately execute arbitrary code as the Dimension WebUI process user by sending specially crafted network requests.	N/A	More Details
CVE-2026-93911	A SQL injection vulnerability in WatchGuard Dimension's log collection service allows an authenticated managed device, or an attacker in possession of a valid device logging key, to achieve arbitrary code execution on the Dimension appliance.	N/A	More Details
CVE-2026-93912	A code injection vulnerability in WatchGuard Dimension's log notification processing allows an authenticated managed device, or an attacker in possession of a valid device logging key, to achieve arbitrary code execution on the Dimension appliance.	N/A	More Details
CVE-2026-93914	An improper input validation vulnerability in WatchGuard Dimension's database snapshot restore feature allows an authenticated Dimension Administrator to execute arbitrary operating system commands as root by uploading a specially crafted database snapshot archive.	N/A	More Details
CVE-2026-93915	A missing authorization vulnerability in WatchGuard Dimension's legacy system settings configuration handler allows a low-privileged authenticated user to execute arbitrary commands with elevated privileges on the underlying OS.	N/A	More Details
CVE-	An insufficient session expiration vulnerability in WatchGuard Dimension's administrator session		

2026-93916	management allows an attacker holding a valid Super Administrator WebUI session to retain full Super Administrator privileges after that account has been downgraded or deleted by another administrator.	N/A	More Details
CVE-2026-93917	An improper authorization vulnerability in WatchGuard Dimension's device management registration process allows an attacker who controls one enrolled managed device's credentials to take over another managed device's identity. This allows the attacker to receive the victim device's queued configuration and apply attacker-supplied configurations under the victim device's identity.	N/A	More Details
CVE-2026-103025	bep/imagemeta is a Go library for reading EXIF, IPTC and XMP image meta data from JPEG, TIFF, PNG, and WebP files. Prior to 0.17.3, imagemeta.Decode could exhaust memory when parsing PNG zTXt chunks whose keyword is "Raw profile type iptc": the declared chunk length can drive an allocation far larger than the input file, and the embedded zlib stream is decompressed without an output limit. The advisory reports a 42-byte input causing a 256 MB allocation and a 127 KB input inflating to 460 MB. The oversized-allocation path's memory effect depends on platform overcommit and configured GOMEMLIMIT or cgroup limits, while the decompression path reliably creates resident memory. Applications that decode untrusted PNGs with IPTC metadata processing can therefore suffer denial of service through memory exhaustion. This issue is fixed in 0.17.3.	N/A	More Details
CVE-2026-104199	aMQTT is an open source MQTT broker and client. From 0.11.0 until 0.11.4, add_subscription() in amqtt/broker.py accepts malformed single-level + wildcard placement because it checks substrings instead of validating each slash-delimited level. The _matches() function converts each + into a greedy regular-expression fragment and evaluates it with re.fullmatch(), allowing a filter containing many adjacent + characters to trigger catastrophic backtracking when compared with a topic that partially matches but ultimately fails. Because matching runs synchronously in _run_broadcast(), one subscription followed by a publish can block the broker event loop, stopping message delivery and new connection processing until restart. Anonymous clients can trigger the issue with the default configuration. Any authenticated client permitted to create arbitrary subscriptions and publish messages may still exploit it when authentication is enabled. This issue is fixed in version 0.11.4.	N/A	More Details
CVE-2026-104885	OpenRemote is an open-source internet-of-things platform. Prior to 1.31.0, an authenticated user who obtains another user's MQTT session ID can disconnect that session across realm boundaries through the GET route GET {realm}/disconnect/{sessionID}. In manager/src/main/java/org/openremote/manager/security/UserResourceImpl.java, disconnectUserSession ignores the requested realm and calls mqttBrokerService.disconnectSession(sessionID) without checking session ownership; model/src/main/java/org/openremote/model/security/UserResource.java exposes the route without an @RolesAllowed restriction. The earlier cross-realm read-method guards did not cover this disconnect method, and same-realm users can also disconnect sessions they do not own. Session IDs are not trivially brute-forceable but may be exposed through monitoring, logs, or correlated API calls, and the unauthorized disconnect can disrupt MQTT connectivity across tenant boundaries without directly disclosing session data. This issue is fixed in 1.31.0.	N/A	More Details
CVE-2026-103579	A CRLF injection vulnerability in WatchGuard Dimension's outgoing email server configuration allows an authenticated Dimension Administrator to execute arbitrary operating system commands on the Dimension appliance with the privileges of the local mail-handling account.	N/A	More Details
CVE-2026-129415	Joomla Extension - vi-solutions - Missing core.create check in example data generation and BreezingForms import in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Example data generation and the BreezingForms import did not check the core.create permission for the affected forms, allowing backend users to create records without permission.	N/A	More Details
CVE-2026-129414	Joomla Extension - vi-solutions - Improper per-form authorization for field management in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Field creation, editing and ordering checked permissions against a user-supplied form ID instead of the form the field belongs to, allowing backend users to manage fields of other forms.	N/A	More Details
CVE-2026-129412	Joomla Extension - vi-solutions - SQL injection in backend data search and export in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Search filters and form IDs in the backend data view and export were not properly sanitized, allowing SQL injection by backend users.	N/A	More Details
CVE-2026-129404	Joomla Extension - vi-solutions - Missing SQL statement permission check in form import and field creation in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - The core.create.sql.statement permission was not enforced when importing forms or creating new fields, allowing backend users without this permission to store arbitrary SQL statements.	N/A	More Details
CVE-2026-129409	Joomla Extension - vi-solutions - Open redirects in data view controller and after form submission in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Return and redirect URLs taken from user input in the data/PDF controller and after a successful form submission were not validated, allowing redirects to external sites.	N/A	More Details
CVE-2026-129410	Joomla Extension - vi-solutions - Code injection via calculation formulas in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Calculation field formulas were executed using PHP eval() without validating their content, allowing backend users with field edit permissions to execute arbitrary PHP code.	N/A	More Details
CVE-	Joomla Extension - vi-solutions - Unauthenticated SQL injection in frontend data view filters in Visforms		

2026-129411	extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Sorting, location radius and published filters of the frontend data view were not properly sanitized, allowing SQL injection, partly without authentication.	N/A	More Details
CVE-2026-129413	Joomla Extension - vi-solutions - IDOR in backend data export in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - The backend export did not correctly verify permissions for the exported form, allowing backend users to export data of forms they are not authorized to access.	N/A	More Details
CVE-2026-129416	Joomla Extension - vi-solutions - Access restriction bypass for form submissions in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Form submissions did not check the access level of the form, allowing guests to submit forms restricted to specific user groups.	N/A	More Details
CVE-2026-129417	Joomla Extension - vi-solutions - Path traversal in file upload in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - File names of uploads were not reduced to their base name when the option to retain original file names was enabled, potentially allowing files to be written outside the upload directory.	N/A	More Details
CVE-2026-133276	When extra data is allowed (extra=Extra.allow), providing a __pydantic_initialised__ key when validating a Pydantic V1 dataclass makes it considered as already validated, resulting in field values not being validated. The __init__() method of Pydantic V1 dataclasses stores all the provided keyword arguments in the instance __dict__ before validation, so a __pydantic_initialised__ keyword argument overrides the internal flag used to skip validation of already validated instances. As a result, data returned by validation doesn't match the declared type. Affected are applications validating untrusted data into Pydantic V1 dataclasses allowing extra data. The issue was fixed in Pydantic 1.10.27. Users relying on the pydantic.v1 namespace in Pydantic V2 should upgrade to 2.14.2 or later.	N/A	More Details
CVE-2026-47278	Frappe is a full-stack web application framework. Prior to 15.108.0 and 16.18.3, the add_comment handler in frappe/templates/includes/comments/comments.py trusts the client-provided comment_by and comment_email fields. A caller can submit chosen identity metadata that is stored with a comment, causing the comment to appear to originate from another name or email address. This issue is fixed in versions 15.108.0 and 16.18.3.	N/A	More Details
CVE-2026-69260	Git LFS is a Git extension for versioning large files. From 0.5.0 until 3.8.1, Git LFS does not format the SSH URL path component as a single argument when constructing remote object-transfer or authentication commands. Shell metacharacters in the SSH URL can therefore escape the intended path argument and cause arbitrary command execution when the remote Git LFS service passes received commands to a shell without adequate sanitization. Git LFS also incorrectly encodes default HTTPS URLs when an SSH URL is configured and the remote host either does not support Git LFS remote commands over SSH or supports HTTP authentication over SSH but git-lfs-authenticate does not return an explicit HTTP URL. This issue is fixed in version 3.8.1.	N/A	More Details
CVE-2026-69261	Git LFS is a Git extension for versioning large files. From 0.5.0 until 3.8.1, Git LFS may forward a Cookie header after an HTTP redirect to a host other than the one for which the header was intended. That unintended host can therefore receive cookies containing credentials or other confidential values. Authorization headers are never forwarded to inappropriate hosts and are not the affected secret in this vulnerability. This issue is fixed in version 3.8.1.	N/A	More Details
CVE-2026-77318	Jupyter MCP Server is a Model Context Protocol server for Jupyter. Prior to 1.0.3, standalone streamable-http mode can enforce MCP_TOKEN on /mcp while leaving the state-changing /api/connect and /api/stop management routes unauthenticated. An unauthenticated network caller can use /api/connect to replace runtime_url and document_url values and their tokens, reset ServerContext, and start a kernel client against attacker-controlled infrastructure, causing the server to send Jupyter API and WebSocket traffic to that backend. The caller can also use /api/stop to stop the current notebook kernel. This boundary failure does not require explicit insecure no-auth mode and occurs outside normal MCP JSON-RPC tool dispatch. This issue is fixed in version 1.0.3.	N/A	More Details
CVE-2026-77338	Sprout Track is a self-hosted baby tracking application. Prior to 1.3.5, Sprout Track uses a publicly known hardcoded signing secret when JWT_SECRET is unset, and standard setup does not generate a unique secret. On a network-reachable affected deployment, an unauthenticated attacker can sign JWTs with arbitrary claims, including system-administrator claims, to access administrative functionality and read or modify family, child activity, account, and email data. Deployments with a strong unique JWT_SECRET are not affected. Managed and hosted instances were reported remediated. This issue is fixed in version 1.3.5.	N/A	More Details
CVE-2026-77359	Jupyter MCP Server is a Model Context Protocol server for Jupyter. Prior to 1.0.3, standalone streamable-http mode validates Host and Origin values for /mcp but does not apply the same browser-origin controls to the state-changing /api/connect and /api/stop management routes. A malicious website can use DNS rebinding so that its hostname resolves to a loopback, LAN, VPN, or published-container MCP listener, then send a browser-shaped request that /api/connect accepts to replace runtime_url and document_url values and start a kernel client against attacker-controlled infrastructure. The attacker can thereby cause Jupyter API and WebSocket traffic to reach that infrastructure and can use /api/stop to stop the current notebook kernel. The default broad bind can increase exposure, but the vulnerability is the missing management-route origin boundary, not a lack of Host validation on /mcp. This issue is fixed in version 1.0.3.	N/A	More Details
	Improper certificate validation vulnerability in Estonian Information System Authority (RIA) libdigidocpp,		

CVE-2026-128978	Estonian Information System Authority (RIA) DigiDoc4, and Estonian Information System Authority (RIA) DigiDoc on Android and iOS. This issue affects libdigidocpp: from 3.13.0 before 4.5.2; DigiDoc4: from 4.0.0 before 4.11.2; DigiDoc: from 3.0.0 before 3.2.0.	N/A	More Details
CVE-2026-129399	Joomla Extension - vi-solutions - Missing core.create check in form creation in Visforms extension 1.0.0-5.6.3, 6.0.0-6.2.0 - The createForm task of the form creator did not check the core.create permission, allowing backend users without create permission to create forms.	N/A	More Details
CVE-2026-129403	Joomla Extension - vi-solutions - Missing access check in reload task in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - The public reload task did not check form and field access levels, allowing unauthenticated users to retrieve results of SQL field statements of restricted forms.	N/A	More Details
CVE-2026-129405	Joomla Extension - vi-solutions - Predictable and reusable email verification codes in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Email verification codes were derived from the email address, did not expire after use and could be guessed without attempt limits, allowing bypass of the email verification.	N/A	More Details
CVE-2026-129406	Joomla Extension - vi-solutions - Verification mail rate limit reset in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - The rate limit for verification mails was reset on every form render, rendering the limit ineffective.	N/A	More Details
CVE-2026-129407	Joomla Extension - vi-solutions - Missing authorization and CSRF checks in reset and check-in tasks in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - The tasks to reset submitted data and to check in forms lacked authorization checks and CSRF protection.	N/A	More Details
CVE-2026-129408	Joomla Extension - vi-solutions - Stored XSS via SQL field labels in Visforms extension 1.0.0-5.6.4, 6.0.0-6.2.1 - Labels and values of SQL field options were returned unescaped by the reload task and inserted into the page as HTML, allowing stored XSS if the statement reads user-supplied data.	N/A	More Details
CVE-2026-99532	In the Linux kernel, the following vulnerability has been resolved: ksmbd: validate security descriptor offsets before pointer arithmetic parse_sec_desc() constructs SID and DACL pointers by adding untrusted 32-bit offsets to the security descriptor before checking their bounds. The existing checks reject offsets into the descriptor header, but do not reject offsets beyond the descriptor's length. On a 32-bit kernel, a large offset can wrap the pointer before the descriptor. The end-only checks in parse_sid() and parse_dacl() can then accept it, allowing an authenticated SMB2 SET_INFO request to access memory outside the security descriptor and fault the kernel. Check each nonzero offset against acl_len before constructing its pointer. Initialize the optional owner and group SID pointers to NULL so an omitted SID remains absent when parsing the DACL.	N/A	More Details
CVE-2026-99531	In the Linux kernel, the following vulnerability has been resolved: ksmbd: verify transform SessionId matches the decrypted header When an encrypted request comes in, the decryption key is picked by the SessionId in the transform header, but the request is then authorized under whatever session the decrypted inner header names. Nothing ever compares the two, so on a connection with more than one session a client can get a request decrypted with one session's key and run it as another session. Encrypted requests are exempt from signing too, so a valid AEAD tag is the only proof of session identity, and it is checked against the wrong session. MS-SMB2 says the server must check that the transform SessionId matches the one in the decrypted SMB2 header and treat a mismatch as a protocol error. Add that check in __handle_ksmbd_work() after decryption, before dispatching: the message has to be at least one fixed SMB2 header, the first operation must not set SMB2_FLAGS_RELATED_OPERATIONS and its SessionId must match the transform SessionId, and every following non-related operation in a compound chain must carry the same SessionId. NextCommand offsets must be 8-byte aligned and stay inside the message; the bounds check is done by subtraction so it cannot wrap. When the decrypted payload is a compression transform, the existing decompression step runs first and the decompressed SMB2 message is checked instead, since a compression transform header carries no SessionId.	N/A	More Details
CVE-2026-99530	In the Linux kernel, the following vulnerability has been resolved: cgroup: serialize migrations of newly forked processes With favordynmods enabled, a cgroup.procs migration takes the target threadgroup's cgroup_threadgroup_rwsem for writing. A process fork holds the forking threadgroup's semaphore for reading. These are different semaphores when fork creates a new process. copy_process() publishes the child's PID before cgroup_post_fork() links the child to its css_set. A writer can find the child in that interval, take its uncontended semaphore, and migrate it while child->cgroups still points to init_css_set and child->cg_list is empty. Migration links the node into init_css_set.mg_tasks. cgroup_post_fork() then adds the same node to the intended css_set as though it were unlinked, corrupting both task lists and their css_set ownership. This regresses the fork/migration race fixed by commit b12b533fa523 ("cgroups: add lock for child->cgroups in cgroup_post_fork()") and restored again in 2012 after the lock had been removed. The old task_lock protection was later superseded by global css_set and cgroup threadgroup locking. The change identified by the Fixes tag narrowed only the migration side back to a per-threadgroup lock, recreating the different-process exclusion failure. Take the new threadgroup's semaphore for reading before the child becomes visible and hold it through cgroup_post_fork(). Use a nested lock subclass because fork already holds the parent's semaphore. CLONE_THREAD children share the parent's signal_struct and remain covered by the existing lock. Migrations of unrelated, fully initialized threadgroups remain concurrent. An AI coding	N/A	More Details

	assistant was used to identify and reproduce the regression and to develop and test this fix.		
CVE-2026-99446	In the Linux kernel, the following vulnerability has been resolved: phy: core: Fix race-condition between _of_phy_get() and try_module_get() Sashiko asked during a patch review if the existing usage pattern had a race condition; specifically in of_phy_get() if it was possible between returning from _of_phy_get() and running try_module_get() that a module might be unbound leading to use-after-free. Looking at the code this appears to be so, there is no linkage between the phy and module under a synchronisation primitive. Using the phy_provider_mutex in phy_get() will ensure there is a link between the returned phy pointer and the module_get() bumping the module reference count. Amend phy_get(), of_phy_get() and devm_of_phy_get_by_index() to fix the same usage pattern. phy_provider_unregister() must take the phy_provider_mutex so amending phy_get()/of_phy_get() to take that same mutex guarantees there is no use-after-free.	N/A	More Details
CVE-2026-99436	In the Linux kernel, the following vulnerability has been resolved: media: rc: After rc_unregister_device() timers can be re-armed If rc_keydown() is called after rc_register_device(), then the timers can be re-armed. Delete them again for a final time.	N/A	More Details
CVE-2026-99437	In the Linux kernel, the following vulnerability has been resolved: media: rc: Use after free in ir_raw_event_handle() If rc_unregister_device() is called while IR is being processed, then ir_raw_event_handle() could call wake_up_process(dev->raw->thread) after kthread_stop(dev->raw->thread). Avoid use after free by using get_task_struct().	N/A	More Details
CVE-2026-99438	In the Linux kernel, the following vulnerability has been resolved: media: rc: mce_kbd: Fix inconsistent locking of keylock data->keylock is taken with spin_lock() in ir_mce_kbd_decode(), which runs in the raw IR decode kthread with local interrupts enabled, but with spin_lock_irqsave() in mce_kbd_rx_timeout(), the rx_timeout timer callback, which runs in softirq context. If the decode kthread is holding keylock via spin_lock() when a timer interrupt fires on the same CPU and the softirq runs mce_kbd_rx_timeout(), the softirq spins forever waiting for a lock that only the now-preempted kthread can release, while the kthread cannot run again until the softirq gives up the CPU. This deadlocks the CPU.	N/A	More Details
CVE-2026-99439	In the Linux kernel, the following vulnerability has been resolved: media: rc: Ensure registered is cleared in error path If rc_register_device() fails, ensure that registered is not set to true. If lirc_register() succeeded, then userspace could have an open file descriptor open. This leads to a use-after-free. lirc_unregister() must be called with registered = false, else there is a possibility that an fd is open and polling, which would never return POLLERR unless registered is false. Note that registered must be changed under dev->lock, many functions read registered under lock (e.g. lirc_ioctl).	N/A	More Details
CVE-2026-99440	In the Linux kernel, the following vulnerability has been resolved: media: rc: Ensure that rc_unregister_device() does not free input device During or after rc_unregister_device(), IR may still be reported which results in a input event being reported. This could result in a null pointer deref in rc_keydown() or a use-after-free of the input device if the pointer was read before it is set to NULL.	N/A	More Details
CVE-2026-99318	In the Linux kernel, the following vulnerability has been resolved: ALSA: bcd2000: Fix use-after-free of MIDI URBs on file close bcd2000_disconnect() calls snd_card_disconnect(), which only prevents new opens, and then frees the MIDI URBs. A rawmidi file that was already open can still be closed afterwards, and its close callback, bcd2000_midi_output_close(), calls usb_kill_urb() on midi_out_urb. The URB pointers are cleared on disconnect and their accesses are serialized with midi_lock, but bcd2000_midi_output_close() reads midi_out_urb without that lock. The check can therefore pass just before the pointer is cleared, and usb_kill_urb() then runs on the URB that has already been freed: BUG: KASAN: slab-use-after-free in usb_kill_urb+0x74/0x80 Read of size 8 at addr ffff8880078e7d40 usb_kill_urb bcd2000_midi_output_close close_substream rawmidi_release_priv snd_rawmidi_release Allocated by: usb_alloc_urb <- bcd2000_init_midi <- bcd2000_probe Freed by: usb_free_urb <- bcd2000_midi_free <- bcd2000_disconnect Take the lock in bcd2000_midi_output_close() and reference the URB with usb_get_urb() before releasing it. The URB then stays valid until usb_put_urb() after usb_kill_urb() has returned, even if a disconnect frees it in the meantime. usb_kill_urb() itself cannot run under midi_lock, as the completion handler takes the same lock.	N/A	More Details
CVE-2026-99442	In the Linux kernel, the following vulnerability has been resolved: nvmem: core: Fix OOB read for bit offsets of more than one byte When the bit offset is BITS_PER_BYTE or larger the read position is advanced by `bytes_offset`. This is not taken into account in the per-byte read loop which still reads `cell->bytes` resulting in an out of bounds read of `bytes_offset` bytes. The information read OOB does not leak directly as the erroneously read bits are cleared. Detected by KASAN while looking for a use-after-free in simplefb.c.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: pps: generators: don't use the driver's info after unregister A file that was opened before pps_gen_unregister_source() keeps pps_gen, and with it the pointer to the driver's pps_gen_source_info. PPS_GEN_SETENABLE and PPS_GEN_USESYSTEMCLOCK keep using that pointer after the driver has gone: - pps_gen_tio allocates its info with devm_kzalloc(), so after an unbind both ioctls read freed memory, and SETENABLE calls the driver's enable() with its freed private data. - pps_gen-dummy does not set info->owner, so it can be unloaded while the file is open, and the ioctls then read the unloaded module's data and call its code. TIO can't probe in a VM, as it needs ART. A test driver that registers its info in devm memory like TIO does, unbound while /dev/pps-gen0 is open, gives: BUG: KASAN: slab-use-after-free in pps_gen_cdev_ioctl+0x4c1/0x580 Read of size 1 at addr ffff88800e647f38 by task	N/A	More

99443	ppsgen64/145 ... Freed by task 146: kfree+0x25a/0x6d0 release_nodes+0xd1/0x140 devres_release_all+0x10e/0x1a0 device_unbind_cleanup+0x71/0x250 device_release_driver_internal+0x41b/0x570 unbind_store+0xd9/0x100 and unloading pps_gen-dummy while its file is open oopses: BUG: unable to handle page fault for address: ffffffffa0203040 Oops: Oops: 0000 [#1] SMP KASAN NOPTI RIP: 0010:pps_gen_cdev_ioctl+0x372/0x580 Protect info in the ioctl handler with a mutex, and clear it under the mutex on unregister, so that unregister waits for the ioctls using it and later ones fail with -ENODEV. PPS_KC_BIND does the same for a removed PPS device since commit 3649f9a6b897 ("pps: don't allow PPS_KC_BIND on removed devices"). The sysfs attributes need nothing new: they are removed, and their callbacks drained, before info is cleared.		Details
CVE-2026-99444	In the Linux kernel, the following vulnerability has been resolved: power: supply: sbs-battery: disable polling before supply teardown The managed poll-work cancellation is registered before the power supply, so cleanup unregisters the supply before draining polling. A pending or running poller can then call power_supply_changed() on the released supply. Commit 8d59cf3887fb ("power: supply: sbs-battery: Fix use-after-free in power_supply_changed()") moved the IRQ request after supply registration, but did not change the poll-work cancellation order. Add a later managed action that disables and drains polling before supply teardown. Its retained disabled state prevents external-power callbacks from rearming the poll work after it has been drained. Keep the original autocancel to initialize work before supply callbacks can run and to cover failed supply registration. The teardown ordering issue was found by static analysis.	N/A	More Details
CVE-2026-99445	In the Linux kernel, the following vulnerability has been resolved: power: supply: pf1550: drain IRQ work before unregistering supplies The VBUS work calls power_supply_changed(), but its managed cancellation is registered before the supplies. Cleanup can therefore unregister a supply while IRQ-triggered work is still pending. Commit 838767f50747 ("power: supply: pf1550: Fix use-after-free in power_supply_changed()") moved IRQ requests after supply registration, but left work cancellation before the supplies. That change protects the direct IRQ callbacks, while this fix drains the work they queue. Register the three work items after both supplies and before the IRQs. Cleanup then frees the IRQ producers, drains the work, and unregisters the supplies. The teardown ordering issue was found by static analysis.	N/A	More Details
CVE-2026-99447	In the Linux kernel, the following vulnerability has been resolved: spi: spi-pci1xxxx: quiesce DMA engines on transfer timeout If a DMA-mapped transfer times out waiting for DMA_TERM and neither DMA abort interrupt arrived, pci1xxxx_spi_transfer_with_dma() returns through the error label while both DMA engines are still enabled, the done/abort interrupt status is still armed and the transfer state (p->xfer, p->tx_sgl, p->rx_sgl) is still registered. The ISRs keep firing afterwards: they reprogram the DMA engines from p->tx_sgl / p->rx_sgl and signal completion through p->xfer and p->spi_xfer_done, all of which the SPI core is free to unmap or reuse once transfer_one() has returned, which is a use-after-free. A late done interrupt can also start the next transfer queued behind the failed one. Stop both DMA engines, drain the pending done/abort interrupt status and synchronize the DMA IRQs before dropping the in-progress state and the transfer pointers, and re-enable the engines afterwards so that subsequent transfers keep working. The engines are only turned back on once both were confirmed to have stopped: re-enabling an engine that failed to stop would cancel the pending stop request and let it keep running into buffers the SPI core is about to unmap or free, which is exactly the window this patch closes. An engine that failed to stop is left disabled, as the abort paths already do. The DMA interrupt status and mask registers are shared between the two SPI instances, so the cleanup drains only this instance's done/abort status bits instead of masking the interrupt lines, which would also suppress the peer instance's completion interrupts. Also make the DMA ISRs ignore interrupts while no transfer is in progress so that a stray MSI after the cleanup can no longer act on the cleared state. No PCI1XXXX hardware was available for testing and QEMU has no model of this PCIe part, so the timeout path was exercised with a userspace replica of the timeout branch and the DMA ISRs, with the register and interrupt accesses mocked: on the unfixed code both a wedged engine that clears after the timeout and a late done interrupt touch the freed transfer state, while with this fix applied neither does and an engine that failed to stop is left disabled. The module was also loaded and unloaded in a QEMU x86_64 guest running a kernel with this fix applied.	N/A	More Details
CVE-2026-99529	In the Linux kernel, the following vulnerability has been resolved: wifi: b43: fix use-after-free in b43_bcma_remove() b43_one_core_detach() frees the struct b43_wldev it is given the bus device of. b43_bcma_remove() then reads wldev->dev again to free the bus device, which is a read of freed memory. b43_ssb_remove() does not have the problem: it takes a copy of wldev->dev at the start. Do the same here.	N/A	More Details
CVE-2026-99449	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: fix refcount underflow in gadgetfs_bind() error path gadgetfs_bind() calls get_dev() only on the success path, but gadgetfs_unbind() called from the enomem error label unconditionally calls put_dev(). When bind fails (e.g. ENOMEM), put_dev() fires without a matching get_dev(), leaving the refcount unbalanced. A subsequent close of the ep0 file descriptor calls dev_release() -> put_dev() which hits zero and frees the object; then gadgetfs_kill_sb() calls put_dev(the_device) again on the already-freed pointer, triggering a refcount underflow and use-after-free. Fix by calling get_dev() at the start of gadgetfs_bind(), before any error path that invokes gadgetfs_unbind(), so the reference is always balanced regardless of whether bind succeeds or fails.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: vfio/nvgrace-gpu: Use a core-only open_device in fallback mode When the ACPI memory properties are missing or the C2C link is down, the device is exposed through nvgrace_gpu_pci_core_ops. That table opens through nvgrace_gpu_open_device() but closes through vfio_pci_core_close_device(), so open and close are not symmetric. nvgrace_gpu_open_device() registers nvdev->usemem for poison handling unconditionally. In fallback mode		

CVE-2026-99450	usemem.memlength is zero, and register_pfn_address_space() inserts a node covering [0, ULONG_MAX]. Nothing removes it: a second open fails with -EBUSY, and once the device is unbound the node outlives the nvdev that contains it. The next open after a rebind reports: BUG: KASAN: slab-use-after-free in interval_tree_iter_first+0x209/0x310 Read of size 8 at addr ffff88800894aa30 by task nvgrace_repro/203 Call Trace: <TASK> dump_stack_lvl+0x5d/0x80 print_report+0x153/0x4c6 kasan_report+0xe4/0x1a0 interval_tree_iter_first+0x209/0x310 register_pfn_address_space+0x97/0x100 nvgrace_gpu_open_device+0x21c/0x730 vfio_df_open+0x213/0x520 vfio_df_ioctl_bind_iommufd+0x3e8/0x6e0 vfio_device_fops_unl_ioctl+0x782/0x1570 __x64_sys_ioctl+0x13c/0x1c0 do_syscall_64+0x106/0x630 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 188: ... _vfio_alloc_device+0x3a/0x500 nvgrace_gpu_probe+0xcb/0x660 local_pci_probe+0xcf/0x170 ... Freed by task 188: ... kfree+0x20d/0x600 device_release+0xcf/0x250 kobject_put+0x152/0x2f0 pci_device_remove+0xb0/0x1c0 device_release_driver_internal+0x38d/0x540 unbind_store+0xde/0x100 ... Give the fallback table its own open_device that only calls vfio_pci_core_enable() and vfio_pci_core_finish_enable(), the counterpart of vfio_pci_core_close_device(). Fallback mode then never touches the nvgrace memory regions, as Alex suggested. This should be safe because the fallback table uses the vfio-pci core read, write and mmap handlers, so nothing there uses the state nvgrace_gpu_open_device() sets up. The one visible difference is that a fallback open no longer fails when BAR0 cannot be mapped, which matches vfio-pci.	N/A	More Details
CVE-2026-99452	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: delete channel-context timers on deinit ath9k_deinit_channel_context() cancels chanctx_work but does not delete the offchannel and sched timers set up by ath9k_init_channel_context(). If either timer fires after deinit (e.g. during driver unload or suspend), it accesses sc->sc_ah which may already be freed by ath9k_hw_deinit(), causing a use-after-free. Delete both timers with timer_shutdown_sync() before cancelling the work item.	N/A	More Details
CVE-2026-99453	In the Linux kernel, the following vulnerability has been resolved: xen-blkfront: unbind irq before tearing down ring and shadow requests In blkif_free_ring(), the driver tears down the ring's persistent grants, shadow request arrays, and shared ring structure (xenbus_teardown_ring), and only calls unbind_from_irqhandler() at the very end. While blkif_free_ring() is freeing persistent grants and clearing the shadow array, the event channel interrupt (blkif_interrupt) is still registered and active. If an interrupt arrives from the backend during this teardown window, blkif_interrupt() reads rinfo->ring.sring and, via blkif_completion(), accesses rinfo->shadow[id].grants_used and rinfo->shadow[id].sg. blkif_free_ring() tears these structures down without holding rinfo->ring_lock, and the handler only checks info->connected at entry, so this is a real race resulting in a use-after-free or NULL pointer dereference. Fix this by moving unbind_from_irqhandler() to the beginning of blkif_free_ring(). Calling unbind_from_irqhandler() first frees the IRQ and synchronizes with any in-flight interrupt handlers on other CPUs before ring memory and shadow request structures are deallocated, matching the teardown order in drivers/net/xen-netfront.c.	N/A	More Details
CVE-2026-99455	In the Linux kernel, the following vulnerability has been resolved: serial: core: shut down initialized port on removal When a serial port configured as a console is opened and subsequently closed, tty_port_shutdown() skips invoking port->ops->shutdown() because port->console is true. As a result, tty_port_initialized() remains true and the port's interrupt handler stays registered in the irq subsystem. If the underlying device is later unbound or removed (e.g. via sysfs driver unbind or device hot-unplug), serial_core_remove_one_port() unregisters the console, frees uport->name via kfree(), and tears down the port without shutting it down or freeing its IRQ. Consequently, the irqaction descriptor remains registered in genirq with action->name pointing to freed memory. This triggers a use-after-free read whenever genirq accesses action->name, such as when reading /proc/interrupts: BUG: KASAN: slab-use-after-free in string+0x476/0x570 Call Trace: string+0x476/0x570 vsnprintf+0x422/0x1300 seq_printf+0x... show_interrupts+... Fix this by acquiring port->mutex in serial_core_remove_one_port() after console unregistration, clearing port->console under the mutex to prevent data races with concurrent tty_port_shutdown(), and invoking uart_shutdown(NULL, state) if the port remains initialized. This ensures that port->ops->shutdown() is called and the IRQ is freed before kfree(uport->name).	N/A	More Details
CVE-2026-99456	In the Linux kernel, the following vulnerability has been resolved: scsi: dc395x: sync the waiting_timer before freeing the host The waiting_timer callback re-arms itself and takes the host lock. adapter_uninit() currently stops it with timer_delete(), which does not wait for a callback already running on another CPU. The final scsi_host_put() that frees the host (and the host lock the callback dereferences) runs right after adapter_uninit(), so a concurrent or just re-armed callback can fire after the free, leaving a potential use-after-free window. Replace timer_delete() with timer_shutdown_sync() at the top of adapter_uninit(), before the host-lock section: the timer is dequeued, a running callback is waited for and rearming is blocked before the chip is halted, so it cannot fire at all once the adapter has ceased to function. Waiting outside the lock avoids the self-deadlock that would result from syncing a callback blocked on the lock we still hold. No shutdown is needed for selto_timer because it has no reachable armer or callback. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: ALSA: us122l: Fix NULL dereference at hwdep open and release usb_stream_hwdep_open() and usb_stream_hwdep_release() pass the result of usb_ifnum_to_if() straight to usb_autopm_get_interface() and usb_autopm_put_interface(), without checking it for NULL. usb_ifnum_to_if() returns NULL once the device is gone, since usb_disable_device() clears udev->actconfig on disconnect. The hwdep file can outlive the disconnect: the card is released with snd_card_free_when_closed(), so a process that keeps the character device open reaches the driver's release	N/A	More

99315	callback after the USB configuration has been dropped, and interface_to_usbdev() then dereferences the NULL interface at offsetof(struct device, parent). The open callback has the same problem when it races with the disconnect. Check both lookups and fail the open with -ENODEV; in release, skip the runtime PM put when the interface is already gone and still run the rest of the per-file cleanup. Move the us122l->first bookkeeping after the checks so that a failed open no longer leaves a stale file pointer behind.		Details
CVE-2026-99458	In the Linux kernel, the following vulnerability has been resolved: mtd: virt-concat: unlink discarded items before freeing mtd_virt_concat_destroy_items() frees each item without removing it from concat_node_list. list_for_each_entry_safe() protects the current traversal, but concat_node_list still points to the freed items afterward. The function also unwinds errors from mtd_virt_concat_node_create(). A later initialization can traverse a freed item in mtd_is_part_concat(), causing a use-after-free. Unlink each item before freeing it.	N/A	More Details
CVE-2026-99459	In the Linux kernel, the following vulnerability has been resolved: mshv_vtl: detach poll waiter on POLLFREE MSHV_SET_POLL_FILE keeps a poll wait entry after the ioctl returns. A file reference does not always pin the registered waitqueue: signalfd_poll() uses current->sighand->signalfd_wqh, whose lifetime is tied to the task. signalfd_cleanup() sends POLLFREE before the sighand is freed. The MSHV callback currently leaves its entry linked, so __wake_up_pollfree() warns and a later replacement calls remove_wait_queue() through a stale wqh pointer. A source-equivalent KASAN run reproduced: BUG: KASAN: slab-use-after-free in _raw_spin_lock_irqsave _raw_spin_lock_irqsave remove_wait_queue harness_ioctl [mshv_pollfree_harness] The object belonged to sighand_cache and had been released through slab_free_after_rcu_debug() after __cleanup_sighand(). On POLLFREE, unlink the entry under the provider waitqueue lock. Preserve the existing CPU cancellation before publishing a NULL wqh, since a NULL load lets a racing replacement reuse the persistent poll state. Use a matching acquire load under RCU when removing a normal registration. The callback takes no mutex and does not sleep, so the existing mutex-to-waitqueue-lock order is unchanged. An LLM assisted the source audit, fix development, and validation of this issue.	N/A	More Details
CVE-2026-99314	In the Linux kernel, the following vulnerability has been resolved: ALSA: us122l: Add missing NULL check in usb_stream_hwdep_mmap() usb_stream_hwdep_mmap() reads us122l->sk.s and dereferences it without checking for NULL. The stream object is released by usb_stream_free(), which is reached through us122l_stop() both when the device is disconnected and when us122l_start() fails, and it leaves us122l->sk.s set to NULL. The hwdep file stays usable after that, so a later mmap() KASAN: null-ptr-deref in range [0x0000000000000010-0x0000000000000017] RIP: 0010:usb_stream_hwdep_mmap+0x107/0x890 sound/usb/usx2y/us122l.c:182 crashes while reading s->read_size, which sits at offset 0x10. That NULL is an expected state for this field: usb_stream_hwdep_vm_fault(), usb_stream_hwdep_poll(), usb_stream_hwdep_ioctl(), usb_stream_stop() and usb_stream_free() all test it before use. Only the mmap callback misses the check, so let it return -ENODEV, matching what the other hwdep mapping in usx2yhwdeppcm.c does when its shared buffer has not been allocated.	N/A	More Details
CVE-2026-99435	In the Linux kernel, the following vulnerability has been resolved: media: rc: Fix use after free in bpf progs Since commit dccc0c3ddf8f ("media: rc: fix race between unregister and urb/irq callbacks"), rcdev->raw is no longer set to NULL after device unregister. raw->progs could point to stale data.	N/A	More Details
CVE-2026-99434	In the Linux kernel, the following vulnerability has been resolved: ocfs2/heartbeat: unlink live slots when their node is missing o2hb_shutdown_slot() returns immediately when its node is no longer in the configfs registry. A slot can still be linked in o2hb_live_slots when that lookup fails. Region teardown or startup-error cleanup then frees the slot array without removing its embedded list member. With the same peer live in another region, the next list update can follow a predecessor into the freed array. Reusing the node number can also make a later insertion encounter that stale member. Always remove the slot under o2hb_live_lock, even when node lookup fails. Preserve aggregate membership accounting and the last-region DOWN callback: the existing callback interface explicitly permits a NULL node for DOWN events. Only drop the node reference when lookup returned one. This closes both normal release and startup-abort paths without changing the publication or callback locking protocol. Also require a configured node before admitting a new live slot in o2hb_check_slot(). The early live-bitmap check is not held through admission: the last live slot in another region can be removed in between. Without a node, admission can then queue an UP event with a NULL argument and hit the event helper's BUG_ON. Existing live slots can still leave and generate NULL-node DOWN events. A controlled probe-instrumented test reproduced this while stopping one of two regions after removing the peer from configfs. Its KASAN report includes: BUG: KASAN: slab-use-after-free in __list_del_entry_valid_or_report+0x1e0/0x210 Read of size 8 at addr ffff888106a4ac60 by task o2hb-pmbd0058_b/497 Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 kasan_report+0xe0/0x110 __list_del_entry_valid_or_report+0x1e0/0x210 o2hb_do_disk_heartbeat+0x1591/0x2c50 o2hb_thread+0x1ed/0xe70 kthread+0x351/0x460 ret_from_fork+0x659/0x940 ret_from_fork_asm+0x1a/0x30 </TASK> Allocated by task 491: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0xaa/0xb0 __kmalloc_noprof+0x2de/0x780 o2hb_region_dev_store+0x7c2/0x1b60 configfs_write_iter+0x2f9/0x4f0 vfs_write+0x639/0x1010 ksys_write+0x111/0x200 do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f Freed by task 491: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 kasan_save_free_info+0x3b/0x60 __kasan_slab_free+0x5f/0x80 kfree+0x308/0x580 o2hb_unmap_slot_data+0x1db/0x330 o2hb_region_release+0x14b/0x4f0 config_item_cleanup+0x14b/0x230 config_item_put+0x79/0x90 configfs_rmdir+0x58a/0x910 vfs_rmdir+0x2e8/0x820 filename_rmdir+0x3b1/0x520 __x64_sys_rmdir+0x4b/0x70	N/A	More Details

	do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f		
CVE-2026-99319	In the Linux kernel, the following vulnerability has been resolved: ALSA: usx2y: Refuse a retried DSP load after a failure snd_usx2y_hwdep_dsp_load() sets up the asynchronous sequence URBs and the interrupt URB when the FPGA firmware image with index 1 is loaded, and then creates the ALSA devices. If usx2y_in04_init() or usx2y_create_alsa_devices() fails, the URBs that were already set up are left in place. The core keeps the index unloaded when the callback fails, so user space may issue the same request again. The second attempt then enters usx2y_async_seq04_init() or usx2y_in04_init() with the URBs still allocated and trips WARN_ON(usx2y->as04.buffer) or WARN_ON(usx2y->in04_urb). The URBs cannot be released at this point: the card may already have registered PCM devices that user space can still open, and the disconnect handler accesses the URBs without holding the DSP lock. Mark the chip as hung instead, so that a retried request is refused with -ENODEV and the assertions keep their meaning. The URBs are released by snd_usx2y_card_private_free() when the card is freed. A failing usx2y_async_seq04_init() still releases its own URBs, so that path remains retryable.	N/A	More Details
CVE-2026-112213	Out-of-bounds read vulnerability in the pixReadMemJp2kCore() function in Leptonica, located in jp2kio.c. The function retrieves the output dimensions only for component 0 and uses a common index to iterate through the buffers of all components. If the remaining components have smaller dimensions, the traversal exceeds the bounds of their sample buffers, which are allocated by OpenJPEG. When processing a specially crafted JPEG 2000 image, data read from adjacent dynamic memory may be incorporated into the pixels of the decoded image and result in information disclosure to anyone who can access the output.	N/A	More Details
CVE-2026-99406	In the Linux kernel, the following vulnerability has been resolved: xfs: fix lost errno returned from xfarray_foliosort LOLLM noticed that we stop the sort early but don't actually pass along the error code, which could enable repair code to continue running with an unsorted array. Fix that.	N/A	More Details
CVE-2026-99407	In the Linux kernel, the following vulnerability has been resolved: xfs: don't repair fs summary counters with garbage LOLLM noticed that if we scan the filesystem and come up with totally garbage numbers for the summary counters, we don't actually mark the scan incomplete. Because we don't do this, repair can install the garbage values into the filesystem, making things worse. Don't do that.	N/A	More Details
CVE-2026-99408	In the Linux kernel, the following vulnerability has been resolved: xfs: don't proceed with xattr repair if the live update fails LOLLM observes that we ought to make one last check for failed live updates before replacing the xattr structure. Let's be cautious here.	N/A	More Details
CVE-2026-99409	In the Linux kernel, the following vulnerability has been resolved: xfs: don't fix the directory tree if live update fails LOLLM complains that we don't double-check dl->aborted one last time to before committing repairs to the filesystem. This is important because the live update handler could have set ->aborted if it encountered a runtime error.	N/A	More Details
CVE-2026-99410	In the Linux kernel, the following vulnerability has been resolved: xfs: always reset incore attr fork buffer when resetting fork LOLLM noticed that the local format attr fork verifier depends on if_bytes being set correctly. However, the fork reset function doesn't reset that correctly, which can cause post-repair verifier failures. Fix that.	N/A	More Details
CVE-2026-99411	In the Linux kernel, the following vulnerability has been resolved: xfs: actually lock the metafile reservation when adjusting after repair LOLLM reported that xrep_reset_metafile_resv performs unlocked loads and stores to the metadata file space reservation counters and thus can race with other threads. Fix this by taking the lock.	N/A	More Details
CVE-2026-99412	In the Linux kernel, the following vulnerability has been resolved: xfs: avoid cross-rtgroup reaping after a repair LOLLM notices that the extents stored in a xrtb_bitmap (aka xfs_rtblock bitmap) can span multiple rtgroups due to two circumstances. The first is that the size of an rtgroup is an exact power of two, which means that rtgroups are adjacent in the segmented xfs_rtblock_t address space. This is fairly common to reduce the amount of multiplication and division needed to handle space on the rt volume. The second is that (unlike in the original rtgroups design), rtgroups do not have fixed-location metadata like AGs do, which means that there's nothing to force a break between rtgroups. Therefore, we must loop through the rtgroups in xreap_rtmeta_extent to avoid running off the end of an rtgroup while reaping.	N/A	More Details
CVE-2026-99413	In the Linux kernel, the following vulnerability has been resolved: xfs: fix buffer overruns in xfs_ioc_attr_list When we converted the al_offset array in struct xfs_attrlist into a VLA, the size of the object shrank by 4 bytes. Unfortunately, the buffer size validation in the attrlist ioctl wasn't updated to notice this, so the al_offset[0] assignment blindly writes off the end of the buffer. LOLLM noticed the omitted check and complained.	N/A	More Details
CVE-2026-99414	In the Linux kernel, the following vulnerability has been resolved: ovl: skip copy up of empty security.capability ovl_copy_xattr() gives up if it can't read an attribute that listxattr() has named. An empty security.capability, which anybody has been able to store on any file for a long time, is such a case: it is listed, and getxattr() has nothing to return for it. A file in a lower layer that carries one can't be copied up: it can't be written to, renamed or have its mode changed. There is nothing to copy if there are no data. Go on with the next one. [amir: treat -EINVAL as no capability too, and do the same in ovl_copy_up_meta_inode_data()]	N/A	More Details

CVE-2026-99415	In the Linux kernel, the following vulnerability has been resolved: ovl: set the revalidate flags after d_instantiate() stat() of a file on an overlayfs mount triggers a kernel warning and fails with ECHILD if another task is creating that file at the same time. This happens on an overlay with a FUSE upper layer, which is the layer where overlayfs creates new files. Write access to a directory on such an overlay is all a user needs to run the two tasks. The revalidate flags are DCACHE_OP_REVALIDATE and DCACHE_OP_WEAK_REVALIDATE. They tell path walk to call ->d_revalidate() or ->d_weak_revalidate() on a dentry, so that the filesystem can check that the cached dentry is still valid. An overlay dentry only has one of these flags if a dentry in one of its layers has it. FUSE dentries have the first one. When a flag is set on an overlay dentry, path walk calls ovl_dentry_revalidate_common(). That function only expects a dentry without an inode in RCU walk, while the dentry is being destroyed. For such a dentry outside of RCU walk, it hits a WARN_ON() and returns -ECHILD. So a negative overlay dentry must never have the flags set. ovl_instantiate() breaks this rule. When a task creates a file, the overlay dentry for the new name is hashed but still negative. ovl_instantiate() then works in this order: - It copies the flags from the new upper dentry to the overlay dentry. - It calls ovl_get_inode(), which can sleep in a getxattr request to the FUSE server. - Only then does it call d_instantiate(). A second task that looks up the name between the first step and the last finds a negative dentry with the flags set and revalidates it. Reproduced with a FUSE server that delays its getxattr reply: stat() of a name that was being created hit the WARN_ON() and failed with ECHILD. Fix this by setting the flags after d_instantiate().	N/A	More Details
CVE-2026-99416	In the Linux kernel, the following vulnerability has been resolved: ovl: read the nlink xattr of every indexed non-dir With nfs_export=on, data written through a file handle to a lower file, i.e. a file that exists only in a lower layer of the overlay, can be lost on the next mount. A lookup of the file by name then returns -EIO. Also, a file with one name reports st_nlink 2 once it has been copied to the upper layer. After such a file is removed, its file handle can still open it until the next mount. Before a lower file is changed for the first time, overlayfs copies it to the writable upper layer. This is called copy up. The index directory is $workdir/index, where $workdir is the directory given by the workdir= option. Each entry is named after the file handle of a lower inode, so that overlayfs can get from the handle to the copied up inode. With nfs_export=on, copy up indexes every non-directory: it copies the file to an entry in the index directory and then hard links that entry into the upper directory. The index entry is an extra link to the upper inode, so overlayfs stores the link count that users should see in the overlay.nlink xattr. A file can also be opened by handle, either by an NFS client or by a process with CAP_DAC_READ_SEARCH calling open_by_handle_at(). If the file is not in the dcache, its overlay dentry is disconnected, i.e. it has no parent yet. Copy up then creates the index entry but no hard link in the upper directory. The index entry has i_nlink 1 and holds the only copy of the data written to the file. ovl_get_nlink() reads the stored count. Its caller passes a default, which the function returns if there is no stored count. If the lower inode has i_nlink 1, the function returns the default without reading the xattr. That was correct as long as only lower files with multiple links were indexed. On the next mount, ovl_verify_index() looks for orphans, i.e. index entries of files whose last name was removed. For an index entry with i_nlink 1, it passes a default of 0 to ovl_get_nlink() and treats a result of 0 as an orphan. Here the lower inode also has i_nlink 1, so the result is the default. ovl_indexdir_cleanup() then replaces the entry with a whiteout, which in the index marks a file as deleted. The written data is lost. A lookup by name then returns -EIO, because a whiteout as the index entry of a lower file that still exists is a bad index. ovl_get_inode() gets the wrong count the same way. Its default is the i_nlink of the upper inode, which includes the index entry, so a copied up file with one name reports st_nlink 2. When that name is removed, the count drops to 1 rather than 0. ovl_nlink_end() replaces the index entry with a whiteout only at a count of 0, so the entry stays, and the handle of the removed file still opens it. Reproduced by writing through open_by_handle_at() to an uncached lower file with one link. After a remount the data was lost, and the name returned EIO. Fix this by adding an argument to ovl_get_nlink() that says whether the caller found an index entry, and by skipping the xattr only if the caller found none. Known gaps: - An index entry that a kernel without this patch left behind for a removed file now stays, and the handle of that file keeps working. The entry has i_nlink 1 and a stored count of 1, just like the entry of a file that was only written by handle. Without this patch, the next mount treated it as an orphan.	N/A	More Details
CVE-2026-99417	In the Linux kernel, the following vulnerability has been resolved: ovl: fix a torn path in ovl_get_redirect() A hard link of a metacopy file that races with a rename of an ancestor stores a redirect to the wrong lower file. Each component of an absolute redirect must be the name that a dentry has in the parent that the walk continues with. ovl_get_redirect() copies the name under d_lock and calls dget_parent() after the unlock, and the callers for a hard link hold no lock that stops a rename across directories. A rename between the two steps combines the old name with the new parent. With delays in ovl_get_redirect(), an unprivileged user hard-linked a metacopy file while renaming its ancestor to another directory, the stored redirect named another lower path, and after a remount the file showed the data of a root-only file. Fix this by taking the parent first and checking it again under d_lock.	N/A	More Details
CVE-2026-99418	In the Linux kernel, the following vulnerability has been resolved: ovl: do not copy up an xattr value that was never read Copy-up stores uninitialized kernel memory in an xattr of the upper file, where the user can read it. ovl_copy_xattr() may only pass bytes to ovl_do_setxattr() that ovl_do_getxattr() has read. When the read fails with -ERANGE the function asks for the size again, and if the new size fits the buffer it stores the buffer although nothing was read into it. The buffer comes from kvmalloc(), or it still holds the value of the previous xattr. A FUSE lower layer that answers every sized GETXATTR with -ERANGE made copy-up store 4096 bytes that the server never sent, with a kernel pointer in them, as an xattr of the upper file. Fix this by failing the copy-up with -EIO when the size that follows -ERANGE fits the buffer.	N/A	More Details

CVE-2026-99419	In the Linux kernel, the following vulnerability has been resolved: ovl: do not copy up a dentry that was removed A change to a removed file through an open descriptor replaces whatever has its old name in the upper layer. Copy up must only create a name in the upper directory, never replace one. After an unlink, an rmdir or a rename over a lower dentry, a dentry that an open file keeps alive still has its old parent and name, but in the upper directory that name is now a whiteout or a new object. ovl_copy_up_workdir() renames the copy over it, which brings back the removed object and destroys whatever took its name. An unprivileged user opened a lower FIFO and unlinked it, root created a file at that name, and fchmod on the old descriptor replaced the file of root with the FIFO. Fix this by looking up the target name of the rename in ovl_copy_up_workdir() with LOOKUP_EXCL, so that copy up fails with -EEXIST if the name is taken, as copy up with O_TMPFILE already does when it links the copy to the name. [amir: use LOOKUP_EXCL instead of relying on d_unhashed()]	N/A	More Details
CVE-2026-99420	In the Linux kernel, the following vulnerability has been resolved: ovl: check that the parents of volatile/dirty are dirs An unprivileged user crashes the kernel by mounting overlayfs with the volatile option. A dentry that serves as the parent of a lookup must be a directory. ovl_create_volatile_dirty() walks work/incompat/volatile/dirty by name, and ovl_lookup_or_create() returns an existing entry whatever its type. A file or a symlink that another task creates during the mount becomes the parent of the next lookup, which calls a NULL ->lookup. An unprivileged user mounted overlayfs with "volatile" in a user namespace while a second process created a symlink named work/incompat, and the mount task oopsed on a NULL ->lookup call. Fix this by returning -ENOTDIR from ovl_lookup_or_create() when the caller asked for a directory and the existing entry is a non-dir.	N/A	More Details
CVE-2026-99421	In the Linux kernel, the following vulnerability has been resolved: ovl: always encode the origin fh when index is enabled On an overlay with index=on,nfs_export=on, copy up of a directory crashes the kernel once an LSM has started to deny CAP_DAC_READ_SEARCH to the mounter, the task that mounted the overlay. The mount and the LSM policy change require root. After that, any user who can write to a directory that exists only in a lower layer triggers the oops by doing so. Before a file or directory that exists only in a lower layer is changed for the first time, overlayfs copies it to the writable upper layer. This is called copy up. On copy up, ovl_get_origin_fh() encodes the file handle of the lower inode, called the origin. It returns NULL if ovl_can_decode_fh() says that the lower filesystem cannot decode handles. Whether the callers can handle NULL depends on the index option: - Without an index, NULL is fine: the callers store an empty origin xattr. - With index enabled, overlayfs keeps entries for copied up inodes in an internal index directory. The name of an index entry is derived from the handle, so ovl_get_origin_fh() must return a handle or an error. The mount only enables the index if ovl_can_decode_fh() succeeds for all layers, but the result can change later. The function checks the export operations of the filesystem and capable(CAP_DAC_READ_SEARCH). Copy up runs with the mounter's credentials, and an LSM can start to deny the capability to those credentials after the mount. ovl_get_origin_fh() then returns NULL with index enabled. This goes wrong in two callers: - With nfs_export=on, copy up of a directory creates an index entry for the directory. ovl_create_index() passes the NULL handle to ovl_get_index_name_fh(), which dereferences it. - If the result changes while mount(2) is still running, ovl_get_indexdir() passes a NULL from ovl_get_origin_fh() to ovl_verify_set_fh() in the same way. Before the commit in the Fixes tag, these two callers encoded the handle themselves, without consulting ovl_can_decode_fh(). Reproduced with a BPF LSM program that started to deny the capability after the mount. An unprivileged touch in a directory that had not been copied up yet oopsed in ovl_get_index_name_fh(). With this patch the touch succeeds and the directory gets its index entry. Fix this by calling ovl_can_decode_fh() in ovl_get_origin_fh() only when index is disabled.	N/A	More Details
CVE-2026-99423	In the Linux kernel, the following vulnerability has been resolved: ovl: take a snapshot of the name for copy up The copy up of a directory reads a freed dentry name when a lookup moves the directory at the same time. A pointer to the name of a dentry is only valid while something stops a move of the dentry. ovl_do_copy_up() keeps a plain copy of c->dentry->d_name in c->destname, and ovl_copy_up_start() only keeps ovl_rename() away. When a lookup finds the same directory under another name after a change to the lower layer, d_splice_alias() moves the dentry and frees its external name. With a delay in ovl_do_copy_up(), a chmod of a lower-only directory with a 60 byte name raced with a rename of that directory in the lower layer and a stat of the new name, and KASAN reported a use-after-free in lookup_noperm_common(). Fix this by using take_dentry_name_snapshot() for c->destname.	N/A	More Details
CVE-2026-99426	In the Linux kernel, the following vulnerability has been resolved: smb: client: hold server info until cifsd exits The last session reference can be dropped while the cifsd demultiplex thread is still running clean_demultiplex_info(). That thread frees TCP_Server_Info there, racing with cifs_put_tcp_session() as it releases the server crypto state. syzbot reported a KASAN use-after-free in cifs_crypto_secmech_release(). Keep an object reference for the cifsd thread. Defer final crypto, session-key, network namespace, string, and structure cleanup until both the last session reference and the thread reference are dropped. Drop the thread reference after clean_demultiplex_info() has finished using the server. Build-tested with ARCH=arm64 and W=1 for fs/smb/client/connect.o. The KASAN reproducer was not run on the patched kernel.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix OOB last_entry pointer from unbounded LastNameOffset CIFSFindFirst() and CIFSFindNext() locate the last entry of a TRANS2_FIND_FIRST2 or TRANS2_FIND_NEXT2 response by adding the LastNameOffset of the response parameters to the start of the data area, and store the result in psrch_inf->last_entry. MS-CIFS 2.2.6.2.2 defines that offset as the offset of the entry's file name; the client follows the reading of NT and Samba servers, which point it at the entry. LastNameOffset is only checked against CIFSMaxBufSize, which says		

CVE-2026-99427	nothing about the response actually received: validate_t2() lets the data area start up to 1024 bytes into the buffer, and SendReceive() copies at most CIFSMaxBufSize + MAX_CIFS_HDR_SIZE bytes of a response, 16468 by default, into a cifs_request allocation of CIFSMaxBufSize + MAX_SMB2_HDR_SIZE bytes, 16588 by default. The pointer itself can therefore be placed anywhere from the end of the received data to 820 bytes past the end of the allocation. The read through that pointer is far larger than its displacement. find_cifs_entry() hands last_entry to cifs_save_resume_key() after every FindNext and after the FindFirst of a search rewind, and cifs_fill_dirent() takes the length of the name from the entry itself: a __le32 FileNameLength for most levels, a __u8 for SMB_FIND_FILE_INFO_STANDARD, and for SMB_FIND_FILE_UNIX the result of a scan for the terminating NUL over up to PATH_MAX + 1 UTF-16 units. The name is recorded as the resume name of the search, and the next CIFSFindNext() rejects only a recorded length of PATH_MAX or more before copying the name into the ResumeFileName of its request and sending it with the resume key taken from the same entry: up to 4 KiB of memory from past the end of the allocation leaks to the server, and for SMB_FIND_FILE_UNIX the NUL scan runs about twice as far before the length is even known. The entry does not have to lie outside the response for that either: a maximal response ends about 120 bytes before the end of the allocation, so an entry in its last bytes with a name just under PATH_MAX has nearly all of that name copied from past the allocation. It needs nothing unusual from the client: a FindNext response that reports no entries and does not end the search makes find_cifs_entry() send the next FindNext at once, carrying whatever was parsed at LastNameOffset, and a search rewind can do the same after its FindFirst. SMB1 is not negotiated by default: this takes a malicious or compromised server, a mount of it with an explicit vers=1.0, which CONFIG_CIFS_ALLOW_INSECURE_LEGACY (default y) permits, and a directory listing on that mount. With KASAN enabled, a response that places the last entry past the buffer, with a large DataOffset and a LastNameOffset near the end of its range, gives: == BUG: KASAN: slab-out-of-bounds in cifs_save_resume_key.isra.0+0x7aa/0x7f0 Read of size 4 at addr ffff88807b55c33b by task ls/83 CPU: 1 UID: 0 PID: 83 Comm: ls Not tainted 7.3.0-rc4-00457-gf14572c203d5 #69 PREEMPT(lazy) Call Trace: <TASK> kasan_report+0xdf/0x1a0 cifs_save_resume_key.isra.0+0x7aa/0x7f0 cifs_readdir+0x1ed6/0x29c0 iterate_dir+0x1c0/0x570 __x64_sys_getdents64+0x133/0x270 do_syscall_64+0x109/0x5d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 83 on cpu 0 at 19.317556s: cifs_buf_get+0x36/0x90 smb_init+0x4f/0x110 CIFSFindNext+0xf7/0x14c0 cifs_readdir+0xf51/0x29c0 The buggy address belongs to the object at ffff88807b558000 which belongs to the cache cifs_request of size 16588 The buggy address is located 623 bytes to the right of allocated 16588-byte region [ffff88807b558000, ffff88807b55c0cc) == The four bytes read there are the FileNameLength of the entry, which cifs_fill_dirent() tak ---truncated---	N/A	More Details
CVE-2026-99428	In the Linux kernel, the following vulnerability has been resolved: ocfs2: snapshot dentry names in file and setattr traces Patch series "ocfs2: stabilize dentry names in file and setattr traces", v4. Snapshot names at all six users of ocfs2__file_ops and at the separate setattr event. Patch 1 handles the five ordinary file-operation callers and setattr. Patch 2 handles release separately to preserve its locked open-count sample and take the snapshot outside that critical section. The release explanation distinguishes its existing RCU lifetime coverage from name/length consistency. It does not claim that the ordinary callers' external-name UAF interleaving crosses the held inode spinlock. The release source change is unchanged from v3. Patch 1 also removes the read callback's redundant NULL-inode check. The two patches apply independently and leave trace formats unchanged. This patch (of 2): The open, fsync, read, write and splice-read trace calls pass the live dentry name and length to the shared ocfs2__file_ops event class. ocfs2_setattr() passes the same unprotected arguments to its own event. Both events size and copy a __string field separately. A reference to the dentry does not pin the external storage backing an individual name. With an event enabled, a caller can be preempted after loading the name pointer. A long-to-short rename can then replace the name and reclaim its old external allocation before the trace callback sizes or copies it, provided no other snapshot retains that allocation. The setattr inode lock does not exclude every rename: a same-directory rename of a source directory need not take that directory's inode lock. It can therefore change the dentry name while setattr records its event. Inline names and separately loaded lengths can also change during trace record sizing and copying. Take a dentry name snapshot around each of these six trace calls and pass both the string and length from the snapshot. Check event enablement before taking a snapshot. Keep the event classes and formats unchanged. Also remove the redundant NULL-inode check in the read callback. file_inode() returns the open file's inode, which is already used by the trace call above the check. The historical report below is from the read path. KASAN report as below: BUG: KASAN: slab-use-after-free in strlen+0x85/0xa0 Read of size 1 at addr ffff88810a5b2b18 by task python3/496 [CPU and host details omitted.] Call Trace: <TASK> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 ? strlen+0x85/0xa0 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x20e/0x420 ? strlen+0x85/0xa0 kasan_report+0xe0/0x110 ? strlen+0x85/0xa0 strlen+0x85/0xa0 trace_event_raw_event_ocfs2__file_ops+0xf0/0x480 ? __pfx_trace_event_raw_event_ocfs2__file_ops+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? pmbd_probe_hit_cookie+0xe7/0x1c0 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_pmbd_probe_hit_cookie+0x10/0x10 ocfs2_file_read_iter+0x6cd/0xb30 ? __pfx_ocfs2_file_read_iter+0x10/0x10 ? selinux_file_permission+0x11d/0x540 ? srso_alias_return_thunk+0x5/0xfbef5 ? security_file_permission+0x26/0x80 vfs_read+0x7d2/0xc20 ? __pfx_vfs_read+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __pfx_do_sys_openat2+0x10/0x10 ksys_read+0x111/0x200 ? __pfx_ksys_read+0x10/0x10 ? irqentry_exit+0xbe/0x820 ? srso_alias_return_thunk+0x5/0xfbef5 ? trace_hardirqs_on_prepare+0x11d/0x150 ?	N/A	More Details

	srso_alias_return_thunk+0x5/0xfbef5 do_syscall_64+0x114/0x620 entry_SYSCALL_64_after_hwframe+0x77/0x7f [User register dump omitted.] </TASK> Allocated by task 489: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0xaa/0xb0 __kmalloc_noprof+0x2de/0x780 __d_alloc+0x5fc/0x8c0 d_alloc_parallel+0x106/0x11d0 lookup_open+0x3ff/0 ---truncated---		
CVE-2026-99431	In the Linux kernel, the following vulnerability has been resolved: ocfs2: cancel quota sync work before freeing quota info OCFS2 quota info owns dqi_sync_work and must outlive both its timer and qsync_work_fn(). ocfs2_local_free_info() frees that owner without cancelling the work, relying on the quota suspend/disable wrappers. Loading quotas on a read-only to read-write remount can successfully queue the work in ocfs2_global_read_info() and then fail while attaching a dquot to an existing inode in add_dquot_ref(). The quota core unwinds that failure directly through dquot_disable(), bypassing the OCFS2 wrapper: Remount task (s_umount write) Timer/workqueue ocfs2_global_read_info() schedules dqi_sync_work add_dquot_ref() returns an error dquot_disable() ocfs2_local_free_info() frees oinfo delayed-work timer expires qsync_work_fn() reads oinfo The timer core can then write to the freed embedded timer, or qsync_work_fn() can read oinfo->dqi_gqinode before its s_umount trylock. The outer ocfs2_disable_quotas() cleanup skips the failed type because dquot_disable() has already cleared its loaded state. Cancel dqi_sync_work synchronously at the start of ocfs2_local_free_info(), before dropping the global inode or freeing any owner state. The work is initialized whenever read_file_info() succeeds. dquot_disable() holds s_umount for write, so the callback skips its quota scan via down_read_trylock() and can finish without waiting for teardown. The synchronous cancel also handles its self-rearm. KASAN report as below: BUG: KASAN: slab-use-after-free in __run_timers+0x9e9/0xa10 Write of size 8 at addr ffff888111b0aad0 by task swapper/3/0 Call Trace: <IRQ> dump_stack_lvl+0x93/0xd0 print_report+0xce/0x630 ? __run_timers+0x9e9/0xa10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x20e/0x420 ? __run_timers+0x9e9/0xa10 kasan_report+0xe0/0x110 ? __run_timers+0x9e9/0xa10 __run_timers+0x9e9/0xa10 ? __pfx___run_timers+0x10/0x10 ? __lock_acquire+0x5a3/0x2140 ? __pfx_do_raw_spin_lock+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 timer_expire_remote+0xfa/0x160 ? __pfx_timer_expire_remote+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? srso_alias_return_thunk+0x5/0xfbef5 ? trace_hardirqs_on+0x18/0x160 tmigr_handle_remote_up+0x4ba/0x9d0 ? __pfx_tmigr_handle_remote_up+0x10/0x10 ? tmigr_handle_remote+0x21a/0x340 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 ? __pfx_tmigr_handle_remote_up+0x10/0x10 __walk_groups_from+0x5a/0x160 tmigr_handle_remote+0x24f/0x340 ? __pfx_tmigr_handle_remote+0x10/0x10 ? lock_acquire+0x190/0x300 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 run_timer_softirq+0x222/0x270 ? __pfx_run_timer_softirq+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? srso_alias_return_thunk+0x5/0xfbef5 ? lockdep_hardirqs_on_prepare+0xb8/0x1a0 ? srso_alias_return_thunk+0x5/0xfbef5 handle_softirqs+0x1fa/0x8d0 ? __pfx_handle_softirqs+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? _raw_spin_unlock+0x23/0x40 ? srso_alias_return_thunk+0x5/0xfbef5 ? __hrtimer_rearm_deferred+0x1fb/0x5d0 __irq_exit_rcu+0x113/0x1e0 irq_exit_rcu+0xe/0x20 sysvec_apic_timer_interrupt+0x6c/0x80 </IRQ> <TASK> asm_sysvec_apic_timer_interrupt+0x1a/0x20 ? ct_kernel_exit.constprop.0+0xc5/0xf0 ? do_idle+0x3ae/0x650 default_idle+0x9/0x10 default_idle_call+0x6c/0xb0 do_idle+0x3ae/0x650 ? cpuhp_online_idle+0x9c/0xb0 ? __pfx_do_idle+0x10/0x10 ? __entry_text_end+0x1020f6/0x1020f9 ? lockdep_hardirqs_on_prepare+0xeb/0x1a0 ? srso_alias_return_thunk+0x5/0xfbef5 ? trace_hardirqs_on+0x18/0x160 ? srso_a ---truncated---	N/A	More Details
CVE-2026-99461	In the Linux kernel, the following vulnerability has been resolved: mm: khugepaged: fix swap entry value to folio_pfn() Patch series "mm: khugepaged: fix tracepoint UAF", v5. The khugepaged tracepoints take a folio pointer and call folio_pfn(), but by then the folio may no longer be valid: freed after folio_put(), folio_unlock() or pte_unmap_unlock(), or not a folio at all but an xarray-encoded swap entry. On classic SPARSEMEM, dereferencing it oopses khugepaged as soon as the trace event is enabled; on other memory models it merely prints a bogus pfn. Pass the pfn to the tracepoints directly, captured while the folio is still pinned, closing the use-after-free windows in mm_khugepaged_scan_file(), mm_khugepaged_scan_pmd() and mm_khugepaged_collapse_file(). This patch (of 3): When the swap entries found exceed max_ptes_swap, the loop is left via break with folio still holding the xarray value that encodes the swap entry, not valid folio pointer. That value is passed to trace_mm_khugepaged_scan_file(), which feeds it to folio_pfn(). On FLATMEM and SPARSEMEM_VMEMMAP, the page_to_pfn() is plain pointer arithmetic, so the trace event merely prints bogus scan_pfn. On classic SPARSEMEM, the page_to_pfn() reads page->flags, dereferencing the tiny encoded integer and oopsing khugepaged whenever the trace event is enabled. So when folio is the swap entry value, simply set pfn to -1, just like exhausted scan naturally. And the folio_put() has maybe dropped the last reference of folio. The trace_mm_khugepaged_scan_file() is left with a dangling folio pointer. so using the folio_pfn() before dropping the reference, closing use-after-free window.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: media: redrat3: fix UAF in probe error path leaving rc device registered redrat3 stores its device name and phys path inside the main rr3 structure, and the rc device's device_name and input_phys point into rr3. When the RC device is registered, it keeps those pointers. KASAN: slab-use-after-free in string_nocheck lib/vsprintf.c:648 Read of size 1 at addr ffff888051fda758 by task udevd/7464 Call Trace: string_nocheck lib/vsprintf.c:648 [inline] string+0x216/0x2d0 lib/vsprintf.c:730 vsnprintf+0x74a/0xef0 lib/vsprintf.c:2945 vscnprintf+0x41/0x90 lib/vsprintf.c:3014 sysfs_emit+0x10e/0x180 fs/sysfs/file.c:761 input_dev_show_name+0x58/0x70		

CVE-2026-99464	drivers/input/input.c:1282 Allocated by task 10: redrat3_dev_probe+0x477/0x2570 drivers/media/rc/redrat3.c:1023 Freed by task 10: redrat3_delete drivers/media/rc/redrat3.c:466 [inline] redrat3_dev_probe+0x1bf4/0x2570 drivers/media/rc/redrat3.c:1124 Syzkaller triggers this via usb probing. It probes the RedRat3 USB interface, which calls redrat3_dev_probe() in redrat3.c. That function allocates rr3, then builds an rc device whose name points into rr3 via redrat3_init_rc_dev() in redrat3.c, and registers it with rc_register_device(). When the detector enable fails after the RC device is registered, the probe jumps to the led_free path. That path unregisters the LED but does not unregister the RC device before freeing rr3 via redrat3_delete() in redrat3.c. The RC device still holds device_name = rr3->name which is now freed. Later udevd reads /sys/.../input device name via sysfs_emit() in file.c, which calls input_dev_show_name() in input.c, which emits dev->name which is the freed rr3->name. When the detector enable fails after the RC device is registered, the error path must not leave the RC device registered. Make the probe error path clean up the same way redrat3_dev_disconnect() does. This change should be safe because the RC device is fully registered at this point and its teardown via rc_unregister_device() is protected by the input device mutex, and rr3 is still alive during unregister so device_name remains valid until after unregister. The URBs are killed before rc_free, so no completion can run after rc is freed. No new lock ordering is introduced.	N/A	More Details
CVE-2026-99465	In the Linux kernel, the following vulnerability has been resolved: media: microchip-isc: take a reference on the parsed endpoints for_each_endpoint_of_node() drops the reference on the current node as it advances. xisc_parse_dt() and isc_parse_dt() store the node in subdev_entity->epn and release it later with of_node_put(), but never took their own reference, so the stored pointer refers to an already-released node. This underflows the refcount and can use-after-free, reachable through the camera device tree overlay. Take a reference with of_node_get() when storing the node, and drop it in microchip_isc_subdev_cleanup() so the entities the bind loop never reaches on an early exit do not leak it.	N/A	More Details
CVE-2026-99499	In the Linux kernel, the following vulnerability has been resolved: s390/pci: Fix use-after-free race in zpci floating interrupt cleanup zpci_clear_irq() stops the adapter from raising new interrupts for the function, but a zpci_floating_irq_handler() already running on another CPU can still be scanning zdev->aibv when zpci_msi_teardown_floating() releases it. Clear the zpci_ibv[] entry so no further handler picks the vector up, then wait for a grace period before releasing it. The handler runs inside the rcu_read_lock() section that do_airq_interrupt() holds across airq->handler(), so synchronize_rcu() drains any handler still in flight. Release the vector and free the summary bit only after the grace period, so neither can be reused while a reader still holds the old pointer. zpci_ibv served both delivery modes, indexed by function under FLOATING and by cpu under DIRECTED. Only the floating vectors are published to and torn down under the interrupt handler, so split the directed vectors out into zpci_dibv and annotate zpci_ibv __rcu, which lets sparse check the accessors above.	N/A	More Details
CVE-2026-99501	In the Linux kernel, the following vulnerability has been resolved: s390/pci: Fix MSI directed-mode teardown IRQ bit count On s390 with directed interrupts enabled, zpci_msi_teardown_directed() frees the platform's maximum number of MSI bits (zdev->max_msi) instead of the actual allocated count (zdev->msi_nr_irqs). This corrupts the shared IRQ bitmap used by all PCI functions, causing lost interrupts and heap corruption. Fix zpci_msi_teardown_directed() to only free the actual allocated IRQ bit count.	N/A	More Details
CVE-2026-99502	In the Linux kernel, the following vulnerability has been resolved: s390/pci: Fix double-free and NULL deref in zpci MSI domain cleanup zpci_remove_parent_msi_domain() dereferences zbus->msi_parent_domain unconditionally and does not clear it afterwards. zpci_bus_create_pci_bus() removes the domain when pci_create_root_bus() fails, then zpci_bus_release() removes it again on the last kref_put(), reading ->fwnode from the freed irq_domain and freeing it twice. If zpci_alloc_domain() or zpci_create_parent_msi_domain() fails, no domain is created at all; zbus is kzalloc'd, so the same release path dereferences NULL. Return early when there is no domain, and clear the pointer after removing one.	N/A	More Details
CVE-2026-99503	In the Linux kernel, the following vulnerability has been resolved: ata: sata_fsl: Fix use-after-free of host_priv on probe() failure sata_fsl_host_stop() releases hcr_base and host_priv: static void sata_fsl_host_stop(struct ata_host *host) { struct sata_fsl_host_priv *host_priv = host->private_data; iounmap(host_priv->hcr_base); kfree(host_priv); } ->host_stop() is called by the driver core through the ata_host_stop() devres action which ata_host_start() registers, so it is called both when the device is unbound and when probe() fails after the host has been started. The error_exit_with_cleanup label of sata_fsl_probe() releases hcr_base and host_priv as well, and it is reachable from the two device_create_file() calls which are done after the host has been activated. In that case sata_fsl_host_stop() runs on an already freed host_priv, resulting in a use-after-free and a double iounmap()/kfree(). Add a separate error label for the failures happening after the host has been activated, which only detaches the host and leaves hcr_base and host_priv to sata_fsl_host_stop(). Note that ata_host_detach() is dropped from error_exit_with_cleanup, as that label is now only reachable while host is still NULL.	N/A	More Details
CVE-2026-99504	In the Linux kernel, the following vulnerability has been resolved: ASoC: audio-graph-card2: Fix use-after-free in audio_graph2_link_c2c() of_graph_get_next_port() consumes the reference of its prev argument. port0 is only borrowed from the caller's iterator, so the call drops a reference the function does not own, and port0 is dereferenced again afterwards. Pass an extra reference, as graph_count_c2c() does, so the reference the callee consumes is the one taken here.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: apparmor: fix a use-after-free in aa_lookupn_ns() aa_lookupn_ns() takes a reference on the namespace it finds with aa_get_ns(), which is		

CVE-2026-99505	kref_get() on the reference count that the namespace is released through, without checking whether the namespace is still in use. A concurrent removal of a ":"-prefixed namespace looks the same namespace up and only then takes the parent lock to remove it, so the last reference can be dropped between the lookup and the aa_get_ns() call: the reference count then reports the increment on a zero count and the kernel classifies the following use of the object as a use-after-free. A user with CAP_MAC_ADMIN can trigger this by writing a ":<name>" removal and a ":<name>/<profile>" removal to /sys/kernel/security/apparmor/.remove concurrently. refcount_t: addition on 0; use-after-free. WARNING: lib/refcount.c:25 at refcount_warn_saturate+0xc0/0xe0, CPU#6 Call Trace: <TASK> aa_lookupn_ns+0xa9/0xc0 aa_remove_profiles+0x4fb/0x730 ? __kvmalloc_node_noprof+0x219/0x5c0 ? __pfx_aa_remove_profiles+0x10/0x10 ? aa_loaddata_alloc+0x65/0x1f0 ? _copy_from_user+0x2d/0x80 ? aa_loaddata_alloc+0x111/0x1f0 profile_remove+0x2ee/0x4a0 vfs_write+0x21e/0xcf0 ? rcu_core+0x29c/0x1740 ? timerqueue_linked_add+0x1f4/0x3d0 ? __pfx_vfs_write+0x10/0x10 ? mutex_lock+0x81/0xe0 ? __pfx_mutex_lock+0x10/0x10 ? __pfx_rcu_core+0x10/0x10 ? __pfx_hrtimer_update_next_event+0x10/0x10 ? mutex_unlock+0x7b/0xd0 ? fdget_pos+0x24d/0x4b0 ksys_write+0xf7/0x1c0 ? __pfx_ksys_write+0x10/0x10 ? restore_fpregs_from_fpstate+0x37/0xc0 do_syscall_64+0xf9/0x540 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Fix by rejecting a namespace whose reference count is zero.	N/A	More Details
CVE-2026-99508	In the Linux kernel, the following vulnerability has been resolved: media: ipu-bridge: Fix use-after-free in IVSC error path ipu_bridge_get_ivsc_acpi_dev() returns an ACPI device with a reference held. If no matching CSI device is found, ipu_bridge_check_ivsc_dev() drops that reference before passing the device to dev_err(). Concurrent CSI device and ACPI namespace removal can drop the other references, making this the final put. dev_err() then reads the freed struct acpi_device and its kobject name. This was reproduced without kprobes by racing normal CSI platform-device unregistration and ACPI configfs table removal against the production ipu_bridge_init() entry. KASAN reported the final free through device_release(), followed by invalid reads through _dev_err(). Emit the error message while the lookup reference is still held, then drop the reference as before. Source reproducers and vulnerable/fixed serial logs, with only the build hostname redacted, are available at: https://github.com/kimaiden1984-boop/linux-ipu-bridge-acpi-uaf-reproducer The fixed target builds with W=1 on x86-64. This was not tested on physical IPU hardware. The tests establish runtime reachability but do not measure occurrence frequency on shipping systems or establish exploitability.	N/A	More Details
CVE-2026-99509	In the Linux kernel, the following vulnerability has been resolved: media: imon: fix use-after-free in display_close via dev_dbg ictx->dev is a raw pointer to the usb_interface's embedded device. On disconnect, usb_disconnect() can drop the last reference on the interface and free it while a userspace fd for /dev/lcd0 is still open. When that fd is later closed, display_close() dereferences the now-freed ictx->dev via dev_dbg(), causing a use-after-free. Remove the dev_dbg() calls that dereference ictx->dev in display_close() and free_imon_context(), rather than taking a reference on the device.	N/A	More Details
CVE-2026-99511	In the Linux kernel, the following vulnerability has been resolved: media: cx23885: Fix use-after-free in cx23885_finidev due to race condition In cx23885_v4l2_dev_notify_init, &dev->cx25840_work is bound with cx23885_av_work_handler, and &dev->ir_rx_work and &dev->ir_tx_work are bound with cx23885_ir_rx_work_handler and cx23885_ir_tx_work_handler. cx23885_irq can schedule these works on system_wq when an AV_CORE interrupt is received, and the IR subdevice can also schedule the IR works from its interrupt service routine via the v4l2_device notify callback. If we remove the device, cx23885_finidev makes cleanup and the memory allocated for dev is released by kfree(dev), while the works mentioned above may still be pending or running. The sequence of operations that may lead to a UAF bug is as follows: CPU0 CPU1 | cx23885_irq | schedule_work(&dev->cx25840_work) cx23885_finidev | cx23885_input_fini(dev) | cx23885_ir_fini(dev) | cx23885_shutdown(dev) | free_irq(pci_dev->irq, dev) | pci_disable_device(pci_dev) | cx23885_dev_unregister(dev) | v4l2_device_unregister(v4l2_dev) | kfree(dev) | // dev is freed | | cx23885_av_work_handler | // use dev (use-after-free) Fix it by canceling the works after the IRQ handler that can schedule them has been stopped, and before proceeding with the remaining cleanup in cx23885_finidev. Note that the flush_work() calls in cx23885_input_ir_stop() do not close this race: they only wait for works that are already queued or running at that moment, they do not prevent the IRQ handler, which is still registered at that point, from scheduling the works again afterwards, and they are skipped entirely when dev->sd_ir is NULL. The cancel_work_sync() calls are therefore placed after free_irq(), the only point at which no new work can be scheduled.	N/A	More Details
CVE-2026-99512	In the Linux kernel, the following vulnerability has been resolved: media: au0828: fix use-after-free in bulk_timeout timer on disconnect start_urb_transfer() arms dev->bulk_timeout, whose callback au0828_bulk_timeout() queues dev->restart_streaming; that work in turn calls start_urb_transfer() and re-arms the timer. au0828_dvb_unregister() cancels restart_streaming on disconnect but does not sync bulk_timeout. au0828_usb_release() then frees dev with kfree(), so a timer still armed or pending at that point can fire after the free and dereference dev through timer_container_of(). Sync the timer with timer_shutdown_sync() before cancel_work_sync(). The timer callback queues the work, so the timer must be stopped first; the shutdown variant is needed because the work re-arms the timer through start_urb_transfer(). The dvb->frontend == NULL early return guards the register-failure path where the timer was never set up. This bug was found by static analysis.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: landlock: Fix superblock use-after-free in release_inode() Without this fix, release_inode() (task 1) can race with a umount operation in hook_sb_delete() (task 2). The order of operations in which the race occurs is: * task 2: hook_sb_delete()		

CVE-2026-99514	arrives at the end of the function, waiting for landlock_superblock(sb)->inode_refs to become zero. * task 1: release_inode(): does atomic_long_dec_and_test(), decrementing inode_refs to zero. * task 2: gets woken up through some other means, sees that inode_refs is zero and returns. The caller is the unmount operation and proceeds to release the superblock. * task 1: Runs wake_up_var(&landlock_superblock(sb)->inode_refs), *dereferencing the sb_security on the freed-up superblock*. Notably, wake_up_var() does not dereference the pointer passed to it, and only uses it to determine an overapproximated set of tasks to wake up. So while it's potentially possible for an attacker to provide a different value for that pointer, this only results in the wrong set of tasks getting woken up, and in this scenario, no further tasks are actually waiting. Nevertheless, KASAN detects it as a slab-use-after-free read on sb->s_security: BUG: KASAN: slab-use-after-free in release_inode+0x196/0x290 Read of size 8 at addr ffff888103b6c0a8 by task repro/217 release_inode+0x196/0x290 landlock_put_object+0x71/0x90 landlock_free_rules+0x38/0xc0 landlock_put_ruleset+0x47/0x130 fop_ruleset_release+0x33/0x40 __fput+0x305/0x780 Freed by task 54: kfree+0x11c/0x390 process_scheduled_works+0x766/0xe50 worker_thread+0x7bc/0xb20 The fix is to determine &landlock_superblock(inode->i_sb)->inode_refs only once earlier in release_inode(), and later use that precalculated address for the wake_up_var() call, without having to dereference through the superblock.	N/A	More Details
CVE-2026-99515	In the Linux kernel, the following vulnerability has been resolved: iommufd: Fix event queue release use-after-free iommufd_eventq_fops_release() drops the file's users reference before reading eventq->ictx. Once that reference is dropped, IOMMU_DESTROY in another thread can remove and free the event queue, so the later load accesses freed memory. Read ictx while the users reference still pins the event queue, then drop the reference and put the saved context. The context reference taken by iommufd_eventq_init() keeps ictx valid until the put. Fault and vEVENTQ file descriptors share this release path.	N/A	More Details
CVE-2026-99517	In the Linux kernel, the following vulnerability has been resolved: nvmet: fix use-after-free of the old registrant in nvmet_pr_preempt() nvmet_pr_preempt() handles a preempt sent by the reservation holder by changing the reservation type with nvmet_pr_update_reg_attr(), unless the reservation is of an All Registrants type. For the holder that helper does not modify the registrant in place: it puts a copy on the registrant list, points pr->holder to the copy and passes the old registrant to kfree_rcu(). If the reservation type changed, nvmet_pr_preempt() then hands &reg->hostid of that old registrant to nvmet_pr_resv_released(). Unless the reservation released notification is masked, that function takes subsys->lock and compares the host id against every controller of the subsystem. The caller only holds pr->pr_sem and is not in an RCU read-side critical section, so nothing keeps the old registrant alive while it is read. The result is a read of 16 freed bytes that decides which controllers get the notification. Found by code inspection. kfree_rcu() defers the free, so it usually happens after the read; when the pending free is forced to run before the read, KASAN reports the host id as freed memory. Pass the host id of the controller that sent the command instead. The registrant was looked up by that host id, and the controller outlives the request.	N/A	More Details
CVE-2026-99518	In the Linux kernel, the following vulnerability has been resolved: platform/x86: asus-wmi: fix use-after-free unregistering kbd_led led kbd_led_update_all() registers asus->kbd_led with devm_led_classdev_register() on the platform device, but kbd_led is embedded in struct asus_wmi, which asus_wmi_remove() frees at its end. The driver core calls .remove() before devres_release_all(), so the deferred devres release runs led_classdev_unregister() on memory that has already been freed. Whether that faults depends on whether the slab has been reused in between, which makes the crash on module unload intermittent and silently corrupts the led classdev state in the remaining runs. Unregister the led classdev explicitly in asus_wmi_led_exit(), after the led workqueue has been destroyed: kbd_led_work can then neither run nor be queued anymore, since asus_ref.asus is NULL too, so kbd_led_registered is stable and the classdev can be released while the memory embedding it is still alive. devm_led_classdev_unregister() also drops the devres entry, leaving nothing to run at unbind cleanup time; it warns when no devres entry exists, hence the kbd_led_registered check.	N/A	More Details
CVE-2026-99519	In the Linux kernel, the following vulnerability has been resolved: PM: EM: Fix use-after-free of perf domain in netlink doit handlers The doit handlers dropped em_pd_list_mutex before dereferencing the perf domain, so a concurrent device unregister could free it while the reply was being built. Encode the reply from em_perf_domain_for_id() callbacks, which run under the mutex, and allocate the skb outside of it, like the dumpit path already does. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-99520	In the Linux kernel, the following vulnerability has been resolved: pps: generators: fix use-after-free when closing a removed device The cdev of a PPS generator is embedded in struct pps_gen_device, but nothing ties the lifetime of that structure to the cdev: pps_gen is freed by the release function of its device, and an open file holds a device reference only until pps_gen_cdev_release() drops it. When the generator is unregistered while /dev/pps-genN is open, that put_device() drops the last reference and frees pps_gen, and __fput() then calls cdev_put() on the freed cdev: BUG: KASAN: slab-use-after-free in cdev_put+0x53/0x60 Read of size 8 at addr ffff88801383e138 by task ppsgen64/149 Call Trace: cdev_put+0x53/0x60 __fput+0x745/0xad0 fput_close_sync+0xd9/0x1b0 __x64_sys_close+0x86/0xf0 ... Freed by task 149: kfree+0x25a/0x6d0 device_release+0xca/0x3c0 kobject_put+0x169/0x320 pps_gen_cdev_release+0x51/0x80 __fput+0x36a/0xad0 pps.c had the same bug, fixed in commit c79a39dc8d06 ("pps: Fix a use-after-free"). Fix it the usual way: embed the struct device in pps_gen_device and register both with cdev_device_add(). This makes the device the parent of the cdev, so the cdev holds a device reference until the last file is closed.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: rapidio: mport_cdev: fix use-after-free in		

CVE-2026-99521	mport_mm_close() A use-after-free vulnerability was identified in mport_mm_close() in drivers/rapidio/devices/rio_mport_cdev.c, identical to the pattern previously addressed in dma_req_free(). This is observable from userspace when an application creates an mmap mapping via the RapidIO character device and subsequently unmaps it (or terminates, triggering exit_mmap()). During munmap, mport_mm_close() is invoked and drops the mapping reference via kref_put(). If kref_put() drops the last reference, mport_release_mapping() is called, which frees the underlying rio_mport_mapping structure. The subsequent mutex_unlock() then dereferences map->md to unlock buf_mutex, leading to a use-after-free: mport_mm_close() -> mutex_lock(&map->md->buf_mutex); ... -> kref_put(&map->ref, mport_release_mapping); /* map is freed */ -> mutex_unlock(&map->md->buf_mutex); /* UAF: map used */ Fix this by caching map->md before kref_put() and using the cached pointer for mutex unlocking, ensuring that freed memory is not accessed.	N/A	More Details
CVE-2026-99522	In the Linux kernel, the following vulnerability has been resolved: RDMA/irdma: avoid use-after-free in icrdma_remove icrdma_remove() dereferences iwdev after calling irdma_ib_unregister_device(): irdma_ib_unregister_device(iwdev) -> ib_unregister_device(&iwdev->ibdev) -> __ib_unregister_device() -> ib_dealloc_device() irdma registers irdma_ib_dealloc_device() as the dealloc_driver callback in its ib_device_ops. The RDMA core explicitly documents ib_unregister_device() that when ops.dealloc_driver is used, ib_dev will be freed upon return from the function. The ib_device is embedded in struct irdma_device and iwdev itself was allocated by ib_alloc_device(irdma_device, ibdev). Consequently, once irdma_ib_unregister_device() returns, iwdev must no longer be dereferenced. However, icrdma_remove() currently accesses iwdev->rf afterwards when deinitializing interrupts, destroying ah_tbl_lock, and freeing rf. Although the final ib_device storage is released with kfree_rcu(), the object lifetime has already ended and callers must not rely on the RCU grace period to continue dereferencing iwdev. Save iwdev->rf before unregistering the RDMA device and use the saved pointer for the remaining cleanup. This issue was found by manual code inspection.	N/A	More Details
CVE-2026-99524	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Fix use-after-free of a local reenq node left on an inactive cpu A local DSQ reenq is queued on the rq through a per-sched node and run from deferred work. When the reenq is requested from dispatch, the work is a balance callback queued once dispatch is done. A cpu going down installs balance push early in its teardown, and queue_balance_callback() refuses new callbacks from then on. ops.dispatch() keeps running on that cpu until the rq goes offline after an RCU grace period, so a reenq it requests in that window may never run. The IMMED handling in dispatch_one() requests one the same way. A reenq requested from elsewhere for a cpu that is already offline goes through irq_work_queue_on(), which warns, and its work doesn't run until the cpu comes back. Either way the node stays on the rq's list. When the scheduler is freed, scx_sched_free_rcu_work() warns and frees the pcpu with the node still linked, and the list is walked again once the cpu comes back. An inactive cpu's local DSQ holds only tasks that are bound to it or are being moved off, so the reenq is pointless. Ignore reenqs of an inactive cpu's local DSQ. The test is lockless. The hotplug thread of a cpu going down runs on that cpu, and it marks the cpu inactive before it installs balance push. A dispatch on the cpu runs to completion before that thread can run there, so a dispatch that saw the cpu active has queued its callback by then. Requests from outside dispatch run from the wakeup path or through irq_work, which balance push doesn't affect. The grace period after the cpu goes inactive waits for any requester that saw it active, so that irq_work is queued while the cpu is online.	N/A	More Details
CVE-2026-99525	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix use-after-free in cifs_ses_add_channel() The error path of cifs_ses_add_channel() calls cifs_put_tcp_session() before cifs_chan_clear_need_reconnect(). The latter reaches cifs_ses_get_chan_index(), which dereferences server->terminate; if the put was the final reference, TCP_Server_Info is already freed and this is a use-after-free. Same put-then-use pattern as the recently fixed sibling in commit 717e0a25036b ("cifs: Fix server use-after-free in cifs_chan_skip_or_disable()"). Move the put last, but do not read chan->server after dropping chan_lock: chan points into ses->chans[], and once chan_count is decremented a concurrent add_channel can reuse the slot and overwrite the pointer, so the put would act on the new channel's server. Cache the pointer under chan_lock, clear the reconnect bit and decrement chan_count while still holding the reference, and release it after the unlock. This issue was discovered by Tencent CodeBuddy Security.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix use-after-free infoleak via the readdir resume name CIFSFindNext() copies the resume name that cifs_save_resume_key() recorded from the last entry of the previous response into its request; psrch_inf->presume_name points into the response buffer that entry came from, and nothing ever clears it. find_cifs_entry() records a new name only from a response with a usable last entry, and cifs_readdir() only from an entry it walks to, so a response that carries no entries and leaves last_entry NULL, because CIFSFindFirst() or CIFSFindNext() rejected LastNameOffset, refreshes neither and presume_name keeps pointing into the previous response. That buffer has already been released: by CIFSFindNext() before it installs the new response, or by find_cifs_entry() when it rewinds the search, which does not drop the name either and can then fail in initiate_cifs_search() before a new response replaces it. The next CIFSFindNext() then copies up to PATH_MAX - 1 bytes of the resume name out of the released buffer into its request and sends them to the server: a use-after-free read that leaks kernel memory to a malicious or compromised server, which can provoke it with a single FindNext response that carries no entries and a LastNameOffset above CIFSMaxBufSize, which is all the existing check rejects. With KASAN enabled, a server that answers a listing that way gives: CIFS: VFS: ignoring corrupt resume name == BUG: KASAN: slab-use-after-free in CIFSFindNext+0x8ca/0x14c0 Read of size 4080 at addr ffff88807c823f98		

CVE-2026-99526	by task ls/83 CPU: 0 UID: 0 PID: 83 Comm: ls Tainted: G B 7.3.0-rc4-00457-gf14572c203d5 #69 PREEMPT(lazy) Tainted: [B]=BAD_PAGE Call Trace: <TASK> kasan_report+0xdf/0x1a0 kasan_check_range+0x10f/0x1e0 __asan_memcpy+0x23/0x60 CIFSFindNext+0x8ca/0x14c0 cifs_readdir+0xf51/0x29c0 iterate_dir+0x1c0/0x570 __x64_sys_getdents64+0x133/0x270 do_syscall_64+0x109/0x5d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 83 on cpu 1 at 13.497119s: cifs_buf_get+0x36/0x90 smb_init+0x4f/0x110 CIFSFindNext+0xf7/0x14c0 cifs_readdir+0xf51/0x29c0 Freed by task 83 on cpu 0 at 13.584018s: cifs_buf_release+0x41/0x80 CIFSFindNext+0xfce/0x14c0 cifs_readdir+0xf51/0x29c0 The buggy address is located 16280 bytes inside of freed 16588-byte region [ffff88807c820000, ffff88807c8240cc) == The same memcpy() first reads past the end of that object while it is still live, a slab-out-of-bounds read that the following patches bound; that is the taint the report above carries. SMB1 is not negotiated by default; reaching this code requires an explicit vers=1.0 mount. Clear presume_name and resume_name_len whenever a new response is installed and when the rewind path releases the search buffer, so that the name never outlives the buffer it points into, and reset resume_key with them, since it was taken from the same entry and would otherwise be sent with an empty name. A response without a usable last entry then continues the search with an empty resume name rather than a stale one; CIFSFindNext() skips the copy of a zero-length name so that it never runs on the NULL pointer. The next patch in the series leaves last_entry NULL for a response that carries no entries, and bounds LastNameOffset against the received response rather than against CIFSMaxBufSize. Responses that are given a last entry today then take the path fixed here, so that patch must not be applied without this one.	N/A	More Details
CVE-2026-99527	In the Linux kernel, the following vulnerability has been resolved: spmi: Fix potential use-after-free by grabbing of_node reference As noticed by Sashiko during a review run of an unrelated patch, in of_spmi_register_devices(), for_each_available_child_of_node() is used to loop through children, and to also assign a node to a newly created SPMI child device. Problem is that the refcount is dropped at every iteration so, in the specific case of DT overlays, a use-after-free may occur when an overlay is dynamically unloaded! To resolve this, increase the of_node refcount when assigning (in function of_spmi_register_devices) and release the reference in spmi_device_remove().	N/A	More Details
CVE-2026-99500	In the Linux kernel, the following vulnerability has been resolved: s390/pci: Fix resource leak in zpci MSI setup If airq_iv_create() fails in __alloc_airq(), the zpci_sbv bit allocated by airq_iv_alloc_bit() is never freed. This permanently leaks one of the ZPCI_NR_DEVICES summary bits, reducing system capacity with each failed device hotplug. In systems with repeated device insertion failures or under memory pressure, all summary bits can be exhausted, preventing new PCI devices from being added until reboot. Add proper error handling to free the zpci_sbv bit and reset zdev->aisb if the AIBV creation fails.	N/A	More Details
CVE-2026-99498	In the Linux kernel, the following vulnerability has been resolved: s390/kprobes: Prevent kprobes on instructions with exception table entry The mvcos exception handler ex_handler_ua_mvcos() decodes the faulting instruction, assuming it is an mvcos instruction. In case the instruction is kprobed the decoded instruction is a breakpoint instruction instead, which leads to incorrect instruction decoding and potential register corruption. Fix this by simply preventing to set a kprobe on such instructions, similar like arm64 is doing it.	N/A	More Details
CVE-2026-99466	In the Linux kernel, the following vulnerability has been resolved: media: microchip-csi2dc: fix async notifier leak on probe error path csi2dc_of_parse() ends up calling csi2dc_prepare_notifier(), which registers the V4L2 async notifier with v4l2_async_nf_register(). The probe error label csi2dc_probe_cleanup_notifier is reached after csi2dc_of_parse() has already succeeded (e.g. when media entity init, csi2dc_power() or v4l2_async_register_subdev() fail), but it calls only v4l2_async_nf_cleanup() and never v4l2_async_nf_unregister(). The notifier therefore stays chained in the global notifier_list while the enclosing struct csi2dc_device is freed, leading to list corruption and a use-after-free when the list is next walked. Unregister the notifier before cleaning it up on the error path, matching the teardown already done in csi2dc_remove().	N/A	More Details
CVE-2026-99497	In the Linux kernel, the following vulnerability has been resolved: s390/probes: Fix wrong opcode for cgrb fixup The last byte of the cgrb instruction is 0xe4, however the probes code compares against 0xe6. Fix it.	N/A	More Details
CVE-2026-99467	In the Linux kernel, the following vulnerability has been resolved: media: marvell: mcam: stop DMA and cancel s_bh_work before freeing DMA buffers In vmalloc mode the frame-completion IRQ queues mcam_frame_work() on cam->s_bh_work, which memcpy()s from the dma_bufs[] that mcam_free_dma_bufs() frees with dma_free_coherent(). mccic_shutdown() frees those buffers without stopping the controller, so while streaming a late frame IRQ can re-arm the work after the buffers are gone, causing a use-after-free. Stop the controller in mccic_shutdown() when streaming is still active (gated on an open fd, which holds a runtime-PM reference, so the device is powered), cancel s_bh_work in mcam_free_dma_bufs(), and move INIT_WORK() into mccic_register() before the device can be published.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: media: gspca: use vb2_video_unregister_device() on disconnect gspca uses vb2_fop_release(), but gspca_disconnect() only unregisters the video device. If the queue owner closes the device after USB disconnect, vb2_fop_release() releases the queue and invokes gspca_stop_streaming() at that point. gspca_stream_off() calls subdriver stop		More

2026-99469	callbacks and usb_set_interface() through gspca_dev->dev. By the time the late close runs, the USB core can have freed the usb_device, resulting in a use-after-free. This was observed as a KASAN use-after-free in sd_stopN(). Use vb2_video_unregister_device() so the queue is released and streaming is stopped synchronously during disconnect, while the usb_device is still valid. Since the helper takes the queue lock, which is usb_lock, call it after dropping that lock.	N/A	Details
CVE-2026-99470	In the Linux kernel, the following vulnerability has been resolved: media: cx23885: unregister DVB bus when SP2 CI setup fails The DVBSKY S950C, DVBSKY T980C and TechnoTrend CT2-4500 CI paths register the DVB bus before attaching their SP2 CI client. If that client cannot be created, bound, or pinned, the current error path removes the I2C frontend components and only deallocates the frontend list, leaving the registered DVB adapter behind. Removing an I2C demod before unregistering the DVB bus can also expose the use-after-free ordering fixed for the normal remove path. Unregister the DVB bus immediately when SP2 CI setup fails, before the existing I2C client cleanup. The bus helper empties the frontend list, so the later frontend deallocation has nothing left to release. Keep failures from vb2_dvb_register_bus() on the existing cleanup path. Also release an unbound SP2 client and propagate the actual I2C creation or CI registration error. This issue was identified during our ongoing static-analysis research while reviewing kernel code.	N/A	More Details
CVE-2026-99474	In the Linux kernel, the following vulnerability has been resolved: HID: hiddev: keep state alive through disconnect unlock mutex_unlock() clears the mutex owner before taking its wait lock. A thread spinning in the final hiddev file release can acquire existancelock after hiddev_disconnect() clears the owner, then free hiddev before the disconnecting thread reaches the mutex wait lock. This causes a use-after-free in the mutex slow unlock path. Give the connection and each open file an explicit reference. Drop each reference only after its existancelock critical section has completed, so neither unlock path can free the mutex while the other is still using it.	N/A	More Details
CVE-2026-99475	In the Linux kernel, the following vulnerability has been resolved: genirq: Defer kfree(irqaction) until irq_thread_dtor() finishes When an interrupt thread exits abnormally via do_exit() (e.g. the interrupt handler oopses) without returning from irq_thread(), a concurrent call to free_irq() could free the irqaction before the irq_thread_dtor() exit task work runs. irq thread (oopsed) free_irq() ---------- ---------- do_exit() exit_mm() complete_vfork_done() -------> kthread_stop() returns exit_files() ... kfree(action) exit_task_work() irq_thread_dtor() action->... <- use-after-free Track the pending exit work with the IRQTF_EXIT_WORK thread flag, and have __free_irq() wait for the bit to clear before kfree(). irq_thread_dtor() only runs on irq_thread abnormal exit. If that dtor also faults, free_irq() will hang indefinitely, but double IRQ thread faults require a reboot anyways. [tglx: Add missing include]	N/A	More Details
CVE-2026-99479	In the Linux kernel, the following vulnerability has been resolved: fbdev: atmel_lcdfb: Prevent FIFO work after removal The LCD controller IRQ schedules sinfo->task to recover from a FIFO underflow. Probe currently requests the IRQ before initializing the work, so an early interrupt can schedule an uninitialized work item. The probe error path and remove path cancel the work before freeing the IRQ. A FIFO-underflow interrupt can schedule the work again after that cancellation. framebuffer_release() may then free sinfo while the work is still pending. KASAN reported a use-after-free in atmel_lcdfb_task() when an underflow interrupt occurred during removal. With the IRQ freed before the work was cancelled, the same test completed without a kernel diagnostic. Initialize the work before requesting the IRQ. In both cleanup paths, free the IRQ before cancelling the work so that it cannot be scheduled again after cancellation.	N/A	More Details
CVE-2026-131920	In OpenStack Cinder before 29.1.1, by making an os-initialize_connection request with an IQN that includes a path traversal as part of the connector's initiator element, an authenticated user with a 'member' role may create or overwrite files in the volume host filesystem. Only cinder-volume hosts using the LIO iSCSI target driver are affected. This target driver is typically used with the LVM volume driver.	N/A	More Details
CVE-2026-99481	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_vhci: Pin module for debugfs file operations An open debugfs file retains its file_operations pointer in the debugfs proxy state. However, force_suspend_fops, force_wakeup_fops, aosp_capable_fops, and force_devcoredump_fops have no .owner set to THIS_MODULE. As a result, full_proxy_open_regular() calls fops_get() which succeeds without taking a reference to the hci_vhci module. When userspace opens any of these debugfs files and later closes the /dev/vhci file descriptor (triggering vhci_release() and debugfs file removal), the open debugfs file descriptor remains valid in userspace. Because the module refcount was not incremented, the hci_vhci module can be unloaded via delete_module (rmmod hci_vhci). When the userspace process subsequently closes the debugfs file descriptor, full_proxy_release() accesses the already unmapped file_operations structure, resulting in a kernel page fault / oops: BUG: unable to handle page fault for address: ffffffbfff846f18b #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page Oops: Oops: 0000 [#1] SMP KASAN NOPTI RIP: 0010:full_proxy_release+0x8e/0x140 Call Trace: <TASK> __fput+0x39f/0xa60 fput_close_sync+0xff/0x200 __x64_sys_close+0x8c/0xf0 do_syscall_64+0xc3/0x590 entry_SYSCALL_64_after_hwframe+0x76/0x7e </TASK> Modules linked in: ... [last unloaded: hci_vhci(E)] The issue was discovered via manual code audit of drivers/bluetooth/hci_vhci.c and verified on a Linux 7.3-rc6 test VM with CONFIG_BT_VHCI=m by keeping a debugfs file open across controller teardown and module unload. Set .owner = THIS_MODULE on all four file_operations structures so that debugfs pins the module until all file descriptors are closed.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ASoC: amd: ps: fix snd_acp63_remove()		

CVE-2026-99482	teardown ordering The ACP IRQ handlers dereference the SoundWire, SoundWire DMA, and PDM child platform devices, but snd_acp63_remove() unregisters them while the interrupt is still held by devres, which frees it only after remove returns. An interrupt in this window is a use-after-free. Mask the ACP interrupt sources and call devm_free_irq() before the first child device is unregistered. The masking goes through the new acp_hw_ops->disable_interrupts callback, keeping the remove path platform-agnostic. The interrupt line is shared, and acp_hw_deinit() clears the sources only after the children are gone, which would leave the line raised with no handler left to ack it. The window predates the tagged refactor, which only reshaped the dereferences. devm_free_irq() waits for in-flight handlers, but not for work the hardirq has already queued: acp63_irq_handler() schedules amd_sdw_irq_thread on the SoundWire manager. Draining that work in the manager's remove path is a soundwire-side change and follows separately. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-99483	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcmtest: handle multiple wraps in random capture fill The random capture helpers assume that the data generated for one timer tick wraps the DMA ring at most once. However, b_rw describes 200 ms of audio because the timer runs at 5 Hz, so it can be several times larger than a small ring. For 48 kHz, S16_LE, four-channel capture with a 512-frame buffer, the DMA ring is 4096 bytes while one tick generates 76800 bytes. After filling the 4096-byte tail, the existing code issues a second 72704-byte write from the ring base, writing 68608 bytes beyond the ring. Both interleaved and non-interleaved capture are affected. KASAN reports the first write into the adjacent free page as a use-after-free in the random-byte generator. The underlying defect is an out-of-bounds write from the live DMA ring, not a stale DMA allocation. Split each random write repeatedly at the interleaved ring or per-channel block boundary. Keep the hardware-pointer update exactly once after the complete logical write. The issue reproduced on three of three clean boots. With this change, the same vulnerable geometry completed six times without a sanitizer finding. Boundary, pattern-fill, no-timer, open/close, and module-reload controls also completed normally. snd-pcmtest is a test driver, and selecting random fill requires an administrator-controlled module parameter. The written values come from the kernel random generator; no privilege escalation or attacker-controlled write primitive was demonstrated.	N/A	More Details
CVE-2026-99486	In the Linux kernel, the following vulnerability has been resolved: rust: configfs: fix data offset calculation for subsystem callbacks `get_group_data()` chooses between `Group::<Parent>::container_of()` and `Subsystem::<Parent>::container_of()` based on whether `this` represents the root group of a configfs subsystem. It detects this by checking `(*this).cg_subsys.is_null()`, but `link_group()` in `fs/configfs/dir.c` unconditionally sets `cg_subsys` for every `config_group` attached anywhere in a registered subsystem, including the subsystem's own `su_group`. The only `config_group` with a NULL `cg_subsys` is the configfs root, on which userspace cannot trigger callbacks. The check is therefore always false at runtime, and the `Subsystem` branch is dead. Subsystem-level callbacks reach the `Group` branch and may read `data` at the wrong offset. Thus change the `is_root` check to correctly identify whether a group is a root group by comparing `this` to `subsys.su_group`. [Andreas: Fix clippy warning on 1.85.0 by moving safety comment.]	N/A	More Details
CVE-2026-99487	In the Linux kernel, the following vulnerability has been resolved: rust: bug: pass TAINT_WARN to warn_slowpath_fmt() on UML On UML, `warn_flags!` calls `warn_slowpath_fmt()` directly. Its third argument is a taint number, but `warn_on!` passes its bug flags there. As a result, `add_taint()` gets 2304 instead of 9 (`TAINT_WARN`). The warning is not recorded: `/proc/sys/kernel/tainted` stays 0 after `warn_on!` fires, where it should be 512. It also causes an out-of-bounds write in `add_taint()`. Pass `TAINT_WARN` directly, as C does on UML.	N/A	More Details
CVE-2026-99488	In the Linux kernel, the following vulnerability has been resolved: klist: avoid accesses after waking klist_remove() klist_remove() waits until a node is unreferenced so that its caller can free the containing object. klist_release() currently publishes waiter->woken and wakes the waiter before its final accesses to the waiter and node. klist_remove() can then return, allowing its stack waiter and the containing object to be freed or reused while klist_release() is still running. In particular, bus_remove_driver() can free drv->p while __device_attach() walks the same bus klist with bus_for_each_drv(). On an arm64 Cortex-A72 system, an unpatched 7.2.0-rc3 kernel with CONFIG_PREEMPT_RT=y and CONFIG_KASAN=y reproduced the bug through the in-tree I2C/at24 path. KASAN reported a use-after-free in klist_dec_and_del() reached from klist_next()/bus_for_each_drv() while at24 was being unregistered. Clear n_klist and take a task reference before publishing woken. Use release/acquire accesses for that publication and wake the referenced task. The task reference keeps the waiter task alive if it returns and exits before wake_up_process().	N/A	More Details
CVE-2026-99489	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: Redesign shared sibling data lifecycle The Wacom driver coordinates state between sibling interfaces of the same physical device using a shared structure 'wacom_shared' inside 'wacom_hdev_data'. The driver kept a volatile representative pointer 'data->dev' pointing to a sibling 'hid_device' for physical path comparisons during sibling matching. This pointer management is fragile. When the representative device is disconnected, wacom_remove_shared_data() failed to clear/update 'data->dev', leading to a Use-After-Free vulnerability when subsequent sibling probes dereference the dangling 'data->dev' pointer. Resolve this issue by redesigning the sibling data lifecycle: - Eliminate the volatile 'data->dev' representative pointer completely - Redesign 'wacom_hdev_data' to store stable static copies of the required attributes upon first allocation: 'phys' path string, 'vendor', 'product' IDs and the sibling's 'device_type' - Use these static attributes for stable sibling matching in wacom_are_sibling() and wacom_get_hdev_data() This ensures sibling matching remains safe and stable even if individual siblings are dynamically added or removed. To secure the lifecycle against concurrent probe/disconnect races: - Switch kref_put() to kref_put_mutex() in wacom_remove_shared_data()	N/A	More Details

	to serialize refcount drops with list traversal and lookup - Modify wacom_release_shared_data() to assume the list lock is already held Also, do not accumulate the 'device_type' capability flag during subsequent sibling probes. Keeping only the first probed sibling's device_type exactly preserves the original sibling matching behavior without introducing side effects.		
CVE-2026-99490	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: Fix Use-After-Free in wacom_bamboo_pad wacom_bamboo_pad_pen_event() accesses wacom->shared->pen locklessly relative to wacom_remove_shared_data() which nullifies it. This can lead to a Use-After-Free if the sibling device is removed while events are being processed. Resolve this by introducing RCU protection for pen and touch pointers: - Annotate 'pen' and 'touch' in wacom_shared struct with __rcu. - Wrap lockless readers in wacom_bamboo_pad_pen_event() with rcu_read_lock() and rcu_dereference(). - Update writers in wacom_sys.c using rcu_assign_pointer(). - Use rcu_dereference_protected for comparisons under wacom_udev_list_lock. - Also use rcu_access_pointer in wacom_sibling_pending() and wacom_mode_change_work() to avoid warnings (while lockless access in the latter remains a pre-existing issue).	N/A	More Details
CVE-2026-99491	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: Fix Use-After-Free in wacom_intuos_pad wacom_intuos_pad() accesses wacom->shared->touch_input locklessly inside the interrupt handler context. If the Touch sibling device is disconnected, wacom_remove_shared_data() clears 'touch_input' outside any lock, creating a Time-of-Check to Time-of-Use (TOCTOU) race condition where a preempted reader in interrupt context dereferences the freed pointer, leading to a Use-After-Free. Resolve this by introducing RCU protection for the touch_input pointer: - Annotate 'touch_input' in wacom_shared struct with __rcu - Wrap all lockless readers in wacom_wac.c with guard(rcu)() and rcu_dereference() using a unified wacom_report_touch_mute() helper - Update writers in wacom_sys.c using rcu_assign_pointer() - Call synchronize_rcu() in wacom_remove_shared_data() to ensure all active RCU readers have finished before the input device is freed Also wrap wacom_set_shared_values() and touch/pen assignments in wacom_add_shared_data() inside the wacom_udev_list_lock to serialize concurrent probe assignments, and verify that 'shared->touch == hdev' before setting touch_input to prevent concurrent sibling probe state desynchronization.	N/A	More Details
CVE-2026-99492	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: Advertise SW_MUTE_DEVICE capability prior to registration Input subsystem guidelines require that device capabilities are advertised before the input device is registered. The Wacom driver was violating this by advertising the SW_MUTE_DEVICE capability post-registration in wacom_set_shared_values() (and duplicating it in device-specific setup cases). Resolve this by moving the SW_MUTE_DEVICE capability setup to wacom_setup_touch_input_capabilities() for all touch devices that support it. For generic touch devices whose capabilities depend on mute switch usages parsed from a sibling Pen/Pad interface, defer registration with -EPROBE_DEFER until the sibling has parsed its descriptors and initialized shared capabilities.	N/A	More Details
CVE-2026-99493	In the Linux kernel, the following vulnerability has been resolved: s390/zcrypt: Guard domain index uses against speculative bypass The previous array_index_nospec() placement only covered the admin permission check, leaving ap_test_config_usage_domain() and ap_test_config_ctrl_domain() in the CCA path exposed to speculative execution. Move the sanitization to before all permission and config checks in both the CCA and EP11 paths. Hold the sanitized domain value in another variable so that the later use can choose to use either the original value (subject to speculative execution) or the sanitized value (which may in case of AUTOSEL_DOM reflect the wrong value). Before that, check the domain value to be either AUTOSEL_DOM or in range 0...AP_DOMAINS-1 and return with -EINVAL in case this check does not pass.	N/A	More Details
CVE-2026-99494	In the Linux kernel, the following vulnerability has been resolved: s390/zcrypt: Fix out-of-bounds ptr advance in cca_query_crypto_facility() The FQ reply parser code blindly advanced the walk pointer by a length value read directly from the hardware reply payload without checking that the advance stayed within the allocated reply buffer. A corrupt or malicious device response could push ptr beyond the cprbmem region, causing an out-of-bounds dereference or kernel memory exposure via the subsequent memcpy(). Fix this by tracking the remaining reply buffer space in a variable and validating each device-supplied length field against it before advancing or dereferencing the pointer.	N/A	More Details
CVE-2026-99495	In the Linux kernel, the following vulnerability has been resolved: s390/zcrypt: Fix and improve zcrypt reply message verification checks Add or improve checks related to buffer sizes and reply sizes to the handling of replies from the crypto cards for CCA, EP11 (AP message type 6) and ICA (AP type 50) messages. The verification code related to reply field length was not designed well and thus firmware deficiencies could lead to unexpected behavior in the zcrypt device driver. Thus improve the code to more closely inspect especially length fields at message replies. Add length validation before accessing reply message structures in zcrypt_msgtype6_receive(), zcrypt_msgtype6_receive_ep11(), and zcrypt_msgtype50_receive() to prevent out-of-bounds reads and potential kernel memory disclosure. Also rework these three functions to validate reply lengths more carefully before copying data back into the request buffer. Use size_t for length calculations, reject inconsistent reply sizes, and add defensive handling for short invalid replies. For XCRB replies, validate both reply segments and derive the effective message length from the covered range instead of trusting only the second segment. Additionally improve some of the later invoked evaluation functions which further check the payload. Add comments about length assumptions and for the type50 processing rework the payload processing completely with removing a misplaced BUG_ON().	N/A	More Details

CVE-2026-99496	In the Linux kernel, the following vulnerability has been resolved: s390/vmcp: Initialize full session info to zero The cma_alloc bit as well as the response can contain old data. cma_release will not free the kmalloced memory in that case and the response code might leak kernel stack data to CAP_SYS_ADMIN users.	N/A	More Details
CVE-2026-118630	Winmate HotTab Exposed Dangerous Function Local Privilege Escalation Vulnerability. This vulnerability allows local attackers to escalate privileges on affected installations of Winmate devices. An attacker must first obtain the ability to execute low-privileged code on the target system in order to exploit this vulnerability. The specific flaw exists within the HotTab driver. The issue results from an exposed dangerous function. An attacker can leverage this vulnerability to escalate privileges and execute arbitrary code in the context of the kernel. Was ZDI-CAN-29948.	N/A	More Details