

Security Bulletin 28 October 2026

Generated on 28 October 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-2026-110063	Improper link resolution before file access ('link following') in VM Inspector for Azure Virtual Machines allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-79848	Vulnerabilities in the affected interface of AOS-8 could allow an unauthenticated remote attacker to conduct SQL injection attacks against an affected AOS-8 instance. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	10.0	More Details
CVE-2026-114383	A weakness has been identified in TOTOLINK EX200 5.2.3c.7200. This impacts the function setWiFiAcIAddConfig of the file /cgi-bin/cstecgi.cgi of the component wireless.so. This manipulation of the argument macAddress causes stack-based buffer overflow. The attack can be initiated remotely. The exploit has been made available to the public and could be used for attacks.	10.0	More Details
CVE-2026-114382	A security flaw has been discovered in TOTOLINK EX200 5.2.3c.7200. This affects the function setStaticDhcpConfig of the file /cgi-bin/cstecgi.cgi of the component lan.so. The manipulation of the argument Comment results in stack-based buffer overflow. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks.	10.0	More Details
CVE-2026-111815	Apache Traffic Server can resume a TLS session across SNI virtual hosts without re-verifying a client certificate required by the resumed host's verify_client STRICT policy. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	10.0	More Details
CVE-2026-111814	Apache Traffic Server's default outbound session-sharing policy can reuse a pooled TLS session authenticated with one client certificate for a transaction configured to present a different one. This issue affects Apache Traffic Server: from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	10.0	More Details
CVE-2026-111812	Apache Traffic Server's origin_server_auth plugin lets a client force signing of attacker-chosen AWS request headers under the proxy's own credentials. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	10.0	More Details
	A flaw has been found in TOTOLINK A950RG 4.1.2cu.5268_B20221104. Affected by this issue is the		

CVE-2026-114575	function UploadCustomModule of the file lib/cste_modules/product.so. Executing a manipulation of the argument FileMd5 can lead to command injection. The attack can be launched remotely. The exploit has been published and may be used.	10.0	More Details
CVE-2026-111813	Apache Traffic Server's experimental magick plugin lets an unauthenticated client control ImageMagick convert arguments in its default (unsigned) configuration. This issue affects Apache Traffic Server: from 8.0.1 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	10.0	More Details
CVE-2026-79846	Command injection vulnerabilities exist in the affected interface of AOS-8 that could allow an unauthenticated remote attacker to run arbitrary commands on the underlying host if certain preconditions outside of the attacker's control are met. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	10.0	More Details
CVE-2026-79847	Buffer overflow vulnerabilities in the affected interface of AOS-10 could lead to unauthenticated remote code execution by sending specially crafted packets to the affected interface. Successful exploitation of this vulnerability results in the ability to execute arbitrary code as a privileged user on the underlying operating system.	10.0	More Details
CVE-2026-79844	Authentication bypass vulnerabilities in the affected interface of AOS-10 could allow an unauthenticated remote attacker to run arbitrary commands on the underlying host if certain preconditions outside of the attacker's control are met. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	10.0	More Details
CVE-2026-79849	Vulnerabilities in the affected interface of AOS-8 could allow an unauthenticated remote attacker to conduct SQL injection attacks against an affected AOS-8 instance. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	10.0	More Details
CVE-2026-111832	Unrestricted Upload of File with Dangerous Type vulnerability in WP Mailster WP Mailster wp-mailster allows Upload a Web Shell to a Web Server.This issue affects WP Mailster: from n/a through 1.9.1.0.	10.0	More Details
CVE-2026-83938	Missing authentication for critical function in Azure Kubernetes Configuration Service allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-113744	A security vulnerability has been detected in Tenda AC19 16.03.10.07. The impacted element is the function formSetFirewallCfg of the file /goform/SetFirewallCfg of the component httpd Web Management Service. Such manipulation of the argument firewallEn leads to stack-based buffer overflow. The attack may be launched remotely. The exploit has been disclosed publicly and may be used.	10.0	More Details
CVE-2026-79850	Vulnerabilities in the affected interface of AOS-8 could allow an unauthenticated remote attacker to bypass existing authentication controls and run arbitrary commands on the underlying host. Successful exploitation could allow an attacker to execute arbitrary commands as root on the underlying operating system leading to complete system compromise.	10.0	More Details
CVE-2026-97785	Improper authentication in AAD App Proxy allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-79851	Buffer overflow vulnerabilities exist in the affected interface of AOS-10 that could lead to unauthenticated remote code execution. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system, leading to complete system compromise.	10.0	More Details
CVE-2026-110264	An issue has been discovered in GitLab CE/EE affecting all versions from 7.0 before 18.11.13, 19.0 before 19.0.10, 19.1 before 19.1.9, 19.2 before 19.2.9, 19.3 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, under certain conditions, could have allowed an unauthenticated user to execute arbitrary code on the GitLab server due to integer overflow when compiling a specially crafted regular expression.	10.0	More Details
CVE-2026-115375	Incus 7.0.0 before 7.0.2 and 7.1.0 before 7.5.0 contains a path traversal vulnerability that allows clients able to create instances to write files on the host. The flaw stems from unvalidated dependent custom volume names in instance backup index.yaml files. Attackers can import a crafted instance backup with ../ sequences in dependent volume names to create root-owned directories with controlled content, leading to full host compromise.	9.9	More Details
CVE-2026-115377	Incus before 6.0.7, 6.1.0 before 7.0.2, and 7.1.0 before 7.5.0 contains an arbitrary file write vulnerability that allows malicious migration sources to write host files by planting symlinks in transferred volumes. Attackers can craft rsync or btrfs receive streams with a symlink such as rootfs	9.9	More Details

	pointing to the host root, writing files as root for full host compromise.		
CVE-2026-107861	The system control service on affected Kiloview E3 NDI encoder firmware fails to sanitize the hostname parameter accepted by its device naming endpoint, passing it unquoted into a shell command. An authenticated web user with access to this function can embed characters to execute arbitrary operating system commands as root, resulting in full device compromise.	9.9	More Details
CVE-2026-115384	Wasmtime 39.0.0 before 48.0.4 and 49.0.0 before 49.0.2 contains a stack buffer overflow vulnerability that allows malicious guests to write ~16KB to the native stack via unvalidated async callback result types. Attackers can supply components with mistyped component-model-async callback functions to overwrite return addresses, crashing the host or escaping the sandbox.	9.9	More Details
CVE-2026-114612	LiteLLM is a proxy server (AI Gateway) to call LLM APIs in OpenAI (or native) format. From 1.87.0 until 1.100.4, 1.101.3, 1.102.2, 1.103.1, and 1.104.0-rc.2, the proxy reuses one cryptographic context for at-rest secrets and session tokens, allowing an authenticated internal user to have crafted metadata containing a forged administrator credential encrypted and returned by the API-key flow. Replaying that ciphertext as a bearer token causes the proxy to trust the forged proxy administrator identity, permitting administrative access and arbitrary commands through the MCP stdio endpoint. Versions 1.87.0 through 1.90.x require EXPERIMENTAL_UI_LOGIN to be enabled, while versions after 1.91.0 are exploitable in the default configuration. This issue was fixed in LiteLLM 1.100.4, 1.101.3, 1.102.2, 1.103.1, and 1.104.0-rc.2.	9.9	More Details
CVE-2026-110243	The network management service on affected Kiloview E3 NDI encoder firmware fails to validate the network interface name accepted by its link maintenance address configuration endpoint before passing it to a shell command, unlike sibling handlers that perform the same validation. An authenticated web user with permission to configure network settings can exploit this to execute arbitrary operating system commands as root.	9.9	More Details
CVE-2026-114625	VeloCloud Orchestrator (VCO) on-prem has a security issue where this issue may allow a remote attacker to access privileged internal functionality and impact the VCO host. Successful exploitation may compromise the confidentiality, integrity, and availability of the orchestrator and data managed by the orchestrator. On an affected system, any authenticated user can gain operator privileges and masquerade as a local "superuser operator" user. This includes access to functionality that is intended for internal use only and which is not intended to be remotely accessible. Hosted, including Dedicated, versions of VCO were impacted and have already been patched.	9.9	More Details
CVE-2026-115889	Wasmtime is a runtime for WebAssembly. Prior to 36.0.19 and 43.0.0, the Cranelift compiler backend bit-packs SSA value indices in a way that can silently overflow in release builds when compiling a sufficiently large or specially crafted Wasm function with many SSA values. The overflow can make Cranelift reference incorrect IR values, causing arbitrary miscompilation, type confusion, and other attacker-controlled effects that may escape the Wasm sandbox on any architecture using the default Cranelift backend. This issue is fixed in versions 36.0.19 and 43.0.0.	9.9	More Details
CVE-2026-115893	Wasmtime is a runtime for WebAssembly. Prior to 36.0.19, 48.0.7, and 49.0.4, the Cranelift compiler backend on aarch64 can compile a Wasm br_table instruction with more than approximately 256,000 targets so that it has an incorrect conditional branch target. The branch can then jump to arbitrary other code at a user-controlled offset, allowing a user-controlled Wasm module to escape the sandbox on aarch64 engines using the default Cranelift backend. This issue is fixed in versions 36.0.19, 48.0.7, and 49.0.4.	9.9	More Details
CVE-2026-115931	Postiz lets a signed-in user ask the backend to run a function on one of their connected channels. The backend did not restrict which functions can be requested, so a user could call internal helper functions that were never meant to be reachable from the web interface. One of those helpers reads a file from the server's disk. A user who connects a channel whose stored token is a file path can have the backend return the contents of any file it can read, including its environment, which holds the secret used to sign login sessions, the database address and the credentials of connected services. Another helper makes the backend process exit with a single request. Self-registration is open by default, so no existing account and no action by another user is needed.	9.9	More Details
CVE-2026-53942	Conda Organization Infrastructure is a centralized location for infrastructure codes. Prior to commit 9c0f238, a GitHub Actions Pwn Request vulnerability appears to remain present in conda/conda. The affected workflow is .github/workflows/update.yml. It is triggered by issue_comment events containing @conda-bot render or @conda-bot recreate. For the render path, the workflow obtains the pull request head repository and branch from the PR API, checks out that PR-controlled repository/ref using secrets.SYNC_TOKEN, and then runs conda/actions/template-files@7f6830b1428a9bd47f0b068892c77eae95207037. The external template-files action reads its default configuration and templates from the checked-out working directory. In the confirmed finding, that action uses an unsandboxed Jinja2 environment to render attacker-controlled template	9.9	More Details

	content. This creates a Pwn Request chain in which an untrusted PR/comment path can reach template rendering in a privileged token context. This issue has been patched via commit 9c0f238.		
CVE-2026-115933	Postiz signs its session, invitation and password reset tokens with one shared secret and without any marker of what a token is for. A route that refreshes the credentials of a browser-extension channel signed data taken from the request itself and stored the result where the requesting user can read it back. A user could therefore obtain a validly signed token containing any claims they choose, which the application accepts as a session, as an organization invitation with a chosen role, as a password reset or as an enterprise provisioning token. A second, narrower route that creates team invitations signs extra fields from the request body in the same way.	9.9	More Details
CVE-2026-116231	A security flaw has been discovered in Oray Pugongying X5 Pro 7.6.1. The impacted element is an unknown function of the component Ser2net. Performing a manipulation of the argument regmsg results in crlf injection. It is possible to initiate the attack remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	9.9	More Details
CVE-2026-76531	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76531 are related to files or directories accessible to external parties and are grouped under the Common Weakness Enumeration (CWE) CWE-552.	9.9	More Details
CVE-2026-76302	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76302 are related to improper neutralization and are grouped under the Common Weakness Enumeration (CWE) CWE-707.	9.9	More Details
CVE-2026-76299	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76299 are related to improper access control and are grouped under the Common Weakness Enumeration (CWE) CWE-284.	9.9	More Details
CVE-2026-71195	In OpenStack Manila before 23.0.1, multiple Manila API operations on public shares do not verify that the caller owns the share. The API controllers authorize operations against the caller's own project context, then load a foreign public share without re-authorizing against the share's actual owner. This affects the following operations: allow_access: grant data-plane read/write access to another tenant's share; deny_access: revoke the owner's existing access rules; update_access: escalate access level (e.g., ro to rw); create_share_backup: exfiltrate share data via a backup created under the attacker's project; soft_delete: move a foreign share to the recycle bin; restore: interfere with the owner's deletion workflow; export location metadata: create, update, and delete metadata on foreign export locations; access rule listing: disclose IPs, usernames, and CephFS keys. Any authenticated project member can exploit these operations on any public share. The share UUID is visible since the share is public.	9.9	More Details
CVE-2026-111834	Unrestricted Upload of File with Dangerous Type vulnerability in Kitae Park Mang Board WP mangboard allows Upload a Web Shell to a Web Server.This issue affects Mang Board WP: from n/a through 2.4.4.	9.9	More Details
CVE-2026-115188	An unrestricted file upload vulnerability exists in the campaign import feature of Mautic 7. When processing an uploaded campaign archive, the application copies the bundled asset files into a publicly accessible media directory without restricting their file names or types. An authenticated user with permission to create campaign imports can exploit this through the user interface or the API (/api/campaigns/import) to place arbitrary files, including server-side scripts, in the web root. Administrator privileges are not required. On web servers that honor per-directory configuration, such as the default Apache setup of the official Mautic Docker image, this leads to remote code execution with the privileges of the web server user, including access to the Mautic configuration, application secrets, and contact data.	9.9	More Details
CVE-2026-116232	A weakness has been identified in Oray Pugongying X5 Pro 7.6.1. This affects the function snmp_set of the component SNMP Handler. Executing a manipulation of the argument info.community can lead to crlf injection. It is possible to launch 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 but did not respond in any way.	9.9	More Details
	A flaw has been found in Tenda BE3L Pro 16.03.60.62. Affected by this vulnerability is the function		

CVE-2026-114268	formWifiBasicSet of the file /goform/WifiBasicSet. This manipulation of the argument security_5g causes buffer overflow. It is possible to initiate the attack remotely. The exploit has been published and may be used.	9.9	More Details
CVE-2026-112856	Deserialization of Untrusted Data vulnerability in ThemeREX Group Vagabonds vagabonds allows Object Injection.This issue affects Vagabonds: from n/a through 1.3.10.	9.8	More Details
CVE-2026-112855	Deserialization of Untrusted Data vulnerability in axiomthemes Redy redy allows Object Injection.This issue affects Redy: from n/a through 1.3.5.	9.8	More Details
CVE-2026-112854	Deserialization of Untrusted Data vulnerability in axiomthemes Holistic Center holistic-center allows Object Injection.This issue affects Holistic Center: from n/a through 1.2.6.	9.8	More Details
CVE-2026-112853	Deserialization of Untrusted Data vulnerability in axiomthemes Hallelujah hallelujah allows Object Injection.This issue affects Hallelujah: from n/a through 1.1.1.	9.8	More Details
CVE-2026-112851	Deserialization of Untrusted Data vulnerability in ThemeREX Group HotLock hotlock allows Object Injection.This issue affects HotLock: from n/a through 1.3.9.	9.8	More Details
CVE-2026-112848	Deserialization of Untrusted Data vulnerability in AivahThemes Hostme v2 hostmev2 allows Object Injection.This issue affects Hostme v2: from n/a through 7.0.	9.8	More Details
CVE-2026-110956	The asVereinbase WordPress plugin before 1.1.8 does not require a valid password or a genuine second factor before authenticating a user, treating an empty stored value as a match, so an unauthenticated attacker can log in as any user, including an administrator.	9.8	More Details
CVE-2026-112847	Deserialization of Untrusted Data vulnerability in ThemeREX ZenSoul zensoul allows Object Injection.This issue affects ZenSoul: from n/a through 1.13.	9.8	More Details
CVE-2026-112858	Deserialization of Untrusted Data vulnerability in axiomthemes Rhodos rhodos allows Object Injection.This issue affects Rhodos: from n/a through 1.3.3.	9.8	More Details
CVE-2026-112846	Deserialization of Untrusted Data vulnerability in ThemeREX Group Save Life save-life allows Object Injection.This issue affects Save Life: from n/a through 1.2.14.	9.8	More Details
CVE-2026-25784	Initialization of a Resource with an Insecure Default vulnerability in Hallo Welt! GmbH BlueSpice (container image bluespice/wiki) allows privilege escalation. The default value of \$mwsgTokenAuthenticatorServiceCIDR in the container's LocalSettings.php was intended to restrict internal service authentication to the container network, but evaluated to an empty string when the environment variable WIKI_SERVICE_TOKEN_AUTH_ALLOWED was not set. An empty value disables the IP restriction in the token authenticator component (mwstake/mediawiki-component-token-authenticator). As a result, any user with read access to the wiki (anonymous users on publicly readable wikis) could obtain an application token from the REST API and use it to get a session with sysop privileges as the internal service user. This issue affects BlueSpice: from 5.2.3 before 5.2.7, and the 5.2.7 container images published before 2026-10-19 15:37 UTC.	9.8	More Details
CVE-2026-112862	Deserialization of Untrusted Data vulnerability in xnau webdesign Participants Database participants-database allows Object Injection.This issue affects Participants Database: from n/a through 2.7.8.5.	9.8	More Details
CVE-2026-63244	An authentication bypass vulnerability in WizLearn Technologies' Learning Management System (LMS) could allow an unauthenticated remote attacker to bypass password authentication and the OAuth authentication flow by sending a specially crafted POST request, potentially resulting in unauthorised access to user accounts and full administrative control of the LMS platform.	9.8	More Details
CVE-2026-59107	A security vulnerability has been identified in the vSolution Matrix file sharing feature. A shared cryptographic trust configuration could allow an attacker who obtains the matching private key to authenticate to an affected device via SSH as root. This vulnerability affects only devices configured as a vSolution Matrix main/master station with the file sharing feature enabled.	9.8	More Details
CVE-2026-65184	The root-privileged `wsserver` terminal-bridge daemon builds a shell command from client-supplied WebSocket fields and passes it to `system()` and `popen()` without quoting, escaping, or character-set validation. The daemon's public-key session path performs no appliance-session check.	9.8	More Details
CVE-2026-112857	Deserialization of Untrusted Data vulnerability in axiomthemes Reprizo reprizo allows Object Injection.This issue affects Reprizo: from n/a through 1.0.8.	9.8	More Details
CVE-2026-59106	A security vulnerability has been identified in the WolfProt interface implementation. Insufficient neutralization of user-controlled input could allow an unauthenticated attacker with network access to a vulnerable service to execute operating-system commands with the privileges of the affected	9.8	More Details

	application service.		
CVE-2026-112859	Deserialization of Untrusted Data vulnerability in ThemeREX Group Little Birdies little-birdies allows Object Injection.This issue affects Little Birdies: from n/a through 1.3.16.	9.8	More Details
CVE-2026-99078	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Prevent stale XSK buffer release on MPWQE refill retry With AF_XDP on a striding RQ, mlx5e defers releasing XSK buffers until an MPWQE is refilled. If XSK allocation then returns -ENOMEM, actual_wq_head is not advanced and a later NAPI poll retries the same WQE. mlx5e_free_rx_mpwqe() leaves each released slot marked as releasable. On retry it can therefore call xsk_buff_free() again through stale pointers after the frames have returned to the XSK pool and been reallocated. Set all skip_release_bitmap bits in the common error path of mlx5e_xsk_alloc_rx_mpwqe(). This matches mlx5e_alloc_rx_mpwqe(). A successful allocation already clears the bitmap after replacing every buffer, so retries become idempotent without changing the success path. Fault injection forced three consecutive failures for one selected MPWQE. Both an early allocation failure and a partial 8-of-16-buffer unwind released the original 16 XSK buffers only once. Each error left a full bitmap, the following NAPI retry skipped the release, and a later successful allocation cleared it. A 20-second AF_XDP zero-copy pressure run exercised 1,575,262 buffer allocation failures without invalid descriptors, WQE errors, or kernel warnings.	9.8	More Details
CVE-2026-115812	Relative path traversal during model archive extraction in Amazon Web Services Multi Model Server 1.1.11 and earlier might allow remote actors to execute arbitrary code via a crafted model archive. Multi Model Server is end of life and no fixed version is available. Users should discontinue use of Multi Model Server and migrate to a supported model serving solution.	9.8	More Details
CVE-2026-114196	Deserialization of Untrusted Data vulnerability in paytiumsupport Paytium paytium allows Object Injection.This issue affects Paytium: from n/a through 5.0.5.	9.8	More Details
CVE-2026-46652	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 24.03 to before version 26.05, ExternalFilesDownloader.getFilename() extracts the destination filename from a server-supplied Content-Disposition response header and passes it directly to Path.resolve() without stripping directory components. This allows for a malicious HTTP server that the viewer fetches from to supply an absolute path or ../ traversal sequences as the filename, writing attacker-controlled content to an arbitrary path writable by the Tomcat process. When targeting the deployed web-application directory, this yields remote code execution as the tomcat OS user. This issue has been patched in version 26.05.	9.8	More Details
CVE-2026-46651	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 24.03 to before version 26.05, the DownloadTaskEndpoint WebSocket endpoint (/tasks/download/monitor.socket) in goobi-viewer-core accepts download requests from unauthenticated users. This issue has been patched in version 26.05.	9.8	More Details
CVE-2026-53921	odhcpd is the DHCPv6, DHCPv4, and router-advertisement server used by OpenWrt. Prior to version 25.12, odhcpd's DHCPv6 IA response handler writes past the end of a fixed 512-byte stack buffer (pdbuf) when a single DHCPv6 REQUEST carries enough IA options to nearly exhaust the reply space. Two independent overflow sites exist in the same call chain, both reachable from the network via a single unauthenticated UDP packet. odhcpd runs as root on OpenWrt; embedded targets often lack stack canaries and ASLR, making code execution realistic. This issue has been patched in version 25.12.	9.8	More Details
CVE-2026-99079	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Prevent stale XSK buffer release on refill retry When an XDP redirect to an AF_XDP socket fails because its RX ring is full, the XSK core frees the buffer. During the subsequent batched refill of a legacy cyclic RQ, mlx5e also releases the WQE's XSK buffer before allocating a replacement. If that refill succeeds only partially, a WQE left without a replacement retains its old buffer pointer. The buffer can meanwhile be allocated to another WQE. A later refill retry can then free the live buffer through the stale pointer and publish the same UMEM frame twice. Mark the WQE as released immediately after the driver-side free. The flag is already cleared when a replacement buffer is assigned, so refill retries no longer release stale pointers. The failure is silent and produces no kernel warning or splat. A standalone legacy cyclic-RQ zero-copy libxsk reproducer, using 64-byte UDP traffic offered at 12 Mpps, detected it: stock stopped after 2,854,914 packets in 4.094 seconds, with 4,542 xdp_rx_ring_full events and 64 ownership/double-publication errors. With this change it processed 356,904,225 packets in 30 seconds despite 571,405 xdp_rx_ring_full events, with no ownership or data errors.	9.8	More Details
CVE-2026-94075	Deserialization of Untrusted Data vulnerability in Tribulant Software Newsletters newsletters-lite allows Object Injection.This issue affects Newsletters: from n/a through 4.18.1.	9.8	More Details
	PraisonAI is a multi-agent teams system. From version 4.2.1 to before version 4.6.59, PraisonAI's		

CVE-2026-57116	AgentOS FastAPI deployment surface remains unauthenticated in current main and in releases after the published patched version for GHSA-pm96-6xpr-978x / CVE-2026-40151. The chat route directly calls agent.chat(request.message). No-auth and wrong-bearer requests both execute the deployed agent. This is broader than passive metadata disclosure. In any deployment where AgentOS wraps agents with tools, private context, memory, API integrations, or cost-bearing model calls, an unauthenticated reachable client can drive those agents. This issue has been patched in version 4.6.59.	9.8	More Details
CVE-2026-115214	The Datawiza Proxy Auth Plugin WordPress plugin before 1.1.1 does not verify that the user-identity and role HTTP request headers it trusts originate from a trusted reverse proxy, allowing unauthenticated attackers to create and authenticate as a WordPress administrator.	9.8	More Details
CVE-2026-107092	NexoPOS v6.2.4 is vulnerable to Directory Traversal via the api/update endpoint.	9.8	More Details
CVE-2026-107656	An authentication bypass using an alternate path or channel vulnerability in Fortinet FortiSandbox 5.0.0 through 5.0.6, FortiSandbox 4.4.0 through 4.4.9, FortiSandbox 4.2 all versions, FortiSandbox PaaS 5.0.0 through 5.0.6, FortiSandbox PaaS 4.4 all versions, FortiSandbox PaaS 4.2 all versions may allow attacker to execute unauthorized code or commands via crafted HTTP requests	9.8	More Details
CVE-2026-116485	danmu_api before 1.21.7 contains an authentication bypass vulnerability that allows unauthenticated attackers to modify server environment configuration when the default TOKEN is used and ADMIN_TOKEN is unset. Attackers can send POST requests to /api/env/set, /api/env/add, or /api/env/del to set ADMIN_TOKEN, lock out operators, redirect OTHER_SERVER or PROXY_URL for SSRF, and trigger redeployment.	9.8	More Details
CVE-2026-115806	The Prime Mover – Backup and Migration plugin for WordPress is vulnerable to PHP Object Injection in all versions up to, and including, 2.2.1 via deserialization of untrusted input in the migration search-and-replace routine (recursiveUnserializeReplace()). An unauthenticated attacker can store a crafted, base64-encoded serialized object in wp_comments.comment_content through the standard comment-submission endpoint (wp-comments-post.php). When an administrator subsequently performs a routine Prime Mover export or restore, the plugin base64-decodes that stored value and passes it to unserialize() with no allowed_classes restriction. A Property-Oriented Programming (POP) chain is present in the plugin's own bundled Freemius library (FS_Clone_Manager::__set → FS_Option_Manager::set_option → store() → update_option()), so no additional plugin or theme is required — the plugin ships the gadget itself.	9.8	More Details
CVE-2026-98959	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix xattr replace when multiple xattrs are packed in the same item If we have a btrfs_dir_item item that packs multiple xattrs and then we replace the value of one of them (with the setxattr(2) family of syscalls) with another value of a different size, we end up not having a fully initialized btrfs_dir_item, resulting in a corruption that the tree checker will detect at extent buffer writeback time. This is because in btrfs_setxattr() when we find a btrfs_dir_item with multiple xattrs (due to the crc32c hash of their name being the same) we delete one of the xattr items (btrfs_dir_item) and then insert a new one, but the deletion and insertion results in shifting existing data in the leaf and therefore when the new value of a xattr has a different size, the new btrfs_dir_item is placed in a leaf section that was not initialized and we only copy the value's data and set the value's length in the new btrfs_dir_item, without setting the name, the name's length, the key (which must be all zeroes for xattrs), flags (BTRFS_FT_XATTR) and transaction ID. The following script reproduces the issue: \$ cat test.sh #!/bin/bash DEV=/dev/sdi MNT=/mnt/sdi mkfs.btrfs -f \$DEV mount \$DEV \$MNT touch \$MNT/testfile # Add two xattrs that, on btrfs, have the same hash (crc32c) for their # name and therefore are packed into the same btrfs_dir_item. setfattr -n user.foobar -v 123 \$MNT/testfile setfattr -n user.WvG1c1Td -v qwerty \$MNT/testfile # Verify the xattrs are present. echo "xattrs before:" getfattr --absolute-names --dump \$MNT/testfile # Now replace the value of the foobar xattr with a significantly larger # value. setfattr -n user.foobar -v abcdefghijklmnopqrstuvwxyz \$MNT/testfile # Check the xattrs have the expected values. echo "xattrs after:" getfattr --absolute-names --dump \$MNT/testfile umount \$MNT Running it: \$./test.sh (...) xattrs before: # file: /mnt/sdi/testfile user.WvG1c1Td="qwerty" user.foobar="123" xattrs after: # file: /mnt/sdi/testfile user.WvG1c1Td="qwerty" So the "user.foobar" xattr is missing and there was a transaction abort when unmounting the fs with the following traces in dmesg: \$ dmesg [869800.159271] BTRFS warning (device sdi): access to eb bytenr 30474240 len 16384 out of range start 16015 len 25964 [869800.159293] -----[cut here]----- [869800.159296] WARNING: fs/btrfs/extent_io.c:4408 at report_eb_range+0x44/0x60 [btrfs], CPU#8: getfattr/2605179 [869800.168961] Modules linked in: btrfs dm_thin_pool (...) [869800.190527] CPU: 8 UID: 0 PID: 2605179 Comm: getfattr Tainted: G W 7.3.0-rc3-btrfs-next-244+ #1 PREEMPT(full) [869800.193726] Tainted: [W]=WARN [869800.194502] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.2-0-gea1b7a073390-prebuilt.qemu.org 04/01/2014 [869800.197576] RIP: 0010:report_eb_range+0x44/0x60 [btrfs] [869800.198985] Code: 48 8b 7b 18 (...) [869800.203531] RSP: 0018:ffffce4541937d58 EFLAGS: 00010246 [869800.204609] RAX: 0000000000000000 RBX:	9.8	More Details

	ffff8dde054b8738 RCX: 0000000000000000 [869800.206137] RDX: 0000000000000000 RSI: 0000000000000001 RDI: ffffffff04c92a0 [869800.207645] RBP: 0000000000003e8f R08: 0000000000000000 R09: 3fffffffefefff [869800.226175] R10: ffffce4541937a88 R11: 0000000000000003 R12: 000000000000656c [869800.227594] R13: ffff8dde14be800f R14: 000000000000656c R15: 0000000000000069 [869800.229097] FS: 00007f59023bb780(0000) GS:ffff8de5788ed000(0000) knlGS:0000000000000000 [869800.231137] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [869800.232321] CR2: 0000558dc45e6a78 CR3: 0000000765b16004 CR4: 0000000000370ef0 [869800.233783] Call Trace: [869800.234330] <TASK> [869800.234791] read_extent_buffer+0x4d/0x100 [btrfs] [869800.235898] btrfs_listxattr+0x199/0x240 [btrfs] [869800.236912] vfs_listxattr+0x51/0 ---truncated---		
CVE-2026-113736	A vulnerability has been found in D-Link DHP-W610AV 1.02EUb03. The impacted element is the function login_prepare_hnap of the file etc/defnodes/defaultvalue.xml of the component HNAP Login. The manipulation of the argument Password leads to empty password in configuration file. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used.	9.8	More Details
CVE-2026-112346	OpenHands Software Agent SDK before 1.52.0 contains an unauthenticated remote code execution vulnerability where the DockerWorkspace and Apptainer workspace launchers bind the sandbox agent server to all network interfaces (0.0.0.0) with an empty session_api_keys list, disabling authentication by default. Attackers who can reach the published container port can send unauthenticated HTTP and WebSocket requests to execute arbitrary shell commands as root in the sandbox, access the mounted workspace, and read plaintext secrets from the secrets store.	9.8	More Details
CVE-2026-79861	There is a command injection vulnerability in AOS-10 that could lead to unauthenticated remote code execution by sending specially crafted PAPI packets. Successful exploitation of this vulnerability results in the ability to execute arbitrary code as a privileged user on the underlying operating system.	9.8	More Details
CVE-2026-76474	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76474 are related to issues concerning improper authentication that are grouped under the Common Weakness Enumeration (CWE) Class CWE-287.	9.8	More Details
CVE-2026-79860	Buffer overflow vulnerabilities exist in an underlying service of AOS-8 that could lead to unauthenticated remote code execution by sending specially crafted packets destined to the PAPI (Aruba's access point management protocol) UDP port 8211. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	9.8	More Details
CVE-2026-79857	There is a command injection vulnerability in the PAPI protocol (Aruba's access point management protocol) of AOS-8 that could lead to unauthenticated remote code execution by sending specially crafted packets to the affected UDP port. Successful exploitation of this vulnerability results in the ability to execute arbitrary code as a privileged user on the underlying operating system.	9.8	More Details
CVE-2026-79856	A buffer overflow vulnerability exists in the PAPI protocol of AOS-8 that 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.	9.8	More Details
CVE-2026-79855	An authentication bypass vulnerability in the SSH daemon of AOS-8 allows an unauthenticated remote attacker to circumvent existing authentication controls. Successful exploitation allows an attacker to gain administrative privileges leading to complete compromise of the AOS-8 host.	9.8	More Details
CVE-2026-79854	A vulnerability in the affected interface of AOS-10 could allow an unauthenticated remote attacker to conduct SQL injection attacks against the AOS-10 instance. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	9.8	More Details
CVE-2026-79853	Sensitive information disclosure vulnerabilities in the affected interface of AOS-10 could allow an unauthenticated remote attacker to obtain sensitive information from the affected system. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	9.8	More Details
CVE-2026-79852	A vulnerability in the web-based management interface of AOS could allow an unauthenticated remote attacker to conduct SQL injection attacks against the AOS instance. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	9.8	More Details

CVE-2022-51026	Bytebase 0.12.0 before 1.0.1 contains an authentication bypass vulnerability that allows unauthenticated attackers to log in as any user. The flaw stems from trusting client-supplied GitLab provider configuration in the POST /api/auth/login/GITLAB_SELF_HOST endpoint. Attackers can set instanceUrl to a server they control that returns an active user with a chosen email. Bytebase then issues session tokens for that principal, letting attackers impersonate existing users including the Owner or create new members.	9.8	More Details
CVE-2026-78846	Multiple AiDot devices are vulnerable to a heap-based buffer overflow via SSDP location header parsing.	9.8	More Details
CVE-2026-112832	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in CozyThemes Rootblox rootblox allows PHP Local File Inclusion.This issue affects Rootblox: from n/a through 1.0.6.	9.8	More Details
CVE-2026-79859	Buffer overflow vulnerabilities exist in the affected interface of AOS-8 that could lead to unauthenticated remote code execution. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	9.8	More Details
CVE-2026-79858	An unauthenticated SQL injection vulnerability exists in the PAPI protocol in AOS-8. Successful exploitation could allow an unauthenticated remote attacker to obtain and modify sensitive information in the underlying database and to interrupt the normal operation of the affected service.	9.8	More Details
CVE-2026-76473	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76473 are related to issues concerning improper neutralization of special elements in output used by a downstream component that are grouped under the Common Weakness Enumeration (CWE) Class CWE-74.	9.8	More Details
CVE-2021-48016	DedeCMS through 5.8.1-beta contains a PHP code injection vulnerability in ShowMsg() where the HTTP Referer header is written into a DedeTemplate string and included. Unauthenticated attackers can send a crafted Referer to /plus/flink.php?dopost=save with an empty validate code to execute arbitrary PHP and system commands. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2026-114108	HKUDS Vibe-Trading from 0.1.0 before 0.1.7 contains a missing authentication vulnerability that allows unauthenticated attackers to access the FastAPI surface because require_auth() returns early when API_AUTH_KEY is unset. Attackers can post messages to /sessions/{id}/messages to trigger root shell commands via BashTool, upload .py or .sh files via /upload, read session histories, and obtain partial API keys.	9.8	More Details
CVE-2021-48014	FineReport 9.0 contains a path traversal vulnerability in the /WebReport/ReportServer design_save_svg operation that allows unauthenticated attackers to write arbitrary files via the filePath parameter. Attackers can send a POST request with directory traversal sequences in filePath and a __CONTENT__ payload to write a JSP webshell and achieve remote code execution. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	9.8	More Details
CVE-2026-96639	Pillow is a Python imaging library. From 11.0.0 until 13.0.0, multi-frame loaders can retain an ImagingCore allocation created for a narrow frame mode while a later frame of the same dimensions changes the logical mode and decoder tiles to wider RGB or CMYK data. ImageFile.load_prepare() reuses the stale native storage because _im is non-null without verifying that its mode and size still match. Crafted PSD, MPO, DCX, or MIC frame transitions can then write wider pixels into the narrow allocation, causing native out-of-bounds writes and crashes. The MPO path can also cause a native over-read and, in demonstrated runtime-specific profiles, command execution. This issue is fixed in version 13.0.0.	9.8	More Details
CVE-2026-76466	A vulnerability in the database service feature of Cisco Expressway Series could allow an unauthenticated, remote attacker with network access to a Cisco Expressway device to execute arbitrary operating system commands as root. This vulnerability is due to the use of persistent cookies. Additionally, there is insufficient validation of client certificates when the TLS Verify Mode is set to Enabled for clustering, which means that client certificates are not validated. An attacker with possession of the cookies could exploit this vulnerability by connecting to the database service and invoking specific remote procedure call (RPC) commands. A successful exploit could allow the attacker to execute arbitrary commands as root on the underlying operating system of an affected device. Note: Cisco Expressway Series refers to Cisco Expressway Core (Expressway-	9.8	More Details

	C) devices and Cisco Expressway Edge (Expressway-E) devices.		
CVE-2026-112241	Authentication Bypass by Spoofing vulnerability in Brainstorm Force SureCart surecart allows Authentication Abuse.This issue affects SureCart: from n/a through 4.9.5.	9.8	More Details
CVE-2026-109618	PHP uses longjmp() for out-of-memory (OOM) errors to exit the current context and abort handling of the request. Some contexts may not properly account for this scenario and leave request memory in an uninitialized state before performing an allocation. The subsequent request cleanup procedure may read or branch on this memory. Multiple OOM strategies can consistently trigger this condition and lead to remote code execution (RCE).	9.8	More Details
CVE-2026-114119	Best Practice Software Bp Premier before 1.14.1.1128 uses a hardcoded Microsoft SQL Server administrative password that is shared across all installations and cannot be changed without breaking the application. Attackers with network access to the database instance can authenticate as sysadmin, bypass application authentication, and execute operating system commands as SYSTEM.	9.8	More Details
CVE-2026-115932	When a post is published, Postiz treats a media path that is not a web address as a file on the server's disk, but it does not check that the file lies inside the upload directory. A registered user can store a crafted media path with a post, get it past the checks that normally restrict it, and have the publishing worker read an arbitrary file and send it as an attachment to a Discord server the user controls. The same unconfined reads can be reached through two further routes. One tells the caller whether a file exists and can exhaust the backend's memory, the other delivers the file the same way as the main route. Self-registration is open by default, so the attacker needs no existing account.	9.6	More Details
CVE-2026-115995	A security vulnerability has been detected in Unitree Robotics Unitree Go2 up to 1.1.11.4. The impacted element is the function set_volume of the file gpt_thread.py of the component chat_go. Such manipulation of the argument Value leads to os command injection. The attack can be launched remotely. Upgrading to version 1.1.11.5 is sufficient to resolve this issue. You should upgrade the affected component. The vendor confirmed early: "Submission 990502 corresponds to our internal ticket #474 (Go2 chat_go set_volume OS command injection, CWE-78)."	9.6	More Details
CVE-2026-79870	An authentication bypass vulnerability in the PAPI protocol of AOS-8 could allow an unauthenticated adjacent attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to gain administrative access, modify system configurations, access or manipulate sensitive data, and disrupt network operations.	9.6	More Details
CVE-2026-79867	Authentication bypass vulnerabilities in the affected interface of AOS-8 could allow an unauthenticated adjacent attacker to circumvent existing authentication controls and run arbitrary commands on the underlying host. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system leading to complete system compromise.	9.6	More Details
CVE-2026-79869	Authentication bypass vulnerabilities in the affected interface of AOS-10 could allow an unauthenticated adjacent attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to execute arbitrary commands as a privileged user on the underlying operating system.	9.6	More Details
CVE-2026-79863	Buffer overflow vulnerabilities in the affected interface of AOS-10 or AOS-8 could allow an unauthenticated adjacent attacker to execute arbitrary code by sending specially crafted data to the affected interface. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	9.6	More Details
CVE-2026-79864	Command injection vulnerabilities exist in the affected interface of AOS-10 that could allow an unauthenticated adjacent attacker to perform command injection. Successful exploitation could allow an attacker to execute arbitrary commands as a privileged user on the underlying operating system.	9.6	More Details
CVE-2026-79865	A memory corruption vulnerability in the affected interface of AOS-10 could allow an unauthenticated adjacent attacker to achieve remote code execution. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	9.6	More Details
CVE-2026-79866	An unauthenticated path traversal vulnerability exists in the PAPI protocol in AOS-8. Successful exploitation of this vulnerability could allow an unauthenticated adjacent attacker to read, modify, and delete files on the underlying operating system.	9.6	More Details
CVE-2026-79871	A sensitive information disclosure vulnerability in the command line interface of AOS-10 could allow an unauthenticated remote attacker to access sensitive information. Successful exploitation could allow an attacker to retrieve information which could be used to potentially gain further access to network services supported by AOS-10.	9.6	More Details

CVE-2026-79862	Buffer overflow vulnerabilities in the affected interface of AOS-10 or AOS-8 could allow an unauthenticated adjacent attacker to execute arbitrary code by sending specially crafted data to the affected interface. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	9.6	More Details
CVE-2026-55981	A firmware verification vulnerability in Creative Sound Blaster Katana could allow an attacker to modify a legitimate firmware image and upload the modified firmware to an affected device.	9.6	More Details
CVE-2026-54630	Part-DB is an open source inventory management system for electronic components. Prior to version 2.12.1, the part attachment upload handler stores uploaded files under the web-served media directory using the client-supplied file extension, and its extension blacklist does not include phar. AttachmentSubmitHandler::renameBlacklistedExtensions() rewrites dangerous extensions (php, phtml, sh, cgi, py, js, html, htaccess, ...) to .txt, but phar is absent from the list, so a .phar file keeps its extension. Public attachments (the default) are written into public/media/, which is the Apache document root. The shipped Apache configuration routes .phar to PHP-FPM, so the stored file is executed as PHP on request. The per-attachment-type extension filter is empty by default and is not consulted by the upload handler in a way that blocks this, and a non-admin user in the default users group can create an attachment type with no filter and attach files to parts. Any authenticated user in the default users group uploads a .phar attachment to a part and fetches it under /media/, executing arbitrary PHP as the www-data web user. Ability to exploit depends on the webserver config, docker image contained the vulnerability by default. This issue has been patched in version 2.12.1.	9.6	More Details
CVE-2026-79868	Authentication bypass vulnerabilities in the affected interface of AOS-8 could allow an unauthenticated adjacent 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, leading to complete compromise of the affected system.	9.6	More Details
CVE-2026-79872	A buffer overflow vulnerability exists in the underlying operating system of AOS that could lead to unauthenticated remote code execution by an unauthenticated remote attacker. Successful exploitation of this vulnerability results in the ability to execute arbitrary code as a privileged user on the underlying operating system.	9.4	More Details
CVE-2026-111852	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in WPdreams Ajax Search Lite ajax-search-lite allows Blind SQL Injection.This issue affects Ajax Search Lite: from n/a through 4.14.6.	9.3	More Details
CVE-2026-111402	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Tribulant Software Newsletters newsletters-lite allows SQL Injection.This issue affects Newsletters: from n/a through 4.18.1.	9.3	More Details
CVE-2026-111809	Apache Traffic Server allows request smuggling via HEAD requests that carry Content-Length or Transfer-Encoding. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	9.3	More Details
CVE-2026-111816	Apache Traffic Server forwarded a malformed 302 response that should have been rejected as a parse error, because the Transfer-Encoding: chunked carve-out bypassed the rejection path. A response that two peers can read differently is then stored and served, allowing cache poisoning and response-framing desync. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	9.3	More Details
CVE-2026-76475	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76475 are related to issues concerning improper neutralization of input during web page generation that are grouped under the Common Weakness Enumeration (CWE) Base CWE-79.	9.3	More Details
CVE-2026-112866	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in WPFactory Product Filter by WBW woo-product-filter allows Blind SQL Injection.This issue affects Product Filter by WBW: from n/a through 3.4.6.	9.3	More Details
CVE-2026-	An authentication bypass in the API of AOS-10 allows an unauthenticated adjacent attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to gain	9.3	More

79874	administrative access, modify system configurations, and access or manipulate sensitive data.		Details
CVE-2026-79876	A sensitive information disclosure vulnerability in the Wi-Fi frame processing of AOS-10 could allow an unauthenticated adjacent attacker to access sensitive information. Successful exploitation could allow an attacker to retrieve information which could be used to potentially gain further access to network services supported by AOS-10, or to interrupt the normal operation of the affected service.	9.3	More Details
CVE-2026-112327	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in TMS Amelia ameliabooking allows SQL Injection.This issue affects Amelia: from n/a through 2.4.11.	9.3	More Details
CVE-2026-55486	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 25.08 to before version 26.05, the IIIF Auth 2.0 endpoints auth/login and auth/token in AuthorizationFlowResource accept an origin query parameter without validation and then (1) embed it into the Content-Security-Policy: frame-ancestors response header and (2) embed it into an HTML response body as a JavaScript single-quoted postMessage argument. Because the attacker controls both the login and token calls that initiate the flow, the unescaped value yields reflected XSS. This issue has been patched in version 26.05.	9.3	More Details
CVE-2026-79873	Authentication bypass vulnerabilities in the PAPI protocol of AOS-10 could allow an unauthenticated adjacent 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 on affected devices.	9.3	More Details
CVE-2026-79875	A vulnerability has been identified in the PAPI protocol of AOS-8 that could potentially allow an unauthenticated adjacent 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.	9.3	More Details
CVE-2026-112860	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in xnau webdesign Participants Database participants-database allows Blind SQL Injection.This issue affects Participants Database: from n/a through 2.7.8.5.	9.3	More Details
CVE-2026-96355	Improper Neutralization of Special Elements Used in a Template Engine vulnerability in Drupal Webform allows Code Injection. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	9.1	More Details
CVE-2026-116234	A vulnerability was detected in Oray Pugongying X5 Pro 7.6.1. Affected is an unknown function of the file /etc/ipsec.conf of the component IPsec Policy Handler. The manipulation of the argument secproto results in crlf injection. The attack can be launched 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.	9.1	More Details
CVE-2026-105014	Improper neutralization of special elements used in a command ('command injection') in Azure Pipelines allows an unauthorized attacker to execute code over a network.	9.1	More Details
CVE-2026-116214	MacCMS v10 before 2026.1000.4054 contains an authentication bypass vulnerability that allows unauthenticated attackers to reset any member's password because check_msg() never compares msg_to. Attackers can request a recovery code to their own mailbox via user/findpass_msg, then submit it to user/findpass_reset with the victim's email or phone to take over the account.	9.1	More Details
CVE-2026-109122	The network management service on affected Kiloview E3 NDI encoder firmware fails to sanitize one field used when configuring static routes via the web management API. An authenticated web user with permission to configure static routes can exploit this to execute arbitrary operating system commands as root, and the malicious configuration persists across reboots. Successful exploitation could result in full device compromise, including exposure of stored credentials and cryptographic material.	9.1	More Details
CVE-2026-115547	An input validation issue in the Okta Access Gateway (OAG) Admin Console allows the processing of specially formatted configuration values. An authenticated administrator can leverage this to execute system commands within the context of the OAG service account.	9.1	More Details
CVE-2026-116235	A flaw has been found in Oray Pugongying X5 Pro 7.6.1. Affected by this vulnerability is an unknown functionality of the component L2TP Authentication. This manipulation of the argument Username causes crlf injection. The attack may be initiated remotely. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	9.1	More Details
CVE-2026-116233	A security vulnerability has been detected in Oray Pugongying X5 Pro 7.6.1. This impacts an unknown function of the component Dnsmasq. The manipulation of the argument rent_time leads to crlf injection. The attack can be initiated remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	9.1	More Details

CVE-2026-114220	NT OneAPP developed by NewType Infotech has an Improper Certificate Validation vulnerability. Unauthenticated network-adjacent attackers can perform a man-in-the-middle attack by failing to validate the server certificate and the hostname, thereby intercepting or modifying communications between the application and the server.	9.1	More Details
CVE-2026-114611	OCO-Server is a self-hosted endpoint management application for computer inventory, software deployment, and mobile device management. Prior to 1.2.3, the Apple MDM enrollment endpoint `frontend/api-mdm.php` verifies the signed enrollment-request plist with `openssl_pkcs7_verify(..., PKCS7_NOVERIFY, ...)` . `PKCS7_NOVERIFY` disables signer certificate chain and trust validation, so the signature only has to be self-consistent with whatever certificate is embedded in the PKCS#7 blob. A self-signed certificate from an unrelated attacker is accepted. Because the endpoint trusts the `SERIAL` field taken from the unverified payload, anyone who knows a registered device's serial number can sign a forged plist (`SERIAL` = victim device, attacker-chosen `UDID`) with an arbitrary self-signed certificate; POST it unauthenticated to `api-mdm.php/profile` (Content-Type: `application/pkcs7-signature`); and receive the victim device's freshly minted MDM identity: a certificate with `CN=<serial>` issued by the server's internal MDM device CA, bundled with its private key in a PKCS#12 payload inside the returned `mobileconfig` enrollment profile. Additionally, the victim's `mobile_device` database row is silently overwritten with the attacker-supplied `UDID` (device-record hijack). With the stolen certificate and key, the attacker can impersonate the victim device on all subsequent MDM traffic (`/checkin`, `/mdm`, `/parameters`), which authenticate callers by the certificate CN (the serial number) via `checkSignature()` in the same file. This yields control over the victim device's MDM relationship: reading its management data, replacing its push token, and disrupting or redirecting its enrollment. Version 1.2.3 fixes the issue.	9.1	More Details
CVE-2026-116207	YesWiki 4.6.0 before 4.7.0 and 5.0.0-alpha1 before 5.0.0-alpha4 contains an unauthenticated SQL injection vulnerability in SearchManager::prepareSearchRequest(). Field names from the query and searchfields parameters of the public Bazar entries search API are placed in a backtick-quoted alias escaped only with mysqli_real_escape_string. Attackers can supply a backtick to inject subqueries, reading data such as user password hashes or exhausting database threads with SLEEP().	9.1	More Details
CVE-2022-51022	YiShaAdmin 2.2 through 3.1, fixed in commit 025b062, contains an unauthenticated path traversal vulnerability in the /admin/File/DownloadFile endpoint that allows attackers to read and delete arbitrary files via the filePath parameter. Attackers can send crafted GET requests with directory traversal sequences in filePath and set the delete parameter to read or remove system files like etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-20.	9.1	More Details
CVE-2026-114820	Vulnerability in the Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: JSSE). Supported versions that are affected are Oracle Java SE: 8u511-b32, 8u511-perf-b32, 11.0.33.0.2, 17.0.21.0.2, 21.0.13.0.2, 25.0.5.0.2, 27.0.1.0.2; Oracle GraalVM for JDK: 17.0.21.0.2 and 21.0.13.0.2; Oracle GraalVM Enterprise Edition: 21.3.20.0.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition accessible data as well as unauthorized access to critical data or complete access to all Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition accessible data. Note: This vulnerability can be exploited by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. This vulnerability also applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. CVSS 3.1 Base Score 9.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N).	9.1	More Details
CVE-2026-91403	Vulnerability in the Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: JSSE). Supported versions that are affected are Oracle Java SE: 8u511, 8u511-perf, 11.0.33, 17.0.21, 21.0.13, 25.0.5, 27.0.1; Oracle GraalVM for JDK: 17.0.21 and 21.0.13; Oracle GraalVM Enterprise Edition: 21.3.20. Easily exploitable vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition accessible data as well as unauthorized access to critical data or complete access to all Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition accessible data. Note: This vulnerability can be exploited by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. This vulnerability also applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that	9.1	More Details

	comes from the internet) and rely on the Java sandbox for security. CVSS 3.1 Base Score 9.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N).		
CVE-2026-114607	Better Auth is an authentication and authorization library for TypeScript. From 1.4.0-beta.18 until 1.7.7, deployments that enable Magic Link with storeToken set to plain, or with storeToken set to hashed and the global verification.storeIdentifier setting also set to hashed, enable a built-in social or Generic OAuth provider, and use account.storeStateStrategy set to database can accept an OAuth sign-in state value as a Magic Link token. These settings can cause Magic Link verification records and OAuth state to resolve to the same stored identifier, allowing an unauthenticated attacker to start OAuth sign-in for a known email address and redeem the returned state through Magic Link without controlling the mailbox or completing the provider callback. The attacker can obtain a normal session for an existing user, or create an account incorrectly marked as email verified when sign-up is enabled. This issue is fixed in version 1.7.7.	9.1	More Details
CVE-2026-94267	The Advanced Custom Fields: Extended WordPress plugin before 0.9.2.8, Advanced Custom Fields: Extended PRO WordPress plugin before 0.9.2.8 does not perform any capability check when updating a post through one of its front-end form actions, allowing unauthenticated visitors to overwrite arbitrary posts, including the ones in which the site stores its own field definitions, which in turn leads to arbitrary code execution on the server. Exploitation requires the site to publish a form whose post-update action maps both its target and its content to visitor-submitted values.	9.0	More Details
CVE-2026-20043	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco owned Open Source Snort 3 Software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20043 are related to issues concerning out-of-bounds writes that are grouped under the Common Weakness Enumeration (CWE) Base CWE-787.	9.0	More Details
CVE-2026-111811	Apache Traffic Server's experimental tsmemcache plugin allows unauthenticated network read, write, and deletion of the shared proxy cache. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2025-31699	Apache Traffic Server can return an origin connection to the shared pool while the origin is still waiting on a request body the client declared but never fully sent, including when the client used Expect: 100-continue. The origin then reads a later request on that connection as leftover body data, which lets an attacker smuggle requests past ATS access controls, disrupt other users' requests, or poison the cache. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	8.9	More Details
CVE-2026-76526	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76526 are related to improper neutralization and are grouped under the Common Weakness Enumeration (CWE) CWE-707.	8.8	More Details
CVE-2026-111841	Cross-Site Request Forgery (CSRF) vulnerability in Aman FunnelKit Automations wp-marketing-automations allows Cross Site Request Forgery.This issue affects FunnelKit Automations: from n/a through 3.8.5.2.	8.8	More Details
CVE-2026-52356	An OS command injection vulnerability in the /htdocs/web/diagnostic.php endpoint of D-Link DAP1160L A1_FW101WWb04 allows attackers to execute arbitrary commands via injecting arbitrary shell metacharacters.	8.8	More Details
CVE-2023-7356	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Copy BASE if service data matches EIR_BAA_SERVICE_UUID Copy the content of a Periodic Advertisement Report to BASE only if the service UUID is Basic Audio Announcement Service UUID.	8.8	More Details
CVE-2026-	Hugging Face Datasets through 5.1.0 contains an insecure deserialization vulnerability that allows attackers to execute arbitrary Python code by crafting a malicious pickle file loaded through the deprecated pandas packaged builder. Attackers can supply a pickle-backed dataset that is processed by Pandas._generate_tables via pandas.read_pickle() without validation, so when a victim calls	8.8	More Details

57845	datasets.load_dataset("pandas", data_files=...), the embedded __reduce__ payload executes code with the victim process's privileges.		
CVE-2026-113779	The pfSense REST API package is an unofficial, open-source REST and GraphQL API for pfSense CE and pfSense Plus firewalls. A privilege escalation vulnerability in pfSense-pkg-RESTAPI prior to version 2.9.0 allows authenticated accounts holding the dedicated settings sync permission to modify REST API configuration secrets. This enables a non-administrative user to elevate their privileges to a full administrator. pfSense-pkg-RESTAPI 2.9.0 patches the issue. Access to the HA settings sync endpoint now strictly requires full administrative privileges (`page-all`). The sub-administrative sync privilege has been removed in versions v2.9.0 and later. In addition to correcting the access control boundary, version 2.9.0 includes the following hardening measures: opt-in default and format and validation modernization. As a workaround, one may mitigate exposure by taking either or all of the following steps. Revoke Scoped Permissions: Remove the `api-v2-system-restapi-settings-sync-post` privilege from all non-administrative users. Apply ACLs: Configure REST API Access Control Lists (ACLs) and/or firewall rules to restrict access to API endpoints exclusively to trusted administrator and HA peer IP addresses.	8.8	More Details
CVE-2026-79982	The Bluetooth HCI USB transport class of the new USB device stack, subsys/usb/device_next/class/bt_hci.c, copies every controller-to-host HCI packet into a USB transfer buffer before enqueueing it on an endpoint. bt_hci_tx_sync_in() allocated that buffer from bt_hci_ep_pool, a fixed pool whose slots are only USBD_MAX_BULK_MPS bytes (64 on full-speed-only builds), and then performed net_buf_add_mem(buf, bt_buf->data, bt_buf->len) with no length check. HCI event and ACL buffers are sized by CONFIG_BT_BUF_EVT_RX_SIZE and CONFIG_BT_BUF_ACL_RX_SIZE and can be far larger, so the copy runs past the end of the 64-byte slot into the adjacent slots and into whatever follows the statically allocated pool. The device runs the local Bluetooth controller in raw HCI mode (CONFIG_BT_HCI_RAW), so the oversized packets originate over the air. With the HCI buffer sizes recommended for full-featured controllers, an unauthenticated attacker within Bluetooth range can supply the overflowing bytes: a connected peer that negotiates Data Length Extension can send ACL fragments up to CONFIG_BT_CTLR_DATA_LENGTH_MAX bytes, and an attacker that merely broadcasts extended advertising data reaches the same path through LE Extended Advertising Report events, which the controller fragments only to CONFIG_BT_BUF_EVT_RX_SIZE. In the stock configuration the overflow is already reached without any radio traffic: the mandatory Read Supported Commands command-complete event is 70 bytes and overruns the 64-byte slot by six bytes on every host initialization. Because net_buf_add_mem() validates tailroom only through an assertion unless CONFIG_NET_BUF_HARDENING is enabled, builds with CONFIG_ASSERT=n — the usual production setting — silently corrupt memory adjacent to the USB buffer pool with attacker-chosen contents, which on a target without memory protection between subsystems can lead to arbitrary code execution. Builds with assertions enabled fail the assertion instead and panic, giving a remotely triggerable denial of service. The fix introduces a separate bt_hci_tx_pool sized to MAX(BT_BUF_ACL_RX_SIZE, BT_BUF_EVT_RX_SIZE) for the controller-to-host direction while leaving the host-to-device pool at the endpoint packet size, which is already bounds-checked during reassembly.	8.8	More Details
CVE-2026-82944	SOEM version 2.0.0 and earlier contains a vulnerability in multi-segment data transfers where a malicious device can overflow memory buffers. The software checks the total size only once at the beginning, but doesn't track the cumulative size as multiple segments arrive. An attacker can declare a small initial size to pass validation, then send additional segments that exceed the allocated buffer. This allows the attacker to write controlled data beyond buffer boundaries, potentially enabling code execution.	8.8	More Details
CVE-2026-115895	Happy DOM is a JavaScript implementation of a web browser without its graphical user interface. From 17.0.0 until 21.0.4, packages/happy-dom/src/module/ECMAScriptModuleCompiler.ts accepts a named ESM import alias without validating that it is a JavaScript binding identifier before inserting it into a generated destructuring binding. With enableJavaScriptEvaluation enabled, attacker-controlled HTML or module source can use an invalid alias that becomes a valid default initializer after compilation, causing the generated JavaScript to execute in the Node.js host process. The demonstrated reachable path uses a JSON module import attribute, which prevents the malicious alias from being emitted again in a later ESM circular-import assignment and keeps the generated function syntactically valid. This incomplete fix for CVE-2026-33943 is a distinct residual path because the earlier remediation validated the local export importName path but not the named-import alias. This issue is fixed in version 21.0.4.	8.8	More Details
CVE-2026-115330	Lotus through 0.10.2 contains an sql injection vulnerability that allows authenticated organization users to execute arbitrary SQL by appending stacked statements to custom metric queries. Attackers can submit semicolon-separated SQL through POST /app/metrics/, bypassing first-statement-only validation, to read or modify data across all tenants.	8.8	More Details
CVE-	The app/process-api-response-headers action accepts fake Craft API response headers from any authenticated control panel user. A non-admin user with only accessCp can place newline characters in X-Craft-Plugin-Licenses. Craft passes this value to its .env writer without removing the newlines. This		More

CVE-2026-115935	lets the user insert a new attacker-known CRAFT_SECURITY_KEY into .env. After Craft loads that key, the user can forge a signed login redirect containing a Twig system() callback. This is different from the previous report. That report reads the existing key through a Twig sandbox bypass. This report replaces the key through fake API headers and CRLF injection. Fixing either source does not fix the other.	8.8	Details
CVE-2026-115748	**UNSUPPORTED WHEN ASSIGNED** Authorization bypass in Apache Vysper, all versions. Apache Vysper is an XMPP (Jabber) server library. Its XEP-0077 (In-Band Registration) implementation is missing an authentication check, allowing any authenticated user to compromise other accounts. As this project is retired, we do not plan to release a version that fixes this issue. Users are recommended to find an alternative.	8.8	More Details
CVE-2026-115938	Craft CMS allows any authenticated control panel user to submit an element-search condition whose element selector is evaluated as a sandboxed Twig object template. Craft's allowed column filter delegates property lookup to Yii ArrayHelper, outside Twig's SecurityPolicy. Repeated single-segment lookups therefore bypass the sandbox's empty property/method allowlists. The vulnerable condition value is ultimately cast to an integer. That does not protect the secret: the template can return author ID 1 when a guessed key character is correct and an impossible ID otherwise. The JSON endpoint returns entries for the first case and an empty array for the second, forming a stable remote character oracle. After recovering GeneralConfig::securityKey, the same user can compute Craft's generic HMAC wire format offline. Signed parameters such as redirect are then accepted and passed to unsandboxed View::renderObjectTemplate(), where Twig's string-callable form can invoke native PHP callables. The regression proof uses the harmless substr() callback and observes abcdef becoming cdef; it does not execute an OS command.	8.8	More Details
CVE-2026-115332	Lotus through 0.10.2 contains an sql injection vulnerability that allows authenticated organization users to inject SQL by registering crafted subscription_filter_keys that bypass keyword-only validation. Attackers can create subscriptions using the malicious key, which is rendered as an unquoted group-by identifier, to inject SQL into usage aggregation queries against the shared multi-tenant database.	8.8	More Details
CVE-2026-114094	Dulwich before 1.2.9 fails to validate DOS drive letter prefixes in Git tree path elements on Windows, allowing arbitrary file write. Attackers can craft a malicious repository with a tree entry named 'C:' that bypasses NTFS validation and writes files outside the worktree to achieve remote code execution.	8.8	More Details
CVE-2026-79877	A remote code execution vulnerability in the underlying operating system of AOS-10 could allow an unauthenticated adjacent attacker to execute arbitrary code on an affected device. Successful exploitation could allow an attacker to execute arbitrary commands as a privileged user on the underlying operating system leading to complete system compromise.	8.8	More Details
CVE-2026-79878	A command injection vulnerability in the AOS-8 command line interface accessed by the PAPI protocol could lead to unauthenticated remote code execution by an adjacent attacker sending specially crafted PAPI packets destined to the PAPI UDP port 8211. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	8.8	More Details
CVE-2026-79879	Memory corruption vulnerabilities in the affected interface of AOS-10 could lead to remote code execution by an unauthenticated adjacent attacker. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	8.8	More Details
CVE-2026-114387	Shopper from 2.0.0 before 2.6.0 contains a missing authorization vulnerability that allows admin-panel users holding only access_dashboard to create administrators and grant permissions. Attackers can call CreateTeamMember::store() or Permissions::togglePermission() directly over /livewire/update to create admin accounts and escalate to full administrator.	8.8	More Details
CVE-2026-115189	A server-side template injection vulnerability exists in the Twig theme rendering sandbox of Mautic. The sandbox introduced to restrict user-uploaded theme templates (the fix for GHSA-9fx4-7cmj-47vg) does not fully enforce its restrictions, allowing a crafted theme template to escape the sandbox. An authenticated user with the core:themes:create permission can upload a malicious theme whose template is later rendered, resulting in execution of arbitrary code on the server with the privileges of the web server process. This allows the attacker to execute operating system commands, read server-side files including configuration files and database credentials, write files, and reach internal services. In default installations the core:themes:create permission is granted to the Administrator role, and in multi-operator setups it may be delegated to other roles. This is an incomplete fix for GHSA-9fx4-7cmj-47vg.	8.8	More Details
CVE-2026-110370	Deserialization of Untrusted Data vulnerability in GS Plugins GS Pins for Pinterest gs-pinterest-portfolio allows Object Injection.This issue affects GS Pins for Pinterest: from n/a through 1.9.2.	8.8	More Details
CVE-2026-110369	Deserialization of Untrusted Data vulnerability in GS Plugins WordPress Team Members – GS Plugins gs-team-members allows Object Injection.This issue affects WordPress Team Members – GS Plugins: from n/a through 2.8.2.	8.8	More Details

CVE-2026-85683	The Jane Menu WordPress plugin before 1.5.1 does not have any capability or nonce checks on one of its form handlers, which is reachable by unauthenticated users, allowing them to bind any published page to a remote document of their choosing. The plugin then fetches that document and outputs its head and body into the page without escaping, so an unauthenticated attacker can run arbitrary JavaScript in the browser of any visitor to that page, including an administrator.	8.8	More Details
CVE-2026-114621	baserCMS is a website development framework. In versions 5.0.0 before 5.3.2 and 5.4.0 before 5.4.2, list and detail actions in the Admin and Api/Admin prefixes pass request query parameters, including contain, conditions, fields, finder, and order, directly to the ORM as options. A user able to log into the admin panel can embed SQL fragments in those options to execute arbitrary SQL against the application's database, including reading password hashes from the users table. Public-page and public-API fixes released in 5.3.1 and 5.4.1 did not protect these administrative paths. Setting USE_CORE_ADMIN_API to false does not prevent access through an authenticated browser session. This issue has been patched in versions 5.3.2 and 5.4.2.	8.8	More Details
CVE-2026-99083	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix heap out-of-bounds write in krb5_authenticate() In krb5_authenticate(), out_len is calculated to determine available headroom in response_buf for the incoming Kerberos AP-REP token. However, when SMB2_SESSION_SETUP is processed as a non-first element of a compounded SMB2 request, the calculation omits work->next_smb2_rsp_hdr_off. This causes out_len to overstate remaining buffer headroom by the cumulative size of prior responses in the compound chain. Consequently, the length check in ksmbd_krb5_authenticate() (*out_len <= resp->spnego_blob_len) passes erroneously, allowing memcpy() to write the AP-REP blob past the end of response_buf into adjacent kernel heap memory. Fix this by subtracting work->next_smb2_rsp_hdr_off when computing out_len, ensuring it accurately reflects physical remaining buffer space.	8.8	More Details
CVE-2026-99082	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Use active memslots for the per-vCPU MMIO cache KVM tags the per-vCPU MMIO cache with a memslot generation so that a memslot update invalidates cached MMIO information. Both the fill and validation paths use kvm_memslots(), which unconditionally selects address space 0, even when the vCPU is running in SMM and using address space 1. Commit 56f17dd3fbc4 ("kvm: x86: fix stale mmio cache bug") added the memslot generation to the cache key so that a memslot update could not leave a stale entry valid. When commit 699023e23965 ("KVM: x86: add SMM to the MMU role, support SMRAM address space") added the SMM address space, these helpers were not converted to use the active memslots. Consequently, an update to the SMM memslots can leave an entry from the old SMM address space apparently valid after the vCPU returns to the normal address space. A guest can then cause an access to valid RAM at the same GFN to be returned to userspace as KVM_EXIT_MMIO. Completing that exit through the VMM's RAM address space writes the backing page without going through KVM's write-tracking path. If the page backs a nested EPT table, KVM can continue using shadow translations derived from the old contents. The WARN_ON_ONCE() in kvm_mmu_write_protect_fault() catches this invalid cache and memslot combination. Use the vCPU's active memslots when caching and validating the entry. Memslot generations are unique across address spaces, so the generation also distinguishes the normal and SMM views. This matches the MMIO SPTE cache, which already uses kvm_vcpu_memslots(). On an unpatched current-mainline kernel, a regression test that fills the cache in SMM and updates only the SMM memslots fails the intended RAM write and triggers the kvm_mmu_write_protect_fault() warning. With this change, the same test completes the RAM write without a warning. The x86/smm_test, set_memory_region_test, and memslot_modification_stress_test selftests also pass.	8.8	More Details
CVE-2026-82814	The Duplicator and Duplicator Pro plugins for WordPress are vulnerable to PHP Object Injection in all versions up to, and including, 1.5.17 and 4.6.10, respectively, via deserialization of untrusted database content in DUPX_UpdateEngine::unserializeTest() during migration. This allows unauthenticated attackers to store a serialized object in public site content before an administrator builds a package and later performs a migration or import that runs database search-and-replace. Duplicator Free bundles a Requests 1.7 POP chain that permits attacker-selected callbacks, resulting in direct operating-system command execution. Duplicator Pro bundles an IndexList destructor gadget that permits deletion of an attacker-selected existing file when its containing directory is writable by the installer PHP process; deletion of wp-config.php can expose the setup wizard and lead to site takeover and code execution if the attacker reaches it first and supplies a reachable database they control. The package must be created after the carrier is stored, and a later administrator migration action is required.	8.8	More Details
CVE-2026-114109	uv before 0.11.15 contains a path traversal vulnerability that allows malicious wheels to write files outside the scripts directory via crafted console_scripts or gui_scripts entry point names. Attackers who induce users to install a malicious wheel can place executables in PATH directories, shadowing or overwriting existing binaries for code execution.	8.8	More Details
	HKUDS Vibe-Trading from 0.1.0 before 0.1.7 contains an indirect prompt injection vulnerability that		

CVE-2026-114107	allows remote attackers to execute shell commands as root by planting instructions in documents or web pages the agent processes. BashTool and BackgroundRunTool pass model-emitted commands to subprocess.run with shell=True and no allow-list or approval step, and read_url fetches arbitrary URLs, enabling command execution and server-side request forgery.	8.8	More Details
CVE-2026-114089	Trigger.dev before 4.5.2 contains a cross-site request forgery vulnerability in the CLI authorization-code flow because loading /account/authorization-code/<code> mints a Personal Access Token without any consent step. Attackers mint an unauthenticated code, lure a logged-in user to that URL, then retrieve the victim's PAT from the unauthenticated /api/v1/token endpoint to take over the account.	8.8	More Details
CVE-2026-98988	In the Linux kernel, the following vulnerability has been resolved: namespace: keep covered mounts covered in OPEN_TREE_NAMESPACE open_tree(OPEN_TREE_NAMESPACE) creates a new mount namespace from a copy of the tree at the given path. The caller needs to be privileged over its current user namespace because the new mount namespace will be owned by it. It doesn't need to be privileged over the mount namespace the tree is copied from. An unprivileged user just needs to create a user namespace first. The copy follows the rules of a detached bind mount. Without AT_RECURSIVE only the mount itself is cloned and its children are left out. With AT_RECURSIVE unbindable mounts are skipped. Both is fine when the caller has privileges over the source mount namespace because it could unmount those mounts anyway. Not so for an unprivileged user in that mount namespace. For them a mount namespace copy is what unshare() does. copy_mnt_ns() copies everything including unbindable mounts and lock_mnt_tree() makes sure nothing can be unmounted in the copy. Nothing that was covered gets revealed. create_new_namespace() only does the locking and so the covered content is right there in the new mount namespace: # mount -t tmpfs none /tmp/otn; mkdir /tmp/otn/covered # echo hidden > /tmp/otn/covered/under.txt # mount -t tmpfs none /tmp/otn/covered uid 1000, after unshare(CLONE_NEWUSER): fd = open_tree(AT_FDCWD, "/tmp/otn", OPEN_TREE_NAMESPACE); setns(fd, CLONE_NEWNS); open("/covered/under.txt", O_RDONLY) -> "hidden" Covering paths with mounts is how container runtimes mask parts of /proc and /sys and how admins hide things. When the caller's user namespace doesn't own the source mount namespace copy the way unshare() does and refuse a non-recursive copy of a mount that has anything mounted below the requested directory and copy unbindable mounts in a recursive copy. A caller that is privileged over the source mount namespace sees no change.	8.8	More Details
CVE-2026-98980	In the Linux kernel, the following vulnerability has been resolved: namei: don't reveal overmounted entries in refwalk In rcuwalk the dentry is validated before it is accepted. step_into() step_into() rechecks d_seq and __follow_mount_rcu() rechecks mount_lock. An entry that gets unlinked in between causes the lookup to retry and miss. A refwalk doesn't do this. lookup_fast() takes a reference on the hashed dentry and simply accepts it. So an unlink that happens after the reference was taken isn't seen by refwalk. That's fine for a simple file. We just happened to open it before it was unlinked, no problem. For a mountpoint and specifically a locked mountpoint it very much isn't. The unlink detaches all mounts and then removes the name. Any refwalk that hasn't traversed the mounts yet simply reveals the underlying entry. It's a very narrow window but it can be hit: reads of a covered file in 60 s, 3 walkers, ~15000 unlinks no widening 27 that step + 200 us 3741 See the appended patch for a more reliable reproducer. So check the entry after step_into(). unlink(), rmdir() and rename() mark the dentry with dont_mount() before they detach the mounts and remove the dentry. So a dentry that is marked with DCACHE_CANT_MOUNT and is unhashed by the time its mounts were looked at is a name that was unlinked under the refwalk. DCACHE_CANT_MOUNT is read after the mounts were looked up and it is set before they are detached by detach_mounts(). A refwalk that missed the mounts will see DCACHE_CANT_MOUNT. Plain d_unlinked() is fine. A rename takes the dentry off its hash chain but __d_drop() leaves d_hash.pprev set. So only __d_drop() and a rename over the dentry unhash it. The only move that flips IS_ROOT splices in a disconnected alias. That doesn't have DCACHE_CANT_MOUNT set. Add the to step_into_slowpath(). "." and LOOKUP_DOWN may legitimately land on an unhashed directory. A refwalk that crossed onto a mount has path.mnt different from nd->path.mnt. A dentry that a filesystem dropped on its own (d_invalidate(), d_drop()) doesn't have the flag set and is treated as before. // SPDX-License-Identifier: GPL-2.0 /* * unlink_covered: unlink a file that is a mountpoint in a detached copy of * its mount, against walkers that open it through that copy. * * A is a tmpfs with the file f1 (content SECRET_F) and the plain file * MARK_A. A' is an open_tree(OPEN_TREE_CLONE) copy of A with MARK_A bound * on A'/f1, held through an O_PATH fd on its root once the tree fd is * closed. Walkers open f1 through that fd while the driver unlinks A/f1, * where nothing is mounted on it. A walker may read MARK_A or get ENOENT. * A read of SECRET_F is a hit: the name was found after its mount was * gone. The lockless walks use openat2(RESOLVE_CACHED). * * usage: unlink_covered [-t seconds] [-w walkers] */ #ifndef _GNU_SOURCE #define _GNU_SOURCE #endif #include <errno.h> #include <fcntl.h> #include <pthread.h> #include <stdatomic.h> #include <stdio.h> #include <stdlib.h> #include <string.h> #include <unistd.h> #include <sys/mount.h> #include <sys/stat.h> #include <sys/syscall.h> #include <linux/openat2.h> #ifndef __NR_open_tree #define __NR_open_tree 428 #endif #ifndef __NR_move_mount #define __NR_move_mount 429 #endif #ifndef __NR_openat2 #define __NR_openat2 437 #endif #ifndef OPEN_TREE_CLONE #define OPEN_TREE_CLONE 1 #endif #ifndef OPEN_TREE_CLOEXEC #define OPEN_TREE_CLOEXEC O_CLOEXEC #endif #ifndef	8.8	More Details

	<pre> MOVE_MOUNT_F_EMPTY_PATH #define MOVE_MOUNT_F_EMPTY_PATH 0x00000004 #endif #define WORK "/tmp/uc" #define ADIR WORK "/"A" static int duration = 60, nwalkers = 3; static atomic_int stop, writer_waiting; static pthread_rwlock_t cur_lock = PTHREAD_RWLOCK_INITIALIZER; static int cur_fd = -1; /* the root of A', -1 while there is none */ static atomic_long n_unlink, n_walk, n_mark, n_enoent, n_othe ---truncated---</pre>		
CVE-2026-56832	PraisonAI is a multi-agent teams system. From version 4.5.2 to before version 4.6.59, praisonai.bots.DiscordApproval approves a pending dangerous tool call when it sees any later non-bot message in the configured Discord channel whose text is classified as approval, such as yes. As a result, any user who can post in the configured approval channel can approve a pending high-risk tool call by sending yes after the approval message appears. the same PoV also shows that the Slack and Telegram messaging approval backends have no explicit approver allowlist parameter, but the primary report-grade issue is the Discord backend's unthreaded channel cross-talk: the approving message does not need to be a reply or otherwise request-bound. This issue has been patched in version 4.6.59.	8.8	More Details
CVE-2026-115331	Lotus through 0.10.2 contains an sql injection vulnerability that allows authenticated organization users to inject SQL by supplying crafted billable metric property_name or categorical filter values. Attackers can embed SQL in values rendered unescaped into Jinja-generated counter, gauge and rate metric queries to access the database shared by all tenants.	8.8	More Details
CVE-2026-115328	JSON-SCADA before 0.65-alpha contains an sql injection vulnerability in opcApi() in src/server_realtime_auth/index.js that allows authenticated users to inject SQL via HistoryReadRequest parameters. Attackers can inject stacked statements through NodeId.Id, StartTime, or EndTime values to read, modify, or delete historian data across all locations regardless of role.	8.8	More Details
CVE-2026-98870	In the Linux kernel, the following vulnerability has been resolved: ksmbd: remove stale channels from all sessions on teardown ksmbd_sessions_deregister() removes a connection's channels from other sessions' channel lists only while conn->binding is still set: if (conn->binding) { hash_for_each_safe(sessions_table, ...) ksmbd_chann_del(conn, sess); } conn->binding is a transient flag: it is cleared once a binding SESSION_SETUP completes, and also by a subsequent non-binding SESSION_SETUP on the same connection (a reauthentication on a bound channel, or a new SessionId==0 setup). A connection that has bound a channel into another session's ksmbd_chann_list and then clears conn->binding leaves that channel behind when it disconnects: the channel, whose chann->conn points at the now freed struct ksmbd_conn, stays on the owner session's list. When the owning connection later tears down, the second loop dereferences the stale channel: xa_for_each(&sess->ksmbd_chann_list, chann_id, chann) if (chann->conn != conn) ksmbd_conn_set_exiting(chann->conn); /* freed */ which is a use-after-free write into the freed ksmbd_conn (the same stale channel is also walked by show_proc_session() through /proc). The session is leaked as well, because its channel list never empties. Remove the conn->binding gate so a connection always removes its channels from every session on teardown.	8.8	More Details
CVE-2026-76530	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76530 are related to use of non-canonical URL paths for authorization decisions and are grouped under the Common Weakness Enumeration (CWE) CWE-647.	8.8	More Details
CVE-2026-76300	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76300 are related to improper control of a resource through its lifetime and are grouped under the Common Weakness Enumeration (CWE) CWE-664.	8.8	More Details
CVE-2026-76298	In Splunk Enterprise versions below 10.4.4 and 10.2.8, a user who does not hold the "admin" or "power" Splunk roles could derive the PostgreSQL sidecar database superuser password and run arbitrary operating system commands as the user account running Splunk Enterprise on search head cluster members. The vulnerability is possible because the PostgreSQL sidecar service derives a privileged database password from the search head cluster identifier that appears in search job logs. For more information see About configuring role-based user access (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/about-configuring-role-based-user-access), server.conf (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.2/configuration-file-reference/10.2.0-configuration-file-reference/server.conf), and Sidecar configuration settings (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.4/splunk-sidecars/sidecar-configuration-settings) in the Splunk documentation.	8.8	More Details
	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user that holds a role with the schedule_search capability could bypass the expected safeguards in the scheduled search workflow. The user could then run risky Search Processing Language (SPL) commands that disclose or alter		

CVE-2026-76291	Splunk Enterprise data or disrupt Splunk Enterprise availability. The vulnerability is possible because the Splunk Enterprise search scheduler does not apply SPL safeguards for risky commands to scheduled searches. 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) in the Splunk documentation. For more information see authorize.conf (https://help.splunk.com/en/data-management/splunk-enterprise-admin-manual/10.6/configuration-file-reference/10.6.0-configuration-file-reference/authorize.conf) in the Splunk documentation.	8.8	More Details
CVE-2026-114658	Grav Login plugin from 3.6.0 before 3.9.13 contains a missing authorization vulnerability in the login.invite task that allows page authors to obtain super user accounts. Attackers can define an invitation form granting super access in page frontmatter, submit it unauthenticated, and register the invited address for full site control.	8.8	More Details
CVE-2026-57117	PraisonAI is a multi-agent teams system. From version 4.6.10 to before version 4.6.59, LocalManagedAgent / SandboxedAgent compute bridging wraps read_file, list_files, and write_file when a compute provider is attached. The bridge converts those file operations into shell command strings using raw path arguments, then sends those strings to shell-backed compute providers. An attacker who can influence a file-tool path argument can break out of the quoted path and execute arbitrary shell commands in the compute environment. With compute="local", commands execute through the local subprocess compute provider on the host. With Docker, commands execute in the container. This issue has been patched in version 4.6.59.	8.8	More Details
CVE-2026-96778	Pillow is a Python imaging library. From 9.4.0 until 13.0.0, ImagingJpegEncode() in src/libImaging/JpegEncode.c can call jpeg_write_marker() for a JPEG COM comment when an ICC profile has exactly filled the 65536-byte encoder buffer. libjpeg writes one byte before checking the exhausted free_in_buffer counter, which then wraps and permits the attacker-controlled comment and subsequent compressed image data to overwrite adjacent heap memory. Re-saving an untrusted image as JPEG while preserving its ICC profile and forwarded comment can crash the process and may enable code execution in suitable runtime and heap layouts. This issue is fixed in version 13.0.0.	8.8	More Details
CVE-2026-116068	Dokploy through 0.30.10 contains an authorization bypass vulnerability that allows owners or admins of any organization to control other tenants' applications by supplying foreign applicationId values. Attackers can call procedures like saveEnvironment, updateTraefikConfig, deploy, or stop to overwrite environment variables and build secrets, repoint sources, and disrupt other tenants' applications.	8.8	More Details
CVE-2026-76476	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76476 are related to issues concerning improper input validation that are grouped under the Common Weakness Enumeration (CWE) Class CWE-20.	8.8	More Details
CVE-2026-116201	Snipe-IT before 8.5.0 contains a SAML assertion replay vulnerability that allows attackers to reuse captured assertions by exploiting a non-atomic nonce check in the SAML acs handler. Attackers can send concurrent POST requests to /saml/acs with one captured assertion to obtain multiple authenticated victim sessions that evade IdP-driven revocation.	8.8	More Details
CVE-2026-114135	Jupyter Server is the backend for Jupyter web applications. From 2.0.0 until 2.6.0, enabling GatewayClient.url replaced the normal KernelWebsocketHandler with WebSocketChannelsHandler, which authenticated the user but did not call authorizer.is_authorized(...) for the execute action on the kernels resource. A restricted authenticated user could therefore open an upstream kernel channel and obtain kernel execution despite the configured Authorizer denying that action, exposing data and credentials available to the kernel and permitting modification or disruption of the kernel environment. The default gateway configuration is fixed in version 2.6.0. In 2.6.0 and later, the same authorization bypass remains possible when a deployment manually registers the deprecated jupyter_server.gateway.handlers module, because its handler still omits the configured kernel-execute authorization decision.	8.8	More Details
CVE-2026-57144	PraisonAI is a multi-agent teams system. From version 4.5.110 to before version 4.6.61, praisonai.sandbox.SandlockSandbox is documented and implemented as the kernel-enforced sandbox backend for untrusted code. Its SandboxConfig.native() path lets callers configure allowed filesystem paths and network=False. On systems where the optional sandlock module imports but reports that Landlock is unavailable, SandlockSandbox.execute() and run_command() do not fail closed. They silently fall back to SubprocessSandbox(self.config). That fallback keeps the same high-level native policy object but does not enforce the native filesystem or network boundary during code execution. A sandboxed payload can read files outside the configured allowed path and open network connections	8.8	More Details

	despite network=False. This issue has been patched in version 4.6.61.		
CVE-2026-116019	Grav 2.0.15 before 2.2.13 contains an allowlist bypass vulnerability that allows page editors to invoke arbitrary public static methods by relocating data-*@ directives with unset-name@ or integer keys. Attackers with page-edit permission can embed crafted fields in Flex form overrides in page frontmatter to call methods like Folder::delete or Utils::download, deleting directories, overwriting files, or reading files outside the page tree.	8.8	More Details
CVE-2026-114045	FreeRDP before 3.32.1 contains a heap-based buffer overflow in the server-side rdpsnd channel because an integer wraparound in rdpsnd_server_select_format() DVI_ADPCM block-alignment arithmetic bypasses validation. A malicious RDP client can send two SNDC_FORMATS PDUs with an nBlockAlign too small for nChannels, causing an unbounded heap write during server audio playback, potentially enabling code execution.	8.8	More Details
CVE-2026-76477	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76477 are related to issues concerning improper neutralization of special elements used in an SQL command that are grouped under the Common Weakness Enumeration (CWE) Base CWE-89.	8.8	More Details
CVE-2026-76478	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76478 are related to improper access control issues that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-284.	8.8	More Details
CVE-2026-112242	Incorrect Privilege Assignment vulnerability in Cozmoslabs Profile Builder profile-builder allows Privilege Escalation.This issue affects Profile Builder: from n/a through 4.0.4.	8.8	More Details
CVE-2026-114177	Cross-Site Request Forgery (CSRF) vulnerability in Xpro Xpro Elementor Addons xpro-elementor-addons allows Cross Site Request Forgery.This issue affects Xpro Elementor Addons: from n/a through 1.7.10.	8.8	More Details
CVE-2026-113870	MongoDB Compass improperly escapes certain BSON values when converting a stored view's aggregation pipeline into editable source. When a user connects to a MongoDB deployment containing a specially crafted view and opens that view for editing, the values can cause arbitrary code to execute with the privileges of the user running Compass.	8.8	More Details
CVE-2026-98981	In the Linux kernel, the following vulnerability has been resolved: namespace: unhash a dentry before detaching the mounts on it unlink(), rmdir() and rename() remove the entry from the filesystem, call detach_mounts() on the entry with the inode locked and then call d_delete() once the inode is unlocked. The thing is that between detach_mounts() and d_delete() the dentry is still hashed and positive. But after detach_mounts() nothing covers the dentry anymore. A lookup that finds this dentry in the dcache can uncover what the mounts hid. That's a problem when unlinking files or directories that are mountpoints in other mount namespaces. Everybody who had the underlying entry covered can race the detach_mount() call until d_delete() has run. The race window isn't all that small. It encompasses namespace_unlock() with a full synchronize_rcu_expedited() grace period and the inode_unlock() of the entry. In my experiments three walkers that kept trying read an overmounted file in 11938 of the 16140 unlinks that removed its mountpoint from a bind mount of the filesystem within a minute. Fun fact, d_invalidate() has the same ordering problem but gets it right. It unhashes the dentry first and detaches the mounts afterwards. Let's do the same in __detach_mounts(): - A lookup that hasn't found the dentry yet misses it in the dcache and waits for the directory lock that the caller holds until the name is gone for good. - A lockless lookup that found it already fails the mount_lock check while the dentry still counts as a mountpoint and the d_seq check once it doesn't, and retries. This fixes the lockless path. We still need to fix the reference count lookup in a follow-up patch. So d_drop() the dentry. The dentry stays positive and held but can't be found anymore. Then proceed with the detach and unlink. Reproducer: The reads were counted with the program below, three walkers for 60 s in a VM with 4 CPUs. The race window is stretched with the debug patch pasted here. fs.detach_race_walk_us sleeps in lookup_fast() once it found a mountpoint dentry and before its mounts are crossed. fs.detach_race_unlink_us sleeps in vfs_unlink() between detach_mounts() and d_delete(). // SPDX-License-Identifier: GPL-2.0 /* * unlink_covered: unlink a file that is a mountpoint in a detached copy of *	8.8	More Details

	its mount, against lookups that open it through that copy. * * A is a tmpfs with the file f1 (content SECRET_F) and the plain file * MARK_A. A' is an open_tree(OPEN_TREE_CLONE) copy of A with MARK_A bound * on A'/f1, held through an O_PATH fd on its root once the tree fd is * closed. Walkers open f1 through that fd while the driver unlinks A/f1, * where nothing is mounted on it. A walker may read MARK_A or get ENOENT. * A read of SECRET_F is a hit: the name was found after its mount was * gone. The lockless walks use openat2(RESOLVE_CACHED). * * usage: unlink_covered [-t seconds] [-w walkers] */ <pre> #ifndef _GNU_SOURCE #define _GNU_SOURCE #endif #include <errno.h> #include <fcntl.h> #include <pthread.h> #include <stdatomic.h> #include <stdio.h> #include <stdlib.h> #include <string.h> #include <unistd.h> #include <sys/mount.h> #include <sys/stat.h> #include <sys/syscall.h> #include <linux/openat2.h> #ifdef __NR_open_tree #define __NR_open_tree 428 #endif #ifdef __NR_move_mount #define __NR_move_mount 429 #endif #ifdef __NR_openat2 #define __NR_openat2 437 #endif #ifdef OPEN_TREE_CLONE #define OPEN_TREE_CLONE 1 #endif #ifdef OPEN_TREE_CLOEXEC #define OPEN_TREE_CLOEXEC O_CLOEXEC #endif #ifdef __NR_move_mount #define MOVE_MOUNT_F_EMPTY_PATH 0x00000004 #endif #define WORK "/tmp/uc" #define ADIR WORK "/A" static int duration = 60, nwalkers = 3; static atomic_int stop, writer_waiting; static pthread_rwlock_t cur_lock = PTHREAD_RWLOCK_INITIALIZER; static int cur_fd = -1; /* the root of A', -1 wh ---truncated---</pre>		
CVE-2026-110018	The aBlocks – Gutenberg Blocks, User Dashboard Builder, Popup Builder, Form Builder & Animation Builder plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 2.17.0 via the `/ablocks/v1/form-builder/registration` REST endpoint. The vulnerability exists because the endpoint is registered with `permission_callback => __return_true` and the `register()` handler reads `roleSlug` directly from attacker-controlled block attributes in `post_content` without ever verifying that the block's `formType` is `registration` — a condition that `RegistrationFormSanitizer::verify()` imposes only at save-time, meaning a Contributor can preserve an arbitrary `roleSlug` by saving a block with `formType='login'` and then triggering the endpoint via a valid `wp_rest` nonce and post-ownership check. This makes it possible for authenticated attackers, with Contributor-level access and above, to escalate their WordPress role to Editor or any other non-administrator role and receive a fully authenticated session at the elevated privilege level via `WP_User::set_role()` and `wp_set_auth_cookie()`. Escalation is bounded below administrator-tier roles because `is_safe_registration_role()` rejects any role holding `manage_options`, `edit_users`, `delete_users`, `create_users`, `promote_users`, or `edit_theme_options`, meaning the highest practically reachable role on a standard WordPress installation is Editor.	8.8	More Details
CVE-2026-116009	Camaleon CMS before 2.9.3 contains a path traversal vulnerability that allows users with only the media permission to write arbitrary files by supplying ../ sequences in the name parameter. Attackers can send a base64 data: URI to the crop_url media action to write an ERB template into a theme's views directory, which executes when rendered.	8.8	More Details
CVE-2026-56840	PraisonAI is a multi-agent teams system. From version 4.5.2 to before version 4.6.59, praisonai.bots.HTTPApproval renders pending tool approval arguments directly into the approval dashboard HTML. An attacker-controlled tool argument can inject JavaScript into that page. When a human opens the approval URL to inspect the risky tool request, the script runs in the dashboard origin and can POST to the same request's /approve/{request_id}/decide endpoint, causing HTTPApproval to return approved=True. This issue has been patched in version 4.6.59.	8.8	More Details
CVE-2026-79880	A vulnerability in the command line interface of AOS-8 could allow an authenticated remote attacker with high privileges to conduct a server-side request forgery (SSRF) attack. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system.	8.7	More Details
CVE-2026-114642	Traefik 2.11.0 before 2.11.58 and 3.0.0 before 3.7.14 stores its NTLM/Kerberos backend round tripper in a client-connection slot not keyed by ServersTransport, so one service's TLS client configuration is reused for another service's backend. An unauthenticated client can arm the slot via a backend issuing an NTLM or Negotiate challenge, then reach another service on the same keep-alive connection, bypassing its root CAs, server name, and certificate verification.	8.7	More Details
CVE-2026-93639	A missing authorization vulnerability in Zimbra Collaboration Suite (ZCS) allows a delegated administrator with no explicit administrative grants to overwrite global Zimlet configuration through ConfigureZimletRequest. The uploaded configuration can also add an attacker-controlled domain to the proxy allowlist of the default class of service. If a signed-in user in that class opens a crafted /service/proxy URL, Zimbra serves attacker-controlled HTML under its own origin. The resulting JavaScript can read mailbox data and perform authenticated actions as the victim. Exploitation requires a delegated-administrator account and access to the admin SOAP service; the browser compromise additionally requires the victim to open the crafted URL. ZCS 10.1.16 and 10.1.20 are affected.	8.7	More Details
CVE-2026-	Traefik from v3.7.0 before v3.7.14 contains an authentication bypass vulnerability in its Kubernetes Ingress-NGINX provider that forwards client-supplied Ssl-Client-* headers on the plaintext HTTP router. When ssl-redirect is disabled alongside auth-tls-pass-certificate-to-upstream, unauthenticated attackers	8.7	More

114640	can send Ssl-Client-Verify: SUCCESS with a chosen certificate subject and issuer to impersonate mTLS identities.		Details
CVE-2026-10632	A reflected cross-site scripting vulnerability in the Classic Web Client of Zimbra Collaboration Suite (ZCS) allows an attacker to inject JavaScript through the packages parameter of the /h/rest document editor. The editor inserts package names into script attributes without HTML encoding. When combined with CVE-2026-50053, an attacker with an ordinary Zimbra account can supply their own authentication token and document identifiers in a URL and execute JavaScript in another signed-in user's Zimbra browser session when that user opens the URL. This permits reading mailbox data and performing authenticated actions as the victim. The issue affects ZCS 10.1.16 and was remediated in 10.1.17.	8.7	More Details
CVE-2026-114639	Traefik 2.11.0 before 2.11.58 and 3.0.0 before 3.7.14 contains an authentication bypass because a first request preauthenticated with an Authorization: Negotiate token authenticates a socket in the globally shared backend connection pool, which is never isolated. Unauthenticated attackers on separate frontend connections can be assigned that NTLM/Kerberos-bound socket and perform reads or writes with the victim's identity.	8.7	More Details
CVE-2026-114637	Traefik from v3.2.0 before v3.7.14 contains an authentication bypass vulnerability in the experimental FastProxy connection pool, which fails to pin NTLM/Negotiate-authenticated backend connections to their frontend connection. Unauthenticated clients can be handed a pooled keep-alive backend connection where a victim completed an NTLM or Negotiate handshake, letting them read victim data and act as that victim.	8.7	More Details
CVE-2026-114636	Traefik v3.7.13 before v3.7.14 contains an access-control bypass in the Kubernetes Ingress NGINX provider, where the HTTP router generated for ssl-passthrough Ingresses on non-TLS entry points never applies the Ingress's access-control annotations. Unauthenticated attackers reaching the HTTP entry point with the Ingress host can access backends protected by auth-type, auth-url, or whitelist-source-range annotations, with nothing logged.	8.7	More Details
CVE-2025-55702	When a client sends a POST with Expect: 100-continue and the origin answers before reading the body, Apache Traffic Server parses the unread body left on the client connection as the next request. An attacker behind a front-end proxy that shares connections to Traffic Server can use this to smuggle requests past the front end and desynchronize responses for other users. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	8.7	More Details
CVE-2026-114635	Traefik before v2.11.58 and v3.7.14 contains an authentication bypass vulnerability in its SNI check, which is skipped when request TLS state is nil. Unauthenticated clients can send HTTP/2 requests with the :scheme http pseudo-header over TLS on Go 1.27 builds to reach routers requiring client certificates without presenting one.	8.7	More Details
CVE-2026-115513	OpenCATS before 0.12.0 contains a stored cross-site scripting vulnerability that allows a low-privileged authenticated user to inject HTML/JavaScript payloads by setting a Candidate, Company, Contact, Job Order, or List name/title field, which is stored unsanitized and later rendered unescaped in the 'Recent' (MRU) navigation panel displayed on nearly every page via lib/MRU.php's getFormatted() function. Attackers can exploit this by crafting a malicious record name that executes arbitrary JavaScript in the session of any other user, including Site Administrators, who views the record, enabling session/cookie theft and privilege escalation to full administrative control.	8.7	More Details
CVE-2026-50053	A request attribute override vulnerability in Zimbra Collaboration Suite (ZCS) allows unauthenticated requests to /h/* to populate servlet attributes intended for trusted internal dispatch. RestFilter copies URL parameters specifying an authentication token, target account, target item and dispatch state into the request attributes consumed by the Classic Web Client. Combined with CVE-2026-10632 on ZCS 10.1.16, this allows an attacker to supply their own valid token and document identifiers and execute JavaScript in another user's signed-in browser session. ZCS 10.1.17 blocks that XSS chain by changing the document renderer, but retains the request attribute override.	8.7	More Details
CVE-2026-114638	Traefik v3.7.0 before v3.7.14 contains an improper authentication vulnerability in the Kubernetes Ingress NGINX provider that merges client CAs when auth-tls-secret names collide after dot-to-dash normalization. Attackers holding certificates from another host's CA can access a host requiring its own client CA, while that host's legitimate clients are refused.	8.7	More Details
CVE-2026-112904	Vvweb is a powerful and easy to use CMS with page builder to build websites, blogs or ecommerce stores. In versions 1.0.8.6 and prior, the nextRestore() and nextBackup() functions in the backup controller accept a file GET parameter and concatenate it directly to DIR_BACKUP without calling sanitizeFileName(). This allows an authenticated administrator to read arbitrary files on the server (via nextRestore) and write SQL dump content to arbitrary locations (via nextBackup). This is inconsistent with the sibling functions restore(), delete(), and download() which all correctly call sanitizeFileName() first. This issue has been patched via commit aa29cb0.	8.7	More Details

CVE-2026-114493	LocalAI v4.6.0 through v4.11.0 contains a server-side request forgery vulnerability in core/gallery/gallery.go that validates only the initial gallery config URL. Attackers can submit a config URL through POST /models/apply that redirects to internal or cloud metadata addresses, receiving those internal responses as gallery content.	8.6	More Details
CVE-2026-56837	PraisonAI is a multi-agent teams system. From version 4.6.56 to before version 4.6.59, PraisonAI's LinearBot starts a public webhook listener on 0.0.0.0 and treats LINEAR_WEBHOOK_SECRET as optional. When the secret is absent, startup only logs a warning and _handle_webhook() skips Linear-Signature verification entirely. An unauthenticated network caller who can reach the webhook endpoint can submit a forged Linear-Event: AgentSession request. The forged request is parsed, scheduled for background processing, dispatched to _handle_agent_session(), and passed into BotSessionManager.chat(). The bot then attempts to post the agent response back to Linear under the configured bot token. The PoV is offline and deterministic. It does not contact Linear. It calls the webhook handler directly, monkey-patches the outbound Linear comment path, and proves both sides of the boundary. This issue has been patched in version 4.6.59.	8.6	More Details
CVE-2026-92998	The Click n Chat (Chat Widget Integration) WordPress plugin before 1.1.3 does not sanitise sort parameters before using them in a SQL ORDER BY clause, allowing unauthenticated attackers to perform SQL injection attacks.	8.6	More Details
CVE-2026-76490	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76490 are related to issues concerning out-of-bounds writes that are grouped under the Common Weakness Enumeration (CWE) Base CWE-787.	8.6	More Details
CVE-2026-111808	Apache Traffic Server can be used to exhaust an HTTP/1.1 origin's worker pool via a malformed padded HTTP/2 DATA frame. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	8.6	More Details
CVE-2026-111805	Apache Traffic Server can be used to exhaust an HTTP/1.1 origin's worker pool via a truncated HTTP/2 Content-Length value. This issue affects Apache Traffic Server: from 8.0.0 through 8.1.11, from 9.0.0 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	8.6	More Details
CVE-2026-92997	The Click n Chat (Chat Widget Integration) WordPress plugin before 1.1.3 does not properly validate a parameter before using it as a column identifier in a SQL query, allowing unauthenticated attackers to perform SQL injection attacks.	8.6	More Details
CVE-2026-115353	Mistune before 3.3.0 contains a path traversal vulnerability in the RST include directive that allows attackers to read arbitrary files by supplying relative or absolute paths. Attackers can submit markdown containing include directives referencing ../ sequences or absolute paths to embed sensitive files like credentials or environment data in rendered HTML.	8.6	More Details
CVE-2026-114661	OpenRemote 1.2.0 before 1.31.0 contains a missing authorization vulnerability in GatewayServiceResource that allows any authenticated user to start, stop and query gateway tunnels without an administrative role. Attackers can send POST or DELETE requests to the gateway tunnel endpoint to open persistent bidirectional TCP circuits into an edge gateway's network or tear down operator tunnels.	8.5	More Details
CVE-2026-115608	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Crocoblock JetEngine jet-engine allows SQL Injection.This issue affects JetEngine: from n/a through 3.9.0.	8.5	More Details
CVE-2026-112822	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in AmentoTech Workreap Core workreap_core allows Blind SQL Injection.This issue affects Workreap Core: from n/a before 3.4.4.	8.5	More Details
CVE-2026-115372	Incus from 7.3.0 before 7.5.0 contains a project restriction bypass that allows restricted project users to set block.create_options through volume update, copy, and move operations. Attackers can set the option on a custom volume and trigger a copy or move onto a formatting pool to inject arguments into mkfs executed as root.	8.5	More Details
CVE-2026-112837	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Ruben Garcia GamiPress gamipress allows Blind SQL Injection.This issue affects GamiPress: from n/a through 8.0.7.	8.5	More Details
	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a mail header injection		

CVE-2026-114716	vulnerability. rmail_sendmail::extract_inline_images() used the media type taken from an inline image data URI without validating it, so a remote authenticated attacker could inject arbitrary MIME part headers into the outgoing message by placing CRLF sequences in that media type. Upstream fixed this in the commit “Fix mail header injection via media type in an inline image data URI”, which validates the media type with rcube_mime::is_mimetype_valid().	8.5	More Details
CVE-2026-115948	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in David Lingren Media Library Assistant media-library-assistant allows Blind SQL Injection.This issue affects Media Library Assistant: from n/a through 3.42.	8.5	More Details
CVE-2026-114663	Vyper 0.4.0b1 through 0.5.0b1 miscompiles unfoldable constants and omitted default arguments from imported modules, because parse_Name resolves their names in the using function's scope rather than at declaration. When a contract's parameter, local, or loop variable shares such a name, attacker-influenced input can replace the intended value, including out-of-range values for narrow types such as uint8, corrupting validation or accounting logic.	8.4	More Details
CVE-2026-98978	In the Linux kernel, the following vulnerability has been resolved: namespace: never expire a locked mount Locked mounts are special. They protect the underlying files and directories from being revealed. do_umount() refuses to unmount locked mounts but shrink_submounts() doesn't. A shrinkable mount can become locked once the owner of a user namespace puts it beneath a locked mount with MOVE_MOUNT_BENEATH. The locked property now moves to the mount at the bottom making it possible to unmount the top mount. So umount() of an unlocked ancestor now expires the bottom mount first and the covered directory is revealed: move_mount(c -> x/hidden, MOVE_MOUNT_BENEATH) = 0 the cover after the lock moved: umount2(x/hidden) = 0 the holder of the lock: umount2(x/hidden) = EINVAL the root of the copy, busy: umount2(x) = EBUSY reads x/hidden/secret: "covered-by-root" So leave a locked mount alone as it dies together with its parent. A plain umount() of an unlocked mount with a locked one below it is EBUSY from now on. It's the same for any other locked child. A lazy umount still takes the whole tree. mark_mounts_for_expiry() never sees a locked mount. lock_mnt_tree() leaves a mount on an expiry list alone and nothing else puts a locked one on such a list. Add an assert for this.	8.4	More Details
CVE-2026-63120	pgbouncer is a PostgreSQL backup daemon and restore tooling that stores backup data in cloud object stores. Prior to version 2.7.2, an incomplete path traversal fix in pgbouncer allows an attacker with write access to the backup object storage (such as S3, GCS, or Azure) to execute arbitrary file writes and directory creation outside the designated PostgreSQL data directory (pgdata). When a database administrator executes a restore operation (via pgbouncer_restore get-basebackup), a maliciously crafted delta backup manifest with path traversal sequences (e.g., ../.././../etc/cron.d/backdoor) will write arbitrary payloads to the host system. This issue has been patched in version 2.7.2.	8.4	More Details
CVE-2026-116237	A vulnerability was found in Oray Pugongying X5 Pro 7.6.1. This affects an unknown part of the component Dnsmasq. Performing a manipulation results in crlf injection. Remote exploitation of the attack is possible. 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.	8.4	More Details
CVE-2026-99042	In the Linux kernel, the following vulnerability has been resolved: io_uring/net: populate CONNECT filter fields for 24-byte sockaddr_in6 io_connect_bpf_populate() fills connect.port and connect.v6_addr for AF_INET6 only when addr_len >= sizeof(struct sockaddr_in6), i.e. 28 bytes. The IPv6 connect paths accept SIN6_LEN_RFC2133 (24 bytes), which already holds sin6_port and sin6_addr. io_connect() passes addr_len to the socket layer unchanged. At addr_len 24 the filter sees family AF_INET6 with port and address zero, while the kernel connects to the real destination. A filter that denies an IPv6 port or address is bypassed. Denying connect.port == P rejects [::1]:P with -EACCES at addr_len 28 but connects at addr_len 24. Use SIN6_LEN_RFC2133 as the threshold. sin6_scope_id stays zero when absent. The threshold mismatch was found by an AI coding assistant (Claude Code) comparing this hook against the net/ipv6 length checks on my instructions; it also drafted the reproducer and this patch. I verified the bypass and the fix in a VM on 7.3-rc6 and that 7.2 has the same code.	8.4	More Details
CVE-2026-114727	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 passes potentially unsafe content into ImageMagick when resizing images. SVG and EPS were accepted as valid image types for contact photo upload and conversion, and both formats can reference external resources and invoke ImageMagick delegates when processed. Upstream fixed this in the commit “Fix passing potentially unsafe attachment content into ImageMagick on image resize”, which removes svg and eps from the accepted types.	8.3	More Details
CVE-2026-79883	A buffer overflow vulnerability in the PAPI protocol of HPE Networking AOS-10 could allow an unauthenticated adjacent attacker to access sensitive information. Successful exploitation could allow an attacker to retrieve information which could be used to potentially gain further access to network services supported by AOS-10, and to interrupt the normal operation of the affected service.	8.3	More Details
CVE-	There is a command injection vulnerability in the PAPI protocol of AOS-10 that could lead to		

2026-79882	unauthenticated remote code execution by an adjacent attacker sending specially crafted packets. Successful exploitation of this vulnerability could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	8.3	More Details
CVE-2026-114068	Trigger.dev before 4.6.0 contains a missing authorization vulnerability in the waitpoint token wait endpoint that allows authenticated attackers to block runs belonging to other organizations. Attackers holding any secret API key can supply a victim run ID with their own waitpoint, causing that run to block indefinitely on a waitpoint the victim cannot complete, and can inject resume data into it.	8.3	More Details
CVE-2026-79881	An authentication bypass vulnerability in the underlying operating system of AOS-8 could allow an unauthenticated remote attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system leading to complete system compromise.	8.3	More Details
CVE-2026-58202	OpenReception's appointment booking software provides an end-to-end encrypted appointment booking platform. Prior to version 2.0.0, the account-confirmation link in emails is built from the request Host header, with no allowlist. The unauthenticated POST /api/auth/resend-confirmation rotates a pending invitee's confirmation token and emails a link whose host comes from the request. An attacker who sends one resend-confirmation with Host: attacker.example and a pending invitee's email causes that invitee to receive a confirmation email whose only link points to the attacker. If the invitee clicks, the one-time confirmation token reaches the attacker, who then confirms the account on the real tenant domain and registers a passkey they control, taking over the account before the legitimate invitee ever uses it. All released versions through v1.1.0 are affected; no patched version exists yet. Reaching the app with a spoofed Host depends on the front proxy: the stock Caddy catch-all forwards it, while an operator who configures explicit per-tenant proxy blocks is not affected (see Details 4). The invite path builds links from the trusted tenant.domain (invite/+server.ts:204-210); the resend-confirmation path builds them from the request Host. This issue has been patched in version 2.0.0.	8.3	More Details
CVE-2026-70388	SOEM version 2.0.0 and earlier contains vulnerabilities where a malicious device on the network can send specially crafted messages that cause the master to write data outside allocated memory boundaries. The software performs arithmetic on message length values without proper validation, allowing negative or extremely large values to bypass safety checks. This can corrupt memory and potentially allow an attacker to execute arbitrary code on the master system.	8.3	More Details
CVE-2026-56835	PraisonAI is a multi-agent teams system. From version 3.11.0 to before version 4.6.59, PraisonAI's Slack bot applies its configured allowed_users, allowed_channels, and unknown-user pairing policy in the normal Slack message event handler, but not in the adjacent Slack app_mention event handler. A Slack workspace user who can mention the bot in a channel where the Slack app is present can trigger the configured PraisonAI agent even when: the sender is not in BotConfig.allowed_users; the channel is not in BotConfig.allowed_channels; unknown_user_policy="deny" is configured; and the same event content is correctly dropped by the normal message handler. This issue has been patched in version 4.6.59.	8.3	More Details
CVE-2026-115325	Netmaker through 1.7.0 contains an authorization bypass vulnerability in the updateExtClient handler that allows network-scoped users to modify external clients on other networks. Attackers with external client permissions on one network can send PUT requests to /api/extclients/{network}/{clientid} to replace public keys, disable clients, and obtain WireGuard private keys.	8.3	More Details
CVE-2026-116067	Paisa before 0.7.8 contains a path traversal vulnerability in the GetFile, GetSheet, DeleteBackups, and DeleteSheetBackups handlers that lets unauthenticated attackers read files and delete backups through ../ sequences in file names. Attackers can send crafted names to the editor file and delete_backups endpoints to read any server-readable file and remove backup files outside the journal or sheets directory.	8.2	More Details
CVE-2024-58393	i-Netsolution Readymade Unilevel Ecommerce MLM (2024-03-15 build) contains an SQL injection vulnerability that allows unauthenticated attackers to manipulate database queries through the 'id' parameter. Attackers can send GET requests to /product-details.php with boolean-based, stacked, or time-based SLEEP payloads in 'id' to extract database information. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-09-14.	8.2	More Details
CVE-2022-51023	PHPOK 4.5.000 before 5.4.081 contains an SQL injection vulnerability that allows unauthenticated attackers to manipulate ORDER BY clauses via the sort parameter. Attackers can send GET requests to /api.php?c=project&f=index with extractvalue payloads in sort to extract database information through error messages. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	8.2	More Details
CVE-2026-79884	A vulnerability in the web-based management interface of AOS-8 could allow an unauthenticated adjacent attacker to bypass existing authentication controls and access sensitive information. Successful exploitation could allow an attacker to gain access to some data in a cleartext format	8.2	More Details

	possibly exposing other network infrastructure to further compromise.		
CVE-2022-51021	NotificationX plugin for WordPress 2.0.0 before 2.3.12 contains an unauthenticated SQL injection vulnerability that allows remote attackers to manipulate database queries by supplying an array-form 'id' parameter to /wp-json/notificationx/v1/notification/{id}. Attackers can pass id[1] with crafted SQL, such as time-based SLEEP payloads, along with an api_key derived from md5(home_url()), which flows unsanitized into Database::get_where_query to extract database contents. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-18.	8.2	More Details
CVE-2026-79885	An unauthenticated file write vulnerability exists in the PAPI protocol in AOS-8. Successful exploitation could allow an unauthenticated remote attacker to make limited modifications to files on the affected system and to disrupt the availability of the affected service.	8.2	More Details
CVE-2026-79886	An unauthenticated buffer overflow vulnerability exists in the PAPI protocol of AOS-8. Successful exploitation could allow an unauthenticated remote attacker to disrupt the availability of the affected service on the controller, with a limited impact to its integrity.	8.2	More Details
CVE-2026-115193	A stored Cross-Site Scripting (XSS) vulnerability exists in the contact detail view of Mautic 7. Values of HTML-type custom contact fields are rendered without output sanitization. A user or integration able to write a value to such a field (for example, through contact editing, the contact API, or a form submission mapped to the field) can store a script payload in it. When another user, including an administrator, opens the affected contact's detail page, the injected script executes within the context of their active browser session. This could allow an attacker to perform administrative actions on behalf of the victim, such as creating new administrator accounts, altering system configuration, or accessing contact and campaign data.	8.2	More Details
CVE-2026-116211	YesWiki 4.5.0 before 4.7.0 contains an SQL injection vulnerability in PageManager::save(), where the date_maj_fiche field of a Bazar entry-creation request is concatenated unescaped into an INSERT statement. Unauthenticated attackers on default installs can inject a sub-SELECT against wiki_users to exfiltrate password hashes.	8.2	More Details
CVE-2019-25778	SeaCMS through 8.7 contains an unauthenticated SQL injection vulnerability that allows remote attackers to manipulate database queries by injecting SQL through the rlist[] array parameter. Attackers can send GET requests to /comment/api/index.php with crafted rlist[] values using backtick-quoted variables and extractvalue to extract database information via error messages. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	8.2	More Details
CVE-2026-111804	Apache Traffic Server's experimental wasm plugin has a use-after-free in its transaction-close handler and forwards unvalidated guest-supplied metric IDs to the internal stats API. This issue affects Apache Traffic Server: from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 10.1.5 or 10.2.1, which fix the issue.	8.2	More Details
CVE-2026-112085	webpack-dev-server is a development server that provides live reloading for webpack projects. In versions before 5.2.7 and in version 6.0.0, when a proxy entry is configured to forward WebSockets, the server forwards non-HMR WebSocket upgrade requests to the proxy target without validating the Host or Origin headers, so the allowedHosts protection does not apply to those proxied upgrades. A malicious web page visited by a developer can open a cross-origin WebSocket through the dev server and read from and write to the proxied backend, even when the server is bound to loopback. The issue is fixed in webpack-dev-server 5.2.7 and 6.1.0, where proxied WebSocket upgrades pass the same Host and Origin validation as the HMR WebSocket, and users should upgrade to one of those versions or later. As a workaround, omit WebSocket forwarding from proxy entries that do not need it and scope the proxy to specific paths.	8.2	More Details
CVE-2026-116116	A vulnerability was detected in Yoosee GM8135S Camera 21.00.01.36. Affected is an unknown function of the component Secure Boot. Performing a manipulation results in improper verification of cryptographic signature. The attack needs to be approached locally. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	8.2	More Details
CVE-2026-107237	The Backup and Staging by WP Time Capsule WordPress plugin before 1.22.28 does not perform capability or nonce checks on the handler that deploys a staging site to the live site, allowing any authenticated user with a Subscriber account on a WP Time Capsule staging site to overwrite the separate live site's database and configuration and take it offline.	8.2	More Details
CVE-2026-56836	PraisonAI is a multi-agent teams system. From version 4.5.112 to before version 4.6.59, PraisonAI's installed console endpoint is Typer-first. In current releases, the recipe command is registered in the Typer app and praisonai recipe serve dispatches to the deprecated Typer command in src/praisonai/praisonai/cli/commands/recipe.py. That Typer command can start the Recipe HTTP server on a non-localhost interface with no authentication. It prints a deprecation warning, and then launches the server. Because config.auth is absent, create_app() does not attach the API-key or JWT middleware. Unauthenticated requests can then reach the recipe API and, when enabled, /admin/reload. This issue	8.2	More Details

	has been patched in version 4.6.59.		
CVE-2026-15723	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 10.1 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, when OmniAuth SAML and LDAP auto-linking were both enabled, an authenticated user could have taken over an existing LDAP-backed account without any interaction from the targeted user, due to improper input sanitization in the LDAP email lookup path used during OmniAuth identity linking.	8.2	More Details
CVE-2025-71434	Weaver E-cology 8.0, 9.0, and 10.0 prior to 10.76 contains an unauthenticated SQL injection vulnerability in the /api/doc/out/more/list endpoint via the elementmore parameter with srcType=2. Attackers can inject SQL through doc_datefrom/doc_dateto fields and read results via /api/ec/dev/table/counts to extract database information. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-07-19.	8.2	More Details
CVE-2026-116176	vLLM from 0.9.2 before 0.24.0 contains a server-side request forgery vulnerability that allows API callers to force outbound connections by embedding a decode address in the X-Request-Id header. Attackers can make a P2pNcclConnector prefill worker ZMQ-connect to an attacker host, leaking the NCCL unique_id and KV-cache tensors or blocking the worker.	8.2	More Details
CVE-2026-79977	The WireGuard VPN service in subsys/net/lib/wireguard/wg.c updated the peer endpoint from a message it could not authenticate. In handle_transport_data(), when the 32-bit receiver index of an incoming transport-data message matched no session, the code both scheduled a new handshake and copied the datagram's source address into peer->endpoint with memcpy(&peer->endpoint, peer_addr, sizeof(peer->endpoint)). Because the receiver index is unknown, there is no key with which such a message could be verified, so the address written into peer->endpoint is entirely attacker-supplied. Every other endpoint update in the implementation happens only after authentication — after the Noise handshake in handle_handshake_init(), or after AEAD decryption plus the anti-replay check in wg_process_data_message(). The WireGuard listener is registered with net_udp_register() using no remote address or port filter and defaults to UDP port 51820 (CONFIG_WIREGUARD_PORT), and wg_ctrl_recv() resolves the peer by interface rather than by any cryptographic identity. Any host able to send a UDP datagram to that port can therefore reach the defect off-path, with no need to observe legitimate traffic: a 16-byte datagram carrying the transport-data type, zeroed reserved bytes and an arbitrary receiver index is sufficient. The vulnerable branch runs whenever the peer has neither a valid handshake nor a current keypair — the state at boot and after a session expires. A single spoofed datagram redirects peer->endpoint to an address of the attacker's choosing, and all subsequent handshake initiations, keepalives and encrypted data packets are sent there. The legitimate peer never receives the initiations and so never re-authenticates, and nothing resets the endpoint afterwards, leaving the tunnel persistently unusable on devices whose peer does not initiate handshakes itself. The attacker can also relay the redirected initiations to the real peer to place itself in the path of the tunnel, observing ciphertext and traffic metadata and selectively dropping packets. WireGuard's end-to-end cryptography still prevents decryption or forgery of tunnel content, so the impact is loss of availability plus limited confidentiality, not key compromise. The fix removes the endpoint learning from this path entirely: the stale-session recovery is kept, but the handshake is sent only to an endpoint learnt from an authenticated message, falling back to the configured cfg_endpoint, and the message is dropped when neither is known. CONFIG_WIREGUARD is marked experimental and is not enabled by default.	8.2	More Details
CVE-2026-94300	Replaying GSS-TSIG signed UDP messages can result in memory corruption in "`named`". An attacker who can capture a previously signed UDP message can replay it while the negotiated TSIG key is still valid. By also sending crafted queries to a zone the attacker controls, the attacker may potentially achieve remote code execution. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.29, 9.21.0 through 9.21.26, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.29-S1.	8.1	More Details
CVE-2026-111594	InterSystems IRIS for Health has TLS/SSL encryption turned off by default, which may enable an attacker to view transmitted data and use that data to log into the device.	8.1	More Details
CVE-2026-112408	Missing verification of source of a communication channel in the default_bucket() function in Amazon SageMaker Python SDK V3 before 3.10.0 and V2 before 2.257.3 might allow remote actors to achieve code execution in the target SageMaker execution environment and obtain IAM credentials via pre-creation of a predictably named S3 bucket that the SDK uses without ownership validation. To remediate this issue, users should upgrade to Amazon SageMaker Python SDK 3.10.0 (V3) or 2.257.3 (V2).	8.1	More Details
CVE-2026-115419	The A2A server component in Amazon Strands Agents SDK for Python before 1.43.0 and Amazon Strands Agents SDK for TypeScript before 1.5.0 does not properly isolate conversation state between callers. An authenticated remote user could read other tenants' conversation history and add persistent instructions that alter model behavior for subsequent callers via HTTP requests to the A2A endpoint using a different context_id. To mitigate this issue, users should upgrade to strands-agents 1.43.0	8.1	More Details

	(Python) or @strands-agents/sdk 1.5.0 (TypeScript).		
CVE-2026-115727	AgenticSeek before commit 409840c contains a remote code execution vulnerability because llm_server/app.py hardcodes Flask debug mode while binding to all interfaces. Attackers who reach the server port and obtain the derivable Werkzeug debugger PIN can execute Python code and operating system commands with llm_server privileges.	8.1	More Details
CVE-2026-113874	SmileCDR HAPI FHIR allows an untrusted bundle body to replace a partition identity previously established by a trusted source, and does not enforce storage partition entitlement by default. An attacker authorized for one tenant can read from or write to another tenant's partition.	8.1	More Details
CVE-2026-115324	mem0 0.1.117 before 1.0.8 contains an sql injection vulnerability in mem0/vector_stores/databricks.py that allows attackers to inject SQL by influencing stored memory content. Attackers can supply crafted memory text, metadata, or message names that reach unescaped insert() and update() string interpolation to execute injected SQL on the configured Databricks SQL warehouse.	8.1	More Details
CVE-2026-79887	A vulnerability in the web interface of AOS-8 could allow an unauthenticated remote attacker to achieve remote code execution. Successful exploitation could lead to complete device control, enabling interception of network traffic and potential compromise of connected systems.	8.1	More Details
CVE-2026-97635	The Agent Zero Web UI before 1.15 does not correctly sanitize user-provided input in chat messages or file uploads. When Agent Zero reproduces content in the chat interface, any injected JavaScript is executed in the context of the victim's authenticated Agent Zero Web UI session.	8.1	More Details
CVE-2026-114069	Trigger.dev before 4.6.0 contains an improper authorization vulnerability in the GET /api/v1/projects/:projectRef/envvars endpoint, which returns the environment's root API key as TRIGGER_SECRET_KEY for staging, production, and preview environments. Attackers holding a restricted read:envvars key can obtain the root secret and plaintext secrets, then trigger tasks and mint admin public JWTs.	8.1	More Details
CVE-2026-112852	Deserialization of Untrusted Data vulnerability in ThemeREX Group Helion helion allows Object Injection. This issue affects Helion: from n/a through 1.1.12.	8.1	More Details
CVE-2026-114403	PHPOffice PHPPresentation through 1.2.0 contains an insecure deserialization vulnerability that allows attackers to instantiate arbitrary PHP objects by supplying a crafted presentation file. Attackers can upload a ZIP containing PhpPresentation.xml that IOFactory::load passes to unserialize(), enabling gadget-chain code execution, file writes, or denial of service.	8.1	More Details
CVE-2026-114087	Trigger.dev before 4.5.6 self-hosted Docker Compose deployments using the default secrets from hosting/docker/.env.example allow authentication bypass because the hardcoded MAGIC_LINK_SECRET lets anyone forge valid magic link tokens. Unauthenticated attackers can log in as any email address, then reach Postgres, Redis, and ClickHouse with default credentials from runner containers on infrastructure networks, compromising all tenant data.	8.1	More Details
CVE-2026-116014	open62541 1.5.0 before 1.5.7 contains an integer overflow in UA_DataType_fromStructureDescription() that wraps the 16-bit memSize field for large custom structure definitions. A malicious OPC UA server can supply DataType definitions via UA_Client_getRemoteDataTypes(), causing decodeBinaryStructure() to overflow the heap, crashing clients or enabling code execution.	8.1	More Details
CVE-2026-114066	Trigger.dev before 4.7.0 contains a missing authorization vulnerability that allows holders of unrestricted environment API keys to block runs they do not own by supplying a foreign parentRunId to batch-creation endpoints. Attackers can send requests to POST /api/v3/batches or POST /api/v2/tasks/batch with resumeParentOnCompletion to suspend, indefinitely block, or inject errors into production or cross-organization runs.	8.1	More Details
CVE-2026-79888	A vulnerability in a certain web interface endpoint of AOS-8 devices could allow an unauthenticated remote attacker to trigger a buffer overflow. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system, leading to complete system compromise.	8.1	More Details
CVE-2026-79889	Vulnerabilities in multiple web interface endpoints of AOS-8 devices could allow an unauthenticated remote attacker to trigger a buffer overflow by using specially crafted requests. Successful exploitation could allow an attacker to execute arbitrary commands on the underlying operating system, leading to complete system compromise.	8.1	More Details
CVE-2026-114078	Trigger.dev before 4.5.6 contains an authorization bypass vulnerability in DeleteTaskScheduleService and SetActiveOnTaskScheduleService that allows authenticated users to delete or disable other organizations' schedules. Attackers can send delete or disable actions through their own project's schedule route with a victim schedule friendlyId to destroy other organizations' IMPERATIVE cron schedules.	8.1	More Details

CVE-2026-116208	YesWiki from 4.6.2 before 4.7.0 contains a missing authorization vulnerability in ActivityPubService::syncActorPosts that lets malicious ActivityPub servers manipulate other actors' mirrored entries. Attackers can serve a poisoned outbox with cross-actor Delete, Update, or Create activities during admin-initiated sync to delete, overwrite, forge, or mass-delete entries.	8.1	More Details
CVE-2026-98968	In the Linux kernel, the following vulnerability has been resolved: amt: send the relay's General Query directly from the receive path An skb queued in a qdisc can outlive the tunnel it references through a raw pointer in skb->cb. For example, with igmp_qrv set to 1 on the relay a tunnel lives for 135s, so a netem delay of 180s on the amt device outlives it; when the tunnel expires and is freed, the subsequent dequeue triggers a use-after-free in amt_dev_xmit(). BUG: KASAN: slab-use-after-free in amt_dev_xmit+0x2763/0x2e20 Call Trace: amt_dev_xmit+0x2763/0x2e20 [drivers/net/amt.c:1262] dev_hard_start_xmit+0x22f/0x620 sch_direct_xmit+0x12e/0xac0 netem_dequeue+0x333/0xc50 net_tx_action+0x35c/0xa60 amt_send_igmp_gg() and amt_send_mld_gg() are only called from amt_request_handler(), inside the rcu_read_lock_bh() section of amt_rcv(). amt_request_handler() already has the tunnel the query is for: it found or created it inside that section. Queuing the query with dev_queue_xmit() only leads back into amt_dev_xmit(), which strips the Ethernet header and calls amt_send_membership_query() for that tunnel. Make that call directly from the two senders instead, the same way amt_send_advertisement() transmits from the receive path. The query never waits in a qdisc, the tunnel is only dereferenced inside the RCU section that found or created it, and nothing is stored in skb->cb, so no lookup or refcount is needed. Remove the query branch of amt_dev_xmit(), amt_skb_cb() and struct amt_skb_cb, which have no users left. Behaviour changes: - The relay's own General Queries no longer pass through the amt device's egress path: its qdisc, tc egress (clsact/tcx), the netfilter egress hook and packet taps. They are still visible as UDP on the underlay. - A query that is sent successfully is no longer counted as tx_dropped. The old query branch left through the unlock label, which counted every sent query as dropped. - A query that reaches amt_dev_xmit() on a relay from elsewhere, such as a userspace querier, is now dropped at the IGMP/MLD type switch. Before, it trusted whatever skb->cb held, and a NULL tunnel hit the WARN_ON(1).	8.1	More Details
CVE-2026-79890	There are command injection vulnerabilities in AOS-8 that could lead to unauthenticated remote code execution by sending specially crafted PAPI packets. Successful exploitation of this vulnerability results in the ability to execute arbitrary code as a privileged user on the underlying operating system.	8.1	More Details
CVE-2026-114489	Missing authorization in the REST and COM administration APIs of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote authenticated user whose credential is confined to named domains to change mail routing and domain configuration beyond those domains on a multi-tenant server. A domain administrator, an account or API key given another role that manages accounts in named domains, or a full-scope API key issued for named domains could create a domain alias for any name (so that other tenants' mail for that name is diverted to an address the attacker controls, or denied), rename a domain to any name the server does not host, set a rule's bind-to-address, send-using-route or script-function action to reach one of the server's own addresses, a configured route's stored credentials or a function of the server's event script, and set the domain's relay host - which delivery then connects to with no address rule - to an address on the server's own network; and over the COM API any account could set the three rule actions on its own rules, where the script-function action runs at delivery. The authorization gate did not refuse these routes to a confined credential and the handlers were not given the caller; over COM the domain-alias save performed no authorization check at all, the domain rename required only a domain administrator, and a rule save required at most that the caller could change the rule's owner.	8.1	More Details
CVE-2026-79891	A buffer overflow vulnerability exists in the command line interface of AOS-8 that could allow an unauthenticated remote attacker to run arbitrary commands on the underlying host if certain preconditions outside of the attacker's control are met. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system leading to complete system compromise.	8.1	More Details
CVE-2026-94380	Weak Password Recovery Mechanism for Forgotten Password vulnerability in HAVELSAN Inc. Liman MYS allows Password Recovery Exploitation. This issue affects Liman MYS: from 2.3.4-1124 before 2.3.5-1131.	8.1	More Details
CVE-2026-112343	Syncplify Server! before 8.8.0 contains an access control bypass vulnerability that allows authenticated SFTP users to evade per-path permission overrides by exploiting case-insensitive filesystem path resolution and 8.3 short names. Attackers can address a protected directory using alternate capitalization or its 8.3 short name to perform file uploads and downloads that the configured override is intended to deny, resulting in unauthorized file access.	8.1	More Details
CVE-	Rsyslog is a rocket-fast system for log processing. From versions 1.15.0 to 8.2610.0, rsyslog's omfile dynaFile facility intentionally allows templates to generate flexible output paths. This supports valid deployments that route logs into different directories or use dynamically generated file names. A deployment becomes unsafe when a dynaFile template uses a value influenced by untrusted input as a		More

2026-92709	path component without explicit secure-path handling. In that configuration, the rendered path can escape the operator's intended log directory. This is not limited to the RFC5424 HOSTNAME field. Any untrusted value that reaches a dynaFile path component can create the condition. At time of publication, there are no publicly available patches.	8.1	Details
CVE-2026-116199	Snipe-IT before 8.6.3 contains an authentication bypass in GoogleAuthController::handleGoogleCallback, which logs in the user whose username matches the Google account email without enforcing the activated flag. Deactivated users with email-shaped usernames and still-live Google identities can sign in via Google SSO and regain sessions at their original permission level, including admin.	8.1	More Details
CVE-2026-116204	YesWiki before 4.7.0 contains a cross-site request forgery vulnerability in the despam action (tools/security/actions/despam.php) that lets unauthenticated attackers delete or restore pages through an administrator's session. Attackers can plant an auto-submitting raw HTML form on a wiki page, so an administrator who views it posts clean, suppr[], or rev[] parameters that delete PagePrincipale, configuration pages, or overwrite current content.	8.1	More Details
CVE-2026-114083	Trigger.dev before 4.5.0-rc.3 contains a path traversal vulnerability in the api/v1/packets and api/v2/packets endpoints that allows authenticated tenants to escape their object key prefix through an unsanitized filename splat segment. Attackers holding one valid API key can supply ../ sequences to obtain presigned GET or PUT URLs and read or overwrite other tenants' run payloads and outputs.	8.1	More Details
CVE-2026-115546	Okta Access Gateway (OAG) does not validate the selected authentication source during sign-on. An authenticated user with an active session can modify the authentication source parameter, which may generate a SAML assertion granting access to the Admin Console.	8.1	More Details
CVE-2026-112345	TTLock Android application before v9.8.260718 and associated smart-lock firmware contain a hardcoded static AES key vulnerability that allows an attacker within Bluetooth range to decrypt the per-pairing session key transmitted during BLE pairing. Attackers can observe the pairing exchange, recover the persistent session key using the embedded defaultKey, and forge authenticated BLE requests to unlock the device, add fingerprints or access codes, and perform administrative actions without authentication or user interaction.	8.1	More Details
CVE-2026-114662	Obot before 0.26.0 contains an authentication bypass vulnerability in the CLI login flow that allows unauthenticated attackers to steal API keys by choosing token request IDs. Attackers can create a token request via POST /api/token-request, send the login link to a signed-in user, and fetch the resulting non-expiring API key to take over the victim's account.	8.1	More Details
CVE-2026-57121	PraisonAI is a multi-agent teams system. From version 0.1.4 to before version 0.1.5 in praisonai-platform, a workspace member can permanently delete any resource — projects, agents, issues, labels, issue dependencies, and issue-label attachments — created by the workspace owner or other members. All six content DELETE endpoints enforce workspace membership but perform no ownership or role check. A single malicious or compromised member account can wipe an entire workspace's content irreversibly. This issue has been patched in version 0.1.5.	8.1	More Details
CVE-2026-99064	In the Linux kernel, the following vulnerability has been resolved: cipso: adjust cached option offsets when removing CIPSO cipso_v4_skbuff_delattr() removes the CIPSO bytes but does not adjust cached offsets for options that follow them. For example, with a valid 10-byte CIPSO option followed by a seven-byte RR option, parsing records rr at offset 30. Removing CIPSO moves RR to offset 20, while the cached offset remains 30. Consumers such as ip_forward_options() and __ip_options_echo() then access the wrong bytes; the latter may interpret packet data as the option length and copy it into fixed-size option storage. Mirror cipso_v4_delopt() and subtract cipso_len from the srr, rr, ts and router_alert offsets when they follow CIPSO. cipso_len is the distance the first memmove() shifts those options. The later header move and network-header reset relocate the bytes and their offset base together.	8.1	More Details
CVE-2026-96644	Pillow is a Python imaging library. Prior to 13.0.0 on Windows, src/PIL/WmfImagePlugin.py passes unchecked 32-bit WMF or EMF dimensions to the native drawwmf handler in src/display.c. The handler sizes a DIB allocation from 16-bit-truncated bcWidth and bcHeight fields but copies pixel data using the original full dimensions. A crafted wide, short metafile can pass the pixel-count decompression-bomb check and make the Windows-only renderer read beyond the heap buffer, potentially disclosing memory or crashing the process. This issue is fixed in version 13.0.0.	8.1	More Details
CVE-2026-114082	Trigger.dev before 4.6.2 contains an authorization vulnerability that allows public-token holders with session write access to obtain environment-wide S2 stream credentials via session input initialization. Attackers can use the credential directly against the storage API to append, create, or trim other session and run streams, forging output or permanently deleting records.	8.1	More Details
CVE-	vLLM from 0.7.3 before 0.31.0 forwards the request-supplied mm_processor_kwargs.code_revision to Transformers AutoProcessor.from_pretrained, letting clients select which revision of a model's dynamic		

2026-114053	processor code is imported when trust_remote_code is enabled. An attacker controlling a revision in the served model repository can send a multimodal request to the unauthenticated /tokenize endpoint to execute arbitrary Python in the vLLM API process.	8.1	More Details
CVE-2026-57113	PraisonAI is a multi-agent teams system. From version 2.6.0 to before version 4.6.59, PraisonAI's template loader accepts GitHub template URIs with refs, for example github:owner/repo/template@v1.0.0. The resolver stores the user-controlled template path and ref verbatim, and the cache layer later joins those values into ~/.praison/cache/templates/github/<owner>/<repo>/<template>/<ref> without normalizing each segment or checking that the final path remains inside the template cache root. A crafted ref such as ../../../../outside-delete-target therefore escapes the cache directory. The first load can write .cache_meta.json outside the cache. If the normal cache hierarchy for the same owner/repo/template has already been created, the same path reaches shutil.rmtree(cache_path) and removes an attacker-selected outside directory before replacing it with cache metadata. This issue has been patched in version 4.6.59.	8.1	More Details
CVE-2026-116149	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a stack-based buffer overflow in the handling of bzip2 compressed images. Attackers can supply a crafted image using bzip2 compression to overwrite stack memory during decoding, causing crashes or memory corruption.	8.1	More Details
CVE-2026-112820	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in Elated-Themes Academist Core academist-core allows PHP Local File Inclusion.This issue affects Academist Core: from n/a through 1.4.1.	8.1	More Details
CVE-2026-112850	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in AncoraThemes Your Dress your-dress allows PHP Local File Inclusion.This issue affects Your Dress: from n/a through 1.4.1.	8.1	More Details
CVE-2026-109619	PHP has various global state that temporarily points into request memory. In some scenarios, global state may not be correctly reset, or may be repopulated after being reset. The global state may then refer to request memory even after it has been freed. Subsequent requests may read or branch on this stale state before it is fully initialized. Multiple strategies can trigger this condition and produce remote code execution (RCE).	8.1	More Details
CVE-2026-94262	The Advanced Custom Fields: Extended WordPress plugin before 0.9.2.8, Advanced Custom Fields: Extended PRO WordPress plugin before 0.9.2.8 does not verify that the requester is authorised to edit the targeted user account in the update-user action of its front-end Forms module, performing that check only when the target account holds an administrator role, which allows unauthenticated visitors to overwrite any non-administrator account's credentials and obtain a session for that account. Exploitation requires the site to publish a form whose user-update action maps its target to a visitor-submitted value rather than to the default current user.	8.1	More Details
CVE-2026-110523	The WP User Manager WordPress plugin before 2.9.26 does not keep its paid-registration capability restrictions in force once the payment gateway is disconnected, while a published paid registration form continues to assign its configured role, allowing unauthenticated users to register and immediately obtain a privileged administrator-equivalent role without completing the payment meant to gate it.	8.1	More Details
CVE-2026-94268	The Advanced Custom Fields: Extended WordPress plugin before 0.9.2.8, Advanced Custom Fields: Extended PRO WordPress plugin before 0.9.2.8 does not resolve the target of the update-user action in its front-end Forms module consistently between the stage that authorises the request and the stage that carries it out, which allows unauthenticated visitors to bypass the guard protecting administrator accounts, take over the site administrator's account and be issued a valid session for it. Exploitation requires the site to publish two separate front-end form configurations whose targets are mapped to visitor-submitted values.	8.1	More Details
CVE-2026-105262	The Falang multilanguage for WordPress plugin before 1.4.6 does not perform a capability check when saving term translations, and exposes a valid nonce to any logged-in user through an unprotected handler, allowing users with Subscriber-level access and above to store content that is later output without proper escaping, leading to Stored Cross-Site Scripting that executes in the context of higher-privileged users such as administrators.	8.0	More Details
CVE-2026-115190	An insecure deserialization (PHP object injection) vulnerability exists in Mautic's handling of segment filters and custom field properties. Several code paths that read these stored values from the database deserialize them without restricting which PHP classes may be instantiated, instead of using Mautic's safe deserialization helper. An attacker who is able to store a crafted serialized value in these segment filter or custom field properties can cause arbitrary PHP objects to be instantiated when Mautic later processes the data, for example during segment evaluation, contact export, or custom field API requests. Combined with gadget chains available in Mautic's dependencies, this could allow an attacker to execute arbitrary code on the server, read, modify, or delete data stored in the Mautic instance,	8.0	More Details

	including contact data, and affect its availability. Exploitation requires high privileges or an existing ability to write to the affected stored data.		
CVE-2026-115392	Wasmtime before 24.0.13, 25.0.0 before 36.0.14, 37.0.0 before 46.0.3, and 47.0.0 before 47.0.4 allows a filesystem sandbox escape in wasmtime-wasi because its cap-std dependency mishandles symlinks with trailing slashes on macOS and Linux kernels lacking openat2. A guest with write access to any sandboxed directory can read and write any host-accessible file, and read-only guests may read host files via preexisting symlinks.	7.9	More Details
CVE-2026-114085	Trigger.dev through 4.7.6 contains a missing authentication vulnerability in the supervisor workload server that allows workloads to manipulate runs belonging to other deployments without ownership verification. Attackers with network access and a valid run/snapshot pair can call the attempts/start endpoint to obtain resolved environment variables and run JWTs, or start, complete, and suspend other tenants' runs.	7.9	More Details
CVE-2026-115726	Aider through 0.86.2 contains a command execution vulnerability that allows attackers to run shell commands by committing a malicious .aider.conf.yml that is auto-loaded without a trust prompt. Attackers can set notifications-command or editor values that are executed via subprocess with shell=True after model replies or /editor use, running commands with the victim's privileges.	7.8	More Details
CVE-2026-79892	A memory corruption vulnerability exists in AOS-8 that could allow an authenticated local attacker to execute arbitrary code. Successful exploitation could allow the attacker to execute arbitrary code as a privileged user on the underlying operating system.	7.8	More Details
CVE-2026-76495	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76495 are related to issues concerning incorrect permission assignments that are grouped under the Common Weakness Enumeration (CWE) Variant CWE-732.	7.8	More Details
CVE-2026-98982	In the Linux kernel, the following vulnerability has been resolved: namespace: don't let a pseudo dentry become the root of a mount d_alloc_pseudo() hands out dentries for pipes, sockets and other files that are never anyone's child or parent. They carry DCACHE_NORCU and dentry_free() frees them right away because no lockless path walk can ever reach them. That holds as long as such a dentry isn't the root of a mount. But bind mounting /proc/self/fd/<fd> of such a file does exactly that. Most callers of alloc_file_pseudo() put their files on kernel-internal mounts and may_copy_tree() refuses those. bpf_token_create() doesn't. It places the token file on the bpf's mount the caller handed it and that mount is in the caller's mount namespace so the clone goes through: mount --bind /proc/self/fd/<token> <file> open_tree(tokfd, "", AT_EMPTY_PATH OPEN_TREE_CLONE) + move_mount() __follow_mount_rcu() then loads ->mnt_root of that mount, reads d_seq and d_flags of the dentry and only then checks mount_lock. The dentry can be gone by then. The final mntput() of an unmounted parent unhooks a child that stayed attached to it under mount_lock alone. When the child's holder does the final put right after that the root is dput() and freed immediately while a walker that found the mount hashed still looks at it. Refuse to clone a mount with a DCACHE_NORCU dentry as its root. Nothing sensible can be done with a bind mount of a bpf token anyway.	7.8	More Details
CVE-2026-114490	Incorrect authorization in the COM administration API of Progressive Robot hMailServer 6.3.4 through 6.3.6 allows an account assigned the serveradmin role with a list of domains - a credential an administrator may create meaning to confine it to those domains - to act as the whole-server administrator over COM. Administrator roles are meant to confine such a credential to the domains its list names, and the REST API did not take it for the server administrator, but the COM API's server-administrator check read the administer_server permission as the whole server whatever list the grant carried. So the credential signed in through the Control Panel, or through a script run on the server, as the server administrator of every domain: it could read and change every domain's accounts and mail, rename any domain, map a domain alias onto any, write the server's settings and secrets, and create further server-administrator accounts. The COM API's domain-administrator check, asked with no domain, likewise let an account assigned the domainadmin role with a list of domains that leaves out the account's own domain administer that domain's ordinary accounts over COM. An account holding serveradmin with no list is the server administrator by design and is unaffected, as are the help desk, auditor and custom roles.	7.8	More Details
CVE-2026-76494	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76494 are related to issues concerning code injections that are grouped under the Common Weakness Enumeration (CWE) Base CWE-94.	7.8	More Details
	A vulnerability in the control channel socket implementation of Open Source Snort 3 could allow		

CVE-2026-76524	an authenticated, local attacker with low privileges to execute arbitrary commands as the Snort 3 process user. This vulnerability is due to a lack of authentication when a local user connects to the Snort control channel socket. An attacker could exploit this vulnerability by connecting to the Snort 3 control channel socket as a local user. A successful exploit could allow the attacker to execute arbitrary commands as well as read and write Snort 3 process memory.	7.8	More Details
CVE-2026-98945	In the Linux kernel, the following vulnerability has been resolved: HID: steam: Rearrange teardown sequence This fixes a narrow window during the teardown sequence where callbacks could still be scheduled during cleanup that would then have a dangling pointer to the now-freed steam struct. This also puts work canceling for rumble and mode switch in steam_unregister, as that shouldn't persist while the client hdev is open.	7.8	More Details
CVE-2026-93640	A local privilege escalation vulnerability in the optional dnscache package of Zimbra Collaboration Suite (ZCS) allows the Unix zimbra account to access files with root privileges. The passwordless sudo policy permits arbitrary arguments to zmunbound, whose wrapper forwards them to Unbound without restricting the configuration path. An attacker can supply a configuration that retains root privileges and disables the chroot. Include directives expose root-readable file contents through parser errors, while an attacker-selected logfile path creates a root-owned file or appends Unbound log records to an existing file. This can disclose password hashes and private keys and corrupt privileged configuration files, preventing affected services from starting with the modified configuration. The affected configuration is present in ZCS 10.1.16 and 10.1.21 dnscache packages.	7.8	More Details
CVE-2026-76479	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76479 are related to issues concerning improper limitation of a pathname to a restricted directory that are grouped under the Common Weakness Enumeration (CWE) Base CWE-22.	7.8	More Details
CVE-2026-107896	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds read vulnerability in VS6ComFile!load_gr_item. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-98869	In the Linux kernel, the following vulnerability has been resolved: bpf: Infer zext_dst based on static register liveness analysis As reported in the thread [1], the verifier's 32-bit operations zero extension logic is broken. This logic is responsible for correct semantics of 32-bit operations on s390 architecture. According to BPF semantics, operation `w1 += 1` is supposed to zero extend the upper half of the register `r1`. On s390 the JIT relies on the verifier emitting explicit zero extension before such operations. The verifier attempts to minimize the amount of zero extensions inserted by tracking whether upper halves of the 64-bit registers are ever used. Previously such tracking worked as follows: - bpf_reg_state->subreg_def field was set by do_check_insn() for each operation defining lower but not the upper halves of the register. - Whenever an operation reading the whole register was verified, the verifier checked register's subreg_def and set bpf_insn_aux_data->zext_dst flag as true via a call to mark_insn_zext() function. - After the verification was complete, a special pass bpf_opt_subreg_zext_lo32_rnd_hi32() extended 32-bit operations with bpf_insn_aux_data->zext_dst set as true by adding explicit zero extension. Note that the logic above relies on bpf_reg_state->subreg_def, which is a property of a current verifier state. Before the commit [2] two additional steps happened: - The verifier tracked upper and lower register halves' liveness as flags REG_LIVE_READ{32,64} in bpf_reg_state->live. - The function propagate_liveness() called mark_insn_zext() in order to transfer the knowledge about which registers have their upper halves alive (and thus might require zero extension). The commit [2] removed the two steps described above, hence making possible a situation like below: - The register's upper half is set and is used on some verification path P1 and the register happens not to be marked as precise. - The checkpoint C is created while processing some instruction between register initialization and usage. - On some other verification path P2 the register's upper half is not initialized and that path ends hitting the checkpoint C. - In such a case the register's initialization on path P2 would lack zext_dst mark, making it possible for the program to inject an arbitrary value in the register's upper half. This commit replaces subreg_def based logic with computing zext_dst statically, as a part of the bpf_compute_live_registers() analysis: - The analysis now tracks usage of upper and lower halves of the registers separately. - If some instruction defines a 32-bit subregister, but not the whole register, *and* the upper half of the register is alive after that instruction, the instruction is marked as zext_dst. There is one notable drop in precision: whenever a BPF subprogram is called, all 64 bits of parameter registers are presumed to be used. The assumption is that such a drop in precision would not inflict a noticeable performance penalty. [1] https://lore.kernel.org/bpf/CAGKGUv=sOuqQtA1Ub-5JXfA4FPosfYKAQE4B79cK+P1erxqtg@mail.gmail.com/ [2] commit 107e16979905 ("bpf: disable and	7.8	More Details

	remove registers chain based liveness")		
CVE-2026-107900	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds read vulnerability in VS6ComFile!CCheck::ItemCheck. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-107899	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds read vulnerability in VS6EditData!CDataRomErrorCheck::CheckSymbolMacroParam. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-20291	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco owned Open Source Snort 3 Software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20291 are related to issues concerning improper privilege management that are grouped under the Common Weakness Enumeration (CWE) Base CWE-269.	7.8	More Details
CVE-2026-107898	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds read vulnerability in VS6EditData!CDataRomErrorCheck::MacroCommandCheck. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-13599	Missing authentication for critical function vulnerability in ABB Ability Edgenius. This issue affects Ability Edgenius: through 3.2.3.0.	7.8	More Details
CVE-2026-98944	In the Linux kernel, the following vulnerability has been resolved: HID: steam: Refactor registration This refactors and simplifies the registration/unregistration flow. Since we now only perform registration when the client isn't opened anymore, the logic for handling that can be removed, and the rest of the function streamlined. We also remove the previous assumption that we have a serial number to show we're registered, replacing it with a single purpose boolean. In a previous refactor the code for unregistering a battery if later registration steps failed was accidentally left out. As a result, a lingering power_supply object could get left over after the steam object was torn down. This is also fixed here.	7.8	More Details
CVE-2026-107894	V-SFT provided by FUJI ELECTRIC CO., LTD. contains a stack-based buffer overflow vulnerability. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-107897	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds read vulnerability in VS6MemInIF!set_temp_type_default. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-107895	V-SFT provided by FUJI ELECTRIC CO., LTD. contains an out-of-bounds write vulnerability. Opening a crafted V7 file may lead to information exposure or arbitrary code execution on the affected product.	7.8	More Details
CVE-2026-115414	OS command injection in the Terraform hook in AWS SAM CLI before 1.166.2 might allow an actor who can influence the Terraform configuration or plan file consumed by a build to execute arbitrary commands as the build user on the host running the build via a crafted Terraform resource address. To remediate this issue, users should upgrade to version 1.166.2 or later.	7.8	More Details
CVE-2026-79979	The arm64 MMU backend invalidates a single virtual address from the TLB with invalidate_tlb_page() in arch/arm64/core/mmu.c, which issued TLBI VAE1{IS} with an operand of virt >> 12. That operand encodes the ASID in bits [63:48], so the invalidation only ever matched ASID 0 (the kernel page tables) and global entries. Memory domains, however, are allocated ASIDs starting at 1 in arch_mem_domain_init(), and every domain-private leaf created by private_map() is mapped with MT_NG, i.e. non-global. A domain's cached translation therefore survived the invalidation. All callers of invalidate_tlb_page() are demand-paging paths. When the eviction algorithm reclaims a frame, page_frame_prepare_locked() in kernel/mmu.c calls arch_mem_page_out(), which marks the kernel leaf invalid, propagates the invalid descriptor into every domain via sync_domains(), and then flushes only ASID 0. do_page_fault() immediately reuses the very same page frame for the incoming mapping through arch_mem_page_in(). A user thread that had already touched the evicted page — a k_mem_partition added with arch_mem_domain_partition_add(), or its own stack mapped by map_thread_stack() — retains a valid EL0 read/write translation to that physical frame under its own ASID. Unprivileged code running in a user thread can drive memory pressure to force eviction and then keep accessing the stale address, obtaining read and write access to whatever the kernel subsequently placed in that frame: supervisor data, another memory domain's partition, or another thread's stack. This defeats the isolation that CONFIG_USERSPACE is meant to enforce, allowing information disclosure and memory corruption up to full compromise of the system. The fix switches the implementation to TLBI VAAE1{IS}, which matches every ASID and both global and non-global entries, so no domain-tagged translation survives a page out, dirty-bit update or access-flag clear. Exploitation requires an	7.8	More Details

	arm64 build with both CONFIG_USERSPACE and CONFIG_DEMAND_PAGING enabled, an active backing store, and at least one demand-paged page that is also privately mapped into a user memory domain. The same commit additionally stops globalize_table() from stripping ELO permission and non-global attributes off domain-private leaves during paging synchronization, which previously caused user threads to lose access to their own partitions and stacks.		
CVE-2026-99032	In the Linux kernel, the following vulnerability has been resolved: net/smc: protect clcsock lifetime in smc_getname smc_getname() dereferences smc->clcsock without holding clcsock_release_lock. Link-group termination can release the CLC socket through smc_close_active_abort() while the SMC socket is still open, for example after shutdown(SHUT_WR). A getsockname() caller can load smc->clcsock, then the termination worker can clear the pointer and call sock_release() before the caller accesses clcsock->ops or invokes getname(). This causes a use-after-free; if the worker clears the pointer before the load, it causes a NULL dereference. The syscall's file reference keeps the SMC socket alive but does not prevent asynchronous release of its CLC socket. KASAN reported the following with test-only timing instrumentation: BUG: KASAN: slab-use-after-free in smc_getname+0x19e/0x1b0 Read of size 8 at addr ffff888109abb4e0 by task poc/103 Call Trace: smc_getname+0x19e/0x1b0 do_getsockname+0xe5/0x170 __sys_getsockname+0x8c/0x100 Allocated by task 95: sock_alloc_inode+0x1e/0x280 sock_alloc+0x3d/0x240 __sock_create+0x7e/0x430 smc_create+0x121/0x240 Freed by task 0: kmem_cache_free+0xcc/0x340 rcu_core+0x50a/0x1850 Last potentially related work creation: evict+0x446/0x6c0 smc_clcsock_release+0xa8/0xd0 smc_close_active_abort+0x26a/0x3a0 __smc_lgr_terminate.part.0+0x137/0x2e0 Hold clcsock_release_lock across the pointer check and the getname() callback to serialize with smc_clcsock_release(). Return -EBADF if the CLC socket has already been released, preserving the existing peer state check and the callback's return value otherwise.	7.8	More Details
CVE-2026-99016	In the Linux kernel, the following vulnerability has been resolved: riscv: futex: untag the user pointer before the atomic access access_ok() checks untagged_addr(uaddr), but arch_futex_atomic_op_inuser() and futex_atomic_cmpxchg_inatomic() hand the raw uaddr to the inline asm amoswap/amoadd/.../lr.w+sc.w through the "+m" (*uaddr) operand. When the tagged address ABI is enabled (CONFIG_RISCV_ISA_SUPM, prctl PR_SET_TAGGED_ADDR_CTRL with PMLen != 0), those two addresses differ: a user pointer whose top PMLen bits hold a tag passes access_ok() after untagging, but the atomic is then performed on the still-tagged raw address. Supervisor-mode data accesses are not subject to the U-mode pointer masking (that is governed by menvcfg.PMM, which the kernel does not set), so hardware does not strip the tag for the kernel's own access. With Sv57 and PMLen=16 the tag bits overlap the canonical-address bits, so a tagged pointer can name a canonical kernel virtual address (e.g. in the linear map) whose untagged form is a valid user address. An unprivileged process can thus make FUTEX_WAKE_OP perform an atomic read-modify-write on an arbitrary kernel address, with the matching FUTEX_OP_CMP_* result serving as a read oracle. get_user()/put_user()/raw_copy_{to,from}_user() already untag the pointer after the access_ok() check; do the same in the futex helpers so the atomic operates on the address that was actually validated.	7.8	More Details
CVE-2026-99053	In the Linux kernel, the following vulnerability has been resolved: tipc: hold a reference to nodes found by link name tipc_node_find_by_name() returns a node after dropping its RCU read lock without taking a reference. The LINK_SET, LINK_GET and LINK_RESET_STATS handlers then lock and access the node, racing with timer-driven cleanup of a down peer. Generic netlink serialization does not exclude the node timer. The following interleaving can leave a handler using a freed node: CPU 0: find the node under RCU and release the node read lock CPU 1: tipc_node_timeout() clears the links and unlinks the down node CPU 1: drop the list and timer references, queuing tipc_node_free() CPU 0: leave the RCU read-side critical section CPU 1: complete the grace period and free the node CPU 0: acquire the node lock through the stale pointer LINK_SET also uses the node's media address after releasing the node lock, when passing queued packets to tipc_bearer_xmit(). KASAN on v7.3-rc5 reported: BUG: KASAN: slab-use-after-free in _raw_read_lock_bh Write of size 4 at addr ffff88807e06f008 by task poc/92 Call Trace: _raw_read_lock_bh kernel/locking/spinlock.c:287 tipc_nl_node_set_link net/tipc/node.c:2475 genl_family_rcv_msg_doit net/netlink/genetlink.c:1114 genl_rcv_msg net/netlink/genetlink.c:1209 netlink_rcv_skb net/netlink/af_netlink.c:2575 Allocated by task 0: tipc_node_create net/tipc/node.c:539 tipc_node_check_dest net/tipc/node.c:1196 tipc_disc_rcv net/tipc/discover.c:252 Freed by task 92: kfree mm/slab.c:6801 rcu_core kernel/rcu/tree.c:2919 Last potentially related work creation: __call_rcu_common.constprop.0 kernel/rcu/tree.c:3181 tipc_node_timeout net/tipc/node.c:814 Acquire a reference to the selected node with kref_get_unless_zero() before leaving RCU, returning NULL if the node has already been released. Release that reference on every caller exit after the last node access, including transmission in LINK_SET. Keep the existing link lookup order and locking so concurrent link removal still takes the existing error paths.	7.8	More Details
CVE-	PraisonAI is a multi-agent teams system. From version 4.5.87 to before version 4.6.61, PraisonAI recipe execution has a dangerous-tool policy that is supposed to block default-denied tools unless the caller explicitly passes allow_dangerous_tools=True. That policy only checks tools declared in TEMPLATE.yaml requires.tools. For steps-based recipes, the actual execution path loads workflow.yaml with		

2026-57142	YAMLWorkflowParser. That parser resolves agent-level tools: declarations and preserves top-level approve:. Workflow.start() then installs those YAML-approved tools into the approval context. As a result, an untrusted recipe can omit execute_command from TEMPLATE.yaml requires.tools, declare it in workflow.yaml agents.*.tools, and add top-level approve: [execute_command]. The caller did not set allow_dangerous_tools=True, but the recipe policy allows the recipe and the workflow approval path self-approves the critical shell tool. This issue has been patched in version 4.6.61.	7.8	More Details
CVE-2026-98889	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mvm: fix sched scan IE sizing Scheduled scan built the probe request before iwl_mvm_scan_fits(), so oversized IEs could be copied into the fixed preq buffer before length validation. Move iwl_mvm_build_scan_probe() after the fits check. Also advertise max_sched_scan_ie_len using iwl_mvm_max_scan_ie_len() so userspace limits account for driver-inserted DS/TPC bytes.	7.8	More Details
CVE-2026-98993	In the Linux kernel, the following vulnerability has been resolved: namespace: don't inherit MNT_UMOUNT in clone_mnt() MNT_UMOUNT gets raised on a mount that umount_tree() or umount_one() has taken out of its namespace. It's never cleared. __detach_mounts() unhooks a mount that carries it instead of unmounting it. disconnect_mount() lets an unmounted child stay attached only if its parent carries it. propagate_umount() takes it as the sign that a mount is going to be unmounted when it decides what else can be unmounted, where a chain of locked mounts ends and where a for reparenting an overmount. clone_mnt() copies the flags of the source and only masks MNT_INTERNAL_FLAGS which doesn't include MNT_UMOUNT. A mount that was lazily unmounted but is kept alive by a file descriptor can still be copied when its root is an nsfs or pidfs file because may_copy_tree() accepts those before checking whether the source is mounted at all: mount --bind /proc/self/ns/net /tmp/x exec 3< /tmp/x umount -l /tmp/x mount --bind /proc/self/fd/3 /tmp/y The mount on /tmp/y is attached to the caller's namespace with MNT_UMOUNT already set and every copy of it inherits the flag again. That causes a number of bugs: - Unlink its mountpoint from a mount namespace where it isn't one. __detach_mounts() takes it for a mount that was unmounted earlier, unhooks it from its parent and drops the reference the namespace holds on it. It stays in the namespace with one reference too few and as its own parent, so it can't be unmounted through the file descriptor anymore and the namespace never finds it again. When the namespace goes away it's leaked and points to a freed struct mnt_namespace that may_change_propagation() and anon_ns_root() read on the next propagation change or mount_setattr() through the file descriptor. - Make it shared and unmount a peer. trace_transfers() follows the flag from the peer into it and strips it from its peer group. - Let a propagated umount reach it. gather_candidates() marks it with T_UMOUNT_CANDIDATE but never puts it on the candidate list, so the mark is never cleared and every later propagated umount skips it. If it's the parent of a locked mount handle_locked() takes it as the cutoff and commits the locked child that should have stayed and disconnect_mount() leaves that child attached to it, unmounted under a mounted parent, which is what disconnect_mount() exists to prevent. When the namespace goes away umount_tree() visits the child a second time. And once __detach_mounts() has made it its own parent reparent() never finds the end of the chain of unmounted parents. A clone is a new mount that has never been unmounted. Add MNT_UMOUNT to MNT_INTERNAL_FLAGS so clone_mnt() never copies it. And warn when umount_tree() reaches a mount that already has it set.	7.8	More Details
CVE-2026-99054	In the Linux kernel, the following vulnerability has been resolved: tipc: destroy topsrv workqueues before closing connections tipc_topsrv_stop() closed subscriber connections while the topology server's send and receive workqueues were still running. Socket callbacks and subscription events could then queue more work, and in-flight send/recv work could drop the last connection reference during the conn_idr walk. That race produced several teardown failures: refcount_t addition on 0 from conn_get() on a connection whose release was blocked on idr_lock, a subsequent use-after-free in tipc_conn_close(), queue_work() on an already destroyed workqueue from the listener data-ready callback, and an RCU stall in tipc_topsrv_exit_net() while the walk spun under idr_lock. Clear srv->listener under idr_lock so it acts as a shutdown flag, skip queue_work() once it is NULL, destroy the workqueues to flush in-flight work, and only then close the remaining connections. Refuse tipc_conn_lookup() after that flag is cleared, and drop idr_lock when the teardown walk finds no connection, so an in-flight subscription event cannot pin idr_in_use while the walk holds the lock. v3 supersedes the narrower in-thread diff Tung Quang Nguyen posted on 2026-09-23. It keeps his teardown order and adds the lookup refusal and the empty-idr unlock.	7.8	More Details
CVE-2026-115504	CFITSIO through 4.7.0 contains an out-of-bounds write vulnerability in ffgcls2() in getcols.c that trusts attacker-controlled heap descriptor lengths for variable-length string columns. Attackers can supply a crafted FITS file that overflows a fixed 28,800-byte stack buffer and the caller's string array, enabling arbitrary code execution.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: namespace: queue a mount only once for mount notifications mnt_notify_add() puts a mount on notify_list. It doesn't check whether the mount is on the list already. That's fine as long as every mount is queued at most once per namespace_sem hold. It isn't. propagate_umount() moves a surviving overmount off a stack of mounts that are going away and queues it as moved. shrink_submounts() and mark_mounts_for_expiry() call		

CVE-2026-98991	umount_tree() for several mounts under a single namespace_sem hold. So the second umount_tree() can take down exactly the mount the first one reparented, reparent it once more and so end up queueing it a second time. The second list_add_tail() cuts the mounts queued in between out of the list while the head still points at the last of them. notify_mnt_list() then only visits that one mount and the others are freed after the grace period with notify_list still pointing at them. Every later mount operation on the host walks freed memory: BUG: KASAN: slab-use-after-free in __list_add_valid_or_report Read of size 8 at addr ffff8881003d57e0 by task notify_dq/162 __list_add_valid_or_report+0x15c/0x1a0 mnt_add_to_ns+0x322/0x890 attach_recursive_mnt.isra.0+0xf2a/0x1a00 path_mount+0x139c/0x1d40 This needs a mount with MNT_SHRINKABLE to bind from. That's what finish_automount() creates (submounts of NFS, AFS and CIFS and the tracefs mount below debugfs). Bind mounts inherit the flag. A mount namespace with a FAN_MARK_MNTNS mark does the rest and root in a user namespace can have both. Initialize to_notify in alloc_vfsmnt() and leave a mount alone that is queued already. mnt_notify() looks at the state the mount has when the list is drained so a single entry per mount is enough. The move event for a mount that is unmounted under the same hold is lost. That's fine. The detach is what matters.	7.8	More Details
CVE-2026-98975	In the Linux kernel, the following vulnerability has been resolved: xfs: stop exchanging reflink flags when exchanging mappings When an exchange covers both files from offset 0 to i_disk_size, xmi_can_exchange_reflink_flags() assumes every mapping moved and hands the reflink flag over to the other file. That assumption is wrong. A file can still own shared blocks past i_disk_size, e.g. after EXCHANGE_RANGE_TO_EOF sets a smaller size without unmapping anything. The flag is cleared anyway, and the next write to such a block happens in place instead of CoW, corrupting the other file that shares it. Commit a23eca88448e ("xfs: fix exchange-range reflink flag clearing issue with INO1_WRITTEN") fixed one instance of this. Rather than add another special case, stop moving the flags. xfs_exchmaps_ensure_reflink() already sets the flag on both inodes before the exchange, so both just keep it. Nothing sets CLEAR_INO{1,2}_REFLINK anymore. Intents logged by older kernels may still carry them; recovery already masks them off, so they stay in the log format mask and are ignored. XFS_IOC_SWAPEXT has the same problem: it moves the reflink flag from one inode to the other along with the extents. Just removing that is not enough here: the exchange path flags both inodes up front in xfs_exchmaps_ensure_reflink(), but swapext has no such step, so the inode receiving the shared extents would be left without the flag. Set it on both inodes instead, and make sure both have a CoW fork, since an inode with the reflink flag is expected to have one. The flag now stays until FALLOC_FL_UNSHARE_RANGE, xfs_scrub, or a truncate to zero removes it.	7.8	More Details
CVE-2026-99014	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_sync: Fix command skb lifetime during scan setup During PA sync scan setup, hci_le_set_ext_scan_param_sync() gets the address used for the PA_LINK lookup from hci_sent_cmd_data(). The returned pointer borrows storage from sent_cmd or req_skb, without holding a reference to either skb. The scan worker runs on req_workqueue while hci_cmd_work() runs on the separate device workqueue. After the scan worker loads sent_cmd, the command worker can free and replace it before hci_sent_cmd_data() dereferences the skb. The command payload can also be freed between returning from the helper and comparing the address. The connection lookup's RCU critical section does not protect the command skb. KASAN reported: BUG: KASAN: slab-use-after-free in hci_sent_cmd_data+0x27c/0x2f0 Workqueue: hci0 hci_cmd_sync_work Call Trace: hci_sent_cmd_data+0x27c/0x2f0 hci_le_set_scan_param_sync+0x43a/0x850 hci_passive_scan_sync+0xb09/0x1460 hci_update_passive_scan_sync+0x47c/0x6e0 hci_le_pa_create_sync+0x200/0xa70 hci_cmd_sync_work+0x13c/0x290 Allocated by task 95: skb_clone+0x13f/0x340 hci_cmd_work+0x2a9/0x7e0 Freed by task 95: kmem_cache_free+0xba/0x3a0 hci_cmd_work+0x29c/0x7e0 Use hdev->lock to serialize the address snapshot with sent_cmd replacement. Protect req_skb replacement and completion cleanup with the same lock because the helper falls back to that skb. Copy the address before releasing the lock and use the copy for the existing RCU-protected connection lookup. Release the mutex before sending or waiting for commands, preserving the lookup and completion ordering.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: xfs: fix exchange-range-to-eof file size exchange Norbert Szetei posted a patch containing an incomplete description of a bug in XFS_IOC_EXCHANGE_RANGE's TO_EOF flag. The bug tricks exchange-range into modifying a file so that it has shared blocks starting beyond EOF, which is never allowed because files never have written data beyond EOF, and you can only share written data blocks. This is key to all the incorrect behavior that follows. Eventually I got from Norbert a description of what's going wrong: > Fair. Here it is with physical block numbers, 4k blocks. peer is the file > that gets rewritten and file1 is the one that ends up with the flag > cleared. > > Start, nothing shared: > > peer size 4096 -> [100] > file1 size 12288 -> [200] [201] [202] > donor1 size 4096 -> [300] > donor2 size 4096 -> [400] > > 1. FICLONERANGE peer's block into file1's middle slot: > > peer size 4096 -> [100] > file1 size 12288 -> [200] [100] [202] reflink flag set > ^^^^^ both files own block 100 > > 2. XFS_IOC_EXCHANGE_RANGE file1 against donor1, file1_offset 0, > file2_offset 8192, length 0, TO_EOF. One block is exchanged, and the sizes > are exchanged with it: > > file1 size 4096 -> [200] [100] [300] reflink flag set > ^^^^^^^^^^^^^ still owned, now past the size > donor1 size 12288 -> [202] > > file1 says it is one block long while it still		

CVE-2026-98974	owns three. Nothing was > unmapped. file2_offset is block 2, so this exchange cannot clear a flag. > > 3. XFS_IOC_EXCHANGE_RANGE file1 against donor2, both offsets 0, length 4096. > By i_disk_size this is two one-block files exchanging everything, so > xmi_can_exchange_reflink_flags() agrees and the flag is cleared: > > file1 size 4096 -> [400] [100] [300] reflink flag CLEARED > donor2 size 4096 -> [200] > > Block 100 is still shared with peer. > > 4. pwrite(file1, 4096, 4096), which is file1's second slot, block 100. > xfs_is_cow_inode(file1) is false, so no CoW: > > peer size 4096 -> [100] same block, new contents > > Block 100 is the only one that matters here. PoC below, it reads a file > called peer in the current directory. The filesystem needs reflink and > exchange-range, which mkfs.xfs 7.x turns on by default. So yes, TO_EOF is screwing up the file size even if it's actually exchanging the data block correctly. The following is the intended behavior of XFS_EXCHANGE_RANGE_TO_EOF, as described by its manpage: "Ignore the length parameter. All bytes in file1_fd from file1_offset to EOF are moved to file2_fd, and file2's size is set to (file2_offset+(file1_length-file1_offset)). Meanwhile, all bytes in file2 from file2_offset to EOF are moved to file1 and file1's size is set to (file1_offset+(file2_length-file2_offset))." In other words, if we only exchange a portion of the file data, then we should only exchange the portion of the file size that relates to the exchanged data. Unfortunately, the code swaps the ondisk file size and the incore file sizes, which is incorrect if the caller passes in a nonzero offset. Having confused the file sizes, the PoC uses a second exchange-range call on what looks like a non-reflinked single-block file. Because the file size is set incorrectly, exchange-range thinks it's exchanging the full contents of two files and clears the reflink flag on the broken file. That enables an extending write of the broken file to rewrite the shared block that's just past EOF. This has become known colloquially as refluxfs. Once the file sizes are set correctly, the second exchange-range no longer thinks that it's doing a full-contents swap, so it won't clear the reflink flag on either of its file arguments. Norbert followed up to an earlier version of this patch with a second complaint about us not handling the i_disk_size state correctly: > The offsets are bounded against the incore i_size but subtracted from > ---truncated---	7.8	More Details
CVE-2026-56838	PraisonAI is a multi-agent teams system. From version 4.5.87 to before version 4.6.59, PraisonAI recipe execution blocks default-denied dangerous tools unless the caller explicitly passes allow_dangerous_tools=True. The normal recipe.run() path enforces this with _check_tool_policy(). The streaming path, recipe.run_stream(), loads the same recipe, checks dependencies, and then calls _execute_recipe() without running the dangerous-tool policy check. As a result, a recipe that honestly declares execute_command in TEMPLATE.yaml requires.tools is denied by recipe.run(), but reaches the execution engine through recipe.run_stream() with allow_dangerous_tools=False. This issue has been patched in version 4.6.59.	7.8	More Details
CVE-2026-19676	The user-mode verifier for the peci_transfer() system call, z_vrfy_peci_transfer() in drivers/peci/peci_handlers.c, checked only the device object permission and copied the caller-supplied struct peci_msg with k_usermode_from_copy(). That copy validates the outer structure only; the buf pointers and len values inside the embedded struct peci_buf members tx_buffer and rx_buffer remained fully caller-controlled and were passed unchecked to the driver. Every in-tree PECE driver dereferences those pointers in supervisor mode: peci_npcx_transfer() in drivers/peci/peci_npcx.c and peci_it8xxx2_transfer() in drivers/peci/peci_ite_it8xxx2.c copy hardware bytes into rx_buffer.buf[i] and read command bytes out of tx_buffer.buf[i], and peci_xec_read()/peci_xec_write() in drivers/peci/peci_mchp_xec.c do the same with no bound on rx_buffer.len at all. A user-mode thread that has been granted the PECE device object — the normal way an application exposes a device under CONFIG_USERSPACE — could therefore point rx_buffer.buf at kernel data structures and have the driver write hardware-supplied bytes there, or point tx_buffer.buf at kernel memory and have its contents pushed onto the PECE/eSPI out-of-band bus. The impact is a full user-mode-to-kernel escape: an arbitrary kernel write primitive (integrity and, through it, code execution at supervisor privilege), an arbitrary kernel read disclosed to a PECE bus observer, and a trivially reachable supervisor-mode fault if an unmapped address is supplied. Exploitation requires CONFIG_USERSPACE=y, one of the PECE drivers enabled (CONFIG_PECI_NPCX, CONFIG_PECI_XEC or CONFIG_PECI_ITE_IT8XXX2), and local code execution in a user-mode thread holding the PECE device. The fix adds K_SYSCALL_MEMORY_READ() and K_SYSCALL_MEMORY_WRITE() checks over the two embedded buffers, with a size_add_overflow() guard on the frame-check-sequence byte accounted for in the RX length.	7.8	More Details
CVE-2026-98999	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Fix stale pinned_link entry when fb-pin performs the final unpin xe_bo_pin_external() and xe_bo_unpin_external() maintain the bo's pinned_link list membership in xe->pinned.late.external, based on whether the current call is the outermost pin or the final unpin. However, the same external bo's pin_count can also be raised and lowered directly by __xe_pin_fb_vma()/__xe_unpin_fb_vma(), which pin the bo as a display scanout buffer without going through xe_bo_pin_external()/xe_bo_unpin_external() at all, and have no notion of, or ownership over, pinned_link. If a bo is pinned both externally (e.g. dma-buf export) and as an fb, and the external unpin happens first, xe_bo_unpin_external() correctly observes pin_count > 1 and leaves the bo on pinned_link. When the fb unpin later performs the true last unpin (pin_count 1 -> 0), it never touches pinned_link, leaving the bo linked on xe->pinned.late.external indefinitely. Once the bo is subsequently freed, this stale list entry points into freed memory, corrupting the list and risking a use-after-free the next time the list is walked or spliced. The backup object pin/unpin sites in	7.8	More Details

	xe_bo_notifier_prepare_pinned()/ xe_bo_notifier_unprepare_pinned() have the same bypass characteristic, though they never add their bo to a pinned list, so are not affected by this particular list-corruption issue. Move the pinned_link removal into xe_bo_account_unpin(), which already runs on every unpin path (kernel, external, framebuffer, backup object) right before the true 1 -> 0 pin_count transition. Since list_del_init() only operates on the node itself, this removal is list-agnostic and safe to perform regardless of which list (external or kernel_bo_present) the bo happens to be linked on, or which code path is performing the final unpin. Drop the now-redundant explicit list_del_init() calls in xe_bo_unpin_external() and xe_bo_unpin(). v2: - New patch v3: - Rebased on changes to previous patch. (cherry picked from commit ca58d36a1f93af586ee93b340c9f9fd81312cc01)		
CVE-2026-99006	In the Linux kernel, the following vulnerability has been resolved: drm/i915/gem: Prevent overstepping exec array boundary It is possible for eb_relocate_parse_slow() to call kvfree() on entries outside the original exec array during following scenario inside eb_relocate_parse_slow(): 1) eb_copy_relocations() allocates as many as eb->buffer_count copies for exec entries; 2) eb_parse() appends up to two VMAs, right after incrementing eb->buffer_count; 3) cleanup under "out" label uses this incremented buffer_count to release relocation pointers from exec array via kvfree(), not from VMA array. To amend this problem, ensure that relocation cleanup uses the original exec object count, excluding the VMAs appended by the command parser. (cherry picked from commit 2a0158925fb943c7ba8f3216e2734952128defc9)	7.8	More Details
CVE-2026-50056	A local privilege escalation vulnerability in the zmmailboxdmgr launcher of Zimbra Collaboration Suite (ZCS) allows the Unix zimbra account to execute arbitrary commands as root. The passwordless sudo interface accepts JVM arguments that are not adequately restricted. An attacker can postpone Jetty's privilege drop, register an OnOutOfMemoryError command and force heap exhaustion during startup, causing the command to execute while the JVM still runs as root. Exploitation requires local command execution as zimbra on a mailbox server and interrupts mailboxd during the malicious restart. ZCS 10.1.16 and 10.1.21 are affected.	7.8	More Details
CVE-2026-98976	In the Linux kernel, the following vulnerability has been resolved: namespace: keep the lock on a mount that a propagated copy is moved beneath attach_recursive_mnt() transfers MNT_LOCKED from the top mount to the mount that is moved beneath it with MOVE_MOUNT_BENEATH. This allows the owner of a user namespace to replace its locked /proc or its root. The mount beneath takes on the job of covering the underlying mountpoint allowing the top mount to be unmounted. Consider two mount namespaces: (H) The host H has a shared mount P with a secret in P/d (Z) Z is a user namespace made from M. Its copy of P receives propagation from the host's P and its copy of M's cover on P/d is locked Z's owner must not get to see P/d. Now (H) mounts X on P/d. This propagates into (Z). The copy of X lands beneath (Z)'s locked cover. The locked property is transferred from (Z)'s cover to the copy of X propagated beneath it. The cover is now unlocked. Now (H) unmounts X again. The copy of X in (Z) gets unmounted and the covering mount is left unlocked on top of P/d. (Z) can now unmount it: Z: umount2(P/d) = EINVAL /* the cover is locked */ H: mount X on P/d, umount X /* both propagate into Z */ Z: umount2(P/d) = 0 /* the cover is now unlocked */ Z: read P/d/secret = "covered-by-root" /* secret revealed */ So only transfer the locked property to the mount beneath for mounts the caller has placed. A propagated copy that lands beneath a locked mount is locked as well so that the mount at the bottom of the stack carries a lock the way every check expects. The mount on top of it remains locked to ensure that it keeps covering even if the propagated mount is unmounted again.	7.8	More Details
CVE-2026-99051	In the Linux kernel, the following vulnerability has been resolved: fsverity: RCU-delay the freeing of struct fsverity_info __fsverity_get_info() calls rhashtable_lookup_fast(), which just uses rcu_read_lock() and doesn't directly synchronize with fsverity_remove_info(). For the same inode this isn't a problem: its fsverity_info is removed only at inode eviction time or upon failure to enable verity, when the inode no longer needs its fsverity_info and it will no longer be accessed via that inode. However, this is broken and can cause a use-after-free for concurrent __fsverity_get_info() for *different* inodes. Those rely on following fsverity_info::rhash_head in the rhashtable under rcu_read_lock() only. They also use fsverity_info::inode to do the key comparison. Fix this by RCU-delaying the freeing of 'struct fsverity_info' after it's been removed from the rhashtable.	7.8	More Details
CVE-2026-99002	In the Linux kernel, the following vulnerability has been resolved: drm/vmwgfx: validate pitch coming from userspace Validate the pitch, alongside the box dimensions before trying to copy data from the underlying surface. Fixes possible out of bounds reads with cursor snooping.	7.8	More Details
CVE-2026-98899	In the Linux kernel, the following vulnerability has been resolved: netfilter: ipset: mark the rcu locked areas properly When we bump the uref counter, there's no need to keep the rcu lock because the referred hash table can't disappear. Also, from the same reason in mtype_gc we need the rcu lock and not a spinlock.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: gpio: mpsse: fix race when arming the IRQ poll worker gpio_mpsse_irq_enable() arms the poll worker before publishing it: schedule_work() runs before the worker is added to priv->workers. If gpio_mpsse_disconnect() walks the list in that window it		

CVE-2026-99050	misses the worker, and once disconnect returns, the USB core frees mpssse_priv while the orphaned gpio_mpsse_poll() work keeps accessing it, causing a use-after-free. Fix this by publishing and arming the worker in one irq_spin critical section. Teardown walks the same list under irq_spin, so a worker found on the list is guaranteed to be armed, and cancel_work_sync() handles it whether it is queued or running. A worker armed after the disconnect walk would still be missed, so also set a new priv->dying flag under irq_spin before the teardown walk, and check it in the same critical section, freeing it via kfree_rcu() instead when the device is going away, since kfree() would be called with the IRQ core's raw spinlock held and may sleep on PREEMPT_RT. schedule_work() is safe to call with irq_spin held, and the next probe gets a fresh mpssse_priv, so the flag never needs to be cleared. This issue was found by an in-house static analysis tool.	7.8	More Details
CVE-2026-99018	In the Linux kernel, the following vulnerability has been resolved: riscv: entry: gate shadow call stack reload on sstatus.SPP, not tp Under CONFIG_SHADOW_CALL_STACK, handle_exception() reloads the kernel shadow call stack pointer (gp) only when the trap came from userspace. It made that decision with scs_load_current_if_task_changed, which compares the trap-time tp (read back from sscratch into s5) against the current tp: beq \prev, tp, _skip_scs scs_load_current _skip_scs: tp is a user-writable register. A user task that sets tp to the linear-map address of its own task_struct (&current) before trapping makes the comparison succeed (can achieve by brute force with KASLR enable under our test), so the kernel skips scs_load_current and keeps running with the gp that userspace left in place. Under CONFIG_SHADOW_CALL_STACK gp is the shadow call stack pointer and load_global_pointer is a no-op, so the compiler-emitted "sd ra, -8(gp)" / "ld ra, -8(gp)" in every non-leaf function then write to, and return through, an attacker-controlled address -- an unprivileged kernel arbitrary-write and control-flow-hijack primitive. The poisoned gp is saved to task_struct.scs_sp on return, so the reload is skipped on every subsequent entry as well. Decide "came from userspace" from the hardware sstatus.SPP bit (already in s1 at this point) instead of from tp, mirroring the check that ret_from_exception already uses a few instructions later. SPP is written by hardware on trap entry and cannot be forged from U-mode. Legitimate behaviour is unchanged: user entries (SPP=0) reload the task shadow call stack, nested kernel entries (SPP=1) skip it.	7.8	More Details
CVE-2026-114098	Dulwich from 0.22.5 before 1.2.8 contains a path traversal vulnerability in stash.py pop() that allows attackers to write arbitrary files outside the worktree through intermediate directory symlinks. Attackers can craft a malicious repository whose popped stash writes a payload into .git/hooks/post-checkout, achieving code execution on the next checkout.	7.8	More Details
CVE-2026-99012	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Serialize concurrent connect calls iso_sock_connect() checks the socket state before taking the socket lock and drops the lock again before setting up a connection. Two callers can both pass the admission check while the socket is open or bound. Caller A can copy its destination for route selection, then caller B can replace the socket destination before A binds or connects the CIS. A then uses B's destination with the route selected for its own request. B can also proceed after A attaches a connection and sets BT_CONNECT. For deferred BIS setup, B can bind a second BIS and overwrite iso_pi(sk)->conn, leaving the first connection's reference and conn->sk back-pointer stranded. Concurrent CIS connects can likewise replace the socket's connection. Later teardown only detaches the current connection, so a callback on the old connection can race with socket release and access a freed socket. KASAN reported: BUG: KASAN: slab-use-after-free in iso_sock_hold+0xf7/0x1b0 Call Trace: iso_sock_hold+0xf7/0x1b0 iso_conn_del+0x7b/0x1d0 iso_connect_cfm+0x186/0x16a0 hci_conn_failed+0x154/0x280 hci_abort_conn_sync+0x3dc/0x7d0 hci_cmd_sync_work+0x173/0x300 Allocated by task 121: sk_alloc+0x2b/0x6d0 bt_sock_alloc+0x29/0x370 iso_sock_alloc.constprop.0+0x19/0x300 iso_sock_create+0x94/0x100 Freed by task 121: kfree+0x121/0x3c0 __sk_destruct+0x42b/0x540 iso_sock_release+0x29d/0x340 __sock_release+0xa1/0x260 Check admission under the socket lock and mark the connect operation in progress before publishing its destination. Keep that flag set across route lookup and connection setup, rejecting another connect with -EBADFD before it can change the destination or attach a connection. Clear the flag after either helper returns, including on failure, so a failed setup can be retried. A separate flag is needed because the socket lock must be released before acquiring the HCI device lock, while BT_CONNECT requires an attached connection. Keep the existing state transitions and the deferred CIS completion through iso_sock_recvmmsg(). Reject an existing socket attachment in __iso_chan_add() before taking a new reference or publishing either pointer, so every caller preserves the connection association. Keep the same-socket, same-connection success case for deferred CIS setup. Reject an attached socket before BIS setup so a reusable BIS is not claimed before the attachment is rejected. In CIS setup, reject a different HCI connection before iso_conn_add() and drop the per-attempt HCI hold directly. An unowned iso_conn may still carry another reference after detachment, so its destruction cannot be relied on to release that hold. iso_listen_bis() creates a fresh PA HCI/ISO pair under the device lock, so the temporary iso_conn_put() frees the ISO candidate and drops its HCI hold on rejection. iso_conn_ready() uses a freshly allocated, unpublished child with no connection, making the socket-side rejection unreachable.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Reject concurrent BIS		

CVE-2026-99013	listener setup iso_sock_listen() drops the socket lock before calling iso_listen_bis() to preserve the hdev->lock ordering. Two concurrent listen() calls on the same broadcast socket can therefore both observe BT_BOUND. Caller A can attach connA and release the locks, then caller B can acquire them and attach connB before A publishes BT_LISTEN. __iso_chan_add() checks whether connB already has a socket, but not whether the socket already has a connection. The second attachment overwrites iso_pi(sk)->conn, leaking the connA reference and leaving connA->sk intact. Socket teardown follows connB, so later events on connA can cause a stale socket access. Recheck BT_BOUND and require iso_pi(sk)->conn to be NULL after taking both the HCI device and socket locks in iso_listen_bis(). The connection check also covers the interval between attachment and publication of BT_LISTEN. Reject a competing setup with -EBADFD through the existing unlock path before allocating another connection. A recorded KASAN report from an instrumented kernel shows: BUG: KASAN: slab-use-after-free in iso_sock_hold+0x119/0x1f0 Call Trace: iso_sock_hold+0x119/0x1f0 iso_conn_del+0xe4/0x270 hci_disconn_complete_evt+0x351/0x8c0 hci_event_packet+0x71b/0xb20 hci_rx_work+0x293/0x730 Allocated by task 93: sk_prot_alloc+0x113/0x220 sk_alloc+0x2b/0x6d0 bt_sock_alloc+0x29/0x370 iso_sock_alloc.constprop.0+0x19/0x300 iso_sock_create+0x94/0x100 Freed by task 93: kfree+0x121/0x3c0 __sk_destruct+0x42b/0x540 iso_sock_release+0x29d/0x340 __sock_release+0xa1/0x260 sock_close+0x10/0x20	7.8	More Details
CVE-2026-98995	In the Linux kernel, the following vulnerability has been resolved: namespace: reset the old parent's ->overmount in mnt_change_mountpoint() __umount_mnt() resets ->overmount of the parent when it unmounts the mount that overmounts it. mnt_change_mountpoint() doesn't. It moves the mount to its new parent and leaves the old one pointing at it. For a mount that is moved beneath in attach_recursive_mnt() that doesn't matter. The new mount has already reset ->overmount of the old parent by the time the old top mount is moved onto it. It matters for reparent() though. propagate_umount() moves a surviving overmount onto the first ancestor that isn't going to be unmounted. The mount it was taken from is unmounted immediately afterwards. If a file descriptor keeps that unmounted mount alive ->overmount points at the moved mount. That's a dangling pointer and UAF as the moved mount is unmounted and freed. do_lock_mount() follows ->overmount from the target path under mount_lock. With a file descriptor on such a mount as the target of a mount beneath topmost_overmount() walks into the freed mount and do_lock_mount() takes references on whatever it finds there. So reset ->overmount of the old parent in mnt_change_mountpoint() the same way __umount_mnt() does. A mount's ->overmount is then always a mount that is attached to it.	7.8	More Details
CVE-2026-114097	Dulwich from 0.23.1 before 1.2.8 contains a path traversal vulnerability in CommitFilter._apply_index_filter() that allows attackers to write files outside the working directory via symlinks persisting between commits. Attackers can craft a repository where an ancestor commit's symlink redirects a descendant commit's file writes during filter_branch, enabling code execution by writing into .git/hooks.	7.8	More Details
CVE-2026-114096	Dulwich from 0.24.0 before 1.2.8 contains a symlink following vulnerability in porcelain.checkout(paths=[...]) that allows attackers to write arbitrary files by abusing raw os.open without O_NOFOLLOW. Attackers can craft a repository where a worktree symlink points to .git/hooks/post-checkout, so restoring the file from a tag writes a malicious hook for code execution.	7.8	More Details
CVE-2026-99056	In the Linux kernel, the following vulnerability has been resolved: strparser: make sure __strp_rcv isn't running before tearing down the parser The comment above strp_done() claims that if strp is stopped, strp_rcv will no longer be called. That's only true if the caller has a mechanism (eg locking) to guarantee that strp_rcv() calls that started before strp_stop() have completed by the time we call strp_done(). This adds a dummy lock_sock/release_sock pair to guarantee that any in-flight strp_rcv() (which runs under either bh_lock_sock or lock_sock, depending if it's called from ->sk_data_ready or strp_work) has completed. Only espintcp can be affected by this race condition, but this patch makes sure no future user of strp can have the bug. This could crash on strp->sk ("general mode" of strp), but this mode has been here for 9 years and has never been used. It'll be gone soon, no point worrying about it. Switch the stopped/paused/etc bits to u8's to avoid races between strp_stop() and strp_pause/unpause(). A reproducer has been published and turns the possible UAF into an exploit [1].	7.8	More Details
CVE-2026-50057	A local privilege escalation vulnerability in the Zimbra Collaboration Suite (ZCS) MTA package allows the Unix zimbra account to execute arbitrary commands as root. The installed sudo policy permits passwordless root execution of Postfix administration commands without restricting their arguments. An attacker can select a zimbra-controlled daemon_directory in the Postfix configuration and invoke postfix set-permissions, causing Postfix to execute the attacker's post-install program as root. The Postfix configuration directory is also writable by zimbra, so restricting postconf alone does not remove the execution path. Exploitation requires local command execution as zimbra on an MTA node. ZCS 10.1.16 and 10.1.21 are affected.	7.8	More Details
CVE-2026-	The Modern Events Calendar Lite WordPress plugin before 7.37.0, Modern Event Calendar Pro WordPress plugin before 7.37.0 do not sanitise and escape a parameter before using it in a SQL statement, which allows authenticated users with Subscriber-level access and above to append	7.7	More Details

102499	additional SQL queries into already existing queries and extract sensitive data from the database.		
CVE-2026-116180	vLLM from 0.8.5 before 0.31.0 has a denial of service flaw. LogprobsTensors.empty_cpu() in vllm/v1/outputs.py allocates an uninitialized int32 tensor, and the V1 model runner leaves prompt positions reported as computed by a KV connector unwritten. A client sharing the engine can send prompt_logprobs (or echo with logprobs) against a connector-cached prefix, making tokenizer decode raise OverflowError, which crashes EngineCore and drops all co-tenant requests.	7.7	More Details
CVE-2026-99059	In the Linux kernel, the following vulnerability has been resolved: wireguard: noise: reject response consumption after intermediate initiation Two threads begin processing the identical response message, received twice. The first thread, A, runs. While it's running, the second one, B, gets partway through, and during that slow calculation, or even while blocking on down_write(), A completes and then also a handshake initiation that's already been queued up runs in thread C, which itself takes that same down_write(). The handshake initiation creation succeeds, and sets the state back to waiting-for-response, and calls up_write(), at which point thread B resumes, because either its finished its calculations or was finally allowed to acquire down_write(). Thread B then copies the state back to the peer, and begins a new session, using that state, which is the same session as the one made in thread A. Thread A Thread B Thread C down_read() sA = handshake->state memcpy(cA, handshake->crypto) up_read() if (sA != 1) goto fail slow_crypto(cA) down_read() sB = handshake->state memcpy(cB, handshake->crypto) up_read() if (sB != 1) goto fail slow_crypto(cB) down_write() if (sA != handshake->state) goto fail memcpy(handshake->crypto, cA) handshake->state = 2 up_write() down_write() if (handshake->state != 2) goto fail derive_session(handshake->crypto) up_write() down_write() slow_crypto(handshake->crypto) handshake->state = 1 up_write() down_write() if (sB != handshake->state) goto fail memcpy(handshake->crypto, cB) handshake->state = 2 up_write() down_write() if (handshake->state != 2) goto fail derive_session(handshake->crypto) up_write() This seems basically impossible to hit in a meaningful way in practice, but ensure that it absolutely cannot happen by comparing the ephemeral private key that's on the stack with the latest one that the peer's handshake state has.	7.7	More Details
CVE-2026-114084	Trigger.dev before 4.5.6 contains an SQL injection vulnerability in the /realtime/v1/runs endpoint, which concatenates the unescaped tags query parameter into the ElectricSQL shape where clause and defeats the runtimeEnvironmentId tenant guard. Authenticated attackers holding any project API key with read:runs scope can append boolean conditions to stream TaskRun records, including payloads and metadata, from every organization on the same self-hosted instance.	7.7	More Details
CVE-2026-115376	Incus 7.0.1 before 7.0.2 and 7.2.0 before 7.5.0 contains an authorization bypass vulnerability that allows project-restricted clients to copy volumes by omitting source.type. Attackers can send crafted POST requests to the custom volumes endpoint with a chosen source.project to copy and read other projects' volumes, snapshots, and configuration.	7.7	More Details
CVE-2026-114081	Trigger.dev before 4.5.6 contains an SQL injection vulnerability in the TSQL query compiler, which concatenates window-function names into generated ClickHouse SQL without allowlisting or escaping. Authenticated customers can submit a backtick-quoted identifier to POST /api/v1/query that injects a raw subquery outside the per-tenant WHERE guard, exfiltrating other organizations' analytics data.	7.7	More Details
CVE-2026-108456	Server-side request forgery (ssrf) in Azure AI Foundry allows an authorized attacker to elevate privileges over a network.	7.7	More Details
CVE-2026-115810	Tandoor Recipes before 2.7.0 contains a server-side request forgery vulnerability in the Pestle importer that allows authenticated low-privileged users to force the server to issue GET requests to arbitrary internal or external hosts by crafting a malicious image URL in a Pestle JSON import file. Attackers can exploit the raw requests.get call in cookbook/integration/pestle.py, which bypasses the application's centralized requests-hardened SSRF filter, to reach internal-only endpoints, perform port/service discovery, and exfiltrate internal responses via the stored recipe image.	7.7	More Details
CVE-2026-115416	Incorrect authorization checks in Amazon Athena engine version 3 query processing could have allowed an authenticated user to read AWS account identifiers and SQL statement text query metadata from other AWS accounts. Query results, and data stored in Amazon S3 or other data sources, were not accessible through this issue. AWS remediated the issue on October 7, 2026. No customer action is required.	7.7	More Details
CVE-2026-115717	Frappe Builder versions before 1.36.1 contain a server-side request forgery vulnerability in the convert_to_webp API endpoint, which lacks the has_page_write access control applied to similar asset-ingest handlers and performs an unprotected requests.get call that follows redirects without re-validating resolved IP addresses. A low-privilege authenticated user can supply an image_url pointing to an internal address or use DNS rebinding and open redirects to make the server fetch internal network resources or cloud metadata endpoints, with successful image responses saved to a world-readable /files/ URL enabling data exfiltration or blind probing of internal hosts.	7.7	More Details

CVE-2026-115484	A flaw was found in the iscsiio component of open-iscsi. The DHCPv4 option parser incorrectly calculates the length of the options data, causing it to read beyond the allocated buffer into adjacent heap memory. An attacker on the same network segment can send a crafted DHCP response to crash the daemon or corrupt its network configuration.	7.6	More Details
CVE-2026-116218	Dolibarr ERP CRM 19.0.0 before 23.0.4 contains an SQL injection vulnerability that allows authenticated API users to inject SQL via the type field of the third-party bankaccounts REST endpoint. Attackers with only societe/creer permission can break out of the unescaped type value in CompanyBankAccount::create() to read password hashes and API keys via time-based blind subqueries.	7.6	More Details
CVE-2026-115863	ConnMan through 2.0 has a stack-based buffer overflow via a DHCPv6 Reply message. This occurs in get_domains in gdhcp/client.c. It copies a fixed NS_MAXDNAME (1025) bytes from a DHCPv6 DOMAIN_LIST (option 24) value into a 1026-byte stack buffer dns_name regardless of the actual option length, then walks DNS labels with an unbounded "c += jump + 1" loop. A crafted DHCPv6 REPLY with a DOMAIN_LIST option of >= 1025 non-zero bytes makes the label walk leave the buffer and read/write '.' past its end in attacker-chosen strides. An on-link attacker can trigger this without authentication via a forged Router Advertisement (O=1) and a forged DHCPv6 REPLY matching xid/CLIENTID with a long DOMAIN_LIST. The impact is unauthenticated on-link crash of connmand (runs as root). This was introduced in bb06c47.	7.6	More Details
CVE-2026-114648	The MongoDB C# Driver does not neutralize dot and dollar characters in dictionary keys, field names and array-filter identifiers used in LINQ queries before placing them into the field paths it sends to the server. An application user whose input reaches these positions can redirect a query or update to fields the application did not name, which can expose other users' documents and widen updates beyond the intended data. Certain values can also write malformed fields that leave stored documents unreadable by the application until they are repaired.	7.6	More Details
CVE-2026-111851	Use of Externally-Controlled Input to Select Classes or Code ('Unsafe Reflection') vulnerability in WProyal Royal Elementor Addons royal-elementor-addons allows Code Injection.This issue affects Royal Elementor Addons: from n/a through 1.7.1068.	7.6	More Details
CVE-2026-115505	libsmb2 before 6.3 contains integer overflow vulnerabilities in the DCE/RPC NDR decoder in lib/dcerpc.c that allow heap out-of-bounds reads and writes. A malicious SMB server or machine-in-the-middle on unsigned connections can send crafted NetrShareEnum replies with wrapping actual_count or max_count values when clients call smb2_share_enum_async(), corrupting heap memory.	7.6	More Details
CVE-2026-98911	In the Linux kernel, the following vulnerability has been resolved: dpaa2-switch: fix the error path in dpaa2_switch_rx() In case of an error in dpaa2_switch_rx(), the dpaa2_switch_free_fd() function is called in order to free the FD. This is incorrect since the dpaa2_switch_free_fd() is intended to be used on Tx frame descriptors, meaning that it expects in the software annotation area of the FD data to find a valid skb pointer on which to call dev_kfree_skb(). Fix this by extracting the dma_unmap_page() from dpaa2_switch_build_linear_skb() directly into the dpaa2_switch_rx() function. This allows us to directly use free_pages() in case of an error before an SKB was created and kfree_skb() afterwards.	7.5	More Details
CVE-2026-20356	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco owned Open Source Snort 3 Software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20356 are related to issues concerning uncontrolled resource consumption that are grouped under the Common Weakness Enumeration (CWE) Base CWE-400.	7.5	More Details
CVE-2026-116543	Open5GS 1.3.0 through 2.8.0 contains a heap out-of-bounds write vulnerability in parse_multipart() that allows unauthenticated attackers to corrupt memory by supplying an empty multipart boundary. Attackers can send a single crafted multipart/related request to the SBI interface of any network function, overflowing the lookbehind buffer and crashing it.	7.5	More Details
CVE-2026-116183	phpMyFAQ before 4.1.10 contains an object-level authorization bypass vulnerability that allows unauthenticated attackers to modify or delete any visitor question by supplying an arbitrary openQuestionID to POST /api/faq/create. Attackers solving the CAPTCHA can rewrite question categories, forge answered entries in faqquestion_history, or permanently delete questions when records.enableDeleteQuestion is enabled.	7.5	More Details
CVE-2026-116489	Timescale pgvector before 0.9.1 contains a type confusion vulnerability that allows low-privileged database users to corrupt memory by shadowing pgvector's vector type or operators. Attackers with CREATE privilege on the vectorscale schema can pre-create shadow objects, then build and query DiskANN indexes over crafted values to cause out-of-bounds reads and writes, disclosing memory, crashing the server, or escalating privileges.	7.5	More Details
	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco owned Open		

CVE-2026-20236	Source Snort 3 Software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-20236 are related to issues concerning incorrect calculation that are grouped under the Common Weakness Enumeration (CWE) Base CWE-682.	7.5	More Details
CVE-2026-56833	PraisonAI is a multi-agent teams system. From version 3.8.1 to before version 4.6.59, PraisonAI's Dynamic Context module provides filesystem-backed history and terminal-log storage. The module also exports agent-callable tool factories. Those tools accept run_id and agent_id arguments from the tool caller. The underlying stores join those values into filesystem paths without rejecting absolute paths or .. traversal. Because run_id can be an absolute path and agent_id can contain traversal, a lower-trust prompt/user that can call these tools can read .jsonl and .log files outside the configured Dynamic Context base directory. This issue has been patched in version 4.6.59.	7.5	More Details
CVE-2026-114055	vLLM from 0.7.3 before 0.31.0 contains an uncontrolled resource consumption vulnerability that allows remote clients to control multimodal processor parameters by supplying per-request mm_processor_kwargs. Attackers can set large image size, patch count, crop count, or frame count values in Isaac, InternVL, H2OVL, and other processors to exhaust memory and terminate the shared server.	7.5	More Details
CVE-2026-114123	D365 F&O MCP Server is a Model Context Protocol server for querying Dynamics 365 Finance & Operations X++ code and metadata. Prior to version 1.11.0, the officially documented Azure one-click deployment uses infrastructure/azuredeploy.json and infrastructure/main.bicep to default the apiKey parameter to an empty string and pass it to the App Service API_KEY setting. With API_KEY unset or empty, src/middleware/apiKeyAuth.ts passes requests through without authentication, while src/index.ts binds the HTTP listener to 0.0.0.0 by default. If the deployer leaves the optional key blank and does not add separate network restrictions or upstream authentication, the resulting public App Service allows unauthenticated JSON-RPC tools/call requests to /mcp. The deployment's MCP_SERVER_MODE=read-only setting does not prevent anonymous access to read/search tools, including search, get_object_info, batch_get_info, extension_info, security_info, find_references, and find_event_handlers, exposing indexed proprietary X++ source snippets and metadata. Exposed information can include business logic, extension and event-handler wiring, security-role and privilege structure, and labels originally indexed from PackagesLocalDirectory. Deployments protected by a configured API key, effective upstream authentication, or network restrictions preventing attacker access are outside this unauthenticated exposure scenario. This issue is fixed in version 1.11.0.	7.5	More Details
CVE-2026-114059	ArcadeDB before 26.10.1 contains a remote code execution vulnerability that allows authenticated users, including read-only reader role holders, to execute Groovy code because GremlinQueryEngine.command() lacks scripting authorization checks. Attackers can submit Gremlin commands to the POST /api/v1/command/{database} endpoint when the Groovy engine is enabled to run OS commands as the server process user.	7.5	More Details
CVE-2026-112212	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 18.3 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, under certain conditions, an unauthenticated user could have poisoned a shared cache's OAuth/OIDC discovery document to redirect certain flows to an externally controlled server, due to improper trust of user-supplied values when constructing endpoint URLs in cached discovery responses.	7.5	More Details
CVE-2026-116179	vLLM from 0.18.0 before 0.31.0 contains a resource exhaustion vulnerability that allows API clients to bypass the video/jpeg frame-count cap by overriding num_frames through media_io_kwargs. Attackers can set num_frames to -1 with a large inline video so VideoMediaIO.load_base64 decodes every frame, exhausting memory and crashing the API frontend.	7.5	More Details
CVE-2026-116171	SiYuan 3.5.4 before 3.8.7 on Windows contains a sensitive-path denylist bypass in IsSensitivePath because the Windows extended-length prefix (\\?\) survives filepath.Clean and matches no forbidden prefix. Authenticated API callers can prepend the prefix to localPath in /api/import/importStdMd to copy workspace conf files, including the API token and TLS private keys, into assets served to unauthenticated clients.	7.5	More Details
CVE-2026-76306	In Splunk SOAR versions below 8.7.0, a user that has the "Observer" role could use query filters in the Representational State Transfer (REST) API to infer password hashes and authorization tokens for other users. The information disclosure is possible because Splunk SOAR applies REST API query filters to user records before it protects sensitive fields from query processing. For more information see 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), Query for Data (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/7.1.0/using-the-splunk-soar-rest-api/query-for-data), 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-	7.5	More Details

	permissions-in-splunk-soar-on-premises) in the Splunk documentation.		
CVE-2026-115937	The Mollie webhook endpoint accepts two independent, attacker-controlled parameters and never verifies that they belong together: * commerceTransactionHash — selects which Craft Commerce transaction is completed * id — selects which Mollie payment is queried for its status An unauthenticated attacker holding any paid Mollie payment ID from the merchant's account can pair it with any transaction hash. Commerce queries Mollie, sees status: paid, and marks the target transaction successful for that transaction's full amount, regardless of how much was actually paid.	7.5	More Details
CVE-2026-115721	**UNSUPPORTED WHEN ASSIGNED** Unbounded allocation in XML parsing in Apache Vysper; all versions, may lead to an out of memory condition. This issue affects all versions of Apache Vysper. Apache Vysper is an XMPP (Jabber) server library. As this project is retired, we do not plan to release a version that fixes this issue. Users are recommended to find an alternative.	7.5	More Details
CVE-2026-114476	Incorrect authorization in the REST administration API of Progressive Robot hMailServer 6.3.4 through 6.3.6 allows a remote authenticated user whose credential holds the manage_roles permission in named domains - a full-scope API key issued for those domains, or an account assigned the serveradmin role with a list of domains - to obtain the server administrator's authority. An account's admin level gives it authority, and a role assigned to it replaces that authority, which is how an administrator narrows an account whose admin level is the server administrator's. The route that takes a role away (DELETE /api/v1/role-assignments/{address}) checked only that the credential takes in what the account held at the time, not what it would hold afterwards, so such a credential could take a narrowing role away from an account in its domains whose admin level is the server administrator's, returning it to the server administrator's authority, and, having set that account's password while it held the narrower role, sign in as the server administrator.	7.5	More Details
CVE-2026-115887	GraphQL Scalars is a library of custom GraphQL scalars for creating precise type-safe GraphQL schemas. From 1.24.1 until 2.0.1, GraphQLPhoneNumber in src/scalars/PhoneNumber.ts uses ambiguous nested digit quantifiers that backtrack exponentially when a PhoneNumber value contains a run of digits followed by a character that makes the match fail. GraphQL evaluates parseLiteral during document validation and parseValue during variable coercion before resolver authentication or authorization, so a small unauthenticated request can block the Node.js event loop and stall the server process. serialize is also affected when unvalidated stored values are returned through the scalar. This issue is fixed in version 2.0.1.	7.5	More Details
CVE-2026-111181	body-parser is body parsing middleware for Node.js and Express. In versions before 1.20.9 and in versions from 2.0.0 before 2.4.0, a single request can crash any application using body-parser with default options. When a request carries a Content-Encoding such as gzip, body-parser pipes it into a decompression stream and, after reading the body, leaves that stream with no error listener. A payload of one valid compressed frame followed by trailing bytes decodes and the application responds normally, but the trailing bytes then reach the already-finished decoder, which emits an error that is re-thrown as an uncaught exception and terminates the process. No matching route is required because body parsing runs before routing, so any endpoint that parses a body is affected. The issue is fixed in body-parser 2.4.0 and 1.20.9, and users should upgrade to one of those versions or later. As a workaround, disable request-body inflation (set the body-parser inflate option to false) so encoded requests are rejected before any decompression stream is created.	7.5	More Details
CVE-2026-113872	multer is a middleware for Node.js that handles multipart/form-data file uploads. In versions from 2.0.0 before 2.5.0, when multer is configured with an asynchronous fileFilter, a single request can crash the Node.js process. If the request is aborted while one file's asynchronous fileFilter is still pending, the pending callback later resumes a parsing stream that multer has already torn down, and the underlying parser emits an error on a stream that no longer has an error listener, which Node raises as an unhandled error event so the process exits. Any application using multer with an asynchronous fileFilter is affected regardless of storage engine, while uploads using a synchronous fileFilter are not affected. The issue is fixed in multer 2.5.0, and users should upgrade to 2.5.0 or later. As a workaround, use a synchronous fileFilter.	7.5	More Details
CVE-2026-113582	UNSUPPORTED WHEN ASSIGNED Improper restriction of XML external entity references in Apache Abdera allows an attacker to access sensitive data or make server-side requests when an application parses attacker-controlled XML with Abdera. This issue affects Apache Abdera: all versions. Apache Abdera is retired and no longer supported, so we do not plan to release a version that fixes this issue. Users are recommended to migrate to a maintained alternative. Applications that cannot migrate should ensure that untrusted XML cannot reach the affected Abdera parser or JDBC adapter parsing paths. Actively maintained downstream products may still be impacted if they include an affected Abdera artifact. NOTE: This vulnerability only affects products that are no longer supported by the maintainer.	7.5	More Details
	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway		

CVE-2026-76493	Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76493 are related to issues concerning uncaught exceptions that are grouped under the Common Weakness Enumeration (CWE) Base CWE-248.	7.5	More Details
CVE-2026-114052	vLLM after 0.24.0 and before 0.31.0 contains a memory exhaustion vulnerability in MOSS-Audio model serving that caches a processor for every distinct request-controlled mm_processor_kwargs value in an unbounded dictionary. Unauthenticated attackers can send distinct or oversized mel_config values to the /tokenize endpoint, which bypasses --api-key, to exhaust host memory and crash the API process.	7.5	More Details
CVE-2026-114191	Missing Authorization vulnerability in Knit Pay UPI QR Code Payment Gateway for WooCommerce upi-qr-code-payment-for-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects UPI QR Code Payment Gateway for WooCommerce: from n/a through 1.6.3.	7.5	More Details
CVE-2026-111810	Apache Traffic Server 9.2.15, 10.1.4 and 10.2.0 write one unthrottled diagnostic log line for every internal @ header they strip, so a request carrying many such headers makes the proxy write far more log data than the request contains. This issue affects Apache Traffic Server: 9.2.15, 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	7.5	More Details
CVE-2021-48017	OrbiTeam BSCW Server 5.x before 5.2.5 and 7.x before 7.4.4 contains an unauthenticated path traversal vulnerability in the style_name parameter of the theme operation. Attackers can send GET requests to /pub/bscw.cgi/<id>?op=theme with ../ sequences in style_name to read local files such as /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-22.	7.5	More Details
CVE-2026-116011	Crow before 1.3.4 contains an uncontrolled recursion vulnerability in crow::json::load() that allows remote unauthenticated attackers to crash the server by sending deeply nested JSON. Attackers can send a request body with thousands of nested objects to exhaust the worker thread stack, terminating the process and dropping all connections.	7.5	More Details
CVE-2026-110353	HAProxy 3.2 before 3.4.6 allows a remote unauthenticated attacker to cause a denial of service (worker-thread CPU exhaustion) via a truncated HTTP/2 PING frame. h2c_handle_ping in src/mux_h2.c does not check for a short read, unlike the other H2 frame handlers. Because the incomplete frame cannot be processed, its data is never consumed, and the H2 multiplexer tasklet is repeatedly woken by that pending data without making progress. The CPU spins until timeout http-request closes the connection, and an attacker opening many such connections can sustain high CPU usage.	7.5	More Details
CVE-2021-48015	JeeWMS before commit 0a610e8 contains an unauthenticated path traversal vulnerability that allows remote attackers to read arbitrary files via the dbPath parameter. Attackers can send GET requests to /systemController/showOrDownByUrl.do with ../ sequences in dbPath, which bypasses AuthInterceptor, to retrieve files like win.ini. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-17.	7.5	More Details
CVE-2021-48013	Weaver E-cology 8.0 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files through the fpath parameter of the /weaver/In.FileDownload endpoint. Attackers can send crafted GET requests with directory traversal sequences in fpath to retrieve sensitive files such as WEB-INF/web.xml from the server. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-18.	7.5	More Details
CVE-2026-110354	HAProxy 3 before 3.4.6 allows a remote unauthenticated attacker to cause a denial of service (persistent worker-thread CPU exhaustion) via an HTTP/1.1 chunked request. Although the H1 multiplexer drains the request before closing the stream, the parser repeatedly re-processes the same bytes without making progress. The affected thread stays busy at 100% CPU after timeouts expire and after the client disconnects, and repeated requests can consume all worker threads until the process is restarted.	7.5	More Details
CVE-2019-25779	WooThemes Diarise through 1.5.9 contains a file read vulnerability that allows unauthenticated attackers to read arbitrary files through the calendar parameter of download.php. Attackers can send GET requests with file:// scheme paths to the calendar parameter to disclose sensitive files like /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	7.5	More Details
CVE-2026-114063	SvelteKit (@sveltejs/kit) from 2.10.0 before 3.0.0 contains a path traversal vulnerability in the vite preview server on Windows that allows unauthenticated attackers to read files outside the prerendered output directory. Attackers who can reach the preview server through network access or a tunnel can send crafted requests to expose source code, configuration, or credentials readable by the server process.	7.5	More Details
CVE-2018-25439	MetInfo CMS 6.0.0 before 6.1.1 contains a path traversal vulnerability in /include/thumb.php that allows unauthenticated attackers to read arbitrary files via the dir parameter. Attackers can send crafted requests with backslash sequences bypassing the str_replace filter to read files like config_db.php on Windows. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-14.	7.5	More Details

CVE-2018-25438	Intermesh Group-Office 3.3.0 before 3.5.5 contains an unauthenticated path traversal vulnerability in compress.php that allows remote attackers to read arbitrary files by manipulating the 'file' parameter. Attackers can send GET requests to /compress.php with '../' sequences in the 'file' parameter, which is appended to the cache path and passed to readfile(), to read files such as /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-21.	7.5	More Details
CVE-2017-20289	Broadstreet Ad Widget plugin for WordPress 2.2.0 before 2.12.0 contains a local file inclusion vulnerability that allows authenticated users to execute attacker-chosen PHP files by passing unsanitized input to require in views/modal/index.php. Attackers can supply directory traversal sequences in the 'step' parameter, optionally with a null byte, to include local files and execute PHP code. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	7.5	More Details
CVE-2017-20287	Javo Spot theme for WordPress before 3.0.0 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files via the fn parameter of the jvfrm_spot_get_json AJAX action. Attackers can send GET requests to /wp-admin/admin-ajax.php with fn set to ../wp-config.php to retrieve database credentials and secret keys. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	7.5	More Details
CVE-2026-111958	The Royal MCP WordPress plugin before 1.5.7 does not limit a preview-link token to the single preview it was issued for, so any party that obtains a valid token is authenticated as the token's owning user, an administrator in the default configuration, across the entire REST API with no nonce, leading to full site takeover.	7.5	More Details
CVE-2026-97328	The WebP Express WordPress plugin before 0.25.16 does not escape values that its bundled HTML parser re-emits into element attributes when a non-default image delivery mode is enabled, allowing unauthenticated users to submit content that is stored and later executed as Cross-Site Scripting in the browser of any user viewing an affected page, including administrators.	7.5	More Details
CVE-2026-112312	Missing Authorization vulnerability in ACF Extended ACF Extended acf-extended allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects ACF Extended: from n/a through 0.9.2.7.	7.5	More Details
CVE-2026-114050	vLLM before 0.31.0 uses a client-supplied multimodal uuid verbatim as cache identity with no consistency check on cache hits, allowing engine crashes. Unauthenticated attackers can desynchronize the P0/P1 processor caches to trigger a fatal token-placeholder mismatch, killing the entire serving instance until restarted.	7.5	More Details
CVE-2026-114136	Django REST framework is a toolkit for building Web APIs. From 3.6.4 until 3.18.2, JSONParser.parse in rest_framework/parsers.py uses the client-supplied Content-Type charset value with codecs.getreader, which accepts byte-transform codecs such as bz2_codec rather than only text encodings. A remote unauthenticated client can therefore cause a compressed JSON request body to be expanded after Django has enforced DATA_UPLOAD_MAX_MEMORY_SIZE against only the wire bytes, allowing a very small request to consume substantial per-worker memory. The default parser set is sufficient, and any reachable endpoint that accepts a JSON body is affected. This issue is fixed in version 3.18.2.	7.5	More Details
CVE-2026-96365	Allocation of Resources Without Limits or Throttling vulnerability in Drupal Webform allows Excessive Allocation. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	7.5	More Details
CVE-2026-114051	vLLM before 0.31.0 skips the min_tokens <= max_tokens check for /inference/v1/generate requests with a null max_tokens, so the EngineCore input thread later dies on an uncaught VLLMValidationError while /health still returns 200. Unauthenticated remote attackers can send one request with an oversized min_tokens to hang all subsequent generation requests until manual restart.	7.5	More Details
CVE-2026-114138	Django REST framework is a toolkit for building Web APIs. Prior to 3.18.2, DefaultContentNegotiation.get_accept_list in rest_framework/negotiation.py accepts an unbounded number and aggregate size of Accept tokens, while select_renderer in that file repeatedly invokes _MediaType parsing from rest_framework/utils/mediatypes.py for each attacker-controlled token, every renderer, and forced 406 renegotiation. A remote unauthenticated client can send many distinct media types containing large quoted parameter values with numerous semicolons, causing Django's parameter parser and DRF's repeated parsing structure to consume disproportionate CPU. Requests to any reachable endpoint can pin synchronous workers for prolonged periods, and concurrent requests can make the service unavailable until the requests stop. This issue is fixed in version 3.18.2.	7.5	More Details
CVE-2026-	As part of Cisco's ongoing commitment to proactive security and product quality, engineering teams for Cisco Emergency Responder, Cisco Prime Collaboration Deployment, Cisco Unified Communications Manager (Unified CM), Cisco Unified Communications Manager IM & Presence Service (Unified CM IM&P), Cisco Unified Communications Manager Session Management Edition (Unified CM SME), and Cisco Unity Connection have conducted a comprehensive internal security review. This review resulted	7.5	More

76481	in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76481 are related to issues concerning improper handling of exceptional conditions that are grouped under the Common Weakness Enumeration (CWE) Pillar CWE-703.		Details
CVE-2026-76290	In Splunk Enterprise versions below 10.4.4, 10.2.8, and 10.0.11, an unauthenticated user could send repeated requests to the Edge Processor package download endpoint in the Splunk Pipeline Builders app. The requests could cause the Splunk daemon to consume excessive system resources and potentially stop responding. The vulnerability is possible because the endpoint accepts package download requests without authentication and reads and base64-encodes the complete package in memory for each request. After upgrading, prepare all existing Edge Processor instances, turn on authentication for Edge Processor package downloads, and restart Splunk Enterprise. For more information, see Enable authentication for Edge Processor package downloads (https://help.splunk.com/en/data-management/process-data-at-the-edge/use-edge-processors-for-splunk-enterprise/10.4/administer-edge-processors/enable-authentication-for-edge-processor-package-downloads).	7.5	More Details
CVE-2024-58396	YouDataSum CPAS Audit Management System V4 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files by manipulating the fileld parameter. Attackers can send GET requests to /cpasm4/pluginManController/downPlugs with ../ sequences in fileld to retrieve system files like /etc/passwd. Exploitation evidence was first observed by the Shadowserver Foundation on 2025-01-13.	7.5	More Details
CVE-2026-76288	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, an unauthenticated user could send a single crafted request to the Splunk Web JavaScript logger endpoint and cause Splunk Web to stop responding to requests. Splunk Web remains unavailable until Splunk Enterprise is restarted. The vulnerability is possible because Splunk Web does not limit request-header parsing before processing the request.	7.5	More Details
CVE-2026-76491	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76491 are related to issues concerning reachable assertions that are grouped under the Common Weakness Enumeration (CWE) Base CWE-617.	7.5	More Details
CVE-2026-76492	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76492 are related to issues concerning null pointers that are grouped under the Common Weakness Enumeration (CWE) Base CWE-476.	7.5	More Details
CVE-2026-56834	PraisonAI is a multi-agent teams system. From version 3.8.1 to before version 4.6.59, PraisonAI's Dynamic Context Discovery feature exposes artifact helper tools through ctx.get_tools(). The implemented artifact tools do not enforce that the supplied artifact_path is an artifact created by the configured store or that it lives under the configured artifact base directory. Instead, artifact_head, artifact_tail, artifact_grep, and artifact_chunk wrap the caller-supplied path directly into an ArtifactRef and then read it from the host filesystem. As a result, any prompt/user/tool-caller that can influence those tool arguments can read files readable by the PraisonAI process, such as project .env files, cloud credentials, SSH keys, source files, or other local data. This issue has been patched in version 4.6.59.	7.5	More Details
CVE-2026-116189	rclone from 1.55.0 before 1.75.2 contains a signature verification bypass in stable selfupdate that trusts checksums from unsigned bytes preceding the PGP-clearsigned SHA256SUMS block. Attackers controlling the release distribution path can prepend an unsigned checksum line matching a malicious archive, causing selfupdate to install an attacker-controlled rclone executable.	7.5	More Details
CVE-2022-51024	Kavita 0.5.2 before 0.5.4.1 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files by supplying '../' sequences in the filename parameter. Attackers can send GET requests to /api/image/cover-upload with filename values like ../appsettings.json to retrieve configuration files outside the temp directory. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-18.	7.5	More Details
CVE-2026-76505	A vulnerability in the IPv4 packet processing functionality of Cisco IOS XR Software that affects Cisco ASR 9000 and 9900 Series Aggregation Services Routers as well as Cisco ASR 9902 and 9903 Compact High-Performance Routers could allow an unauthenticated, remote attacker to cause network processor (NP) errors, resulting in a line card reset or shutdown and a denial of service (DoS) condition. This vulnerability is due to the incorrect handling of crafted IPv4 packets that are received on an affected line card interface and processed as for-us traffic. The issue can result in an NP reset, causing traffic loss on the affected line card while it recovers. Under repeated reset conditions, the network	7.5	More Details

	processor may enter a persistent shutdown state, in which case an explicit line card reload may be required to restore traffic. The vulnerability is within the IPv4 packet-processing path. Transit IPv4 traffic does not trigger this issue and IPv6 traffic cannot trigger it.		
CVE-2026-116191	Rclone before 1.75.2 contains a path traversal vulnerability in serve webdav that allows unauthenticated attackers to write above the served root by sending an unvalidated '..' path component. Attackers can send MKCOL ../ to plant a VFS node, then PUT, MOVE, COPY, or MKCOL to create, overwrite, or delete files belonging to other tenants.	7.5	More Details
CVE-2026-115728	@hapi/content from 6.0.1 before 6.0.3 contains a regular expression denial of service vulnerability in Content.disposition() caused by an incomplete fix for CVE-2026-35213. Unauthenticated remote attackers can send multipart uploads with a long Content-Disposition parameter lacking an equals sign, triggering quadratic backtracking that blocks the event loop.	7.5	More Details
CVE-2026-114656	Traefik before 2.11.58 and 3.7.13 contains an authentication bypass vulnerability that allows unauthenticated clients to skip client-certificate checks because negated Host matchers are counted as served domains. Attackers can reach backends protected by RequireAndVerifyClientCert when a router rule with !Host() creates a TLS options conflict that falls back to the default options.	7.5	More Details
CVE-2026-96640	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, PdfParser.linearize_page_tree() in src/PIL/PdfParser.py treats a PDF /Pages graph as a tree without tracking duplicate indirect references. When Image.save(..., format="PDF", append=True) processes an attacker-controlled existing PDF whose /Kids arrays repeatedly reference the same child /Pages node, Pillow expands the shared subtree as distinct pages. A compact malformed PDF can therefore create exponentially many page references and consume disproportionate CPU before the append completes. This issue is fixed in version 13.0.0.	7.5	More Details
CVE-2026-114145	GitPython is a python library used to interact with Git repositories. Prior to 3.2.1, create_submodule() in git/objects/submodule/base.py accepts an attacker-influenced submodule path such as .git/hooks and invokes _clone_repo() before safely rejecting the destination. In 3.1.x, _validate_repo_path() does not exist and the call succeeds silently, while in 3.2.0 and later affected revisions the check in git/index/base.py raises ValueError only after the clone has written to disk and does not remove the planted file. An attacker-controlled submodule can therefore place an executable post-checkout hook that a later Git operation executes. Exploitation requires an application to accept an untrusted submodule path and a clone destination such as .git/hooks that is empty enough for Git to clone into it, which commonly means that default sample hooks are absent. This issue is fixed in version 3.2.1.	7.5	More Details
CVE-2026-114144	toml-node is a TOML parser for Node.js and the browser. Prior to 5.0.1, toml.parse() passes unbounded hexadecimal, octal, binary, and decimal digit sequences from lib/parser.js to BigInt() before reduceInteger() in lib/compiler.js checks whether the resulting value fits the TOML 64-bit integer range. Converting a sufficiently long integer string consumes super-linear CPU time and synchronously blocks the Node.js event loop, allowing one attacker-supplied TOML document to deny service to concurrent work in the process. The later range check cannot prevent the conversion cost because it executes after BigInt(). This issue is fixed in version 5.0.1.	7.5	More Details
CVE-2026-57146	PraisonAI is a multi-agent teams system. From version 4.5.115 to before version 4.6.61, the published A2U advisory GHSA-f292-66h9-fpmf says unauthenticated A2U event streaming was fixed in praisonai 4.5.115. Current head still exposes the same A2U subscription and event routes without authentication when the operator starts the documented CLI entryptpoint. The current CLI wrapper does not expose --api-key, does not install the common API-key middleware, and does not generate a token for A2U. It calls create_a2u_routes(app) directly. That helper only enforces auth if A2U_AUTH_TOKEN is already present; if the variable is missing, _authenticate_request() returns None and treats auth as disabled. This issue has been patched in version 4.6.61.	7.5	More Details
CVE-2026-65825	OpenReception's appointment booking software provides an end-to-end encrypted appointment booking platform. In versions 2.1.0 and prior, a returning client's six-digit booking PIN can be brute-forced online in the victim tenant. A successful guess yields the client's booking access token, granting access to that client's appointment tunnel, which holds their end-to-end-encrypted appointment data, and the ability to act as the client, for example adding or cancelling appointments in the tunnel. The bypass is cross-tenant: state owned by one tenant removes an authentication protection of a different tenant. This crossing of the tenant boundary is the basis for the scope-changed vector. At time of publication, there are no publicly available patches.	7.5	More Details
CVE-2026-115355	Mistune before 3.3.0 contains a path traversal vulnerability in the Include directive that allows attackers to read arbitrary files by supplying ../ sequences in include paths. Attackers can also inject HTML through unescaped error messages and included file contents to perform cross-site scripting even when escape=True is set.	7.5	More Details
CVE-	The Simple Banner WordPress plugin before 3.3.3 does not correctly limit the capabilities it grants to		

2026-107506	the user roles an administrator allows to manage its banners, granting such roles the full administrator capability, which allows users in a low-privileged role to escalate their privileges and fully compromise the site.	7.5	More Details
CVE-2026-79955	An unauthenticated memory corruption vulnerability exists in a service accessed via the PAPI protocol in AOS-8 that could allow an unauthenticated remote attacker to conduct a denial of service attack. Successful exploitation could allow an attacker to disrupt the availability of the affected service.	7.5	More Details
CVE-2026-113873	SmileCDR HAPI FHIR does not properly handle highly compressed data, allowing an unauthenticated attacker to send compressed payloads that exhaust server heap or crash certain services when decompressed.	7.5	More Details
CVE-2026-111833	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Bold Grid weForms weforms allows Path Traversal.This issue affects weForms: from n/a through 1.6.28.	7.5	More Details
CVE-2026-79894	An unauthenticated Denial-of-Service (DoS) vulnerability exists in the PAPI protocol of AOS-8. Successful exploitation of this vulnerability results in the ability to disrupt the availability of the affected interface.	7.5	More Details
CVE-2026-96185	Stack-based buffer overflow in Microsoft Exchange Online allows an authorized attacker to execute code over a network.	7.5	More Details
CVE-2026-114613	MessagePack is the serializer implementation for Python msgpack.org. Prior to 1.2.3, the msgpack-python C extension (msgpack._cmsgpack) leaks a decoded map value when deserialization rejects a key with strict_map_key=True (the default). Repeated malformed inputs can cause unbounded process memory growth and denial of service in applications that deserialize attacker-controlled MessagePack data. Both unpackb() and the streaming Unpacker API are affected. The pure-Python fallback does not exhibit this leak in the reported tests. This issue is fixed in version 1.2.3.	7.5	More Details
CVE-2026-79893	Memory corruption vulnerabilities in AOS-10 could allow an unauthenticated remote attacker to conduct a denial of service attack. Successful exploitation could allow an attacker to disrupt the availability of the affected service.	7.5	More Details
CVE-2026-112330	Motor.NET before 0.19.0 contains a denial of service vulnerability that allows a remote attacker who can publish to a consumed queue to crash the service permanently by sending a message with crafted cloudEvents headers. Attackers can supply an invalid cloudEvents:motorversion, a single double-quote header value, or an unparseable typed attribute such as cloudEvents:time to trigger an unhandled exception in ExtractCloudEvent that is treated as fatal, causing the application to stop and the unacknowledged message to be redelivered on restart, resulting in a permanent crash-loop.	7.5	More Details
CVE-2026-50787	e-SIC Livre (esiclivre) v0.2.2 and earlier contains an uncontrolled resource consumption vulnerability in restrito/inc/captcha.php. The application uses user-supplied parameters (l, a, tf, ql) without validation to generate CAPTCHA images, allowing remote attackers to cause denial of service via excessive memory and CPU consumption.	7.5	More Details
CVE-2026-115347	Mistune before 3.3.0 contains an inefficient algorithmic complexity vulnerability in parse_link_text that allows remote attackers to cause denial of service through unbalanced square brackets. Attackers can submit markdown containing long runs of unmatched brackets to mistune.html(), triggering repeated forward scans that stall worker processes for minutes.	7.5	More Details
CVE-2026-115348	Mistune from 3.0.0 before 3.3.0 contains an algorithmic complexity vulnerability in the inline link parser's parse_link_text function that allows unauthenticated attackers to exhaust CPU with unbalanced brackets. Attackers can submit Markdown containing many unbalanced '[' or '[' characters, forcing repeated unbounded forward scans and tying up worker processes for minutes.	7.5	More Details
CVE-2026-115349	Mistune from 3.0.0 before 3.3.0 contains an algorithmic complexity denial of service vulnerability in the parse_link_text function in src/mistune/helpers.py. Attackers can submit markdown with deeply nested square brackets triggering quadratic parsing time, exhausting server CPU with small crafted requests.	7.5	More Details
CVE-2026-115352	Mistune before 3.3.0 contains an algorithmic complexity vulnerability in the inline link parser that allows unauthenticated attackers to exhaust CPU by submitting long runs of unmatched '[' characters. Attackers can send small markdown payloads that trigger O(N^2) parse_link_text and parse_link_label probing, occupying rendering workers for 30+ seconds per 16 KB input.	7.5	More Details
CVE-2026-116177	vLLM from 0.10.2 before 0.25.0 contains an uncontrolled resource consumption vulnerability that allows unauthenticated attackers to exhaust memory because the Nemotron-VL processors globally set Image.MAX_IMAGE_PIXELS to None. Attackers can submit a pathological decompression-bomb image to /v1/chat/completions to OOM-kill the worker and deny service to co-tenants.	7.5	More Details

CVE-2026-97206	NLnet Labs Unbound, 1.22.0 up to and including version 1.26.1, contains a DNS-over-QUIC (DoQ) lifetime bug reachable by a remote DoQ client under multi-threaded operation. A client-triggered QUIC path migration can enter stream/connection cleanup while a mesh reply still retains a pointer to the DoQ stream. Cleanup frees the DoQ stream and when the pending mesh reply later completes it dereferences a stale pointer leading to a crash.	7.5	More Details
CVE-2026-111949	The miniOrange 2FA WordPress plugin before 6.3.1 does not prevent a pending passkey login from registering an additional passkey on an account that already has one configured, allowing an attacker who knows a user's password to add their own authenticator and bypass the second factor.	7.5	More Details
CVE-2026-112331	Motor.NET before 0.19.0 contains a denial-of-service vulnerability that allows a remote attacker who can produce records to a consumed topic to crash the service permanently by sending a single crafted record. By setting the content-type header to application/cloudevents+json and omitting required CloudEvent attributes (id, source, type), the attacker triggers an ArgumentException in KafkaMessageConsumer.SingleMessageHandlingAsync that is treated as fatal, causing the host to stop and re-process the same poison record on every restart, resulting in a permanent crash-loop.	7.5	More Details
CVE-2026-97313	A recursive BIND server that receives a query with recursion desired (RD=1) followed by a non-recursive query (RD=0) may exit with an assertion failure, resulting in a denial of service. To be vulnerable, the BIND server must allow recursive queries and have at least two RPZ zones configured: one with the default setting of "recursive-only yes", and the other with "recursive-only no". This issue affects BIND 9 versions 9.16.0 through 9.18.50, 9.20.0 through 9.20.29, 9.21.0 through 9.21.26, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.29-S1.	7.5	More Details
CVE-2026-96888	A BIND resolver that is configured with both a "redirect" zone and "qname-minimization strict" may abort with an assertion failure while processing the redirect. This issue affects BIND 9 versions 9.16.0 through 9.18.50, 9.20.0 through 9.20.29, 9.21.0 through 9.21.26, 9.16.8-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.29-S1.	7.5	More Details
CVE-2026-99030	In the Linux kernel, the following vulnerability has been resolved: net: skbuff: don't leave stale bytes in skb_copy_and_csum_bits() When skb_copy_and_csum_bits() reaches unreadable frags it returns 0 after copying only the linear part, and the rest of the caller's buffer is left as it was. The callers copy into a buffer that is about to go out on the wire: an ICMP error quoting the offending packet, or a driver's TX bounce buffer in skb_copy_and_csum_dev(). Neither buffer is zeroed beforehand, so whatever was in memory there gets sent. Zero the part of the buffer we didn't fill. The checksum usually won't match the data any more, so the receiver will usually drop the packet, but either way it no longer carries anything it shouldn't. Only zero for a positive @len, a negative one from a broken caller must not turn into a huge memset().	7.5	More Details
CVE-2026-115417	Memory allocation with excessive size value in the MQTT 3.1.1 compatibility adapter in AWS aws-c-mqtt before 0.16.2 might allow a remote actor to cause a denial of service of the client application. To remediate this issue, users should upgrade to aws-c-mqtt 0.16.2 or later, or to an AWS Common Runtime or AWS IoT Device SDK release that bundles it.	7.5	More Details
CVE-2026-115497	The Music Store – WordPress eCommerce plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 1.4.7 via the 'f' 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 a valid purchase_id token, which any user can obtain by making a single low-cost purchase or receiving a shared download link.	7.5	More Details
CVE-2026-106765	An issue in @rog0x/mcp-docker-tools 1.0.1 allows a remote attacker to execute arbitrary code via the handleImageList function.	7.5	More Details
CVE-2026-106764	An issue in @rog0x/mcp-git-tools 1.0.1 allows a remote attacker to execute arbitrary code via the resolveBaseBranch function.	7.5	More Details
CVE-2026-106763	An issue in @justice8096/sagemath-mcp-server 1.0.1 allows a remote attacker to execute arbitrary code via the escapeShellString and sanitizeCode functions.	7.5	More Details
CVE-2026-	Uncontrolled recursion in the mcpack2pb component in Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to make the server crash via sending deeply nested mcpack data. Root Cause: The bRPC mcpack2pb parser (used to serve nshead_mcpack / ubrpc / public_pbrpc requests) recursively parses nested objects and arrays without any depth limit. For a protobuf message that references itself (e.g. a repeated message field), the generated parser recurses once per nesting level of the incoming mcpack object. A deeply nested request makes the recursion grow unbounded, overflow the stack and crash the process. Affected Scenarios: Use bRPC server with a recursive-schema protobuf message to	7.5	More

82230	serve nshead_mcpack / ubrpc / public_pbrpc requests from untrusted network. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3499 Note: No matter which option you choose, you should know that the fix enforces the documented parsing depth limit MAX_DEPTH (128) on the parse path, the same limit the serializer already applies. If your requests contain mcpack messages nested deeper than 128 levels, the request will be failed after applying the fix.		Details
CVE-2026-19630	In a BIND resolver, inbound DNS protocol messages with a particular structure will result in additional processing, with no limit imposed by BIND. A DNS message with a maliciously large number of such structures will cause BIND to consume disproportionate amounts of CPU time and memory during processing. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.29, 9.21.0 through 9.21.26, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.29-S1.	7.5	More Details
CVE-2026-115506	libexpat through 2.9.0 contains a denial of service vulnerability because billion laughs protection in accountingDiffTolerated() in xmlparse.c detects entity amplification too late. Attackers can supply a moderately sized XML document with a large DTD and nested entity expansions to consume tens of seconds of CPU per parse.	7.5	More Details
CVE-2026-98964	In the Linux kernel, the following vulnerability has been resolved: nvmet: use the local P2P device on metadata allocation failure nvmet_req_alloc_p2pmem_sgls() assigns req->p2p_dev only after the data and metadata SGL allocations succeed. If the data allocation succeeds but the metadata allocation fails, out_free_sg passes req->p2p_dev to pci_p2pmem_free_sgls() before that assignment has taken place. Pass the local p2p_dev argument instead, so cleanup uses the same PCI device as the successful data allocation.	7.5	More Details
CVE-2026-112332	Motor.NET before 0.19.0 contains an uncaught exception vulnerability in JsonSerializer.Deserialize that allows a remote attacker who can publish to a consumed queue or topic to crash the service permanently by sending structurally valid JSON that does not match the handler's expected type. Attackers can exploit the fact that JsonSerializerException is not caught alongside JsonReaderException and InvalidCastException, causing TypedConsumerService.SingleMessageConsumeAsync to treat the error as CriticalFailure and call StopApplication, which combined with the unacknowledged message being redelivered on restart results in a permanent crash-loop.	7.5	More Details
CVE-2026-79895	Vulnerabilities in the affected interface of AOS-10 could allow an unauthenticated remote attacker to conduct a denial of service attack. Successful exploitation could allow an attacker to disrupt the availability of the affected interface.	7.5	More Details
CVE-2026-111574	The Redirection for Contact Form 7 WordPress plugin before 3.2.12 does not sanitise or escape submitted form data stored as lead entries before outputting it in the lead debug metabox, allowing unauthenticated attackers to store a malicious script that executes when a logged-in user views the entry with the Redirection for Contact Form 7 WordPress plugin before 3.2.12's debug mode enabled.	7.5	More Details
CVE-2026-102483	Reachable exit in NTPsec ntpd 1.1.4 through 1.2.5 allows a remote attacker to terminate the daemon via an NTS request with excess NTS Cookie Placeholder extension fields.	7.5	More Details
CVE-2026-75507	GraphQL-core is a Python port of the GraphQL.js reference implementation. From 3.0.0 until 3.2.12 and 3.3.0rc1, the OverlappingFieldsCanBeMerged rule in specified_rules repeatedly compares fields that share response names and recursively re-compares flattened inline fragments during the validation phase. The prior named-fragment memoization does not cover this path, so nested repeated inline fragments can drive quadratic and worst-case $O(n^2 \cdot m^2)$ CPU work before any resolver or AST-based depth or complexity control runs. An unauthenticated remote attacker can submit a shallow, highly compressible GraphQL query that pins a worker process and can saturate a multi-worker service. This issue is fixed in versions 3.2.12 and 3.3.0rc1.	7.5	More Details
CVE-2026-85642	The Formidable Forms WordPress plugin before 7.0 does not verify the identity of the caller before recording the timestamp it uses to rate limit payment gateway verification requests, allowing unauthenticated attackers to consume the daily allowance and cause legitimate verification attempts to be refused for the next 24 hours.	7.5	More Details
CVE-2026-74757	Uncontrolled resource allocation in the mcpack2pb component in Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to make the server allocate a large amount of memory via sending a crafted mcpack request with a huge array item count. Root Cause: The bRPC mcpack2pb parser trusts the item count stored in the mcpack array header without validating it against the actual payload size. The code generated by the mcpack2pb protoc plugin uses this count to preallocate memory for repeated protobuf fields via Reserve(). A malformed request declaring item_count = 0x7fffffff can force the server to preallocate about 16GB of virtual memory, which may abort the process (SIGABRT) on memory-constrained hosts. Affected Scenarios: Use a bRPC server that wires a NsheadService to	7.5	More Details

	UbrpcMcpack2Adaptor (the Baidu ubrpc deployment pattern) and parses requests with code generated by the mcpack2pb protoc plugin. The request body is parsed before authentication. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3451 Note: No matter which option you choose, you should know that the fix caps the array item count to the remaining bytes of the array payload.		
CVE-2026-107127	The Uncanny Toolkit for LearnDash WordPress plugin before 3.8.2 does not validate the origin of a password-reset request, allowing an unauthenticated attacker to set the password of a user who is partway through a password reset to an attacker-chosen value through a forged request, resulting in full account takeover.	7.5	More Details
CVE-2026-115482	In Isquic before 4.9.0, the HTTP/3 implementation can use an excessive amount of memory because the number of buffered header sets is not limited, as demonstrated by a QPACK amplification attack. This is related to header_out_dynamic_entry and es_max_header_list_size.	7.5	More Details
CVE-2026-115778	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, twisted.conch.telnet.Telnet.dataReceived appends every byte received after IAC SB to a persistent subnegotiation buffer until IAC SE is received, without applying a maximum length. Telnet option negotiation occurs before authentication, so an unauthenticated peer can omit the terminator and continuously grow the buffer. The resulting memory exhaustion can terminate the process and every service it hosts. This issue is fixed in version 26.10.0rc1.	7.5	More Details
CVE-2026-88784	The Translate WordPress with Weglot 2.0 WordPress plugin before 6.4.3 does not escape the URLs it re-inserts into a page when rewriting links for a translated language, and reverses the encoding that was applied to them beforehand, allowing unauthenticated users to store a payload which executes in the browser of anyone viewing the page in a translated language.	7.5	More Details
CVE-2026-115776	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, twisted.words.protocols.irc.IRC.dataReceived retains all bytes received without a line feed in its per-connection buffer and applies no maximum line length. The twisted.words.service.IRCUser server subclass inherits this behavior, and buffering occurs before PASS, NICK, or USER authentication commands. An unauthenticated connection can continuously grow the buffer until the shared server process runs out of memory and drops all users. This issue is fixed in version 26.10.0rc1.	7.5	More Details
CVE-2026-92450	The authentication framework used by affected Kiloview E3 NDI encoder firmware signs session JSON Web Tokens using an HMAC-SHA256 secret hardcoded as a static string within the firmware image, identical across all deployed devices. An attacker with access to the publicly distributed firmware can recover this secret and forge arbitrary session tokens, including administrative tokens, without valid credentials, enabling unauthenticated access to privileged device functions.	7.5	More Details
CVE-2026-110825	The ONVIF device service on affected Kiloview E3 NDI encoder firmware accepts hostname, DNS, and NTP configuration write operations without requiring authentication, regardless of the device's configured authentication setting for this service. A remote, unauthenticated attacker can alter the device hostname, redirect DNS resolution, or modify the configured time source, enabling follow on attacks such as traffic interception or certificate validation bypass.	7.5	More Details
CVE-2026-114605	dnspython is a DNS toolkit for Python. Prior to 2.9.0, dns.name.from_wire_parser follows attacker-controlled DNS compression pointer chains without a hop limit. A malicious DNS wire-format message can force excessive pointer traversal during name decompression, causing high CPU use, application degradation, or denial of service in applications that process untrusted DNS messages. This issue is fixed in version 2.9.0.	7.5	More Details
CVE-2026-102486	Out-of-bounds read in NTS extension field parsing in NTPsec ntpd 1.1.4 through 1.2.5 allows a remote attacker to crash the daemon via an NTS Authenticator and Encrypted Extension Fields (AEEF) field with an oversized nonce length.	7.5	More Details
CVE-	Improper Restriction of XML External Entity Reference vulnerability in Apache Camel Saxon component. This issue affects Apache Camel: from 4.0.0 before 4.18.5, from 4.19.0 before 4.22.2. XQueryBuilder in camel-saxon, which evaluates expressions for the xquery language and the xquery component, built its default Saxon Configuration with parse options that set only a whitespace stripping rule, leaving the XML parser at its defaults, under which a DOCTYPE declaration is accepted and external entities are resolved. A message body of type String or byte[], or any body read as an InputStream, is first turned into a javax.xml.transform.Source by Camel's own XML type converters, which configure the parser against external entities, so those bodies were not affected. A body that was already a Source, however, was handed to Saxon unchanged, and when that Source still had to be parsed, as a stream-based Source does, Saxon parsed it with that Configuration, bypassing the protection applied to the other body types. Such a body arises, for example, after a route converts the message with convertBodyTo(Source.class), or when a component or processor earlier in the route sets the body to a stream-based Source. A DOCTYPE with an external subset, or an external entity, declared in that body		More

2026-90885	was therefore resolved and expanded into the document the expression was evaluated against. An external sender whose XML reaches an XQuery expression or an xquery endpoint in that form could read files readable by the Camel process and cause the process to issue requests to hosts the document names. No non-default option is required, since the default Configuration is used whenever the route does not supply one of its own. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The 4.14.x releases stream and earlier, as well as the 4.15.x to 4.17.x and 4.19.x to 4.21.x release lines, reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. For deployments that cannot upgrade immediately, do not evaluate XQuery over a message body that is already a javax.xml.transform.Source when that body carries XML from an untrusted sender, or set the configuration option of the xquery endpoint or language to a Saxon Configuration whose parse options refuse a DOCTYPE declaration and disable external entities and external DTD loading, which is what the fixed releases now apply by default.	7.5	Details
CVE-2026-98965	In the Linux kernel, the following vulnerability has been resolved: ice: Restore Ordered MMIO Writes for Tx Doorbells The DQL accounting state which tracks the number of bytes queued for the NIC to transmit, namely dql->num_queued, is updated by the ICE driver when it invokes __netdev_tx_sent_queue(). Subsequently the driver updates the hardware queue tail allowing the NIC to begin processing the work queue. The queue accounting update should be ordered before the NIC begins transmitting descriptors. The transmit path currently updates the hardware doorbell using writel_relaxed(), which (on arm64) does not provide the same ordering guarantees between writes to normal memory and writes to MMIO registers that writel() does. This introduces a potential race where the NIC begins transmitting descriptors before the dql->num_queued update is globally visible. If the NIC finishes before the update is observed by dql_completed(), it detects an invalid state where more bytes have been completed than have been queued. When this happens the following BUG_ON is triggered. BUG_ON(count > num_queued - dql->num_completed); This has been observed and manifests as the following crash (note that the trace has been trimmed) in an environment with sustained network load that uses an Intel Corporation Ethernet Controller E810-XXV for SFP (rev 02) on an arm64 machine. Replacing writel_relaxed() with writel() in a test kernel eliminated the crash in the user's workload, which previously reproduced the issue reliably. kernel BUG at lib/dynamic_queue_limits.c:99 Internal error: Oops - BUG: 00000000f2000800 [#1] SMP pc : dql_completed+0x268/0x2a0 lr : ice_clean_tx_irq+0x1d4/0x620 [ice] Call trace: dql_completed+0x268/0x2a0 (P) ice_napi_poll+0x94/0x520 [ice] __napi_poll+0x48/0x3f0 net_rx_action+0x194/0x420 This restores the behaviour prior to commit ccde82e90946 ("ice: add E830 Earliest TxTime First Offload support"), which changed the notification mechanism from writel() to writel_relaxed().	7.5	More Details
CVE-2026-75508	GraphQL-core is a Python port of the GraphQL.js reference implementation. From 3.2.3 until 3.2.12 and 3.3.0rc1, the parser's max_tokens limit in Parser.advance_lexer counts only non-comment tokens returned after Lexer.lookahead skips comments. Each skipped comment is still allocated as a Token and retained through token.next and token.prev, allowing a highly compressible comment-padded GraphQL document to pass a low token limit while consuming memory. An unauthenticated remote attacker can exhaust worker memory before AST-based depth or complexity controls run. This issue is fixed in versions 3.2.12 and 3.3.0rc1.	7.5	More Details
CVE-2026-114106	HKUDS Vibe-Trading from 0.1.0 before 0.1.7 contains a path traversal vulnerability that allows remote attackers to read sensitive server files through LLM agent file-read tools. Attackers can prompt the agent to invoke read_document, which lacks sandbox checks, or safe_user_path-gated tools to leak /etc/shadow, /proc/self/environ API keys, and /root credentials.	7.5	More Details
CVE-2026-19741	Zephyr's OpenThread integration defined CONFIG_OPENTHREAD_JOINER_PSKD with default "J01NME" in modules/openthread/Kconfig.thread (previously subsys/net/l2/openthread/Kconfig, then subsys/net/l2/openthread/Kconfig.thread). J01NME is the example Joining Device Credential used throughout the public OpenThread and Thread commissioning documentation, including Zephyr's own samples/net/sockets/echo_client/README.rst, so it is not a secret. A firmware image built with CONFIG_OPENTHREAD_JOINER_AUTOSTART=y that never overrode the PSKd therefore shipped with a publicly known joining credential compiled in, with no build-time or run-time warning. The PSKd is the only value that authenticates Thread commissioning: joiner and commissioner run a DTLS-PSK handshake keyed on it, after which the commissioner supplies the full operational dataset. In ot_configure_instance() (modules/openthread/openthread.c), a device with no stored dataset calls otJoinerStart() with the configured PSKd, and ot_joiner_start_handler() calls otThreadSetEnabled() on success. An unauthenticated attacker within 802.15.4 radio range of an uncommissioned device (factory-fresh, after a factory reset, or after a settings erase) can run any stock Thread commissioner registered with J01NME, complete the handshake, and hand the device an attacker-controlled operational dataset. The device then joins and enables Thread on the attacker's mesh. The attacker becomes its router, seeing, dropping, injecting and spoofing all of the device's IPv6 traffic and impersonating the peers the application expects, while the device is cut off from the network it was	7.5	More Details

	meant to join and stays mis-provisioned until it is reset. Only application-layer security with independently provisioned credentials survives. Exposure is limited to builds that enable joiner autostart and leave the PSKd unset; the option is not enabled by default and no in-tree sample uses the vulnerable combination. The fix removes the default from CONFIG_OPENTHREAD_JOINER_PSKD and adds a BUILD_ASSERT() in modules/openthread/openthread.c, active only under CONFIG_OPENTHREAD_JOINER_AUTOSTART, that fails the build unless a 6-to-32-character PSKd is configured. Products should provision a per-device PSKd from factory data and call otJoinerStart() from the application rather than relying on a credential compiled identically into every unit.		
CVE-2026-113499	The LiteSpeed Cache WordPress plugin before 7.9.2 does not restrict one of its HTML rewriting routines to stylesheet content, so it removes a quote that terminates an attribute in the rendered page, and combined with unescaped comment author output this allows unauthenticated users whose comment is published to inject attributes that execute JavaScript in the browser of any visitor, including a logged-in administrator, who views front-end pages while the affected option is enabled.	7.5	More Details
CVE-2026-113500	The LiteSpeed Cache WordPress plugin before 7.9.2 does not limit one of its HTML optimization routines to stylesheet content, so it can delete the markup that closes an attribute in the rendered page, and combined with unescaped comment author output this allows unauthenticated users whose comments are published to inject event-handler attributes that run JavaScript in the browser of any visitor, including a logged-in administrator, viewing front-end pages while the optimization is enabled.	7.5	More Details
CVE-2026-99046	In the Linux kernel, the following vulnerability has been resolved: gve: fix XSK buffer leak on error descriptor When the error bit is set in the RX completion descriptor, the buf_state and its attached buffer should be freed. In the case of AF_XDP ZC, the XSK buffer was not freed, leading to a leak. This issue was caught by LLM.	7.5	More Details
CVE-2026-112359	Missing Authorization vulnerability in Amauri.IO WPMobile.App wpappninja allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WPMobile.App: from n/a through 11.86.	7.5	More Details
CVE-2026-56386	Integer underflow and out-of-bounds read in the HTTP/2 server DATA frame parser in Apache bRPC allows an unauthenticated remote attacker to cause denial of service by sending a malformed DATA frame with the PADDED flag set and a zero-length payload. Root Cause: The HTTP/2 DATA frame handler computed the fragment size from the attacker-controlled payload length and, when the PADDED flag was present, decremented the fragment size before ensuring that the frame contained the mandatory pad length byte. A zero-length padded DATA frame could therefore underflow the unsigned fragment size and read the pad length byte past the end of the frame buffer. Affected Scenarios: Apache bRPC servers that accept HTTP/2 traffic are affected when a remote client can send crafted HTTP/2 frames to the server. HTTP/2 may be reached through cleartext h2c detection on an HTTP port or through TLS ALPN negotiation. How to Fix: (Choose one of the following options) 1. Upgrade Apache bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3325 Note: The fix validates that a padded DATA frame has enough payload bytes for the pad length field before decrementing the fragment size or reading the pad length byte. Malformed zero-length padded DATA frames are rejected with HTTP/2 FRAME_SIZE_ERROR.	7.5	More Details
CVE-2026-109936	The Otter Blocks WordPress plugin before 3.2.8 does not correctly bound a pattern it uses to rewrite its dynamic-content tokens while rendering blocks, so the pattern can consume the quote that closes an HTML attribute and pull adjacent unescaped data into the markup as new attributes. This allows unauthenticated users to achieve stored Cross-Site Scripting that executes in the browser of any visitor who views an affected page, including a logged-in administrator.	7.5	More Details
CVE-2026-114102	adm-zip before 0.6.1 contains a denial of service vulnerability that allows unauthenticated attackers to crash the Node.js process via a malformed ZIP file, because inflateAsync in methods/inflater.js lacks a zlib stream error listener. Attackers can supply a ZIP with corrupted DEFLATE data processed by readFileAsync, readAsTextAsync, extractAllToAsync, or getDataAsync, triggering an uncatchable exception that terminates all in-flight requests.	7.5	More Details
	In the Linux kernel, the following vulnerability has been resolved: bnxt_en: fix DMA mapping length for padded small packets Stefan Fleischmann reported Intel IOMMU DMA Read faults on BCM57412 NetXtreme-E NICs when transmitting packets on VLAN/macvlan interfaces: DMAR: [DMA Read NO_PASID] Request device [18:00:0] fault addr 0xfc499000 [fault reason 0x06] PTE Read access is not set bnxt_en 0000:18:00:0 eno1np0: Abandoning msg {0xb4 0x41a} len: 0 due to firmware status: 0x2000001 ... NETDEV WATCHDOG: eno1np0 (bnxt_en): transmit queue 0 timed out The fault address (0xfc499000) is on an exact 4KB page boundary, pointing to a DMA read buffer overrun. In bnxt_start_xmit(), packets smaller than BNXT_MIN_PKT_SIZE (52 bytes), such as 42-byte untagged ARP frames, are padded: if (length < BNXT_MIN_PKT_SIZE) { pad = BNXT_MIN_PKT_SIZE - length; if (skb_pad(skb, pad)) goto tx_kick_pending; length = BNXT_MIN_PKT_SIZE; } mapping = dma_map_single(&pdev->dev, skb->data, len, DMA_TO_DEVICE); ... dma_unmap_len_set(tx_buf, len,		

CVE-2026-99081	len); However, 'len' was initialized earlier to skb_headlen(skb) (e.g. 42 bytes) and is left unadjusted after padding. Consequently, dma_map_single() and dma_unmap_len_set() map and track only 42 bytes. Later, the hardware TX buffer descriptor is programmed with the padded length: txbd->tx_bd_len_flags_type = cpu_to_le32(((len + pad) << TX_BD_LEN_SHIFT) flags TX_BD_FLAGS_PACKET_END); The NIC DMA engine is thus instructed to read 52 bytes from a region where only 42 bytes were DMA-mapped. If skb->data ends near the boundary of a 4KB page (within 'pad' bytes of the next page), the hardware DMA read overruns into the unmapped adjacent page, triggering an IOMMU fault. This issue was exposed after commit 447cbe95ebb9 ("vlan: fix skb_under_panic and races when toggling HW VLAN offload") because reserving extra VLAN headroom rounded LL_RESERVED_SPACE from 48 up to 64 bytes, shifting skb->data offsets and potentially causing small frames to land right against page boundaries. Fix this by using skb_put_padto(skb, BNXT_MIN_PKT_SIZE) in bnxt_start_xmit(), before skb_shinfo(skb)->nr_frags is sampled, because padding might linearize the skb. This ensures skb->len and skb_headlen(skb) reflect the padded size so that dma_map_single() maps the full buffer and the descriptor length is consistent. This also removes the temporary 'pad' variable and masking logic. The padding is done after the SW USO branch, otherwise bnxt_sw_udp_gso_xmit() would account the padding as UDP payload and send it in extra segments. Note that SW USO segments are still not padded: with IPv4, a segment carrying less than 10 bytes of UDP payload (small gso_size, or a short last segment) is sent as a frame shorter than BNXT_MIN_PKT_SIZE. This is a separate issue, left for a followup patch.	7.5	More Details
CVE-2026-110491	The ONVIF service on the E3 can be rendered unresponsive by sending malformed or incomplete requests to certain operations. The affected service stops accepting new connections while the rest of the device continues operating normally, and will recover automatically after a period. This can be triggered repeatedly by any unauthenticated party with network access to the device.	7.5	More Details
CVE-2021-48018	Vyper from 0.2.13 before 0.2.15 allocates storage slots for @nonreentrant reentrancy locks at the same locations as contract storage variables. When a function using @nonreentrant is called, setting the lock overwrites the overlapping storage variable, corrupting contract state and breaking the contract's intended logic.	7.4	More Details
CVE-2026-99077	In the Linux kernel, the following vulnerability has been resolved: ipv4: udp: Create exceptions before socket matching Currently, when ICMP Fragmentation Needed and Redirect Message packets are locally delivered and quote a UDP packet, a FIB nexthop exception (FNHE) is only created if the kernel can match the UDP packet to an existing socket. This behavior allows off-path attackers to conduct a side-channel attack on the FNHE cache in order to discover the ephemeral port used by a connected UDP socket. Commit 6457378fe796 ("ipv4: use siphash instead of Jenkins in fnhe_hashfun()") and commit 67d6d681e15b ("ipv4: make exception cache less predictable") tried to mitigate such attacks by making it harder for attackers to discover hash collisions in the FNHE cache and by randomizing the number of exceptions a hash bucket can hold, respectively. Unfortunately, both of the mitigations can be bypassed. Instead, mitigate such attacks by always creating a FNHE, even before trying to find a matching socket. Do that by calling ipv4_update_pmtu() and ipv4_redirect(), the helpers used when the quoted packet did not originate from a socket. This means that guesses (right or wrong) from an off-path attacker will always result in a FNHE being created or updated in the cache that the attacker can observe. Pass an oif of 0, in a similar fashion to icmp_err(). This is also the oif used by the socket path for sockets that are not bound to a device. Note that this does not allow attackers to create FNHEs that they could not create before, as both helpers can already be reached with little to no validation. For example, by sending an ICMP error that quotes an ICMP Echo Reply or one that quotes a UDP source port that matches a wildcard socket. Also note that in the good case (matched socket) the above scheme comes at the cost of an extra route lookup, as the no socket helpers perform their own lookup before the one performed by ipv4_sk_update_pmtu() / ipv4_sk_redirect(). When the two resolve to different nexthops, it also results in two exceptions being created for the same destination IP. One in the FNHE cache of the nexthop resolved by the no socket helpers and another in the FNHE cache of the nexthop used by the socket.	7.4	More Details
CVE-2026-114221	NT OneAPP developed by NewType Infortech has a Credential Disclosure vulnerability. Unauthenticated remote attackers can trigger an HTTP Basic Authentication challenge from an unvalidated source in the WebView under specific conditions, thereby obtaining the NUP Portal passcode of a logged-in user.	7.4	More Details
CVE-2026-89035	MikroTik RouterOS through 7.24.5 contains a session token replay vulnerability in the WebFig interface that allows attackers to bypass authentication by replaying captured session tokens across new ECDH sessions. The makeSessionKey function in libcrypto.so generates tokens using fixed salts without incorporating channel-specific or session-specific input, and the user service validates tokens against the login store without verifying channel context, enabling stolen tokens to grant persistent administrative access until the token store is explicitly cleared. This vulnerability extends the impact of CVE-2025-61481, which identified the underlying transport-layer exposure without a corresponding fix.	7.4	More Details
CVE-2026-	NT OneAPP developed by NewType Infortech has a Hard-coded Credentials vulnerability. Unauthenticated remote attackers can obtain the hard-coded cryptographic key and decrypt login	7.4	More Details

114219	request contents, thereby obtaining users' plaintext account credentials and device information.		
CVE-2026-114397	Trudesk through 1.2.11 contains a missing authentication vulnerability in debug routes enabled whenever NODE_ENV is unset, letting unauthenticated attackers invoke them remotely. Attackers can call /debug/populatedb to delete all groups, tags and tickets, then log in with seeded accounts using the password 'password'.	7.4	More Details
CVE-2026-114724	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 contains an object injection primitive in rcube_session::unserialize(). A session key prefixed with “!” was treated as having no value, and the reconstructed serialized string was emitted with a key but no corresponding value, desynchronizing the remainder of the input so that following session data was interpreted as a serialized value. This yields a PHP object injection primitive for an attacker able to influence stored session data. Upstream fixed this in the commit “Fix object injection primitive in rcube_session::unserialize() via “!”-prefixed session keys”, which rejects such keys.	7.4	More Details
CVE-2026-114615	Authlib is a Python library which builds OAuth and OpenID Connect servers. From 1.3.1 until 1.9.0, OctKey.import_key() checks asymmetric-key markers only at byte zero while Authlib's PEM loaders accept leading whitespace, comments, or a byte-order mark. In deployments that use the documented mixed algorithm registry and already configure verification with PEM or certificate bytes containing such a prefix, an unauthenticated attacker who knows those same public-key bytes can use them as an HMAC secret and forge HS256 tokens with arbitrary identity or authorization claims. This is a bypass of the remediation for CVE-2024-37568. The module-level jwt object reaches OctKey.import_key() in authlib/jose/rfc7518/oct_key.py, where POSSIBLE_UNSAFE_KEYS fails to identify prefixed asymmetric key material. This issue is fixed in version 1.9.0.	7.4	More Details
CVE-2026-116163	Parse Server before 8.6.102 and 9.0.0 before 9.10.5-alpha.2 contains an authentication bypass because the validateSetUp, validateLogin, and validateUpdate methods of custom adapters subclassing AuthAdapter are never called, and an inherited default that accepts any authData is used instead. Attackers can supply arbitrary authData to sign up, link identities, log in as existing users, or skip custom additional authentication factors.	7.4	More Details
CVE-2026-65824	OpenReception's appointment booking software provides an end-to-end encrypted appointment booking platform. Prior to version 2.2.1, an attacker who captures a pending invite token can replay it to obtain a passkey-registration credential and register their own passkey on the invited account, taking control of that account. For a staff or tenant-admin invite this yields access to the tenant context tied to that account. This issue has been patched in version 2.2.1.	7.4	More Details
CVE-2026-61791	OCO-Server is a self-hosted endpoint management application for computer inventory, software deployment, and mobile device management. Prior to 1.2.1, `Html::commandButton()` in `lib/Html.class.php` embeds an action URL containing a stored computer hostname into a JavaScript string in an inline `onclick` handler, using `htmlspecialchars(ENT_QUOTES)` only for HTML escaping. Browsers decode the HTML entities before interpreting the handler as JavaScript, allowing a crafted hostname to escape the string and execute script in an administrator's authenticated session when an affected command button is activated. `frontend/views/computer-details.php` passes `\$computer->hostname` to this helper using `Computer::DEFAULT_COMPUTER_COMMANDS` in `lib/Models/Computer.class.php`, whose VNC, RDP, SSH, and Ping entries use `new_tab=false` and the vulnerable `window.location` branch. A remote attacker who possesses the shared `agent-registration-key` can use the `oco.agent.hello` handler in `frontend/api-agent.php` to store an unchecked hostname when `agent-self-registration-enabled` is enabled; self-registration is disabled by default. The vulnerable default command buttons require no custom command configuration. The injected script can abuse the administrator's privileges, including permission to deploy software jobs to managed computers. This issue is fixed in version 1.2.1.	7.4	More Details
CVE-2026-98932	In the Linux kernel, the following vulnerability has been resolved: EDAC/altera: Fix use-after-free in error paths In both altr_edac_a10_device_add() and altr_portb_setup(), the error path freed the dci structure before releasing the devres group. Since the managed single and double bit IRQ handlers use altdev(dci->pvt_info) as their data, an IRQ firing between freeing dci and unregistering the IRQs could dereference the freed memory. Release the devres group first so the managed IRQs are unregistered before the dci structure is freed.	7.4	More Details
CVE-2026-115570	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to 42.11.16, 43.8.1, 44.9.1, and 45.0.0-beta.6, Electron does not block compromised eMudhra intermediate certificates in Chromium networking, Node.js tls, Node.js https, or Node.js fetch. An attacker who holds a certificate issued through one of these intermediates and can intercept network traffic can impersonate a TLS server to an affected application. This issue is fixed in versions 42.11.16, 43.8.1, 44.9.1, and 45.0.0-beta.6.	7.4	More Details
	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. Prior to version 26.05, Configuration.getWebApiToken returns the empty string when		

CVE-2026-55457	webapi.authorization.token is not configured. AuthorizationFilter compares the supplied token against this default without first checking that the configured token is non-empty. A request to any @AuthorizationBinding-protected endpoint with an empty ?token= value is therefore accepted. This issue has been patched in version 26.05.	7.4	More Details
CVE-2026-114137	Django REST framework is a toolkit for building Web APIs. From 3.12.0 until 3.18.2, TokenProxy.delete() in rest_framework/authtoken/models.py uses the proxy's user-id pk while the concrete Token token-string primary key identifies rows, so Django issues a deletion query that matches no row and reports zero deletions without raising an error. Applications, administrative integrations, or operational scripts that revoke credentials through a TokenProxy instance can therefore believe revocation succeeded while the token remains live and continues authenticating. DRF's own TokenAdmin.delete_model workaround is not affected, but other callers can leave access available indefinitely because DRF tokens do not expire automatically. This issue is fixed in version 3.18.2.	7.4	More Details
CVE-2026-115723	Snap7 through 1.4.3 contains an out-of-bounds write vulnerability in the passive partner that allows a malicious active partner to overflow application buffers by declaring an unvalidated BSend total length. Attackers connecting from the configured peer address can declare up to 65535 bytes and send a short fragment, causing BRecv to memcpy past the buffer and crash or corrupt SCADA, HMI or gateway processes.	7.4	More Details
CVE-2026-111079	A flaw was found in Goose. When the terminal integration feature is enabled, the application captures shell command history and sends it verbatim without sanitization to configured Large Language Model (LLM) provider endpoints during queries. If a user enters sensitive data—such as API tokens, passwords, or environment variables containing secrets—in the terminal, this unredacted information is transmitted across trust boundaries. This can lead to sensitive information disclosure, exposing confidential credentials to external service providers and logs.	7.4	More Details
CVE-2026-110975	A flaw was found in Kruize Autotune. Its REST API ships without an authentication layer, and the endpoints that return datasource configuration — the list-datasources endpoint (GET /datasources) and the datasource-metadata endpoint — serialized datasource objects using a default Gson configuration with no field-exclusion strategy. As a result, the authentication secrets stored for each configured monitoring datasource (basic-auth passwords, bearer tokens, API keys, and OAuth2 client secrets) were returned in cleartext to any client able to reach the service. An attacker with network access to Autotune could harvest these credentials and reuse them to authenticate to the backing monitoring stack.	7.4	More Details
CVE-2026-114616	Authlib is a Python library which builds OAuth and OpenID Connect servers. Prior to 1.9.1, Authlib's OctKey.import_key() and HMAC signing and verification path accept zero-length key material. If an application's secret lookup can return an empty string or empty bytes, a remote attacker can compute the deterministic HS256, HS384, or HS512 signature for arbitrary token claims and have Authlib treat the token as authentic. This issue is fixed in version 1.9.1.	7.4	More Details
CVE-2026-55459	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 25.01 to before version 26.05, AuthenticationEndpoint.verifyOpenIdToken builds its verifier as JWT.require(algorithm).withIssuer(issuer).build() and does not enforce the aud claim during cryptographic validation. The aud check is performed manually afterwards via a single string-equality call, which only handles single-string aud values. In a multi-RP IdP setup, a token whose single-string aud matches the viewer's client ID but was issued for a different relying party passes the check and authenticates the attacker as the token's sub. This issue has been patched in version 26.05.	7.4	More Details
CVE-2026-115999	A vulnerability has been found in mansi-k Restaurant_IMS d7ad49f61ab7cc1ec6eb359ac2b5378bdbde0694. Affected is the function mysqli_connect of the file lsmodal.php. The manipulation of the argument login_id/epsw leads to sql injection. Remote exploitation of the attack is possible. This product follows a rolling release approach for continuous delivery, so version details for affected or updated releases are not provided. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-116000	A security flaw has been discovered in Zylon PrivateGPT 1.0.0/1.0.1. This vulnerability affects the function load_file_from_uri of the file private_gpt/server/ingest/uri_loader.py of the component URI Artifact Loader. Performing a manipulation of the argument Value results in server-side request forgery. It is possible to initiate the attack remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-114256	A weakness has been identified in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. Affected is the function save_folder/delete_folder/delete_file/file_rename/save_files of the file admin_class.php of the component Folder Management. This manipulation of the argument name/parent_id/id/type/folder_id/description causes sql injection. The attack is possible to be carried out remotely. The exploit has been made available to the public and could be used for attacks. This	7.3	More Details

	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.		
CVE-2026-116050	A vulnerability has been found in krupalshah Employee Recruitment System up to 922ff6d208278282e7c8e36d70d9ae356adec9fe. This affects an unknown part of the file /source/registerform3.php of the component Registration. The manipulation of the argument pemail leads to sql injection. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used. Continuous delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-116094	A vulnerability was determined in EyouCMS up to 1.8.3. This impacts the function remote_file_to_local of the file application/function.php. Executing a manipulation can lead to server-side request forgery. It is possible to launch the attack remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-116260	A weakness has been identified in libimobiledevice up to 1.4.0. This issue affects the function service_receive of the file src/afc.c of the component AFC Service. Executing a manipulation can lead to double free. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-116062	A security flaw has been discovered in SourceCodester Performance Indicator System 1.0. This issue affects some unknown processing of the file /opils/staff/save_profile.php. The manipulation of the argument Contact results in sql injection. The attack may be launched remotely. The exploit has been released to the public and may be used for attacks.	7.3	More Details
CVE-2026-116049	A flaw has been found in krupalshah EmployeeRecruitmentSystem up to 922ff6d208278282e7c8e36d70d9ae356adec9fe. Affected by this issue is the function mysql_query of the file /source/registerform2.php of the component Registration Wizard. Executing a manipulation of the argument university/institute/branch/degree/status/cpi/semester/experience can lead to sql injection. It is possible to launch the attack remotely. The exploit has been published and may be used. This product takes the approach of rolling releases to provide continuous delivery. Therefore, version details for affected and updated releases are not available. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-116061	A vulnerability was identified in SourceCodester Performance Indicator System 1.0. This vulnerability affects unknown code of the file /opils/staff/edit_dealer.php. The manipulation of the argument Contact leads to sql injection. The attack may be initiated remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2026-116060	A vulnerability was determined in SourceCodester Performance Indicator System 1.0. This affects an unknown part of the file /opils/staff/adddealer.php. Executing a manipulation of the argument Contact can lead to sql injection. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2026-114439	A vulnerability was determined in oretnom23 php-travel-agency-system. Affected by this issue is some unknown functionality of the file admin/session.php of the component Session Handler. Executing a manipulation of the argument user_id can lead to missing authentication. It is possible to launch the attack remotely. The exploit has been publicly disclosed and may be utilized. This product takes the approach of rolling releases to provide continuous delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-114242	A vulnerability was identified in dhikshitha29 Insurance-Management-System 91de5edda5baa109ad8e9baf03ce6a3ceda84183/125362cc21dca20d74091a4913a5096bb58dd837. Affected is the function mysqli_query of the file /Admin/View/editPlan.php. Such manipulation of the argument Plan_no/Name leads to sql injection. The attack can be executed remotely. The exploit is publicly available and might 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.	7.3	More Details
	A vulnerability was detected in krupalshah EmployeeRecruitmentSystem up to 922ff6d208278282e7c8e36d70d9ae356adec9fe. Affected by this vulnerability is the function mysql_query of the file /source/registerform1.php of the component Registration. Performing a manipulation of the argument		

CVE-2026-116048	firstname/middlename/lastname/gender/birthdate/state/statespecify/city/cityspecify results in sql injection. It is possible to initiate the attack remotely. The exploit is now public and may be used. This product is using a rolling release to provide continuous delivery. Therefore, no version details for affected nor updated releases are available. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-115970	A vulnerability has been found in Travisun Opptrix 0.4.0. This impacts an unknown function of the file /api/system-update/import of the component Import Endpoint. Such manipulation leads to code injection. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used. The project has been completely taken offline and deprecated. The project maintainer explains: "The mentioned security risk was a design flaw that manifested primarily in local environments during earlier stages of version development. We previously addressed and resolved this specific risk; however, because the project was restricted to a small learning scope and the main versions were intended to run exclusively on localhost, the security fix was only pushed directly to our GitHub repository before it went private." This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-116006	A vulnerability was detected in YunaiV/zhijiantianya ruoyi-vue-pro up to 2026.09. The affected element is an unknown function of the file AuthController.java of the component SMS Code Verification. The manipulation of the argument mobile/code results in improper authentication. The attack can be launched 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.	7.3	More Details
CVE-2026-114267	A vulnerability was detected in SourceCodester Food Order Management System 1.0. Affected is an unknown function of the component Admin Endpoint. The manipulation results in missing authorization. The attack may be performed from remote. The exploit is now public and may be used.	7.3	More Details
CVE-2026-115969	A flaw has been found in Travisun Opptrix 0.4.0. This affects an unknown function of the file /api/auth/setup of the component Schedule API. This manipulation causes command injection. It is possible to initiate the attack remotely. The exploit has been published and may be used. The project has been completely taken offline and deprecated. The project maintainer explains: "The mentioned security risk was a design flaw that manifested primarily in local environments during earlier stages of version development. We previously addressed and resolved this specific risk; however, because the project was restricted to a small learning scope and the main versions were intended to run exclusively on localhost, the security fix was only pushed directly to our GitHub repository before it went private." This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2026-115973	A vulnerability was identified in FaveoSuite Faveo Helpdesk 2.0.3. Affected by this issue is the function ForgotPasswordController::sendResetLinkEmail/PasswordController::reset of the file app/Http/Controllers/Auth/ForgotPasswordController.php. The manipulation leads to weak password recovery. The attack may be initiated remotely. The vendor was contacted early about this disclosure.	7.3	More Details
CVE-2026-114255	A security flaw has been discovered in orenom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. This impacts the function save_user of the file admin_class.php of the component User Account Management. The manipulation of the argument name/username/password/type/id results in sql injection. The attack can be executed remotely. The exploit has been released to the public and may be used for attacks. 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.	7.3	More Details
CVE-2026-98952	In the Linux kernel, the following vulnerability has been resolved: pfcpc: make sure the SEID is linear before reading it pfcpc_encap_rcv() only makes sure that the UDP header and the 4 byte PFCP header are in the linear area. When the S flag is set, pfcpc_session_rcv() then reads the 8 byte SEID that follows, which is not covered by the pskb_may_pull() check. A short PFCP packet with the S flag set, or one whose session header lies in a fragment, therefore makes pfcpc_session_rcv() read beyond the end of the packet data, and whatever it finds there is stored in the tunnel metadata that flower later classifies on. Pull up to the end of the SEID before reading it, and drop packets that are too short to contain it. Reload the header pointer afterwards since pskb_may_pull() may reallocate the skb head. Use offsetofend() rather than sizeof(struct pfcphdr_session): the structure is not packed, so its size is 16 bytes because of the alignment of the __be64 member, while the header on the wire is only 12 bytes, and valid session messages would be dropped.	7.3	More Details
CVE-2026-	A vulnerability was determined in file-management-system dc9ce1bf27f94b23d906cb1e86ce374dc94ca5df. Impacted is the function store of the file app/Http/Controllers/FileController.php. This manipulation of the argument File causes unrestricted upload. The attack may be initiated remotely. The exploit has been publicly disclosed and may be	7.3	More

114262	utilized. This product uses a rolling release model to deliver continuous updates. As a result, specific version information for affected or updated releases is not available. The project was informed of the problem early through an issue report but has not responded yet.		Details
CVE-2026-114257	A security vulnerability has been detected in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. Affected by this vulnerability is the function readfile of the file download.php of the component File Download. Such manipulation of the argument ID leads to sql injection. The attack may be performed from remote. The exploit has been disclosed publicly and may be used. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-114254	A vulnerability was identified in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. This affects the function Login of the file admin_class.php. The manipulation of the argument username/password leads to sql injection. Remote exploitation of the attack is possible. The exploit is publicly available and might be used. 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.	7.3	More Details
CVE-2026-114258	A vulnerability was detected in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. Affected by this issue is some unknown functionality of the file files.php of the component File Management. Performing a manipulation of the argument fid/id results in sql injection. It is possible to initiate the attack remotely. The exploit is now public and may be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-114259	A flaw has been found in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. This affects the function save_files of the file admin_class.php. Executing a manipulation of the argument Upload can lead to unrestricted upload. It is possible to launch the attack remotely. The exploit has been published and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116368	A vulnerability was determined in DeepSeek Harness up to 0.2.0-rc.2. This issue affects some unknown processing of the file packages/client/connection/src/browser-auth.ts of the component Browser Auth. This manipulation causes improper authorization. The attack can be initiated remotely. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-114261	A vulnerability was found in MomoTommy file-management-system dc9ce1bf27f94b23d906cb1e86ce374dc94ca5df. This issue affects the function index/store/update/destroy of the file app/Http/Controllers/FileController.php of the component Files API. The manipulation of the argument Request results in missing authentication. The attack can be launched remotely. The exploit has been made public and could be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-115323	EverShop 1.2.2 before 2.2.1 contains a missing authorization vulnerability in the customer address API that allows unauthenticated attackers to manage other customers' addresses. Attackers who know customer and address UUIDs can send POST, PATCH, or DELETE requests to /api/customers/:customer_id/addresses routes to add, modify, or delete saved shipping addresses.	7.3	More Details
CVE-2026-114249	A vulnerability was detected in dhikshitha29 Insurance-Management-System 91de5edda5baa109ad8e9baf03ce6a3ceda84183/125362cc21dca20d74091a4913a5096bb58dd837. This vulnerability affects unknown code of the file /Agent/includes/UpdatePolicyHolder-inc.php of the component Policy Holder Update Form. The manipulation of the argument Policy_no results in improper authorization. It is possible to launch the attack remotely. The exploit is now public and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116339	A vulnerability was determined in iFlytek astron-agent up to 1.1.2. This affects an unknown function of the file core/memory/database/api/v1/exec_ddl.py of the component DDL Whitelist Handler. This manipulation causes incomplete blacklist. The attack may be initiated remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details

CVE-2026-114577	A vulnerability was found in oretnom23 php-travel-agency-system. This vulnerability affects unknown code of the file admin/insert.php of the component Insertion Handler. The manipulation of the argument costomer_id/payment_type/amount/description/user/uploadimage/full_name/visa_no results in sql injection. The attack may be launched remotely. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116253	A vulnerability was determined in DedeCMS 5.7.118. This vulnerability affects the function common.inc.php::_RunMagicQuotes of the file plus/ad_js.php of the component Advertisement Rendering. Executing a manipulation of the argument aid can lead to code injection. The attack can be executed remotely.	7.3	More Details
CVE-2026-114245	A weakness has been identified in dhikshitha29 Insurance-Management-System 91de5edda5baa109ad8e9baf03ce6a3ceda84183/125362cc21dca20d74091a4913a5096bb58dd837. Affected by this issue is some unknown functionality of the file /Admin/includes/Admin-Register-inc.php of the component Add New Admin Form. Executing a manipulation of the argument submit/Admin_id/password can lead to improper authorization. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-114243	A security flaw has been discovered in dhikshitha29 Insurance-Management-System 91de5edda5baa109ad8e9baf03ce6a3ceda84183/125362cc21dca20d74091a4913a5096bb58dd837. Affected by this vulnerability is the function mysqli_query of the file /Agent/ManagePolicy/EditPolicyHolder.php. Performing a manipulation of the argument Policy_no/Name results in sql injection. The attack is possible to be carried out remotely. The exploit has been released to the public and may be used for attacks. This product adopts a rolling release strategy to maintain continuous delivery. Therefore, version details for affected or updated releases cannot be specified. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116117	A flaw has been found in halo-dev halo up to 2.27.0-beta.1. Affected by this vulnerability is the function run.halo.app.theme.DefaultViewNameResolver.resolveViewNameOrDefault of the file application/src/main/java/run/halo/app/theme/DefaultViewNameResolver.java of the component View Name Resolver. Executing a manipulation of the argument template can lead to path traversal. The attack may be launched remotely. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-114263	A vulnerability was identified in ProjectWorlds Advanced Library Management System 1.0. The affected element is the function mysqli_query of the file individual_report_search.php. Such manipulation of the argument roll_number leads to sql injection. The attack may be launched remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2026-116110	A vulnerability was found in GiantappMan LiveWallPaper 2.6.8.0 on Windows. This issue affects some unknown processing of the file /LiveWallpaper/UploadFile of the component SignalR Hub. Performing a manipulation results in missing authentication. The attack is possible to be carried out remotely. The exploit has been made public and could be used. The patch is named b62f878899adbfc2968e3345442f3a3df91b774d. It is suggested to install a patch to address this issue. The vendor was contacted early about this disclosure.	7.3	More Details
CVE-2026-114440	A vulnerability was identified in oretnom23 php-travel-agency-system. This affects an unknown part of the file admin/delete.php. The manipulation of the argument costomer_id/ticket_id/payment_id/id/document_id/salary_id/emp_id/user_id leads to sql injection. The attack can be initiated remotely. The exploit is publicly available and might be used. Continious delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116257	A vulnerability was determined in BillaBear up to 2025.01.03. Affected by this issue is the function processInstall of the file src/BillaBear/Controller/InstallController.php of the component Installation Handler. This manipulation causes missing authentication. Remote exploitation of the attack is possible. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-116106	A security vulnerability has been detected in HeyCyan App 1.0.103 on Android. Affected by this vulnerability is an unknown functionality of the component FileProvider. The manipulation leads to improper access controls. The attack may be initiated remotely. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
	A weakness has been identified in oretnom23 php-travel-agency-system. This issue affects some		

CVE-2026-114441	unknown processing of the file admin/update.php. This manipulation of the argument action/full_name/visa_no/hiden/title/description/page_id/user causes sql injection. The attack may be initiated remotely. The exploit has been made available to the public and could be used for attacks. This product uses a rolling release model to deliver continuous updates. As a result, specific version information for affected or updated releases is not available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2026-116112	A vulnerability was identified in HashenUdara edoc-doctor-appointment-system up to cc258425b915d1af9e298777ccd26d160493e68d. The affected element is an unknown function of the file admin/add-new.php. The manipulation of the argument spec leads to sql injection. It is possible to initiate the attack remotely. The exploit is publicly available and might be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-116256	A vulnerability was found in geesondog RhaPHP up to 7ae639ccf8dcb627faace096d030bd4cb4cfb905. Affected by this vulnerability is the function image of the file application/mp/controller/Show.php of the component Mp Module Media Proxy. The manipulation of the argument url results in server-side request forgery. The attack may be launched remotely. The exploit has been made public and could be used. This product operates on a rolling release basis, ensuring continuous delivery. Consequently, there are no version details for either affected or updated releases. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2026-111982	The User Profile Builder WordPress plugin before 4.0.5 does not verify that a user is authorized to edit the specific account selected on its front-end Edit Profile form, allowing an Administrator of a Multisite subsite to change any other user's account email, including a Network Super Admin's, and take over that account through the password-reset flow.	7.2	More Details
CVE-2026-95706	Command injection vulnerabilities in the affected interface of AOS-8 could allow an authenticated remote attacker with high privileges to perform command injection. Successful exploitation could allow an attacker to execute commands with the privileges of the impacted mechanism on the underlying operating system.	7.2	More Details
CVE-2026-116219	Dolibarr ERP CRM before 24.0.0 contains an OS command injection vulnerability in htdocs/admin/tools/export.php that lets authenticated administrators bypass the restrict_os_commands allowlist. Attackers can submit a mysqldump value chaining commands with & characters around an allowlisted basename to execute operating system commands as the web server user.	7.2	More Details
CVE-2026-95707	A buffer overflow vulnerability in the PAPI protocol of AOS-8 could allow an authenticated remote attacker with high privileges to achieve remote code execution. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	7.2	More Details
CVE-2026-76307	In Splunk SOAR versions below 8.7.0, a user that holds a role with the Users and Roles view permission could use the Representational State Transfer (REST) API to infer authentication token values for other users. If the inferred token grants administrative access, that user could gain administrative access to the Splunk SOAR instance. The vulnerability is possible because the REST API accepts caller-controlled query filters on sensitive user record fields. For more information see Using the REST API reference for Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/7.1.0/using-the-splunk-soar-rest-api/using-the-rest-api-reference-for-splunk-soar-on-premises), Query for Data (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/7.1.0/using-the-splunk-soar-rest-api/query-for-data), and Manage roles and permissions in Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/6.4.0/manage-your-splunk-soar-on-premises-users-and-accounts/manage-roles-and-permissions-in-splunk-soar-on-premises) in the Splunk documentation.	7.2	More Details
CVE-2026-104068	JetBackup plugin for WordPress before 3.1.24.10 contains a symlink-based arbitrary file write vulnerability that allows an authenticated WordPress administrator to write attacker-controlled file contents outside the intended restore destination during backup import and restore. Attackers can craft a backup archive containing a symbolic link that archive extraction follows without validating targets or enforcing destination containment, potentially leading to server-side code execution depending on the writable destination and server configuration.	7.2	More Details
CVE-2026-114704	Deserialization of Untrusted Data vulnerability in Christoph Vielgrader Relevanssi relevanssi allows Object Injection.This issue affects Relevanssi : from n/a through 4.28.4.	7.2	More Details
CVE-2026-	The Duplicator WordPress plugin before 5.0.0 does not restrict its one-click upgrade routine to users allowed to install Duplicator WordPress plugin before 5.0.0, and additionally exposes the routine without authentication, allowing an administrator of a multisite child site who is not a network administrator to	7.2	More Details

115251	install and run arbitrary PHP in the network-wide shared Duplicator WordPress plugin before 5.0.0 directory.		
CVE-2026-116236	A vulnerability has been found in Oray Pugongying X5 Pro 7.6.1. Affected by this issue is the function plugin_init of the file /tmp/upload_files of the component PPPoE WAN Setting Handler. Such manipulation of the argument pppd_options leads to unrestricted upload. The attack may be launched remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	7.2	More Details
CVE-2026-114622	baserCMS is a website development framework. In versions 5.2.3 before 5.3.2 and 5.4.0 before 5.4.2, the optional BcThemeFile plugin's Admin API validates a raw path in ThemeFilesService::update() and getFullpath(), but ThemeFileForm::_execute() subsequently applies rawurldecode() to fullpath before writing. Because normalizePath() resolves only literal dot segments, percent-encoded traversal can pass the containment checks and become a path outside the theme directory at the write sink. An authenticated administrator can write PHP files into writable locations such as webroot and execute them, exceeding the intended theme-template editing boundary. The source also identifies the same check-before-decode ordering in ThemeFilesService::copy(), although that path was not dynamically verified. Exploitation requires the non-default plugin, an administrator session, and a same-origin Origin header; USE_CORE_ADMIN_API=false does not block session-authenticated requests. The Admin web form is not affected by this encoded-path tampering because FormProtection protects fullpath. This is a third bypass of the protections associated with CVE-2026-30940, using encoded traversal decoded after validation. This issue has been patched in versions 5.3.2 and 5.4.2.	7.2	More Details
CVE-2026-29001	TastyIgniter through 4.4.5 contain a PHP object injection vulnerability that allows authenticated administrators to achieve arbitrary code execution by submitting a serialized PHP object as a settings value, which is stored unsanitized and later unserialized by getValueAttribute() when the setting is read. Attackers can submit a serialized gadget chain, such as one generated via phpggc, through the admin settings form (for example the maintenance_message setting loaded by CheckMaintenance.php) to achieve arbitrary operating-system command execution as the PHP application process user, with the impact persisting until the value is overwritten.	7.2	More Details
CVE-2026-114624	baserCMS is a website development framework. In versions 5.0.0 before 5.3.2 and 5.4.0 before 5.4.2, the widget-area entity serializes widgets only when the submitted value is an array, allowing a string submitted through the admin JSON API to be stored unchanged. Entity reads, including front-end rendering and API responses, deserialize the stored value with unserialize() without an allowed_classes restriction, accepting both raw and base64-encoded serialized data. An administrator, or an attacker who obtains widget-write capability through a separate vulnerability such as chained XSS, can plant an object graph that subsequently triggers on unauthenticated page requests. Two demonstrated gadget chains in the default dependency set use composer/composer ProcessExecutor::startJob() with Process::fromShellCommandline(), or logiscope/mcp-sdk-php StdioTransport::connect() with StdioServerParameters, to reach proc_open() and execute commands as the web-server user. Deserialization bypasses constructor-based readonly safeguards, and the stored payload can execute repeatedly until the database row is repaired. Unauthenticated page access is a trigger after privileged payload storage, not an unauthenticated means of initially planting the payload. This issue has been patched in versions 5.3.2 and 5.4.2.	7.2	More Details
CVE-2026-113673	The GeoDirectory WordPress plugin before 2.8.194 does not properly escape administrator-defined text before writing it into a generated PHP file in its own folder, allowing high privilege users such as a subsite administrator on a multisite network, who cannot otherwise run code, to achieve Remote Code Execution that any unauthenticated visitor can trigger.	7.2	More Details
CVE-2026-95709	Command injection vulnerabilities exist in the affected interface of AOS-10 that could allow an authenticated remote attacker with high privileges to perform command injection. Successful exploitation could allow an attacker to execute arbitrary commands as a privileged user on the underlying operating system.	7.2	More Details
CVE-2026-76293	In Splunk Enterprise versions below 10.4.3, 10.2.8, 10.0.11, and 9.4.16, a user that holds a role with the edit_manager_xml capability could execute arbitrary operating-system commands using the permissions of the user account running Splunk Enterprise by opening the affected configuration in the Splunk Web Manager page. The vulnerability is possible because Splunk Web evaluates user-controlled expressions from Manager XML without sufficiently restricting the expression syntax.	7.2	More Details
CVE-2026-39902	Cacti is an open source performance and fault management framework. Prior to version 1.2.31, the prior blacklist ([;& `\$\\n\r]) omitted single-quote, double-quote, redirect operators, subshell parens and braces. An authenticated admin could save an input_string` containing single-quote, redirect operator, or subshell parens, and pass the check. This issue has been patched in version 1.2.31.	7.2	More Details
CVE-	Roundcube Webmail 1.6.14 before 1.6.20 and 1.7.0 before 1.7.5 is affected by a server-side request forgery filter bypass. The network blacklist used by rcube_utils::is_local_url() omitted the RFC 6052		

2026-114730	NAT64 prefix 64:ff9b::/96, the RFC 8215 local-use NAT64 prefix 64:ff9b:1::/48 and the unspecified address ::/128, so URLs such as http://[64:ff9b::7f00:1]/ passed the filter and reached loopback, link-local and cloud metadata addresses. Upstream fixed this in the commit “Fix SSRF filter (is_local_url()) bypass via NAT64 64:ff9b::/96”.	7.2	More Details
CVE-2026-109927	The HIPAA Forms WordPress plugin before 3.3.2 does not perform a capability check when handling a request that rebuilds its own custom user role, allowing a user who already holds that role to grant it administrator-level capabilities and escalate their privileges.	7.2	More Details
CVE-2026-109198	The Anti-Spam by CleanTalk – Spam Protection Without CAPTCHA plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'user_email' parameter in all versions up to, and including, 6.89.1 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 user registration to be open on the site, and script execution is triggered when an Administrator subsequently runs the plugin's user-spam scan and views the resulting Non-checkable users or Spam users tab — the plugin's own documented workflow.	7.2	More Details
CVE-2026-95705	Command injection vulnerabilities in the affected interface of AOS-8 could allow an authenticated remote attacker with high privileges to perform command injection. Successful exploitation could allow an attacker to execute commands with the privileges of the impacted mechanism on the underlying operating system.	7.2	More Details
CVE-2026-95708	Command injection vulnerabilities exist in the affected interface of AOS-10 that could allow an authenticated remote attacker with high privileges to perform command injection. Successful exploitation could allow an attacker to execute arbitrary commands as a privileged user on the underlying operating system.	7.2	More Details
CVE-2026-110744	A flaw was found in multicloud-global-hub-operator-bundle. The operator is configured with overly broad role-based access control (RBAC) permissions that grant cluster-wide impersonation of arbitrary users, groups, and service accounts without restriction. An attacker who compromises the operator or obtains its service account credentials can exploit this flaw to achieve privilege escalation, gaining full administrative control over the cluster.	7.2	More Details
CVE-2026-95704	A vulnerability in the web-based management interface of wireless network management systems could allow an authenticated remote attacker to execute remote code. Successful exploitation could enable an attacker to execute arbitrary commands on the underlying operating system, potentially compromising system integrity and data security.	7.2	More Details
CVE-2026-57114	PraisonAI is a multi-agent teams system. From version 4.5.126 to before version 4.6.59, PraisonAI's Async Jobs API validates webhook_url when a job request is parsed and again when the internal Job object is constructed. That validation blocks direct loopback/private targets, but it is not bound to the later network request. When a job completes, _send_webhook() passes the original hostname to httpx.AsyncClient.post() with no send-time validation, IP pinning, or guarded transport. An attacker-controlled hostname can therefore resolve to a public IP during Pydantic validation and later resolve to loopback/private/cloud-metadata infrastructure during webhook delivery. This bypasses the intended SSRF guard in current supported releases. This appears to be an incomplete fix / patch bypass for GHSA-8frj-8q3m-xhgm. This issue has been patched in version 4.6.59.	7.2	More Details
CVE-2026-109987	The PixelYourSite WordPress plugin before 11.4.3, PixelYourSite Pro WordPress plugin before 12.6.8 do not sanitise and escape a cookie value before printing it into the page output, which, on a site running a full-page cache, allows unauthenticated attackers to store a Cross-Site Scripting payload in the cached page that is then served to other visitors.	7.2	More Details
CVE-2026-79956	A vulnerability in the CLI command handling of the AOS-8 web interface could allow an authenticated remote attacker to execute arbitrary commands on the underlying operating system. Successful exploitation could allow an attacker to run unauthorized commands on the system, potentially leading to system compromise.	7.2	More Details
CVE-2026-114090	Trigger.dev cli-v3 from 3.0.0 before 4.5.4 contains a missing authentication vulnerability in the trigger dev --mcp server, which binds all network interfaces and exposes the /sse and /messages endpoints. LAN attackers or malicious web pages can send MCP tools/call requests to invoke trigger-task, cancel-run, list-runs, and get-run-logs using the developer's stored access token.	7.1	More Details
CVE-2026-111948	The User Profile Builder WordPress plugin before 4.0.1 does not escape the current page URL before outputting it in a JavaScript context on one of its admin pages, leading to Reflected Cross-Site Scripting against high-privilege users.	7.1	More Details
CVE-2026-112863	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Hogash KALLYAS kallyas allows Reflected XSS.This issue affects KALLYAS: from n/a through 4.25.0.	7.1	More Details

CVE-2026-109147	The WP Booking System WordPress plugin before 2.1.0.3 does not validate and escape one of its parameters before reflecting it back in an admin page, leading to Reflected Cross-Site Scripting which could be used against high privilege users such as administrators.	7.1	More Details
CVE-2026-112838	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in e4jvikwp VikRentCar vikrentcar allows Stored XSS.This issue affects VikRentCar: from n/a through 1.4.8.	7.1	More Details
CVE-2026-112849	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in AivahThemes Hostme v2 hostmev2 allows DOM-Based XSS.This issue affects Hostme v2: from n/a through 7.0.	7.1	More Details
CVE-2026-112835	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in e4jvikwp VikRentCar vikrentcar allows Reflected XSS.This issue affects VikRentCar: from n/a through 1.4.8.	7.1	More Details
CVE-2026-79627	The Advanced File Manager WordPress plugin before 6.0.0 does not correctly determine whether a stored file is an SVG before sanitising it, allowing users granted file-manager access by an administrator to store a file containing script that is not stripped and that runs in the browser of anyone who opens it, including an administrator.	7.1	More Details
CVE-2026-114664	Frappe Framework before 16.33.0 contains a blind SQL injection vulnerability in the frappe.desk.doctype.event.event.get_events method, which fails to properly neutralize user-supplied input before including it in SQL queries. Authenticated attackers can send crafted input to get_events and infer query results from application behavior to extract sensitive database information.	7.1	More Details
CVE-2026-99011	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: SCO: serialise sco_conn lifetime against sco_recv_scodata() sco_recv_scodata() upgrades the weak hcon->sco_data back-pointer to a strong reference under hdev->lock: hci_dev_lock(hdev); hcon = hci_conn_hash_lookup_handle(hdev, handle); ... conn = sco_conn_hold_unless_zero(hcon->sco_data); hci_dev_unlock(hdev); but the pointer is cleared from the other side without that lock. When the last sco_conn reference is dropped, sco_conn_free() ran conn->hcon->sco_data = NULL; with no hdev->lock held, so the RX path could read hcon->sco_data and call kref_get_unless_zero() on an sco_conn that was concurrently freed: BUG: KASAN: slab-use-after-free in sco_conn_hold_unless_zero+0xbe/0x160 Write of size 4 by task kworker/u17:0 Workqueue: hci0 hci_rx_work Call Trace: sco_conn_hold_unless_zero+0xbe/0x160 sco_recv_scodata+0x13f/0x490 hci_rx_work+0x3af/0x730 kref_get_unless_zero() only guards against a zero refcount, not against the backing memory already being freed. Give hcon->sco_data an actual reference on the sco_conn it points to, so the object cannot be freed while the pointer is still installed, and only clear and drop it from sco_conn_del(), which runs under hdev->lock (its callers, sco_connect_cfm() and sco_disconn_cfm(), hold it). The reader in sco_recv_scodata() also takes hdev->lock, so the store and the read are now serialised: the RX path either observes NULL or a reference that is guaranteed to stay valid until it drops its own. sco_conn_add() no longer hands its kref_init() reference to the caller as the connection's only reference; that initial reference is the one owned by hcon->sco_data, and callers get their own via sco_conn_hold(). Because the sco_conn can now outlive the socket (the association keeps it alive until the link goes down), the hci_conn can no longer be dropped from sco_conn_free() without regressing socket close: closing a connected SCO socket must still tear the link down. Move hci_conn ownership to the socket instead: __sco_chan_add() takes an hci_conn reference and it is released from sco_chan_del() and sco_sock_destruct(), exactly once. The sco_conn no longer owns an hci_conn reference, so sco_conn_add() stops consuming one and sco_connect() drops the hci_connect_sco() reference on every path; sco_connect_cfm() no longer needs its hci_conn_hold() dance.	7.1	More Details
CVE-2026-114720	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by an IMAP command injection vulnerability. A remote authenticated attacker can exploit this via CRLF sequences injected into the _flag parameter of the mail/mark action to execute arbitrary IMAP commands. Upstream fixed this in the commit “Fix IMAP command injection via CRLF in `_flag` parameter of mail/mark action”, which rejects flags containing control, whitespace and IMAP-special characters.	7.1	More Details
CVE-2026-114697	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in QuarkA QA Analytics qa-heatmap-analytics allows Stored XSS.This issue affects QA Analytics: from n/a through 5.3.0.1.	7.1	More Details
CVE-2026-	baserCMS is a website development framework. In versions 5.0.0 before 5.3.2 and 5.4.0 before 5.4.2, global shortcode expansion runs over the complete rendered page, including HTML-escaped blog comments supplied by unauthenticated visitors. A comment containing the BcBaser.getGoogleMaps shortcode can set apiBaseUrl to an attacker-controlled script URL, which is emitted as a script src and executes JavaScript on the site's origin when another visitor or administrator opens the article displaying that comment. Escaping raw HTML in comments does not prevent this stored cross-site	7.1	More Details

115457	scripting because shortcode processing happens after HTML rendering. Other registered shortcodes also expand in comments; Mail.getForm can embed a working contact form with a valid CSRF token. The flaw can compromise the authenticated browser context of an administrator who views the affected article. This issue has been patched in versions 5.3.2 and 5.4.2.		
CVE-2026-112379	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Codepeople Contact Form Email contact-form-to-email allows Stored XSS.This issue affects Contact Form Email: from n/a through 1.3.68.	7.1	More Details
CVE-2026-114178	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Bold Grid W3 Total Cache w3-total-cache allows Stored XSS.This issue affects W3 Total Cache: from n/a through 2.10.7.	7.1	More Details
CVE-2026-112869	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Noah Hearle We're Open! opening-hours allows Reflected XSS.This issue affects We're Open!: from n/a through 2.6.	7.1	More Details
CVE-2026-112844	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Arraytics WP Event Solution wp-event-solution allows Stored XSS.This issue affects WP Event Solution: from n/a through 4.1.26.	7.1	More Details
CVE-2026-114391	ZenTao PMS before 22.3 contains an authorization bypass vulnerability in commonModel::isOpenMethod() that lets logged-in users skip group-based ACL checks for any method containing 'ajax'. Minimal-privilege users can call methods like task.ajaxGetByID, system.ajaxDBAuthUrl and admin.ajaxChangeTableEngine to read data, obtain database credentials, trigger CI jobs and alter tables.	7.1	More Details
CVE-2026-114396	Trudesk through 1.2.11 contains a server-side JavaScript injection vulnerability that allows authenticated users to inject code into a MongoDB \$where clause via the search parameter. Attackers can close the regular expression in GET /api/v1/tickets/search to loop mongod for denial of service or leak full ticket documents including internal notes.	7.1	More Details
CVE-2026-115374	Incus 7.0.0 before 7.0.2 and 7.1.0 before 7.5.0 contains a signature verification flaw in its built-in S3 server that allows presigned URL holders to inject unsigned x-amz-* headers. Attackers can add an unsigned X-Amz-Copy-Source header to a presigned upload URL to copy other bucket objects onto the granted key, disclosing their contents and ETags.	7.1	More Details
CVE-2026-115371	Incus 6.0.3 before 6.0.7, 6.6.0 before 7.0.2, and 7.1.0 before 7.5.0 contains an improper authorization vulnerability that allows restricted TLS clients to read storage bucket keys in the default project. Attackers can query the bucket keys endpoints to obtain S3 access and secret keys, then read, modify or delete objects in those buckets.	7.1	More Details
CVE-2026-112313	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in ACF Extended ACF Extended acf-extended allows Stored XSS.This issue affects ACF Extended: from n/a through 0.9.2.7.	7.1	More Details
CVE-2026-112342	Tandoor Recipes before 2.7.0 contains an improper cache-key design in cookbook/helper/permission_helper.py that allows authenticated users to retain stale group-based permissions across different spaces by exploiting a cache key keyed only on user id while omitting the active space, with the cached entry's TTL refreshed on every request. Attackers can create their own space to obtain an admin group designation, then switch their active space to a victim space where they hold lesser privileges, causing the stale cached admin decision to authorize elevated access and bypass group revocations or demotions across multiple permission classes and views.	7.1	More Details
CVE-2026-111838	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in ThemeHunk ThemeHunk Customizer themehunk-customizer allows Reflected XSS.This issue affects ThemeHunk Customizer: from n/a through 2.8.7.	7.1	More Details
CVE-2026-111840	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in VideoWhisper.com Paid Videochat Turnkey Site ppv-live-webcams allows Reflected XSS.This issue affects Paid Videochat Turnkey Site: from n/a through 7.5.14.	7.1	More Details
CVE-2026-111848	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Dylan Kuhn Geo Mashup geo-mashup allows Reflected XSS.This issue affects Geo Mashup: from n/a through 1.13.22.	7.1	More Details
CVE-2026-112221	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Brainstorm Force SureCart surecart allows Stored XSS.This issue affects SureCart: from n/a through 4.9.5.	7.1	More Details

CVE-2026-112207	The Download Monitor WordPress plugin before 5.3.4 does not properly validate that a clicked link belongs to the site before fetching it client-side and rendering part of the response in the page, which allows unauthenticated attackers to execute arbitrary JavaScript in the browser of any user who is tricked into clicking a crafted link.	7.1	More Details
CVE-2026-115807	The User Profile Builder – Beautiful User Registration Forms, User Profiles & User Role Editor plugin for WordPress is vulnerable to Limited Privilege Escalation in all versions up to, and including, 4.0.3. This is due to the Gutenberg Register block passing the attacker-controlled Assigned Role attribute directly into the `[wppb-register]` shortcode without server-side authorization. The role guard in the form builder only downgrades roles that carry the manage_options or remove_users capability, leaving the Editor role and other privileged non-Administrator roles completely unrestricted. This makes it possible for authenticated attackers with Author-level access and above to create new WordPress accounts bearing the Editor role (or any other role that lacks the two guarded capabilities).	7.1	More Details
CVE-2026-82947	SOEM version 2.0.0 and earlier contains multiple vulnerabilities in frame processing where a malicious device can cause the master to read data outside allocated buffers. The software uses length and offset values from device messages without validating them against buffer sizes. This can cause the master to access invalid memory locations, potentially disclosing sensitive information or causing crashes.	7.1	More Details
CVE-2026-114510	Flexible GraphRAG is a Python framework for building graph and retrieval-augmented generation applications across multiple graph, vector, and search backends. From 0.6.0 until 0.8.2, the _escape() helper in flexible-graphrag/langchain/graph/pg_store_adapters/falkordb_adapter.py escapes quotes before backslashes and then interpolates the result into Cypher assembled by FalkorDBAdapter.add_graph_documents, allowing an inserted escape to be doubled and a quote to terminate the string early. When GRAPH_BACKEND=langchain and GRAPH_DB=falkordb are configured, attacker-influenced ingested document content can cause the extraction model to produce malicious entity or relationship values that reach these queries, allowing injected Cypher clauses to modify or delete graph data and, when the calling code surfaces query results to the user or LLM, disclose graph data. Dynamic node and relationship type identifiers are also inserted into backtick-quoted query positions without safe parameterization. The LlamaIndex FalkorDB store is not affected. This issue is fixed in version 0.8.2.	7.1	More Details
CVE-2026-115486	A flaw was found in the iscsiui component of open-iscsi. The IPv6 Neighbor Discovery handler accesses link-layer address options at a fixed offset without verifying the packet contains them. Additionally, the ARP table is updated unconditionally with attacker-controlled addresses from the packet headers. An attacker on the same network segment can send a minimal Neighbor Advertisement to poison the ARP table and redirect network traffic.	7.1	More Details
CVE-2026-114179	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in wordplus BP Better Messages bp-better-messages allows DOM-Based XSS.This issue affects BP Better Messages: from n/a through 3.1.3.	7.1	More Details
CVE-2026-111830	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Mailster WP Mailster wp-mailster allows Stored XSS.This issue affects WP Mailster: from n/a through 1.9.1.0.	7.1	More Details
CVE-2026-115192	An SQL Injection vulnerability exists in the dashboard widget data API endpoint (/api/data/{type}). Chart filter parameters supplied by the client are not properly validated before they are used in the database queries that build dashboard chart data. An authenticated user with API access can exploit this to inject SQL into these reporting queries. Using blind SQL injection techniques, the attacker can read arbitrary data from the Mautic database, including contact and other marketing data and user account data with password hashes. The injected SQL can also be used to slow down database queries, which can have a limited impact on availability. Data modification is not possible through this vulnerability. Exploitation requires the Mautic API to be enabled.	7.1	More Details
CVE-2026-107435	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 3.3.12.	7.1	More Details
CVE-2026-110899	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Kevin Stover Ninja Forms ninja-forms allows Stored XSS.This issue affects Ninja Forms: from n/a through 3.15.5.	7.1	More Details
CVE-2026-116016	yudao-cloud through 2026.09 contains an authorization bypass vulnerability that allows authenticated users to evade row-level data scope restrictions by sending a data-permission-enable header set to false. Attackers can add this header to /admin-api requests, which DataPermissionRpcWebFilter trusts, to read or modify other departments' and users' records within the same tenant.	7.1	More Details
	The MCTP-over-I2C+GPIO target binding in subsys/pmci/mctp/mctp_i2c_gpio_target.c emulates a small		

CVE-2026-79976	pseudo-register file, where the bus controller first writes a register number and then writes or reads payload bytes. The selected register (b->reg_addr) and the data-phase flag (b->rxtx) were never cleared when the controller issued an I2C STOP, so after a completed transfer the target still had a data register selected while the packet buffer that register refers to had already been consumed and set to NULL by mctp_i2c_gpio_target_stop(). A controller on the bus can re-enter the same branch of mctp_i2c_gpio_target_stop() with a NULL buffer using two ordinary transactions and no authentication: first a normal packet write to MCTP_I2C_GPIO_RX_MSG_ADDR followed by STOP (which delivers and frees b->rx_pkt), then a bare read transaction followed by STOP. The failing read_requested() call leaves the stale state untouched and the I2C driver still delivers the STOP callback, which then calls mctp_bus_rx() with a NULL packet. The symmetric case exists on the transmit side, where the stale MCTP_I2C_GPIO_TX_MSG_ADDR selection dereferences b->tx_pkt->end after b->tx_pkt has been cleared and additionally signals b->tx_complete a second time, causing the next outbound packet to be dropped. The dereference happens in the I2C target interrupt handler, so on any target that traps accesses to address 0 the result is a fatal exception and a device reset or halt — a denial of service inflicted by any device able to act as controller on the shared management bus (for example compromised firmware on a peer MCTP endpoint or on the BMC). On parts where address 0 is readable, the impact degrades to a malformed packet buffer being handed to the MCTP stack and lost transmit packets. No attacker-controlled value is written and no confidentiality or integrity boundary is crossed. The fix resets b->reg_addr to MCTP_I2C_GPIO_INVALID_ADDR and clears b->rxtx at every STOP, so a new register must be selected before any data phase can run again. Only builds that opt into the experimental CONFIG_MCTP subsystem with CONFIG_MCTP_I2C_GPIO_TARGET=y are affected.	7.1	More Details
CVE-2026-110884	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in geekontheroad Live Summary for Gravity Forms live-summary-for-gravity-forms allows Reflected XSS.This issue affects Live Summary for Gravity Forms: from n/a through 1.2.10.	7.1	More Details
CVE-2026-93631	A cross-site request forgery vulnerability in the OAuth extension of Zimbra Collaboration Suite (ZCS) Network Edition allows an attacker to link an attacker-controlled external account to a victim's mailbox. The OAuth callback does not bind its state value or authorization code to the initiating Zimbra session and instead selects the destination mailbox from the browser's current authentication cookie. On a deployment with the provider configured, an attacker can induce a signed-in victim to load a callback containing a fresh code for the attacker's provider account. This creates a persistent data source that imports attacker-controlled contacts or events. With bidirectional CalDAV synchronization, appointments subsequently created or edited in the linked calendar are also disclosed to the attacker. ZCS 10.1.16 and 10.1.20 are affected.	7.1	More Details
CVE-2026-114694	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Pixelite Events Manager events-manager allows Reflected XSS.This issue affects Events Manager: from n/a through 7.4.8.	7.1	More Details
CVE-2026-25716	Missing Authorization vulnerability in Hallo Welt! GmbH BlueSpice (Extension:BlueSpiceUserManager) allows any authenticated user without special privileges to block or unblock arbitrary user accounts, including administrators. The REST endpoints /bs-usermanager/v1/block/create and /bs-usermanager/v1/block/delete call UserManager::blockUser() and UserManager::unblockUser(), which do not verify the permissions of the acting user. This allows locking out all administrators (denial of service) and lifting blocks imposed by administrators. This issue affects BlueSpice: from 5.0.0 before 5.1.11, from 5.2.0 before 5.2.7.	7.1	More Details
CVE-2026-115191	A SQL injection vulnerability exists in a contact lookup AJAX endpoint of the Mautic administration interface (/s/ajax). A request parameter that selects the contact field to search by is not properly validated before it is used in a database query. Any authenticated Mautic user can exploit this, regardless of the role or permissions assigned to the account. Using blind (boolean-based or time-based) techniques, the attacker can run arbitrary SQL queries against the Mautic database and read any data the Mautic database user can access, including Mautic user accounts with email addresses and password hashes, contact data, and campaign and email content.	7.1	More Details
CVE-2026-114077	Trigger.dev before 4.5.6 contains an authorization bypass vulnerability in createDeploymentBackgroundWorkerV4 that resolves a WorkerDeployment by friendlyId alone, without environment, project, or organization scoping. Attackers holding only their own environment API key can POST to the deployment background-workers endpoint with a victim's deployment ID to re-point that in-progress deployment at a worker in the attacker's environment.	7.1	More Details
CVE-2026-95712	A sensitive information disclosure vulnerability in the PAPI protocol of AOS-8 could allow an unauthenticated adjacent attacker to access limited information handled by the affected service. Successful exploitation could allow an attacker to view some information and to disrupt the availability of the affected service.	7.1	More Details
CVE-	An authentication bypass vulnerability in the web-based management interface of AOS-10 could allow an unauthenticated adjacent attacker to circumvent existing authentication controls and run arbitrary		

2026-95711	commands on the underlying host if certain preconditions outside of the attacker's control are met. Successful exploitation of this vulnerability could allow an attacker to execute arbitrary commands on the underlying operating system.	7.1	More Details
CVE-2026-115975	Broken access control in JiraController in Cervantes version v1.3-beta could allow an attacker to create, edit, or delete Jira tickets for vulnerabilities in other projects	7.1	More Details
CVE-2026-95710	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-2026-114695	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Infomaniak Network VOD Infomaniak vod-infomaniak allows Stored XSS.This issue affects VOD Infomaniak: from n/a through 1.5.14.	7.1	More Details
CVE-2026-95713	An unauthenticated buffer overflow vulnerability exists in the command line interface of AOS-8. Successful exploitation could allow an unauthenticated remote attacker to disrupt the availability of the affected interface, along with limited access to or modification of some system data.	7.0	More Details
CVE-2026-98927	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: nv: Fix life cycle of the nested_mmus array The nested_mmus array holds the shadow page tables that are used when a guest is running a nested context. These structures are allocated on VCPU_INIT for whole guest, which implies that they may have to be relocated as the array grows. Should a VCPU_INIT occur whilst a vcpu is actively running an L2 and that the allocation requires relocation, that vcpu will still be running with a pointer to the previous structure, which will have been freed. Fix this by turning the array of structures to an array of pointers, which is now allocated at VM creation, sized to the absolute maximum that KVM can handle. In turn, each VCPU_INIT contributes S2_MMU_PER_VCPU to the pool. No reallocation is ever performed, and the life cycle of each object is much clearer: - the nested_mmus array is allocated in kvm_init_nested(), and freed in kvm_arch_destroy_vm() - s2_mmu structures are allocated in kvm_vcpu_init_nested(), and freed on kvm_arch_flush_shadow_all() Finally, the freeing of vcpu->arch.vncr_array is made consistent rather than being done on some failure paths, but not others.	7.0	More Details
CVE-2026-98941	In the Linux kernel, the following vulnerability has been resolved: ext4: don't enable large folios on encrypted inodes Commit 709f0f1f1bf5 ("ext4: add checks for large folio incompatibilities when BS > PS") disables large folios at mount time if the "encrypt" feature is set. However, enabling that feature on a mounted filesystem circumvents that check, and writing encrypted data results in a crash: BUG: kernel NULL pointer dereference, address: 0000000000000028 Workqueue: ext4-rsv-conversion ext4_end_io_rsv_work RIP: 0010:ext4_finish_bio+0xf7/0x3e0 Call Trace: ext4_release_io_end+0x51/0x110 ext4_end_io_end+0x4c/0xe0 ext4_end_io_rsv_work+0xb3/0x100 Fix this by disabling large folios on encrypted inodes, as is already done for inodes that use data journaling. This affects all kernels between v6.19 and v7.2, but it is not needed in v7.3, where the code was significantly refactored.	7.0	More Details
CVE-2026-99017	In the Linux kernel, the following vulnerability has been resolved: riscv: errata: select RISC_V_ALTERNATIVE_EARLY for THEAD GHOSTWRITE errata_probe_ghostwrite() only marks the CPU as vulnerable when it is called at the RISC_V_ALTERNATIVES_EARLY_BOOT stage, and the mitigation in ghostwrite_enable_mitigation() only takes effect if the vulnerable state was recorded before the boot alternatives are patched. The early boot stage is only executed when CONFIG_RISC_V_ALTERNATIVE_EARLY is enabled, which currently is only selected by ERRATA_THEAD_MAE. In a kernel built with ERRATA_THEAD_GHOSTWRITE=y but ERRATA_THEAD_MAE=n, the early alternatives never run, so ghostwrite_state is never set to VULNERABLE: ghostwrite_enable_mitigation() returns false, xtheadvector remains enabled, and sysfs reports "Not affected" on a vulnerable T-Head C9xx CPU. Select RISC_V_ALTERNATIVE_EARLY from ERRATA_THEAD_GHOSTWRITE so that the vulnerability is detected in the early boot stage and the default mitigation is applied.	7.0	More Details
CVE-2026-98895	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: HWS, Handle destroying table that has a miss table If a table has a miss table that was created by 'mlx5hws_table_set_default_miss' API function, its miss_tbl keeps the table that points to it in a list. If such table is deleted, we need to also remove it from the miss_tbl list, otherwise the node in miss_tbl list will contain garbage.	7.0	More Details
	The kernel helper z_waitq_head() (kernel/include/wait_q.h) reads the head of a thread wait queue, but on SMP systems it was called from several places without holding the scheduler spinlock _sched_spinlock that serializes wait-queue mutation. The consequential call site is z_timer_expiration_handler() in kernel/timer.c, which held only the timer subsystem lock timer_lock		

CVE-2026-78118	while reading timer->wait_q and then unpending, setting the return value of, and readying the thread it obtained. halt_thread() in kernel/sched.c removes an aborted thread from its wait queue while holding only _sched_spinlock, so the two paths are not mutually exclusive. When the last waiter is removed on another CPU inside that window, sys_dlist_peek_head() can return the list sentinel — a pointer into the k_timer object itself — typed as a struct k_thread * (with CONFIG_WAITQ_SCALABLE, rb_get_min() can return an arbitrary node of a tree being restructured), or it can return a pointer to a thread that has already been unpended and, for dynamically allocated threads, freed. The handler then calls unpend_thread_no_timeout(), which performs a sys_dlist_remove() through pointers read from that object, and arch_thread_return_value_set(), which writes unconditionally through the pointer. The result is out-of-bounds and use-after-free writes in supervisor context, plus scheduler state corruption; a milder outcome is a thread that has re-pended elsewhere being woken with a success return value from an unrelated blocking call. All of the primitives needed to drive the interleaving — k_timer_start(), k_timer_status_sync(), k_thread_create() and k_thread_abort() — are system calls, so with CONFIG_USERSPACE enabled an unprivileged user-mode thread can attempt the race against kernel-owned data. Exploitation requires a multi-CPU build (CONFIG_MP_MAX_NUM_CPUS greater than one; the affected region is a no-op on uniprocessor targets), local code execution, and winning a narrow timing window, so the practical outcome is most often a kernel crash rather than controlled corruption. The fix splits the helper into z_waitq_head_locked() for callers that already hold the scheduler spinlock and a z_waitq_head() that acquires it, and brings the timer expiration handler's head read, unpend and ready under a single LOCK_SCHED_SPINLOCK region.	7.0	More Details
CVE-2026-76496	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. &nbsp; The vulnerabilities tracked by CVE-2026-76496 are related to issues concerning improper link resolutions that are grouped under the Common Weakness Enumeration (CWE) Base CWE-59.	7.0	More Details
CVE-2026-54532	CloudPirates Open Source Helm Charts is a collection of Helm charts. Prior to commit 271d415, pull_request_target workflow is vulnerable to cache poisoning and OIDC token theft. A pull_request_target workflow checks out fork code using github.event.pull_request.head.sha or .head.ref. This gives attacker-controlled code access to base repository secrets and tokens. A pull_request_target workflow writes to a shared Actions cache. A subsequent release or publish workflow restores from the same cache scope, allowing a poisoned cache entry to execute during release. This issue has been patched via commit 271d415.	6.9	More Details
CVE-2026-114429	The WPVR WordPress plugin before 9.0.3 does not properly sanitise and escape some of its tour hotspot content before storing it and outputting it on the front end, which could allow users with the Contributor role and above to perform Stored Cross-Site Scripting attacks that execute in the browser of any user viewing the affected tour.	6.8	More Details
CVE-2026-112344	Synclify Server! before 8.8.0 contains a path traversal vulnerability that allows attackers with a symbolic link placed inside a user's home folder to access files outside the virtual filesystem root because confinement is enforced only against the requested path and not the on-disk destination of the link. Attackers can leverage a symlink or, on Windows, a directory junction that points outside the jail to read or access restricted files beyond the intended home-folder boundary.	6.8	More Details
CVE-2026-115993	Stored XSS vulnerability in the “Notes” field of the Checklists in Cervantes version v1.3-beta could allow an attacker to execute arbitrary code in another user's browser with elevated privileges	6.8	More Details
CVE-2026-78392	The Link Library WordPress plugin before 7.9.6 does not sanitise and escape a link field before outputting it inside an attribute of a link preview, allowing users with the Contributor role and above to perform Stored Cross-Site Scripting attacks. The preview is generated without authentication and without checking whether the link has been published, so an unapproved submission is enough to reach any visitor who opens it.	6.8	More Details
CVE-2026-114618	Authlib is a Python library which builds OAuth and OpenID Connect servers. Prior to 1.9.0, the default RFC 7523 JWTBearerGrant.get_audiences() returns an empty list, causing JWTBearerGrant.verify_claims() to require only that an aud claim exists without matching the current authorization server. An attacker holding a legitimate assertion for another service can present it to a server that shares the signing key and uses the default grant, obtaining an access token across a trust boundary. Implementations that override get_audiences() with the server's valid audience identifiers are not affected. This issue is fixed in version 1.9.0.	6.8	More Details
CVE-2026-111971	The Photo Engine (Media Organizer & Lightroom) WordPress plugin before 6.5.8 does not sanitise and escape collection and gallery names before storing them and outputting them in an administrative dashboard, allowing authenticated users with the author role and above to perform Stored Cross-Site Scripting attacks.	6.8	More Details

CVE-2026-105199	The Pixel Tag Manager for WooCommerce WordPress plugin before 2.2.3 does not perform a capability check in one of its AJAX actions, and exposes the nonce protecting it to any user able to access the admin dashboard, allowing users with a role as low as Contributor to overwrite the Pixel Tag Manager for WooCommerce WordPress plugin before 2.2.3's tracking configuration and have arbitrary JavaScript executed in the browsers of all visitors, including administrators.	6.8	More Details
CVE-2026-115929	MongoDB Compass relies on Electron, a third-party runtime that carries its own TLS root certificate store. This root certificate store has the recoverable eMudhra intermediate CA certificates. Issue is resolved by updating to Compass 1.53.1, which includes Electron 44.9.1.	6.8	More Details
CVE-2026-116194	Rclone from 1.52.0 before 1.75.2 contains a cross-site request forgery vulnerability in the pCloud backend's localhost OAuth callback, which disables the state check and trusts the hostname parameter. Malicious web pages can send a stateless request to port 53682 during authorization, planting attacker tokens and API hosts so later transfers reach the attacker.	6.8	More Details
CVE-2026-114715	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by an incorrect authorization vulnerability in the Modoboa driver of the password plugin. The driver took the first entry of the account search result returned by the Modoboa API instead of matching the authenticated user, so a password change could be applied to a different account. Upstream fixed this in the commit "Fix bug where Password Modoboa driver could change password for another account", which selects the account whose username matches and fails if none does.	6.8	More Details
CVE-2026-116488	NXP MCUXpresso SDK USB middleware (mcuxsdk-middleware-usb) before 2.13.0 contains a heap buffer overflow in host audio, video and CDC class drivers that trust device-supplied bNumEndpoints and bLength values. A malicious USB device can supply crafted descriptors that write endpoint entries past the fixed-size epList array, corrupting heap memory for code execution or denial of service.	6.8	More Details
CVE-2026-114649	The MongoDB C# Driver compares TLS client certificates by issuer and serial number only when deciding whether two clients can share cached connections. In an application that creates clients from certificates supplied by different users, a user who provides a certificate copying another user's issuer and serial number, without that user's private key, can be given connections authenticated with the other user's certificate. Database operations can then run under the other user's identity, affecting the confidentiality and integrity of data that identity can access.	6.8	More Details
CVE-2026-115930	The MongoDB Shell (mongosh) relies on Node.js as a third-party runtime that carries its own TLS root certificate store. This certificate store has the recoverable eMudhra intermediate CA certificates. This issue impacts MongoDB Shell (mongosh) versions prior to 2.14.0 and can be resolved by updating to mongosh 2.14.0 which bundles Node.js 24.22.1 or later. If mongosh is installed through npm, it uses the Node.js runtime installed on the host; update that runtime separately.	6.8	More Details
CVE-2026-94264	The Advanced Custom Fields: Extended WordPress plugin before 0.9.2.8, Advanced Custom Fields: Extended PRO WordPress plugin before 0.9.2.8 does not validate or escape some of its form shortcode attributes, which can be pointed at values taken from the visitor's own request and are then output without escaping, allowing users with a role as low as contributor to have arbitrary script run in the browser of anyone led to open a crafted link or page, including a logged-in administrator. This is an incomplete fix for CVE-2023-5292, whose deny-list covers only the three HTML wrapper attributes.	6.8	More Details
CVE-2026-111966	The User Profile Builder WordPress plugin before 4.0.5 does not sanitise and escape one of its shortcode attributes before outputting it back in a redirect message, which could allow users with a role as low as contributor to perform Stored Cross-Site Scripting attacks against higher-privileged users who submit the affected form.	6.8	More Details
CVE-2026-107790	The HireZoot WordPress plugin before 4.1.1 does not sanitise and escape some of its stored data before outputting it in the admin dashboard, and additionally reverses the HTML-entity encoding that would otherwise neutralise it, allowing users with the Author role to perform Stored Cross-Site Scripting attacks that execute in the context of a logged-in administrator.	6.8	More Details
CVE-2026-81813	The Elementor Website Builder WordPress plugin before 4.3.5 does not sanitize the text content of SVG files, checking only their elements and attributes, allowing users with the Contributor role and above to store a payload which breaks out when the SVG is rendered inline in a page and executes in the browser of any user viewing it. Execution is specific to Chromium-based browsers; Firefox is not affected.	6.8	More Details
CVE-2026-114732	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by an improper authentication vulnerability in multi-host installations. When a user logs in with an alias address, rcube_user::query() matched the address against identity records without qualifying the lookup by IMAP host, so a login could be resolved to a user account belonging to a different IMAP host. Upstream fixed this in the commit "Fix bug where an alias login wasn't qualified by IMAP host".	6.8	More Details

CVE-2026-116066	Firecrawl before 2.11.586 contains a server-side request forgery vulnerability in the Playwright service SSRF proxy that allows attackers to bypass isInternalHost checks via DNS rebinding. Attackers controlling DNS for a scrape target can return a public address during validation and a private address at connect time, retrieving cloud metadata and internal resources through the scrape API.	6.8	More Details
CVE-2026-109933	The Otter Blocks WordPress plugin before 3.2.8 does not properly sanitize and escape user-supplied block styling values before writing them into an inline stylesheet, allowing users with contributor-level access and above to inject arbitrary web scripts that execute when any visitor, including administrators, views the affected post.	6.8	More Details
CVE-2026-110651	The Swift Performance Lite WordPress plugin before 2.3.7.4 does not properly sanitize and escape a value before using it in a SQL statement, which could allow users with the author role and above to perform SQL injection attacks.	6.8	More Details
CVE-2026-114406	The system certificate authority (CA) store in Red Hat products, distributed via ca-certificates RPM packages, included multiple root CA certificates of authorities, which signed intermediate CA certificates with weak private keys issued by the eMudhra certificate authority. A man-in-the-middle attacker could use this flaw to issue certificates for arbitrary host names signed by the weak intermediate CA certificates, which would be considered trusted by a TLS/SSL client using system CA store and cause such clients to connect to malicious sites. Due to inability to only distrust specific intermediate certificates in all clients using the system CA store, root CA certificates of authorities that signed weak intermediate CA certificates are being removed.	6.8	More Details
CVE-2026-115334	Lotus through 0.10.2 contains an authorization bypass vulnerability in ChangeUserOrganizationView.post that allows authenticated users to join any organization by supplying its UUID. Self-registered attackers who know another organization's UUID can send a POST request to /app/switch_organization/ to access that tenant's data and billing configuration.	6.8	More Details
CVE-2026-89209	A vulnerability has been identified in Mendix SAML (Mendix 10 compatible) (All versions < V4.2.4), Mendix SAML (Mendix 11 compatible) (All versions < V4.2.4), Mendix SAML (Mendix 9 compatible) (All versions < V3.6.29). Affected versions of the module do not properly protect against SAML assertion replay in IdP-initiated authentication flows. This could allow unauthenticated remote attackers to gain unauthorized access to the application by replaying a valid SAML response.	6.8	More Details
CVE-2026-113493	The GeoDirectory WordPress plugin before 2.8.194 does not sanitise and escape some of its map shortcode attributes before outputting them within a style attribute, allowing users with the Contributor role and above to perform Stored Cross-Site Scripting attacks that execute in the context of any user viewing the affected content.	6.8	More Details
CVE-2026-116166	SiYuan 3.6.0 before 3.8.7 accepts and persists the networkServeTLS setting but never enforces it, so the kernel port keeps serving full API access over plain HTTP alongside HTTPS. Network-adjacent attackers can sniff the cleartext login flow and session cookie, which lacks the Secure attribute while responses omit HSTS, to hijack kernel API sessions.	6.8	More Details
CVE-2026-115729	@pdfme/common from 5.5.10 before 6.2.6 contains a server-side request forgery vulnerability that allows attackers to bypass the isUrlSafeToFetch denylist in getB64BasePdf. Attackers can supply basePdf URLs using IPv4-mapped IPv6, [::], NAT64 or localhost alias hostnames to read loopback, internal or cloud metadata responses through generated documents.	6.8	More Details
CVE-2026-111979	The GiveWP WordPress plugin before 4.19.0.1 does not properly sanitise and escape one of its donation form template settings before outputting it inside an HTML attribute, allowing users with a custom form-editing GiveWP WordPress plugin before 4.19.0.1 role to inject arbitrary web scripts that execute in the browser of any visitor or administrator who opens the affected donation form.	6.8	More Details
CVE-2026-108013	BigFix Platform uses Mozilla CA certificate bundle (PEM). Versions of the bundle dated before 2026/10/20 still trust certificates whose private keys have been compromised. When BigFix Platform validates servers using such a bundle, an on-path attacker can impersonate remote servers and tamper with files downloaded by the application. The issue is resolved by using bundle version 2026/10/20 or later. The fixed bundle is available starting from BES Support v1520.	6.8	More Details
CVE-2026-116167	SiYuan from 2.9.0 before 3.8.7 contains a cross-site scripting vulnerability in the unauthenticated /check-auth page that assigns the to parameter to window.location.href without scheme validation. Attackers can send a crafted link with a javascript: URI that executes after victim login, invoking kernel APIs to read or modify workspace content and settings.	6.8	More Details
CVE-2026-46347	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. Prior to version 26.05, when a password reset link is visited, the application generates a replacement password in a predictable way and emails it to the account holder. The insufficient entropy allows the password to be brute-forced by an attacker who can approximate the server-side generation time. Exploiting this finding in isolation requires the victim to have triggered the reset themselves and the	6.8	More Details

	attacker to infer the server time from network or email metadata. This issue has been patched in version 26.05.		
CVE-2026-114047	FreeRDP from 2.0.0 before 3.32.0 contains an improper certificate validation vulnerability in its proxy that accepts untrusted target TLS certificates when no verification module is configured. Network-positioned attackers can intercept the proxy's outbound connection and present a self-signed, wrong-hostname certificate to impersonate the RDP target and receive proxied session data.	6.8	More Details
CVE-2026-95715	An unauthenticated command injection vulnerability exists in the command line interface of AOS-10. Successful exploitation could allow an attacker with physical access to execute arbitrary commands as a privileged user on the underlying operating system leading to complete system compromise.	6.8	More Details
CVE-2026-12664	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 18.2 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, under certain conditions, could have allowed an authenticated user with a read_api-scoped token to retrieve a project's incoming email address and create work items, bypassing the token's restrictions, due to insufficient scope enforcement on a GraphQL field.	6.8	More Details
CVE-2026-95714	A buffer overflow vulnerability exists in the API of AOS-10 that could allow an authenticated adjacent attacker with high privileges to achieve remote code execution. Successful exploitation could allow an attacker to execute arbitrary code as a privileged user on the underlying operating system.	6.8	More Details
CVE-2026-115365	OpenBao from 2.3.0 before 2.6.3 contains an incorrect authorization vulnerability in its ACL policy cache that allows tokens to reference policies from arbitrary namespaces using specially crafted policy names. Attackers holding such tokens can gain capabilities of cached policies in other namespaces, including the root namespace, to act across namespace boundaries.	6.6	More Details
CVE-2026-76497	As part of Cisco's ongoing commitment to proactive security and product quality, the Cisco Expressway Series software engineering team has conducted a comprehensive internal security review. This review resulted in software hardening releases that address multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76497 are related to issues concerning insecure temporary files that are grouped under the Common Weakness Enumeration (CWE) Base CWE-377.	6.6	More Details
CVE-2026-114652	A data-query-logic neutralization weakness in the update-statement construction component of the MongoDB C# Driver can cause application-supplied values to be interpreted as query/update logic rather than as literal data. An application user who can influence the content of values that the application passes into the affected builder surface, without any special privilege beyond ordinary read-write application access, may cause stored data to be modified or removed in ways the application never intended. The result is unauthorized integrity impact on the affected application data; there is no direct loss of confidentiality or availability of the database service.	6.5	More Details
CVE-2026-115354	Mistune before 3.3.0 contains a path traversal vulnerability in the opt-in Include directive that allows attackers to read local files by supplying relative or absolute include paths. Attackers who submit Markdown rendered via Markdown.read() can disclose files such as /etc/passwd and inject unescaped HTML or script content into rendered output.	6.5	More Details
CVE-2026-114707	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Crocoblock JetElements For Elementor jet-elements allows Stored XSS.This issue affects JetElements For Elementor: from n/a through 2.9.2.5.	6.5	More Details
CVE-2026-114646	A data-validation flaw exists in the BSON handling and binary serialization path of the MongoDB C# Driver. When an application constructs a raw or lazy BSON document object from externally supplied binary data, the driver does not verify that the size quantity declared inside that data matches the amount of data actually provided, and it copies the data onto the wire unchanged. An unauthenticated party who can supply binary content to such an application can thereby cause additional elements to be inserted into, or existing elements to be removed from, the database command the application issues, without needing any database credentials of their own. This can result in unauthorized modification of stored data and disclosure of data outside the intended scope of the application's request, bounded by the database privileges the application itself holds.	6.5	More Details
CVE-2026-97254	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WProyal Royal Elementor Addons royal-elementor-addons allows Stored XSS.This issue affects Royal Elementor Addons: from n/a through 1.7.1069.	6.5	More Details
CVE-2026-116254	A vulnerability was identified in rathena FluxCP up to 98cd4b8fd83228c11b36fb3d097b8e0f7c105a8d. This issue affects the function verifyiprange of the file lib/Flux/PaymentNotifyRequest.php of the component PayPal IPN Handler. The manipulation of the argument X-Real-IP leads to origin validation error. The attack is possible to be carried out 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.	6.5	More Details

CVE-2026-114192	Missing Authorization vulnerability in Magepeople inc. WpEvently mage-eventpress allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WpEvently: from n/a through 5.7.8.	6.5	More Details
CVE-2026-115949	Missing Authorization vulnerability in WP Smart Widgets Better Post & Filter Widgets for Elementor better-post-filter-widgets-for-elementor allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Better Post & Filter Widgets for Elementor: from n/a through 1.9.1.	6.5	More Details
CVE-2026-90484	The MySQL client response parser in Apache bRPC dereferences a null pointer when the server peer sends a zero-payload-length response packet. A response packet whose declared payload length is zero is accepted by the frame-length check (is_full_package), after which the parser reads the response type byte via an unbounded 5-byte fetch that returns nullptr, and type = p[4] dereferences it, crashing the client process. Root Cause: In brpc::policy::MysqlReply::ConsumePartialIOBuf, buf.fetch(header, 5) returns nullptr when the buffered data only holds the 4-byte header of a zero-length packet, then the code reads p[4] (and, in the fast-auth path, p2[4]) without checking for a null or zero-length packet. Affected Scenarios: A brpc MySQL client connecting to a malicious or compromised MySQL server (or a MITM) can be crashed by a zero-length server response packet, resulting in denial of service of the client. Fix: The fix rejects zero-length response packets (returning PARSE_ERROR_ABSOLUTELY_WRONG) both for the initial response type determination and after a fast-auth marker, before reading the type byte. Upgrade to the Apache bRPC version 1.18.0 or apply the PR patch at https://github.com/apache/brpc/pull/3536 .	6.5	More Details
CVE-2026-115602	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Liquid Web / StellarWP Gutenberg Blocks by Kadence Blocks kadence-blocks allows Stored XSS.This issue affects Gutenberg Blocks by Kadence Blocks: from n/a through 3.8.0.1.	6.5	More Details
CVE-2026-96386	Missing Authorization vulnerability in Drupal Smart Content allows Forceful Browsing. This issue affects Smart Content versions: from 0.0.0 to 3.2.1.	6.5	More Details
CVE-2026-114653	A client-side encryption schema construction component in the MongoDB C# Driver places caller-supplied field identifiers into a position that is interpreted as a pattern expression, without normalizing or validating them. As a result, an application that designates certain fields for client-side field-level encryption may have the protection silently fail to apply, so data intended to be encrypted is stored and returned in plaintext; conversely, the rule may apply more broadly than intended. No user action is required to cause the condition, and any party with ordinary read access to the stored data — without possession of the encryption keys — can observe the unprotected values.	6.5	More Details
CVE-2026-96385	Authentication Bypass Using an Alternate Path or Channel vulnerability in Drupal REST & JSON API Authentication for Drupal allows Authentication Bypass. This issue affects REST & JSON API Authentication for Drupal versions: from 0.0.0 to 3.2.0.	6.5	More Details
CVE-2026-112865	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WProyal Royal Elementor Addons royal-elementor-addons allows DOM-Based XSS.This issue affects Royal Elementor Addons: from n/a through 1.7.1069.	6.5	More Details
CVE-2026-46346	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. Prior to version 26.05, the password reset token is predictably based on the time of the reset request. An attacker who knows the target's email address can request a reset and then brute-force the token. Exploiting this finding alone forces an unauthorized password reset on any account. This issue has been patched in version 26.05.	6.5	More Details
CVE-2026-114709	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Igor Funa Ad Inserter ad-inserter allows Reflected XSS.This issue affects Ad Inserter: from n/a through 2.8.18.	6.5	More Details
CVE-2026-96380	Missing Authorization vulnerability in Drupal CSS Usage Analyzer allows Forceful Browsing. This issue affects CSS Usage Analyzer versions: from 1.0.0 to 1.0.2.	6.5	More Details
CVE-2026-55463	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 4.11.0 to before version 26.05, several Jakarta WebSocket endpoints in goobi-viewer-core (DownloadTaskEndpoint, CampaignEndpoint, AutocompleteEndpoint, UserEndpoint) accept connections without verifying that the associated HttpSession belongs to a logged-in user, and none validate the Origin header against an allowlist. An unauthenticated attacker can therefore connect to these endpoints; through CampaignEndpoint this allows polluting the in-process crowdsourcing page-lock map and broadcasting forged lock messages, disrupting reviewers' work. This issue has been patched in version 26.05.	6.5	More Details

CVE-2026-115613	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in bPlugins Image Gallery block – Create and display photo gallery/photo album. 3d-image-gallery allows Stored XSS.This issue affects Image Gallery block – Create and display photo gallery/photo album.: from n/a through 2.4.1.	6.5	More Details
CVE-2026-114654	A numeric-type conversion flaw in the BSON handling library of the MongoDB C# Driver allows an out-of-range binary subtype value supplied through untrusted textual input to bypass the library's binary-element integrity checks. When such a value is later encoded, it is narrowed to a smaller wire representation, producing binary elements that misdeclare their type or length. An unauthenticated party able to supply text parsed by an affected application can cause it to emit and persist malformed binary data, which subsequent readers either reject with an error or silently interpret as shorter than the data that was validated.	6.5	More Details
CVE-2026-82943	SOEM version 2.0.0 and earlier contains vulnerabilities where a malicious device can send messages that cause the master to read data from memory locations beyond the intended buffer. The software uses length values from device responses without validating them against actual buffer sizes. An attacker can exploit this to extract up to 64KB of data from the master's memory.	6.5	More Details
CVE-2026-114610	fpdf2 is a PDF creation library for Python. From 2.5.0 until 2.8.8, fpdf2's SVGObject.build_xref() passes nested SVG <use> references to GraphicsContext.add_item(), whose default clone=True deep-copy behavior copies the complete referenced graphics tree during FPDF.image() processing. Repeating references at each nesting level makes the in-memory object count and processing work grow exponentially, allowing a small user-controlled SVG to consume excessive CPU and memory, raise MemoryError, fail PDF generation jobs, and reduce or exhaust worker capacity. ResourceAccessPolicy does not prevent the issue because the references are inline and require no external resource, XML entity, local file, or permissive network setting. This issue is fixed in version 2.8.8.	6.5	More Details
CVE-2026-114665	Goploy 1.18.0 before 1.18.1 contains a link following vulnerability in the /deploy/fileDiff and /deploy/fileCompare endpoints that allows authenticated members to read arbitrary files. Attackers with repository write access can commit a symlink pointing outside the project directory to read database configuration and deployment SSH private keys.	6.5	More Details
CVE-2026-114700	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Gloriafood Restaurant Menu – Food Ordering System – Table Reservation menu-ordering-reservations allows Stored XSS.This issue affects Restaurant Menu – Food Ordering System – Table Reservation: from n/a through 2.4.4.	6.5	More Details
CVE-2026-114701	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in xootix Waitlist Woocommerce (Back in stock notifier) waitlist-woocommerce allows Stored XSS.This issue affects Waitlist Woocommerce (Back in stock notifier): from n/a through 2.9.1.	6.5	More Details
CVE-2026-114702	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in nicdark ND Shortcodes For Visual Composer nd-shortcodes allows Stored XSS.This issue affects ND Shortcodes For Visual Composer: from n/a through 7.8.	6.5	More Details
CVE-2026-114703	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in hupe13 Extensions for Leaflet Map extensions-leaflet-map allows Stored XSS.This issue affects Extensions for Leaflet Map: from n/a through 6.0.	6.5	More Details
CVE-2026-114492	gin-vue-admin through 2.9.2 contains a path traversal vulnerability that allows authenticated administrators to recursively delete arbitrary directories by supplying traversal sequences in the pluginName parameter. Attackers can send a crafted POST request to /autoCode/removePlugin with '..' sequences, causing os.RemoveAll to delete any directory writable by the server process.	6.5	More Details
CVE-2026-110892	Missing Authorization vulnerability in Brainstorm Force SureDash suredash allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects SureDash: from n/a through 1.12.2.	6.5	More Details
CVE-2026-114131	Nuxt is an open-source web development framework for Vue.js. From 3.11.0 until 3.21.4 and 4.4.4, Nuxt compares route-rule keys with the raw request path while route resolution decodes dynamic and catch-all parameters, allowing a percent-encoded segment to resolve to a protected page without matching the routeRules entry intended for that path. As a result, appMiddleware can be skipped during server rendering and an unauthenticated requester can receive the gated page and its SSR data. The same mismatch affects app-manifest rules including redirect, appLayout, ssr, prerender, and payload flags derived from cache, isr, and swr. This vulnerability is fixed in 4.4.4 and 3.21.4.	6.5	More Details
CVE-2026-114693	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Plugins360 Labs All-in-One Video Gallery all-in-one-video-gallery allows Stored XSS.This issue affects All-in-One Video Gallery: from n/a through 4.9.7.	6.5	More Details

CVE-2026-114128	Astro is a web framework for content-driven websites. Prior to 7.3.2, Astro's JSX and MDX rendering pipeline treats string children of script and style elements as trusted HTML without distinguishing compiler-generated literal content from dynamic expression values. An application that passes attacker-controlled data directly as a dynamic child can allow a value containing a closing script or style tag to terminate the element and inject arbitrary HTML and JavaScript. Other JSX expression contexts escape dynamic strings automatically; the issue is limited to script and style children rendered through the astro:jsx renderer, including MDX components. Successful exploitation can produce reflected or stored cross-site scripting under the application's origin. This issue is fixed in 7.3.2.	6.5	More Details
CVE-2026-103381	The Mintpay for WooCommerce WordPress plugin through 2.2.3 does not verify the authenticity of the payment callback it processes, allowing unauthenticated attackers to mark arbitrary orders as paid, to cancel arbitrary orders, and to disclose order details including customer contact information.	6.5	More Details
CVE-2026-114184	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WPVibes Elementor Addon Elements addon-elements-for-elementor-page-builder allows Stored XSS.This issue affects Elementor Addon Elements: from n/a through 1.14.4.	6.5	More Details
CVE-2022-51025	Penta Security WAPPLES 4.0.54.1, 5.0, and 6.0 contains a path traversal vulnerability that allows authenticated attackers to download arbitrary files via the name parameter of /webapi/file/transfer. Attackers can send GET requests with type=db_backup and ../ sequences in name to read files like /etc/passwd, and a hardcoded systemi account makes authentication effectively trivial. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	6.5	More Details
CVE-2026-112842	Missing Authorization vulnerability in Cozmoslabs Paid Member Subscriptions paid-member-subscriptions allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Paid Member Subscriptions: from n/a through 3.1.2.	6.5	More Details
CVE-2026-112867	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in H5P H5P h5p allows Stored XSS.This issue affects H5P: from n/a through 1.17.9.	6.5	More Details
CVE-2026-87835	The Front End Users WordPress plugin through 3.2.35 does not restrict which account attributes a visitor may set when registering through its front-end forms, allowing unauthenticated users to assign themselves an arbitrary membership level and to mark their own account as approved, confirmed and paid, bypassing the Front End Users WordPress plugin through 3.2.35's content restrictions and its registration gating.	6.5	More Details
CVE-2026-111835	Missing Authorization vulnerability in bjorntech BjornTech Swish for WooCommerce woo-swish-e-commerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects BjornTech Swish for WooCommerce: from n/a through 3.7.8.	6.5	More Details
CVE-2026-114260	A vulnerability has been found in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. This vulnerability affects the function delete_folder of the file ajax.php of the component AJAX Handler. The manipulation of the argument ID leads to missing authentication. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used. Continuous delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	6.5	More Details
CVE-2026-107907	The MAS Videos WordPress plugin through 1.3.4 does not perform any capability or ownership check when handling its front-end video and playlist deletion requests, and the nonce it verifies is tied only to a static action rather than to the targeted object, allowing any authenticated user with a minimal role to permanently delete arbitrary posts, pages, attachments and other content belonging to other users.	6.5	More Details
CVE-2026-114194	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Igor Funa Ad Inserter ad-inserter allows Reflected XSS.This issue affects Ad Inserter: from n/a through 2.8.18.	6.5	More Details
CVE-2026-114475	Missing authorization in the REST role-management routes of Progressive Robot hMailServer 6.3.4 through 6.3.6 allows a remote authenticated user whose credential is confined to named domains to modify a custom role that accounts in other domains hold, and to create custom roles and delete any that no account holds. The routes that create, update and delete a custom role require the manage_roles permission, which only the serveradmin role holds; but a full-scope API key issued for named domains, and serveradmin assigned with a domain list, are serveradmin confined to those domains and hold manage_roles. The server's authorization gate did not refuse the three role routes to a confined credential and the route named no domain, so a credential issued to one tenant could raise or lower the permissions of a custom role assigned across other tenants - granting another domain's role every permission over its mailboxes and settings, or removing what it needs. A custom role may not hold administer_server, administer_domain or manage_roles, so the confined credential cannot	6.5	More Details

	make itself or another a server administrator by creating or rewriting a custom role.		
CVE-2026-114705	Missing Authorization vulnerability in quadlayers WP Social Feed Gallery insta-gallery allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Social Feed Gallery: from n/a through 5.0.13.	6.5	More Details
CVE-2026-115886	MLflow is an open source AI engineering platform for agents, large language models, and machine learning models. From 3.12.0rc0 until 3.17.0, the scorer invoke, register, and GenAI evaluate endpoints, including unauthenticated POST /ajax-api/3.0/mlflow/scorer/invoke, accept serialized third-party scorer data that Scorer.model_validate_json resolves using a module-path allowlist without constraining the selected class or its constructor arguments. When optional trulens and trulens-providers-litellm packages are installed, an attacker can select mlflow.genai.scorers.trulens, choose a model URI outside MLflow's native provider registry so execution uses the LiteLLM fallback, and supply api_base and api_key values that cause the server to send an outbound request to an attacker-chosen or internal host with attacker-controlled headers and body. The vulnerable deserialization occurs before a tracking-store lookup or job submission. This issue is fixed in version 3.17.0.	6.5	More Details
CVE-2026-116223	A flaw has been found in icret EasyImage up to 2.8.7. The impacted element is an unknown function of the file app/del.php. Executing a manipulation of the argument hash can lead to improper privilege management. The attack can be launched remotely. The exploit has been published and may be used. This patch is called b12cb67d447bcbb67e21c8f011d0f5638229760c. A patch should be applied to remediate this issue.	6.5	More Details
CVE-2026-114515	OCO-Server is a self-hosted endpoint management application for computer inventory, software deployment, and mobile device management. Prior to 1.2.3, a valid Self-Service Portal user with permission to create job containers and read visibility on a computer and package can queue an uninstall job without deploy permission for those resources. Requests to `self-service/ajax-handler/job-containers.php` pass `uninstall_package_assignment_id[]` values to `lib/SelfService/CoreLogic.class.php::uninstallSelfService()`. Its permission-check loop continues after denied `METHOD_DEPLOY` checks without removing rejected IDs from `\$installationIds`, then passes the original list to `lib/CoreLogic.class.php::uninstall()`. The CoreLogic object is constructed with a null SystemUser, causing `lib/CoreLogic.class.php::checkPermission()` to treat the operation as system context and unconditionally permit it. Enumerable `computer_package.id` values allow selection of readable installed-package assignments, including those on other users' computers. The `deploySelfService()` install path instead aborts on denied deployment permission, and resources lacking read permission remain blocked. This permits unauthorized software-removal jobs across resources visible to the acting user's role. This issue is fixed in version 1.2.3.	6.5	More Details
CVE-2026-115594	Insertion of Sensitive Information Into Sent Data vulnerability in David Lingren Media Library Assistant media-library-assistant allows Retrieve Embedded Sensitive Data.This issue affects Media Library Assistant: from n/a through 3.42.	6.5	More Details
CVE-2026-82945	SOEM version 2.0.0 and earlier contains a vulnerability during device configuration where a malicious device can cause memory corruption. The software writes configuration data to a fixed-size array without checking if the array index exceeds its maximum size of EC_MAXFMMU entries. A device presenting more than EC_MAXFMMU configuration blocks can force writes beyond the array boundary, corrupting adjacent memory structures and potentially disrupting subsequent operations.	6.5	More Details
CVE-2026-82946	SOEM version 2.0.0 and earlier contains a vulnerability where a malicious device can crash the master system. When the software receives an error response, it clears an internal pointer but continues processing. If another message arrives before timeout, the software attempts to use the cleared pointer, causing a crash. An attacker can trigger this with a simple sequence of two messages.	6.5	More Details
CVE-2026-111401	Cross-Site Request Forgery (CSRF) vulnerability in Tailored Media Tailored Tools tailored-tools allows Cross Site Request Forgery.This issue affects Tailored Tools: from n/a through 3.0.3.	6.5	More Details
CVE-	The DNS stub resolver in subsys/net/lib/dns/resolve.c binds one UDP socket per configured server when the resolver context is initialised (dns_resolve_init_locked() -> register_dispatcher() -> dns_dispatcher_register()), with the local port left at 0 so the stack picks an ephemeral port exactly once. dns_write() then sends every subsequent query from that same socket without re-binding, so the source port is constant for the lifetime of the resolver context while only the 16-bit transaction identifier (ctx->queries[i].id, drawn with sys_rand16_get()) varies per query. RFC 5452 section 9.2 requires both to vary; with a fixed port, an off-path attacker has to guess only the identifier. The stable port is cheap to learn. Any single observed query reveals it permanently, and it can also be found by probing the ephemeral UDP range: subsys/net/ip/connection.c answers unmatched unicast UDP with an ICMP port-unreachable, so the resolver's socket is the one port in that range that stays silent. Acceptance of a reply is decided solely by get_slot_by_id() matching the transaction identifier and a CRC-16 of the queried name — the reply's source address is never checked (dispatcher_cb()) does		

2026-79981	ARG_UNUSED(addr)), so forged answers need not even be spoofed to appear to come from the configured server; they only need to reach the resolver's port. An attacker with UDP reachability to the device (same LAN, Wi-Fi or mesh segment, or a globally routable device address) can therefore flood answers for a predictable query name across the 16-bit identifier space during a query window and race the legitimate server. A win makes the application resolve an attacker-chosen address, and with CONFIG_DNS_RESOLVER_CACHE enabled the forged record is retained for its attacker-chosen TTL. The attack still requires winning that race, and applications that validate TLS certificates for the resolved endpoint are largely protected. The fix draws a fresh local port before each query to a server that has nothing outstanding: dns_randomize_source_port() unregisters the dispatcher, closes and reopens the socket, and re-registers it bound to port 0 so the system selects a new port, raising the per-query guessing space from roughly 16 to roughly 31 bits. Multicast DNS and LLMNR are excluded because their ports are protocol-fixed. The behaviour is controlled by CONFIG_DNS_RESOLVER_RANDOMIZE_SOURCE_PORT (default y).	6.5	More Details
CVE-2026-111837	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Wasiliy Strecker Contest Gallery contest-gallery allows Stored XSS.This issue affects Contest Gallery: from n/a through 34.0.1.	6.5	More Details
CVE-2026-114708	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Crocoblock JetSmartFilters jet-smart-filters allows Stored XSS.This issue affects JetSmartFilters: from n/a through 3.8.6.1.	6.5	More Details
CVE-2026-115508	Crawl4AI 0.7.0 before 0.9.5 contains an inefficient regular expression complexity vulnerability in the Docker API server that allows authenticated API clients to supply unbounded regular expressions in /crawl request bodies. Attackers can submit catastrophic-backtracking patterns via RegexExtractionStrategy.custom, RegexChunking.patterns or JsonCss, JsonXPath and JsonLxml schema regex fields against a crafted page to freeze the entire server, including /health.	6.5	More Details
CVE-2026-114399	Trudeusk through 1.2.11 contains a path traversal vulnerability that allows administrators to delete arbitrary files by supplying encoded traversal sequences to DELETE /api/v1/backup/:backup. Attackers can also pass unvalidated file values to POST /api/v1/backup/restore to restore archives from, or extract to, paths outside the intended directories.	6.5	More Details
CVE-2026-114659	Grav Flex Objects plugin 1.4.0 before 1.4.17 contains an authorization bypass because its generic /api/v1/flex-objects/pages API routes check only directory-level admin.pages permissions, without enforcing the Pages API's api.pages.read/write permissions or page-level permissions rules. Authenticated non-super users with api.access and admin.pages or admin.configuration.pages permissions can read page source and frontmatter, modify content, or delete pages they are denied.	6.5	More Details
CVE-2026-115679	Arm Keil MDK-Middleware 8.3.0 may allow out-of-bounds reads and writes to the FTP Server RX Buffer	6.5	More Details
CVE-2026-114088	Trigger.dev before 4.5.4 contains a hardcoded default secret vulnerability in the /coordinator Socket.IO namespace that allows unauthenticated attackers to connect using the published coordinator-secret value. Attackers can emit READY_FOR_EXECUTION to retrieve decrypted run environment variables or send TASK_RUN_FAILED_TO_RUN and similar events to tamper with arbitrary runs.	6.5	More Details
CVE-2026-114092	Trigger.dev before 4.5.0 performs no authentication or organization check in the action handler of the dashboard run replay route (resources.taskruns.\$runParam.replay) and looks up the run by friendlyId alone. A remote attacker who knows a victim's run friendlyId can replay it into an environment of their choosing, exposing the victim's task payload or running the task again with the victim's secrets.	6.5	More Details
CVE-2026-114095	Dulwich before 1.2.9 contains a denial of service vulnerability in Pack.resolve_object that allows attackers to cause infinite recursion via self-referential OFS_DELTA objects. Attackers can supply a malicious packfile with an OFS_DELTA delta_offset of 0 to exhaust CPU during clone, fetch, or cat-file operations.	6.5	More Details
CVE-2026-114508	Laravel is a web application framework. From 12.45.0 until 12.68.0 and 13.27.0, crafted relationship inclusion on a read-only endpoint can trigger a destructive database operation on a related table instead of retrieving data. A user with read access but no write permission can empty related-table data, causing irreversible data loss, denial of service, and loss of data integrity. This issue is fixed in versions 12.68.0 and 13.27.0.	6.5	More Details
CVE-2026-112825	Missing Authorization vulnerability in Magepeople inc. Car Rental Manager car-rental-manager allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Car Rental Manager: from n/a through 1.5.3.	6.5	More Details

CVE-2026-112821	Missing Authorization vulnerability in Knit Pay Knit Pay knit-pay allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Knit Pay: from n/a through 9.6.3.0.	6.5	More Details
CVE-2026-93850	Mattermost versions 11.9.x <= 11.9.1, 11.8.x <= 11.8.5, 11.7.x <= 11.7.10, 11.10.x <= 11.10.1 fail to bind the volume of text extracted from an uploaded PDF which allows an authenticated user to exhaust server memory and cause a denial of service via a crafted PDF whose compressed content stream expands into a very large amount of text.. Mattermost Advisory ID: MMSA-2026-00769	6.5	More Details
CVE-2026-115950	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Smart Widgets Better Post & Filter Widgets for Elementor better-post-filter-widgets-for-elementor allows Stored XSS.This issue affects Better Post & Filter Widgets for Elementor: from n/a through 1.9.1.	6.5	More Details
CVE-2026-93849	Mattermost versions 11.9.x <= 11.9.1, 11.8.x <= 11.8.5, 11.7.x <= 11.7.10, 11.10.x <= 11.10.1 Mattermost fail to apply allocation limits to the byte-string decode paths in its WebSocket msgpack decoder, which allows an authenticated remote attacker to cause excessive server memory allocation via a crafted binary WebSocket frame declaring a Bin32 payload length of up to 4,294,967,295 bytes.. Mattermost Advisory ID: MMSA-2026-00768	6.5	More Details
CVE-2026-92171	A flaw was found in the Infinispan Operator for Red Hat Data Grid. When cross-site (xsite) TLS is enabled, the operator launches the Gossip Router container with keystore and truststore passwords passed as plain-text command-line arguments (`-tls_keystore_password`, `-tls_truststore_password`). Container args are stored in the Pod spec and visible to any principal with `get pods` permission in the namespace (no `get secrets` required), to node-level processes via `/proc/<pid>/cmdline`, and frequently captured in audit/event logs.	6.5	More Details
CVE-2026-96643	Pillow is a Python imaging library. From 11.0.0 until 13.0.0, Exif.load() in src/PIL/Image.py repeatedly slices the remaining EXIF data while it begins with the six-byte Exif header. src/PIL/PngImagePlugin.py accepts an uncapped PNG eXif chunk, so a file filled with repeated headers causes a new shrinking copy on every iteration and quadratic CPU work. The condition is reachable through Image.getexif(), Image._getexif(), and ImageOps.exif_transpose(), including common photo-upload rotation workflows. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-96642	Pillow is a Python imaging library. Prior to 13.0.0, the TarIO initializer in src/PIL/TarIO.py accepts a negative octal member size and computes a backward seek that returns to the same 512-byte header. When the requested member name does not match, the parser repeatedly reads that header, consumes a full CPU core, and never returns. TarIO is reached only when an application directly uses PIL.TarIO.TarIO rather than the default Image.open() path. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-96641	Pillow is a Python imaging library. Prior to 13.0.0, PILfont text-width accumulation in src/_imaging.c checks positive glyph advances against INT_MAX but does not check negative advances against INT_MIN. Repeating a signed 16-bit advance of -32768 can therefore execute a signed C integer overflow in xsize += dx, producing a large positive width from a negative mathematical result. Applications that use getlength(), getsize(), or derived layout, quota, serialization, or allocation calculations can experience logic-integrity errors or resource consumption, while the high-level mask path commonly retains a decompression-bomb check. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-100145	The Synchrony Financing WordPress plugin before 1.1.2 does not properly authorise requests to one of its REST routes that modifies WooCommerce orders, allowing unauthenticated attackers to reassign any order to an account they control, tamper with its details, and read the associated customer's personal data.	6.5	More Details
CVE-2026-96638	Pillow is a Python imaging library. From 9.1.0 until 13.0.0, BmpRleDecoder.decode() in src/PIL/BmpImagePlugin.py permits crafted RLE BMP delta escapes and absolute-mode runs to append beyond the remaining dest_length without a per-branch cap. A delta can add roughly 255 image rows in one iteration, so a kilobyte-sized file with a large but permitted width can expand the data bytearray to hundreds of megabytes or more, causing memory-exhaustion denial of service. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-115680	Arm Keil MDK-Middleware 8.0.0-8.3.0 may allow DHCPv6 Option-Walk Over-Reads and Unintended Data Retransmission	6.5	More Details
CVE-2026-76289	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could make the Splunk Enterprise instance unavailable to other users by sending a request to the Representational State Transfer (REST) API endpoint for lookup table files. The affected instance requires a restart to restore service. The vulnerability is possible because the endpoint does not limit resource use when it processes response field selections. For more information see Endpoints reference list (https://help.splunk.com/en/splunk-enterprise/rest-api)	6.5	More Details

	reference/10.4/introduction/endpoints-reference-list) 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.		
CVE-2026-76292	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, and Splunk Secure Gateway versions below 3.10.12, 3.9.26, and 3.8.73, a low-privileged user that does not hold the "admin" Splunk role or the "sc_admin" Splunk Secure Gateway role could read mobile alert notification data for other users through Splunk Secure Gateway App Key Value Store (KV Store) Representational State Transfer (REST) API endpoints. The vulnerability is possible because the affected Splunk Secure Gateway App Key Value Store collections do not enforce user-level authorization for read requests. For more information see About the app key value store (https://help.splunk.com/en/splunk-enterprise/administer/admin-manual/10.2/administer-the-app-key-value-store/about-the-app-key-value-store) and About configuring role-based user access (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/about-configuring-role-based-user-access) in the Splunk documentation.	6.5	More Details
CVE-2026-114067	Trigger.dev 4.5.0 before 4.6.2 contains an improper authorization vulnerability in the PUT snapshot-url session endpoint that allows public session token holders to obtain presigned S3 write URLs. End users can overwrite their own session's durable chat snapshot to forge the agent's conversation history, since the route lacks a secret-key-only check.	6.5	More Details
CVE-2026-116175	vLLM from 0.19.0 before 0.24.0 contains a decompression bomb vulnerability because the AudioMediaIO multimodal path calls load_audio without the max_duration_s decoded-duration limit. Authenticated attackers can send small compressed audio via input_audio or audio_url in /v1/chat/completions requests that expands to huge decoded PCM, exhausting server memory.	6.5	More Details
CVE-2020-37286	Vyper from 0.1.0b4 before 0.3.0 contains a memory corruption vulnerability caused by an incorrect stack tip pointer when internal function calls appear inside tuples or nested calls. Contracts using such calls as tuple elements, call arguments, or builtin arguments like concat or ecrecover can return corrupted values or pass wrong arguments to external calls.	6.5	More Details
CVE-2026-76308	In Splunk SOAR versions below 8.7.0, an authenticated user could change the Enterprise Security edition through the Representational State Transfer (REST) API and cause Splunk SOAR to apply the license change. The vulnerability is possible because Splunk SOAR authenticates requests to manage the Enterprise Security edition but does not authorize the caller before it applies the license change. For more information see Using the REST API reference for Splunk SOAR (On-premises) (https://help.splunk.com/en/splunk-soar/soar-on-premises/rest-api-reference/8.7.0/using-the-splunk-soar-rest-api/using-the-rest-api-reference-for-splunk-soar-on-premises), Obtain and configure a Splunk SOAR (On-premises) license (https://help.splunk.com/en/splunk-soar/soar-on-premises/administer-soar-on-premises/8.7.0/configure-your-companys-settings-in-splunk-soar-on-premises/obtain-and-configure-a-splunk-soar-on-premises-license), and Integration of Splunk SOAR with Splunk Enterprise Security (https://help.splunk.com/en/splunk-enterprise-security-8/administer/8.7/automation-with-playbooks/integration-of-splunk-soar-with-splunk-enterprise-security) in the Splunk documentation.	6.5	More Details
CVE-2026-55190	Event Schedule is an open-source platform to share events, sell tickets and bring communities together. In versions 1.0.0 and prior, App\Utils\UriUtils::isUriSafe() is the SSRF guard applied to user-supplied import URLs. It can be bypassed two ways, both reaching internal/localhost services when the URL is later fetched by ImportCuratorEvents (Http::get(\$url) then the response body is parsed and imported as events, and file_get_contents(\$ogImage)). At time of publication, there are no publicly available patches.	6.5	More Details
CVE-2026-76305	In Splunk SOAR versions below 8.7.0, an authenticated user who does not hold a role with permission to view apps could download an archive of an installed app. This could expose all data in the app archive, including proprietary source code and sensitive information. The vulnerability is possible because Splunk SOAR does not verify permission to view apps before it returns the app archive. For more information see Export or import an app (https://help.splunk.com/en/splunk-soar/soar-on-premises/develop-apps/8.4.0/develop-an-app-using-the-splunk-soar-app-wizard/export-or-import-an-app) 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.	6.5	More Details
CVE-2021-48020	Vyper from 0.1.0b4 before 0.3.0 contains a memory corruption vulnerability caused by an incorrect stack tip pointer when internal function calls are performed inside arrays. Contracts placing internal calls in nested array elements, external call arguments, struct literals, or raw_log topics can have memory corrupted, producing incorrect contract behavior.	6.5	More Details

CVE-2025-71437	Astral uv before 0.8.22 contains a path traversal vulnerability in source distribution (sdist) tar extraction that allows a specially crafted sequence of symlinks to redirect file writes outside the intended extraction prefix. Attackers can publish a malicious sdist that, when installed by a victim, writes arbitrary files to filesystem locations outside the installation prefix.	6.5	More Details
CVE-2026-114041	FreeRDP before 3.32.1 on Windows contains a heap out-of-bounds read in client_cli_present_gateway_message that passes a byte count as wide-character printf precision. A malicious RD Gateway can send pre-authentication consent or service messages, causing the client to read twice the buffer size, leaking heap memory or crashing.	6.5	More Details
CVE-2026-114042	FreeRDP 3.32.0 contains a use-after-free vulnerability in the drdynvc channel that frees a pool-owned stream with the owned-stream deallocator when totalLength is smaller than the chunk size. A malicious RDP server can send a single well-formed Virtual Channel PDU to trigger the free and crash the client during teardown.	6.5	More Details
CVE-2026-114043	FreeRDP before 3.32.1 contains a heap-based out-of-bounds read in WinPR's clipboard PNG-to-CF_DIB synthesis because winpr_read_png_from_buffer() allocates a bit-packed buffer for PNGs with bit depth below 8, while winpr_bitmap_write_buffer() copies rows assuming one byte per pixel. A malicious RDP server can send a crafted low-bit-depth PNG clipboard payload that, when pasted by the victim, discloses megabytes of heap memory.	6.5	More Details
CVE-2026-116181	vLLM from 0.20.1 before 0.29.0 contains an improper input validation vulnerability that allows remote attackers to crash the engine by supplying out-of-vocabulary stop_token_ids. Attackers can send OpenAI chat, completion, generate, or batch requests with invalid IDs to terminate the engine or write outside the V2 runner's logits buffer.	6.5	More Details
CVE-2026-114046	FreeRDP 3.0.0 before 3.32.0 contains a null-pointer dereference in the smartcard emulation backend that accepts a null per-reader name in SCardGetStatusChange requests. A malicious RDP server can send a crafted reader-state entry with a null szReader, causing strcmp or _wcsncmp to crash the client process.	6.5	More Details
CVE-2026-114054	vLLM before 0.31.0 contains a cache aliasing vulnerability that allows authenticated attackers to collide shared cache keys by supplying caller-controlled values such as EXIF ImageID or cache_salt. Attackers can set a UUID EXIF ImageID to receive another request's cached processed media, or match cache_salt to a LoRA name to reuse first-block KV state.	6.5	More Details
CVE-2026-114065	Trigger.dev before 4.7.3 contains a tenant filtering bypass vulnerability that allows authenticated users to read other tenants' data by using placeholder table syntax in POST /api/v1/query. Attackers can submit queries with FROM {database.table} expressions, including within subqueries, to skip table allowlist and tenant filters and expose run metadata and payloads.	6.5	More Details
CVE-2026-114633	WWBN AVideo before 29.4.0 contains an incorrect authorization vulnerability that allows delegated Users Admin accounts to delete or modify administrator accounts. Attackers can submit requests to confirmDeleteUser.json.php or userAddNew.json.php targeting administrators to delete their videos and accounts, or change and demote them.	6.5	More Details
CVE-2026-114722	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a CSS scoping bypass in HTML messages. rcube_utils::mod_css_styles() did not treat an at-rule prologue such as @layer a; as a selector boundary, and did not handle selectors beginning with ::, so selectors following an at-rule were left unscoped and applied to the whole Roundcube interface instead of being confined to the message body. This allows an HTML message to restyle application chrome for UI redressing and spoofing. Upstream fixed this in the commit “Fix CSS scoping bypass in HTML messages using at-rules”.	6.5	More Details
CVE-2026-106932	The Moloni WordPress plugin before 5.0.14 does not perform capability or nonce checks on several of its authenticated AJAX actions, relying only on whether the site holds a stored connection to its external invoicing provider, allowing any authenticated user such as a Subscriber to generate fiscal invoices for arbitrary orders, discard arbitrary orders, and create or modify products and stock.	6.5	More Details
CVE-2026-96637	Pillow is a Python imaging library. Prior to 13.0.0, SpiderImageFile.seek() in src/PIL/SpiderImagePlugin.py reparses a later SPIDER stack frame and replaces self._size without rerunning Image._decompression_bomb_check(). A tiny file can use a safe first frame to pass Image.open() and then declare very large dimensions in a later mode F frame. Loading that frame allocates approximately four bytes per attacker-declared pixel and can consume multiple gigabytes of memory. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-95716	An authentication bypass in the PAPI protocol of AOS-10 allows an unauthenticated adjacent attacker to circumvent existing authentication controls. Successful exploitation allows an attacker to change the state of certain settings of a vulnerable system.	6.5	More Details
	Authentication tokens issued by the E3's web management interface on login are given an expiration		

CVE-2026-96616	time far beyond any reasonable session lifetime, effectively making them non-expiring. A token obtained through interception or device compromise remains valid for authenticated access for an extended period.	6.5	More Details
CVE-2026-17557	A flaw was found in the `Im-evaluation-harness` component of Red Hat OpenShift AI. An authenticated user, able to submit an evaluation job, could exploit this by providing a malicious model URL in the `JobSpec`. This vulnerability allows for the disclosure of the tenant's model authentication secret to an attacker-controlled external listener. Additionally, it enables Server-Side Request Forgery (SSRF) attacks against internal network endpoints reachable from the evaluation job pod.	6.5	More Details
CVE-2026-55606	odhcpd is the DHCPv6, DHCPv4, and router-advertisement server used by OpenWrt. In versions 25.12 and prior, odhcpd's DHCPv6 self-loop detection at dhcpv6.c:720 uses a network-byte-order DUID length value directly as the memcmp() size parameter. On little-endian hosts, this causes memcmp() to read far beyond the stack buffer — a single unauthenticated DHCPv6 packet triggers a 2560-byte overread from a 128-byte buffer. At time of publication, there are no publicly available patches.	6.5	More Details
CVE-2026-87943	Missing Authorization vulnerability in Drupal SAML SSO - Service Provider allows Forceful Browsing. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	6.5	More Details
CVE-2026-50106	When NLnet Labs Unbound up to and including version 1.26.1 is configured with both a root trust anchor and an additional 'trust-anchor:' for an island-of-security zone whose parent is provably insecure under the root chain, a malicious actor who can deliver one forged NXDOMAIN response for a name inside the island, with the NSEC/SOA RRSIG 'SignerName' set to the insecure parent, tricks Unbound's validator to use the root trust anchor instead, and consecutively the insecure parent, thus causing that record to be treated as insecure and denying existence for records in the trust island.	6.5	More Details
CVE-2026-116210	Webasyst Framework through 4.4.0 contains a cross-site request forgery vulnerability in the Photos app that allows remote attackers to perform backend actions because the app does not enable the framework's opt-in CSRF check. Attackers can lure logged-in Photos users to a malicious page that submits forged POST requests to photo/delete, album/save, photo/manageAccessSave or photo/saveField, deleting photos and changing access rights.	6.5	More Details
CVE-2026-115195	A broken access control vulnerability exists in the Mautic Stats API (/api/stats/{table}). The endpoint did not enforce permissions for several internal tables: an intended administrator-only restriction was never applied, and tables without an explicit permission mapping were readable by default. As a result, any authenticated API user, including one whose role grants no permissions, could read data from these tables, including the application audit log with actions performed by all users and change details that may contain contact personal data, stored visitor IP addresses and their geolocation details, and tracked redirect URLs and their channel mappings.	6.5	More Details
CVE-2026-115194	A path traversal vulnerability exists in the Mautic Asset API. When creating a local asset through the API, the supplied file path is not restricted to the asset upload directory. An authenticated API user with permission to create assets can register any file that is readable by the web server process as a Mautic asset. When the asset is published, its public download URL serves the contents of that file without authentication. This allows the attacker to read arbitrary files from the server, including Mautic configuration files containing the database password, secret key, and mail (SMTP) credentials, as well as system files and any other file readable by the web server user. Exposed credentials can be used to access the database or other connected services, and anyone with the public asset URL can retrieve the file contents.	6.5	More Details
CVE-2026-96636	Pillow is a Python imaging library. Prior to 13.0.0, MpoImageFile.seek() and DcxImageFile.seek() reparse later MPO or DCX frames and replace self._size without rerunning Image._decompression_bomb_check(). A crafted multi-frame image can use a tiny first frame to pass Image.open() and then declare an over-limit later frame. Normal frame iteration through seek(), load(), or ImageSequence.Iterator can decode that frame and commit hundreds of megabytes or more, causing memory-exhaustion denial of service. This issue is fixed in version 13.0.0.	6.5	More Details
CVE-2026-105947	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Innotim Software, Telecommunications, and Consulting Co. Ltd. Logsign SIEM allows Path Traversal. This issue affects Logsign SIEM: from 6.4.120 before 6.4.122.	6.5	More Details
CVE-2026-112311	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in ACF Extended ACF Extended acf-extended allows Path Traversal.This issue affects ACF Extended: from n/a through 0.9.2.7.	6.5	More Details
CVE-2026-115961	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in bPlugins B Slider b-slider allows Stored XSS.This issue affects B Slider: from n/a through 2.2.2.	6.5	More Details

CVE-2026-115956	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in bPlugins Panorama Viewer - 360 Degree Image + Video Viewer panorama allows Stored XSS.This issue affects Panorama Viewer - 360 Degree Image + Video Viewer: from n/a through 1.8.3.	6.5	More Details
CVE-2026-115955	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in bPlugins Advanced Post Block - Post Grid for WordPress block editor advanced-post-block allows Stored XSS.This issue affects Advanced Post Block - Post Grid for WordPress block editor: from n/a through 2.5.1.	6.5	More Details
CVE-2026-115953	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Dylan Kuhn Geo Mashup geo-mashup allows Stored XSS.This issue affects Geo Mashup: from n/a through 1.13.22.	6.5	More Details
CVE-2026-115951	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Kevin Weber Lazy Load for Videos lazy-load-for-videos allows Stored XSS.This issue affects Lazy Load for Videos: from n/a through 2.18.11.	6.5	More Details
CVE-2026-87937	URL Redirection to Untrusted Site ('Open Redirect') vulnerability in Drupal Central Authentication System (CAS) Server allows Fake the Source of Data. This issue affects Central Authentication System (CAS) Server versions: from 0.0.0 to 2.0.4, from 2.1.0 to 2.1.3.	6.5	More Details
CVE-2026-87936	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Drupal amaze.ai Private AI Provider allows SQL Injection. This issue affects amaze.ai Private AI Provider versions: from 0.0.0 to 1.3.7, from 1.4.0 to 1.4.3.	6.5	More Details
CVE-2026-115786	A mailbox quota-accounting vulnerability in Zimbra Collaboration Suite (ZCS) allows an authenticated user to consume another user's mailbox quota by repeatedly sharing and revoking a document. Creation of the recipient's shared-file record charges the document size to mailbox-wide usage. Revocation deletes the record and reverses its folder counters without reversing the mailbox-wide charge. Repeated normal share notifications therefore accumulate quota usage after the visible records are gone. On accounts with a finite quota, this can prevent incoming mail and other size-increasing operations without recipient interaction. The issue affects ZCS 10.1.16 and 10.1.20.	6.5	More Details
CVE-2026-84923	Authentication Bypass Using an Alternate Path or Channel vulnerability in Drupal Email Verification / SMS Verification / OTP Verification allows Authentication Bypass. This issue affects Email Verification / SMS Verification / OTP Verification versions: from 0.0.0 to 2.4.0.	6.5	More Details
CVE-2026-116516	Mattermost versions <=11.9 11.1.9 11.5.8 11.10.7.0 fail to validate Zoom chat line length before slicing which allows an authenticated attacker to crash the Agents plugin via summarizing a Zoom chat transcription containing short continuation lines. Mattermost Advisory ID: MMSA-2026-00723	6.5	More Details
CVE-2026-116073	Leon before commit 50df142 contains a path traversal vulnerability in getModelFileNameFromURL() that allows attackers to write files outside the models directory via percent-encoded separators. Attackers can trick a user into running '/download model' with a crafted URL containing encoded '../' sequences to write attacker-controlled content to arbitrary paths writable by Leon.	6.5	More Details
CVE-2026-111955	The WholesaleX WordPress plugin before 3.0.9 does not check that a stored account confirmation code is set before comparing it with the one supplied in a request, allowing unauthenticated users to activate pending accounts and bypass the administrator approval required before login.	6.5	More Details
CVE-2026-116070	Docker-Proxy through 5.1.9 contains a missing authentication vulnerability in the hubcmdui routes/config.js that allows unauthenticated attackers to read and overwrite configuration via /api/config. Attackers can retrieve the WeChat webhook URL and Telegram bot token, and overwrite menuitems links, proxyDomain, and the monitoring webhook destination receiving alerts.	6.5	More Details
CVE-2026-116069	NiubiGEO before 0.3.2 contains an unauthenticated server-side request forgery vulnerability that allows remote attackers to make the server fetch internal addresses by supplying a crafted domain value to POST /audit-plan. Attackers can submit IP literals, localhost, or private and link-local hosts to read loopback, internal-network, or cloud metadata responses returned in the plan's siteEvidence JSON.	6.5	More Details
CVE-2026-112845	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Jory Hogeveen Off-Canvas Sidebars & Menus (Slidebars) off-canvas-sidebars allows Stored XSS.This issue affects Off-Canvas Sidebars & Menus (Slidebars): from n/a through 0.5.9.	6.5	More Details
	baserCMS is a website development framework. In versions 5.0.0 before 5.3.2 and 5.4.0 before 5.4.2, an authenticated user with only theme-file view permission can traverse outside the theme directory through the optional BcThemeFile plugin's Admin web screen. ThemeFilesController::parseArgs() constructs a path containing literal ../ segments, and the view and edit actions pass it to ThemeFilesService::get() without containment validation. Authorization checks the unresolved request path, after which ThemeFile::_getContents() reads the resolved file and exposes its contents in the form		

CVE-2026-114623	textarea. A client must preserve the dot segments rather than normalize them away. Readable files are limited to .php, .css, .js, and .ctp; this can disclose application source and config/install.php database credentials in MySQL or PostgreSQL deployments, while SQLite configurations have an empty database password. Reading config/app_local.php may expose the Security.salt fallback, which is effective only when not overridden by environment configuration. The flaw does not disclose .env, .pem, or SQLite .db files, does not itself permit file writes or code execution, and does not affect the separate API getFullpath() route. The plugin is disabled by default and unauthenticated requests are not sufficient. This issue has been patched in versions 5.3.2 and 5.4.2.	6.5	More Details
CVE-2026-84917	Incorrect Authorization vulnerability in Drupal Jsonapi Role Access allows Forceful Browsing. This issue affects Jsonapi Role Access versions: from 0.0.0 to 2.0.2.	6.5	More Details
CVE-2026-116492	Frappe CRM app before 1.89.1 contains a missing authorization vulnerability that allows authenticated low-privileged users to read CRM Lead records they are not permitted to access by querying the get_contact_by_phone_number API, which performs raw database queries without invoking the permission system. Attackers can supply a known phone number to the crm.integrations.api.get_contact_by_phone_number endpoint to retrieve full lead details otherwise blocked by standard API and UI access controls, resulting in unauthorized disclosure of sensitive CRM data.	6.5	More Details
CVE-2026-87838	The Wava Payment Plugin WordPress plugin before 0.3.10 does not perform any authorisation, nonce or origin verification on one of its public REST endpoints, allowing unauthenticated attackers to overwrite the stored credential that identifies the merchant account the store's payments are created against, to force the integration into a verified state the site owner never completed, and to remove the payment method from the store's checkout.	6.5	More Details
CVE-2026-77156	Incomplete validation of extracted archive entry paths in ast-cli's tar extraction logic on Linux and macOS (used when installing the Vorpall ASCA real-time engine and the SCA Resolver) would allow a malicious tar archive containing absolute or traversal path entries to write files outside the install directory. Windows builds were not affected, as the Windows extractor performs this path validation. Exploitation required substituting the archive served at "download.checkmarx.com". An attacker able to compromise that server and replace the archive could write arbitrary, executable files to arbitrary filesystem paths via this path traversal. Extraction now validates that each entry's destination path resolves within the working directory and rejects absolute paths.	6.4	More Details
CVE-2026-113491	The GeoDirectory WordPress plugin before 2.8.194 does not validate a user-supplied image URL before using it as a local file path when saving a front-end listing, allowing users with Subscriber-level access and above to read files from outside the web root and to delete files on the server, limited to file names that carry an allowed media extension. A disclosed file is re-published at a public URL and its original is deleted as a side effect.	6.4	More Details
CVE-2026-77157	ast-cli downloads two helper binaries, the Vorpall ASCA real-time engine and the SCA Resolver, from "download.checkmarx.com" as part of its auto-update flow, triggered by the Checkmarx's "VSCode/JetBrains" extensions and by cx scan "sca-realtime". Prior to the fix, the downloaded archive was not verified against a checksum before extraction and execution. An attacker able to substitute the served archive could deliver an arbitrary binary that ast-cli would extract and execute the contents of the archive, since no integrity check was performed on the download. Successful exploitation requires compromising "download.checkmarx.com" and replacing the served archive with one containing an attacker-controlled binary. If the served archive had been replaced that way, any developer machine or CI pipeline invoking ASCA (via Checkmarx's VSCode/JetBrains extensions) or cx scan "sca-realtime" on an affected version would have extracted and executed that binary in place of the intended download. The download is now verified against a published SHA-256 checksum before extraction.	6.4	More Details
CVE-2026-108360	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. Prior to 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0, Worksheet\Drawing::setPath validates only the initial external-image URL with setIsWhitelisted when setAllowExternalImages is enabled. The HTTP stream context permits PHP to follow a redirect without applying the callback to the redirected URL, so a permitted image host can redirect the server to an internal address and cause a blind GET request. The response body is not returned to the attacker because the reader validates the downloaded content as an image and removes the temporary file, but internal GET-driven endpoints, network reachability, and cloud metadata endpoints may be reached. This issue is fixed in versions 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0.	6.4	More Details
CVE-2026-115845	An incomplete list of disallowed inputs in the Python code validator of the run_script tool in AWS MCP Server might allow an authenticated remote user to execute arbitrary code within their own isolated session environment via a crafted script. AWS fixed this issue service-side. No customer action is required.	6.4	More Details

CVE-2026-81812	The Elementor Website Builder WordPress plugin before 4.3.5 does not sanitize uploaded SVG files on every upload route, allowing users with the Author role and above to upload a file containing JavaScript, which will run in the browser of any user who opens it. Exploitation requires unfiltered file uploads to be enabled in the Elementor Website Builder WordPress plugin before 4.3.5 settings, which is not the default.	6.4	More Details
CVE-2026-102484	Heap-based buffer overflow in the JJY refclock driver in NTPsec ntpd before 1.2.6 allows an attacker who controls the Telephone JJY time source to corrupt memory or crash ntpd via more than six time replies.	6.4	More Details
CVE-2026-116206	YesWiki from 4.2.0 before 4.7.0 contains a stored cross-site scripting vulnerability that allows authenticated users to inject script into comment bodies using the wakka raw-HTML syntax. Attackers can post comments via the /api/comments endpoint that render unescaped through the Twig raw filter, executing JavaScript in administrators' and visitors' sessions.	6.4	More Details
CVE-2026-96376	Improper Neutralization of Special Elements used in an OS Command ("OS Command Injection") vulnerability in Drupal Cloud allows OS Command Injection. This issue affects Cloud versions: from 0.0.0 to 7.0.1.	6.4	More Details
CVE-2026-116109	A vulnerability has been found in continuedev Continue 2.0.0. This vulnerability affects the function MCPManagerSingleton.setConnections of the file core/config/ConfigHandler.ts of the component MCP Server Transport. Such manipulation leads to os command injection. The attack can be executed remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-114378	A vulnerability was detected in itsourcecode Online Cake Ordering System 1.0. This affects an unknown part of the file /deleteproduct.php. The manipulation of the argument ID results in sql injection. The attack may be launched remotely. The exploit is now public and may be used.	6.3	More Details
CVE-2026-116111	A vulnerability was determined in 2K Games Civilization V 1.0.3.144. Impacted is the function loadstring of the component Lua Sandbox. Executing a manipulation can lead to type confusion. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-115997	A flaw has been found in pig-mesh pig up to 4.1.0. This impacts the function orderItem.getColumn of the file pig-common/pig-common-data/src/main/java/com/pig4cloud/pig/common/data/resolver/SqlFilterArgumentResolver.java. Executing a manipulation of the argument ascs/descs can lead to sql injection. The attack may be launched remotely. The exploit has been published and may be used. Upgrading to version 4.1.1 will fix this issue. This patch is called 9b6311161e346d7626c81d993044072fed2dca07. It is suggested to upgrade the affected component. The vendor was contacted early about this disclosure.	6.3	More Details
CVE-2026-116258	A vulnerability was identified in libimobiledevice up to 1.4.0. This affects an unknown part of the file src/afc.c of the component AFC. Such manipulation leads to heap-based buffer overflow. The attack can be executed remotely. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-116044	A security vulnerability has been detected in prasathmani TinyFileManager up to 2.6. Affected by this issue is some unknown functionality of the file tinyfilemanager.php of the component Remote URL File Upload Module. The manipulation leads to unrestricted upload. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-116275	A vulnerability was found in AllenAI AllenNLP up to 2.10.1. This impacts the function Registrable.resolve_class_name of the file allennlp/common/registrable.py of the component Registrable Configuration Machinery. Performing a manipulation of the argument Type results in code injection. The attack may be initiated remotely. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	6.3	More Details
CVE-2026-114376	A weakness has been identified in leomoon-studios wiki-go up to 1.9.1. Affected by this vulnerability is the function SaveHandler/SourceHandler of the file internal/handlers/editor.go of the component API. Executing a manipulation can lead to path traversal. The attack can be launched remotely. The exploit has been made available to the public and could be used for attacks. Upgrading to version 1.9.2 addresses this issue. This patch is called 94f7e73ef8df7c58091d72d77a5e26e9fdad98b8. Upgrading the affected component is advised.	6.3	More Details
CVE-	A flaw was found in the iscsiui component of open-iscsi. The DHCPv6 option parser calls strlen() on raw network data without ensuring a NUL terminator is present, causing it to read past the buffer into heap		More

2026-115485	memory. This can result in a daemon crash, a 1-byte heap overwrite, or corruption of the iSCSI initiator name. An attacker on the same network segment can trigger this by sending a crafted DHCPv6 Advertise message.	6.3	Details
CVE-2026-116098	A vulnerability was detected in SteelRaven7 Ravenfield 25407073. This vulnerability affects the function cacheCoverPoints of the component RFL Custom Map Handler/RFLD Custom Map Handler. Performing a manipulation results in deserialization. Remote exploitation of the attack is possible. The patch is named EA39 Security Patch. Applying a patch is the recommended action to fix this issue.	6.3	More Details
CVE-2026-54631	Part-DB is an open source inventory management system for electronic components. Prior to version 2.12.1, the file .env committed in the Part-DB repository contains a fixed value for APP_SECRET (a03498528f5a5fc089273ec9ae5b2849). The official Docker image jbttronics/part-db1 does not override or randomize this value at container start-up. Because Symfony uses APP_SECRET (via kernel.secret) as the HMAC key for signing remember-me cookies, any attacker who knows the secret -- or fetches it from the public repository -- can forge an arbitrary-user remember-me cookie and gain a fully authenticated session with no credentials required. This issue has been patched in version 2.12.1.	6.3	More Details
CVE-2026-116099	A flaw has been found in Fusion Creators Studio SCP Containment Breach Multiplayer up to 1.295. This issue affects the function BlitzBanks of the component Skynet++ Virtual Machine. Executing a manipulation can lead to out-of-bounds write. The attack can be executed remotely. The exploit has been published and may be used. Upgrading to version 1.3R is capable of addressing this issue. It is recommended to upgrade the affected component.	6.3	More Details
CVE-2026-114039	A weakness has been identified in projectworlds Advanced Library Management System 1.0. Impacted is the function mysqli_query of the file report_search.php. Executing a manipulation of the argument dateto/datefrom can lead to sql injection. The attack can be executed remotely. The exploit has been made available to the public and could be used for attacks.	6.3	More Details
CVE-2026-116101	A vulnerability was found in itsourcecode Online Cake Ordering System 1.0. The affected element is an unknown function of the file /deletetransac.php. The manipulation of the argument ID results in sql injection. The attack may be performed from remote. The exploit has been made public and could be used.	6.3	More Details
CVE-2026-115373	Incus before 6.0.7, 6.1.0 before 7.0.2, and 7.1.0 before 7.5.0 contains a missing authorization vulnerability in GET /1.0/operations/{uuid} that lets project-restricted clients read operations in other projects. Attackers holding a restricted certificate and a known operation UUID can retrieve command lines, environment secrets, websocket secrets, and trust enrolment tokens.	6.3	More Details
CVE-2026-114375	A security flaw has been discovered in leomoon-studios wiki-go up to 1.9.1. Affected is the function RenameFileHandler of the file internal/handlers/files.go of the component Files Rename API. Performing a manipulation of the argument currentPath/newName results in path traversal. The attack can be initiated remotely. The exploit has been released to the public and may be used for attacks. Upgrading to version 1.9.2 is able to address this issue. The patch is named 81cd12e248314a52605f0f39718c8a76272eac29. It is recommended to upgrade the affected component.	6.3	More Details
CVE-2026-113743	A weakness has been identified in Tenstorrent tt-metal up to 0.80.0. The affected element is the function load_hf_state_dict_filtered of the file models/tt_transformers/tt/load_checkpoints.py. This manipulation of the argument weight_map causes path traversal. The attack may be initiated remotely. The exploit has been made available to the public and could be used for attacks. Upgrading to version 0.81.0-dev20261016 is sufficient to fix this issue. Patch name: 4851ccf347f6903c3d819eed60705600386e52c3. It is suggested to upgrade the affected component.	6.3	More Details
CVE-2026-114362	A flaw has been found in HarnessRouter up to 0.31.12. This impacts an unknown function of the file gateway/app.py of the component Targets Test Endpoint. Executing a manipulation of the argument base_url can lead to server-side request forgery. The attack may be performed from remote. Upgrading to version 0.31.13 will fix this issue. This patch is called 1c4b1fc1f89737dabd089ac1d1c174a57f764e28. It is suggested to upgrade the affected component.	6.3	More Details
CVE-2026-116124	A security vulnerability has been detected in GNOME libgsf up to 1.14.60. The impacted element is the function ole_info_read_metabat of the file gsf/gsf-infile-msole.c of the component OLE2 Parser. The manipulation leads to heap-based buffer overflow. The attack can be initiated remotely. The identifier of the patch is a8e25560b74f84c3f86df6ca0ac59195fa5af68a. It is recommended to apply a patch to fix this issue.	6.3	More Details
CVE-2026-116058	A vulnerability was found in Alfresco Activiti up to 7.21.0-rc.692. Affected by this issue is the function JuelScriptEngine.eval of the file JuelScriptEngine.java. Performing a manipulation results in code injection. The attack can be initiated remotely. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details

CVE-2026-116057	A vulnerability has been found in KodCloud KodExplorer up to 4.55. Affected by this vulnerability is an unknown functionality of the file app/controller/editor.class.php of the component Online Editor Save Endpoint. Such manipulation of the argument path/filestr leads to unrestricted upload. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-78119	The Nuvoton NPCX eSPI driver parses Target-Attached-Flash (TAF) requests in espi_taf_parse() (drivers/espi/espi_npcx.c), taking the 12-bit length field directly out of the request header in the FLASHRXBUF FIFO. For write and RPMC OP1 requests it copied DIV_ROUND_UP(len, 4) words into struct espi_taf_pkt.src, a fixed 16-word (64-byte) member of a static structure, without ever comparing the length against that buffer or against the negotiated eSPI receive payload size (rx-plsize, 64 bytes on all in-tree NPCX/NPCK SoCs). Both the loop counter and the word count were uint8_t, so the copy ran up to 255 iterations. An eSPI controller that emits a flash request whose header claims a payload larger than the negotiated maximum therefore drives up to 239 32-bit stores past the end of the structure — roughly 956 bytes of out-of-bounds write into .bss — from the eSPI flash receive interrupt handler, while the source reads walk past the end of the 18-word FLASHRXBUF register array into neighbouring eSPI registers. The same unvalidated length was then passed on to flash_write() in espi_taf_npcx_flash_write() with the 64-byte src buffer as source, reading up to 4 KB of adjacent controller SRAM into SPI flash at a host-chosen address, where the host can read it back. A related fixed-size stack buffer in taf_npcx_completion_handler() was overflowed by four bytes on the maximum-size completion path on SoCs whose tx-plsize is 64. The defect is reachable only from the eSPI host side, and only by a request that violates the negotiated payload size, so it requires host firmware or bus-level control rather than ordinary host software; a compliant maximum-length write fits the buffer exactly and is handled correctly. Because TAF exists specifically so that the embedded controller arbitrates host access to shared flash (espi_taf_check_read_protect() / espi_taf_check_write_protect()), corrupting controller memory from a host-supplied field crosses the trust boundary the feature asserts; consequences range from controller crash or hang to corruption of driver state that gates the flash protection regions. The fix widens the counters, rejects any request whose word count exceeds both ARRAY_SIZE(FLASHRXBUF) - NPCX_ESPI_FLASH_HEADER_ADDRESS_LEN and NPCX_TAF_RX_PAYLOAD_WORDS (derived from the devicetree rx-plsize), marks such packets with a new invalid_flag so espi_taf_work() answers them with an unsuccessful completion instead of processing them, bounds the transmit payload against FLASHTXBUF, and adds a default case for unsupported cycle types. Only configurations with CONFIG_ESPI_TAF enabled on an NPCX/NPCK target are affected; no in-tree board enables it by default.	6.3	More Details
CVE-2026-116230	A vulnerability was identified in Ndemic Creations Plague Inc Evolved up to 1.24.0.3. The affected element is an unknown function of the component Workshop Scenario Loader. Such manipulation leads to deserialization. The attack may be performed from remote. Upgrading to version 1.24.0.4 is sufficient to fix this issue. It is recommended to upgrade the affected component.	6.3	More Details
CVE-2026-114266	A security vulnerability has been detected in ProjectWorlds Advanced Library Management System 1.0. This impacts the function mysqli_query of the file profile.php. The manipulation of the argument admin_id leads to sql injection. The attack is possible to be carried out remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2026-116055	A vulnerability was detected in SourceCodester Performance Indicator System 1.0. This impacts an unknown function of the file /opiis/admin/deleteproduct.php. The manipulation of the argument ID results in sql injection. The attack may be performed from remote. The exploit is now public and may be used.	6.3	More Details
CVE-2026-114250	A flaw has been found in itsourcecode Sales-and-Inventory-System up to 0a3c45a78aaff3601b8cf9ce8ba82ff5c345caac. This issue affects some unknown processing of the file /pages/cust_edit.php. This manipulation of the argument ID causes sql injection. The attack can be initiated remotely. The exploit has been published and may be used. Continuous delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-114251	A vulnerability has been found in itsourcecode Sales-and-Inventory-System up to 0a3c45a78aaff3601b8cf9ce8ba82ff5c345caac. Impacted is an unknown function of the file /pages/cust_del.php. Such manipulation of the argument ID leads to sql injection. The attack can be launched remotely. The exploit has been disclosed to the public and may be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-	A vulnerability was found in itsourcecode Sales-and-Inventory-System up to 0a3c45a78aaff3601b8cf9ce8ba82ff5c345caac. The affected element is an unknown function of the file /pages/cust_searchfrm.php. Performing a manipulation of the argument ID results in sql injection. The		More

2026-114252	attack may be initiated remotely. The exploit has been made public and could 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 project was informed of the problem early through an issue report but has not responded yet.	6.3	Details
CVE-2026-114253	A vulnerability was determined in oretnom23 file-management-system-php up to 2e3c1ebe60f121e8df4fb3f544f0e9591d9a01e4. The impacted element is the function include of the file index.php. Executing a manipulation of the argument page can lead to file inclusion. The attack may be launched remotely. The exploit has been publicly disclosed and may be utilized. This product operates on a rolling release basis, ensuring continuous delivery. Consequently, there are no version details for either affected or updated releases. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-116102	A vulnerability was determined in QuantGeekDev docker-mcp 0.1.0/0.2.0. The impacted element is the function _build_windows_command of the file docker_executor.py. This manipulation of the argument project_name causes os command injection. It is possible to initiate the attack remotely.	6.3	More Details
CVE-2026-114380	A vulnerability was found in itsourcecode Online Cake Ordering System 1.0. Impacted is an unknown function of the file /inserttrans.php. Performing a manipulation of the argument prcd results in sql injection. The attack is possible to be carried out remotely. The exploit has been made public and could be used.	6.3	More Details
CVE-2026-116056	A flaw has been found in KodCloud KodExplorer up to 4.56. Affected is the function request_url_safe of the file app/function/web.function.php of the component Remote Download. This manipulation causes server-side request forgery. It is possible to initiate the attack remotely. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-116103	A vulnerability was identified in Sony Neural Network Libraries up to 1.39.0. This affects the function load_param_in_txt of the file python/src/nnabla/utlis/cli/encode_decode_param.py of the component Encode Param Command Handler. Such manipulation leads to code injection. It is possible to launch the attack 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. This vulnerability only affects products that are no longer supported by the maintainer.	6.3	More Details
CVE-2026-116118	A vulnerability has been found in Z.ai Zcode up to 3.14.3/3.15.1. Affected by this issue is some unknown functionality of the component Sed Command Filter. The manipulation leads to os command injection. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used.	6.3	More Details
CVE-2026-116367	A vulnerability was found in DeepSeek Harness up to 0.1.0-rc.5. This vulnerability affects the function createContext of the file packages/extensions/cordis-host-runner/src/sandbox.ts. The manipulation results in sandbox issue. It is possible to launch the attack remotely. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-116369	A vulnerability was identified in Sunwood-ai-labs github-kanban-mcp-server up to 0.4.0. Impacted is the function handleListIssues of the file src/handlers/issue-handlers.ts. Such manipulation of the argument state/labels leads to os command injection. The attack can be launched remotely. Upgrading to version 0.4.1 is recommended to address this issue. The name of the patch is 253d93bbdf7ee9fe42bd8fdf15fee93ae59cbae. It is suggested to upgrade the affected component.	6.3	More Details
CVE-2026-114265	A weakness has been identified in ProjectWorlds Advanced Library Management System 1.0. This affects the function mysqli_query of the file returned_book_search.php of the component Search Form. Executing a manipulation of the argument datefrom/dateto can lead to sql injection. The attack can be executed remotely. The exploit has been made available to the public and could be used for attacks.	6.3	More Details
CVE-2026-114660	Grav from 2.0.0 before 2.2.5 contains an improper access control vulnerability in FlexDirectory::dynamicFlexField() that allows page authors to call arbitrary Flex object methods via flex-*@ directives in page frontmatter. When a user with page edit rights opens the page, directives calling setStorageKey() and delete() run in their session, deleting protected pages neither account may delete.	6.3	More Details
CVE-2026-113738	A vulnerability has been found in R6410418 Jackrong-llm-finetuning-guide up to ef2b17f2798cd18cd85375080c35a58397ad0d09. This vulnerability affects the function extract_mtp_heads of the file qwen-mtp-gguf/scripts/qwen_mtp_gguf_pipeline.py of the component MTP weight_map Handler. The manipulation of the argument shard_name leads to path traversal. It is possible to initiate the attack remotely. The exploit has been disclosed to the public and may be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
	A security flaw has been discovered in GouGuOA up to 6.1.2. The impacted element is the function		

CVE-2026-116113	chunkUpload of the file app/disk/controller/Api.php. The manipulation of the argument file_extension results in unrestricted upload. It is possible to launch the attack remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2026-113740	A flaw has been found in itsourcecode Online Cake Ordering System 1.0. Affected by this vulnerability is an unknown functionality of the file /deletecustomer.php. This manipulation of the argument ID causes sql injection. The attack is possible to be carried out remotely. The exploit has been published and may be used.	6.3	More Details
CVE-2026-114381	A vulnerability was determined in itsourcecode Online Cake Ordering System 1.0. The affected element is an unknown function of the file /orderdetail.php. Executing a manipulation of the argument ID can lead to sql injection. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized.	6.3	More Details
CVE-2026-116104	A security flaw has been discovered in Sony Neural Network Libraries up to 1.39.0. This impacts the function func_set_nnabla_version_decorate of the file python/src/nnabla/utlils/converter/utlils.py of the component NNP Converter. Performing a manipulation of the argument function.type results in code injection. The attack can be initiated remotely. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way. This vulnerability only affects products that are no longer supported by the maintainer.	6.3	More Details
CVE-2026-114384	A security vulnerability has been detected in itsourcecode Online Cake Ordering System 1.0. Affected is an unknown function of the file /transactionsave.php. Such manipulation of the argument del leads to sql injection. The attack can be launched remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2026-115388	wasmtime-wasi-http in Wasmtime before 36.0.16, 37.0.0 before 48.0.3, and 49.0.0 before 49.0.1 does not enforce the wasi:io/streams check-write allowance for WASIp2 outgoing HTTP bodies, so oversized writes succeed instead of trapping. A malicious guest can pass large buffers to write on outgoing-body resources, holding host allocations for the store's lifetime and exhausting host memory.	6.2	More Details
CVE-2026-115387	Wasmtime before 36.0.16, 37.0.0 before 48.0.3, and 49.0.0 before 49.0.1 contains a denial of service vulnerability in wasmtime-wasi that allows guest WASI programs to panic the host via overflowing filesystem timestamps. Attackers can pass atim or mtim values whose sub-second field overflows the seconds field to path_filestat_set_times, fd_filestat_set_times, set-times, or set-times-at, crashing the host.	6.2	More Details
CVE-2026-115385	Wasmtime 46.0.0 before 48.0.4 and 49.0.0 before 49.0.2 contains a denial of service vulnerability in the wasmtime-wasi WASIp3 implementation that allows guests to abort the Unix host process. Guests with mutable preopened directory access can call set-times or set-times-at with negative seconds timestamps before the epoch, triggering an assertion failure that crashes the host.	6.2	More Details
CVE-2026-92823	A flaw was found in libmodulemd. In builds with rpmio support enabled, compressed modulemd YAML is decompressed through a custom read handler that does not enforce a cumulative decompressed-size limit before parsing. An attacker who can supply a highly compressible compressed YAML file can cause excessive CPU and memory consumption, resulting in a denial of service.	6.2	More Details
CVE-2026-92824	A flaw was found in libmodulemd. In the module stream v2 rpm-map emission path, an inner loop incorrectly uses the count of digest algorithms as its bound instead of the number of checksums in the current digest table. Crafted modulemd v2 metadata with multiple digest types can trigger an out-of-bounds read during YAML serialization, typically causing a crash and resulting in a denial of service.	6.2	More Details
CVE-2026-114129	stream-json is a micro-library of stream components for processing JSON and JSONC with a minimal memory footprint. From 2.0.0 until 3.7.0, FlexAssembler compiles string and regular-expression path filters into matchers that join the full key stack on every container start. Because the key stack grows with unbounded object or array nesting, a document nested to depth D causes D joins whose cost grows with the current depth, resulting in quadratic CPU consumption on the single-threaded event loop. The join occurs before the rule match decision, so the attacker does not need to know the target schema. Consumers using only function filters or no rules are not affected. A crafted nested document can stall all concurrent work in a process that feeds it to an affected FlexAssembler. This issue is fixed in 3.7.0, which adds a default maximum depth.	6.2	More Details
CVE-2026-114222	NT OneAPP developed by NewType Infotech has an Authentication Bypass vulnerability. Unauthenticated local attackers can bypass the App PIN lock mechanism via a Deep Link, thereby accessing protected resources of a logged-in user.	6.2	More Details
CVE-2026-	Adm-zip before 0.6.1 contains an incorrect permission assignment vulnerability that allows attackers to set setuid, setgid, or sticky bits on extracted files when keepOriginalPermission is enabled. Attackers can craft zip external file attributes that pass unmasked into chmod, letting attacker-controlled	6.2	More Details

114105	executables gain elevated privileges when extracted as root and later run.		
CVE-2026-92822	A flaw was found in libmodulemd. When parsing module metadata, the xmd field is converted into nested GVariant values through recursive YAML mapping and sequence helpers without a recursion-depth or node-count limit. An attacker who can supply crafted modulemd YAML with deeply nested xmd content can cause stack exhaustion or excessive resource consumption, resulting in a denial of service.	6.2	More Details
CVE-2026-116168	SiYuan before 3.8.7 contains a path traversal vulnerability in the MCP asset tool because validateAssetUploadPaths relies on an incomplete sensitive-path blacklist that omits .env and other secret files. Authenticated attackers driving the MCP asset upload action can copy such local files into the workspace and read them back via /assets/ over HTTP.	6.2	More Details
CVE-2026-115350	Mistune before 3.3.0 contains a cross-site scripting vulnerability in the Include directive that emits included .html, .xhtml, or .htm file content as raw block_html tokens. Attackers who can plant an HTML file readable by the renderer can reference it via an include directive, combined with path traversal, to execute stored scripts when escape=False is set.	6.1	More Details
CVE-2024-5139	The Authentication Endpoint accepts user-supplied input without sufficient encoding, allowing for the injection of malicious JavaScript payloads. This vulnerability can be triggered by submitting crafted input to the Authentication Endpoint. Successful exploitation allows an attacker to redirect the user's browser to a malicious website, modify the web page's user interface, or extract sensitive information from the browser. However, due to the `httpOnly` flag on session-related cookies, session hijacking is not possible.	6.1	More Details
CVE-2025-69723	A reflected cross-site scripting (XSS) vulnerability in localstack v4.9.2 allows attackers to execute arbitrary Javascript in the context of the user's browser via a crafted URL.	6.1	More Details
CVE-2026-113787	odf-kit allows users to generate, fill, read, and convert OpenDocument Format files (.odt, .ods) in TypeScript and JavaScript. Starting in version 0.6.0 and prior to version 0.14.2, the renderer has an `escapeHtml()` helper and applies it consistently to text content and to attribute values — `href`, `src`, `alt`, `id`, `data-author`, `data-date`, citation text all go through it. The inline CSS strings do not. In `renderTextSpan` (`dist/odt/to-html/html-renderer.js:195`), the value of e.g. `fo:color` is copied from the document into `style.fontColor` verbatim by the parser (`extractSpanStyle`, read/parser.js:348), then concatenated raw into the attribute. A ` ` in the value terminates the attribute early. The same unescaped `style="{...}"` pattern appears at roughly ten other emission sites: paragraph and heading (`paragraphStyleToCss`), table cells and rows (`cellStyleToCss`, `rowStyleToCss`), ` <col>`, (via="" 0.14.2="" `fo:background-color`,="" `fo:border*`,="" `fo:color`,="" `fo:font-family`="" `fo:letter-spacing`,="" `fo:line-height`,="" `fo:margin-*`,="" `fo:padding-*`,="" `fo:text-shadow`,="" `style:vertical-align`.="" a="" affected="" and="" attribute.="" attributes="" both="" contains="" direct="" font-face="" image="" include="" patch.<="" resolution),="" source="" style="" td="" the="" version=""><td>6.1</td><td>More Details</td></col>`,>	6.1	More Details
CVE-2026-114407	e107 CMS before 2.3.13 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject malicious HTML by manipulating the author GET parameter in search.php, which is parsed with parse_str() and concatenated unescaped into a rendered input tag via the form handler's 'other' option key. Attackers can craft a single GET request containing a crafted author parameter to break out of the HTML input tag and inject an img element with an onerror handler, executing arbitrary JavaScript in the victim's browser session.	6.1	More Details
CVE-2026-115722	RPLY through 1.1.4 contains an improper array index validation vulnerability in ply_read_header_property() that accepts list as a list property's length or value type. Attackers can supply a crafted PLY file that makes ply_read_list_property() call an out-of-bounds handler, crashing ASCII parsing or corrupting memory in binary parsing.	6.1	More Details
CVE-2026-93638	A reflected cross-site scripting vulnerability in the login page of Zimbra Collaboration Suite (ZCS) allows an unauthenticated attacker to execute JavaScript through the zrememberme parameter. A crafted relogin URL selects a two-factor error display that inserts the parameter into a hidden input attribute without HTML encoding. A victim who opens the URL receives attacker-controlled script and page content under the legitimate Zimbra origin, allowing credential capture through a replacement login form. The relogin path logs out an existing session before rendering the payload, so the demonstrated attack does not inherit that session's mailbox access. ZCS 10.1.16 and 10.1.17 are affected.	6.1	More Details
CVE-2026-85518	A flaw was found in OpenSCAP. When decomposing a Source DataStream (SDS) file, the application does not properly validate relative file path attributes, leading to a path traversal vulnerability. An attacker can exploit this flaw by tricking a user or automated system into decomposing a crafted datastream containing directory traversal sequences. Successful exploitation allows unauthorized file creation or modification outside the intended target directory with the permissions of the executing process.	6.1	More Details
	Discourse 1.9.0.beta2 before 1.9.0.beta13 contains a reflected cross-site scripting vulnerability that		

CVE-2017-20290	allows unauthenticated attackers to inject scripts through the email parameter of /email/unsubscribed. Attackers can craft GET requests with SVG onload payloads, which unsubscribed.html.erb renders unescaped via html_safe, executing JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-19.	6.1	More Details
CVE-2024-58394	Retool before 3.88.0 contains a cross-site scripting vulnerability that allows unauthenticated attackers to execute scripts by proxying malicious SVG files through /api/imageProxy. Attackers can supply an attacker-hosted SVG in the url parameter, which is served from the Retool origin without a Content Security Policy, executing JavaScript in victims' sessions. Exploitation evidence was first observed by the Shadowserver Foundation on 2024-10-15.	6.1	More Details
CVE-2020-37284	Sickbeard through build-507 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the pattern parameter of /config/postProcessing/testNaming. Attackers can craft GET requests with malicious SVG onload payloads that the testNaming handler returns unescaped, executing JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-22.	6.1	More Details
CVE-2026-114404	PHPOffice PHPPresentation 1.2.0 contains a cross-site scripting vulnerability in the HTML writer that allows attackers who supply presentation content to inject script into generated HTML. Attackers can craft a PPTX with malicious text runs, hyperlink relationship targets, or shape descr attributes that are emitted unescaped, executing script in viewers' browsers.	6.1	More Details
CVE-2026-116195	Rclone before 1.75.2 contains a path traversal vulnerability in the local backend's translateLink function, which trims the .rclonelink suffix from a path after the containment check, so a name can resolve outside the destination root. Attackers can plant a directory named ...rclonelink in a source tree synced with --links and --metadata to change ownership, timestamps, and xattrs of the destination's parent directory.	6.1	More Details
CVE-2026-116217	RosarioSIS before 12.9.2 contains a reflected cross-site scripting vulnerability in modules/Food_Service/MenuReports.php that allows unauthenticated attackers to inject attributes via unescaped menu_id and modname parameters. Attackers can lure a logged-in administrator to a crafted Modules.php link whose double-quoted menu_id injects an onmouseover handler, stealing session and CSRF tokens to act as the administrator.	6.1	More Details
CVE-2017-20288	QCCubed 2.2.0 through 3.1.1 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject scripts through the installation_path parameter. Attackers can craft GET requests to /assets/php/_devtools/installer/step_2.php with malicious installation_path values echoed unescaped, executing JavaScript in victims' browsers. Exploitation evidence was first observed by the Shadowserver Foundation on 2023-10-20.	6.1	More Details
CVE-2026-116203	YesWiki before 4.7.0 contains a stored cross-site scripting vulnerability in the filtertags action, which echoes attacker-controlled token values without output encoding. Anonymous attackers can inject markup via action-attribute parameters that bypasses the allow_raw_html gate, executing JavaScript for every page viewer.	6.1	More Details
CVE-2023-54413	WPML Multilingual CMS plugin for WordPress before 4.6.1 contains a reflected cross-site scripting vulnerability that allows unauthenticated attackers to inject script by supplying crafted input in the wp_lang parameter. Attackers can lure users to /wp-login.php with a wp_lang value injecting an image input with an obfuscated onfocus handler to execute JavaScript in the victim's browser. Exploitation evidence was first observed by the Shadowserver Foundation on 2025-11-12.	6.1	More Details
CVE-2026-111978	The GiveWP WordPress plugin before 4.19.0.1 does not properly validate a user-supplied filename before reading it from disk, allowing users with a high-privilege custom role to perform path traversal, disclose the contents of JSON-parseable files anywhere on the server, and overwrite the GiveWP WordPress plugin before 4.19.0.1's stored configuration.	6.0	More Details
CVE-2026-116174	vLLM from 0.22.0 before 0.24.0 contains an improper input validation vulnerability in the Rust gRPC Generate and GenerateStream service that forwards unchecked prompt token IDs to EngineCore. Attackers can submit requests with out-of-vocabulary prompt.token_ids values to trigger a fatal embedding lookup failure that terminates EngineCore and leaves the server unhealthy.	5.9	More Details
CVE-2020-37287	Vyper before 0.2.6 contains a call stack corruption vulnerability that causes incorrect function outputs when complex types with non-base members are passed as arguments. Contracts passing multi-dimensional arrays or structs to internal or external functions receive corrupted values because the compiler assumed each subtype size is 32.	5.9	More Details
CVE-2026-116541	Netty netty-tcnative before 2.0.84.Final contains a denial of service vulnerability in the tcn_write_to_bytebuffer BIO write callback that overreports bytes consumed. Remote TLS 1.3 peers can provoke queued post-handshake records like NewSessionTicket or KeyUpdate to truncate flights, fail connections, or abort the JVM process.	5.9	More Details

CVE-2026-116539	Netty netty-tcnative before 2.0.84.Final contains a heap buffer overflow vulnerability in the tcn_write_to_bytebuffer BIO callback caused by failing to wrap the non-application ring buffer write index. Remote TLS peers can drive handshake, renegotiation, or post-handshake messages into this buffer, triggering an out-of-bounds memcopy that crashes the JVM process.	5.9	More Details
CVE-2026-116173	vLLM from 0.22.0 before 0.24.0 contains a denial of service vulnerability that allows remote attackers to terminate EngineCore by bypassing logprob_token_ids limits in the Rust gRPC frontend. Attackers can send Generate or GenerateStream requests with output_logprobs enabled and more than 128 output_candidates.token_ids, triggering a fatal sampler exception that crashes the engine.	5.9	More Details
CVE-2026-116172	vLLM from 0.22.0 before 0.24.0 contains an improper input validation vulnerability in the Rust gRPC converter that allows remote attackers to bypass max_logprobs limits via CandidateTokens.top_n. Attackers can send Generate or GenerateStream requests with oversized or integer-wrapping top_n values to trigger torch.topk failures that terminate EngineCore.	5.9	More Details
CVE-2026-116540	netty-tcnative 2.0.66.Final before 2.0.84.Final built against OpenSSL 3.x contains a use-after-free vulnerability in ssl_tlsexch_ticket_key_cb during session ticket key rotation. Remote clients performing TLS handshakes while setSessionTicketKeys0 frees the shared key array cause OpenSSL to read freed HMAC keys, crashing or corrupting ticket processing.	5.9	More Details
CVE-2026-76406	In Splunk SOAR versions below 8.7.0, an unauthenticated user with access to encrypted credential data from an affected Splunk SOAR instance could submit Basic authentication data to the Representational State Transfer (REST) API to expose and decrypt that encrypted credential data. The vulnerability is possible because Splunk SOAR decrypts encrypted Basic authentication values before completing authentication checks. The attack requires the unauthenticated user to have encrypted credential data from an affected Splunk SOAR instance. The unauthenticated user should not be able to exploit the vulnerability at will. For more information see 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.	5.9	More Details
CVE-2026-114644	The MongoDB C# Driver's BSON library contains an uncontrolled resource consumption weakness in its deserialization handling of one BSON value type. An application that deserializes untrusted content with the library retains data derived from that content for the lifetime of the process, with no limit on how much accumulates. An unauthenticated remote party who can cause such content to be deserialized can drive unbounded memory growth, exhausting the host application's available memory and denying service. There is no impact to confidentiality or integrity, and no MongoDB server component is affected.	5.9	More Details
CVE-2026-111956	The WholesaleX WordPress plugin before 3.0.9 does not protect or remove the user-import file it keeps in a publicly accessible location under a predictable name, allowing unauthenticated users to download the most recently imported user data, including any account passwords and customer contact and address details it contains.	5.9	More Details
CVE-2026-116136	ProjectSend 2.0.0 through 2.7.1 contains an incomplete authorization vulnerability in FileVersions::link() that allows staff users to share files without authorization by linking file versions. Attackers can link their file onto an intermediate revision or assign out-of-scope clients, causing moveAssignmentsToRoot() to expose protected original files to unauthorized clients.	5.9	More Details
CVE-2026-112381	Apache Traffic Server can crash with a null pointer dereference when a client POST request is torn down via half-close while a chunked origin response is being reparsed from a reentrant write-ready callback, a timing-dependent race introduced by the CVE-2021-32564 flow-control backports. This issue affects Apache Traffic Server: 9.2.15, 10.1.4, 10.2.0. Users are recommended to upgrade to version 9.2.16, 10.1.5, or 10.2.1, which fix the issue.	5.9	More Details
CVE-2026-116071	rotki before 1.44.2 contains a weak key derivation vulnerability that allows attackers to brute-force backups because the premium sync key is a single unsalted SHA-256 of the user password. Attackers who obtain an uploaded backup blob can run GPU-parallel offline password guessing and decrypt the database containing exchange API keys, secrets, wallet addresses and transaction history.	5.9	More Details
CVE-2026-114650	A flaw in the client-side encryption support of the MongoDB .NET/C# driver can cause data designated for client-side field-level encryption to be written without encryption. Under certain platform text-encoding configurations, an internal name conversion loses information, so the encryption layer fails to associate the operation with its configured encryption rules and the operation proceeds without an error being raised. An application user with the ability to write data through an affected application can therefore have sensitive fields persisted in plaintext, defeating the confidentiality guarantee that client-side encryption is intended to provide. No integrity or availability impact was observed.	5.9	More Details
	adm-zip before 0.6.1 contains an inconsistent entry resolution vulnerability that allows attackers to		

CVE-2026-114104	bypass content validation by supplying ZIP archives with duplicate entry names. Attackers can craft archives where getEntry() returns the last duplicate while extractAllTo() writes the first, causing unvalidated bytes to be written to disk.	5.9	More Details
CVE-2026-114655	libmongocrypt builds requests to external key management services (AWS, Azure and GCP) by inserting values from key vault documents, KMS provider configuration and KMS responses into request lines, headers and signed request bodies without validating or encoding them for each context. A user with write access to the key vault collection can alter the structure of these requests, including adding members to request bodies and token claims that the library then signs with the application's credentials. This may result in key management operations being performed with parameters the application did not intend.	5.9	More Details
CVE-2026-57207	NetBox MCP Server is a simple read-only Model Context Protocol server for NetBox. Prior to version 1.2.1, when running with TRANSPORT=http, the MCP server has no authentication on the HTTP endpoint. Any client that can reach the port gets a valid MCP session and full access to all four NetBox tools - no token or credential needed. This issue has been patched in version 1.2.1.	5.9	More Details
CVE-2026-16255	GitLab has remediated an issue in GitLab EE affecting all versions from 12.0 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, when the instance-wide SAML or LDAP membership lock was enabled, an authenticated user could have granted or accepted project and group membership outside the identity provider's control, due to incomplete enforcement of the membership lock on requests created before the lock was enabled.	5.9	More Details
CVE-2026-116192	rclone from 1.37 before 1.74.2 contains a server-side request forgery vulnerability that allows unauthenticated attackers to redirect http and netstorage backend requests by supplying a double-slash root on an existing named remote via --rc-serve. Attackers can make HEAD/GET requests to attacker-chosen hosts to read responses, capture configured custom headers, or obtain valid NetStorage request signatures.	5.9	More Details
CVE-2026-116200	Snipe-IT before 8.6.3 ignores the ldap_start_tls() return value in Ldap::connectToLdap(), allowing on-path attackers to force cleartext LDAP binds. Attackers can answer the StartTLS request with a protocol error to capture service-account credentials and end-user login passwords.	5.9	More Details
CVE-2026-114400	Trudesk through 1.2.11 contains a missing authentication vulnerability that allows unauthenticated attackers to download private ticket attachments under /uploads. Because express.static serves the public/ directory before the hasAuth-guarded /uploads mount, attackers who know or guess a ticket ObjectId and filename can retrieve confidential attachments.	5.9	More Details
CVE-2026-114117	Vyper 0.3.8 through 0.5.0b1 contains an optimizer miscompilation in which (msg.value K) or (len(msg.data) K), where K is any nonzero integer literal, is compiled to the constant 1 when used as a ternary expression branch under -O gas (the default) or -O codesize. Affected contracts use the wrong value in storage writes, arithmetic, indexes, return values, or events without warning or revert, which can record deposits incorrectly and lock funds.	5.9	More Details
CVE-2026-116178	vLLM from 0.20.0 before 0.30.0 contains a denial of service vulnerability that allows remote attackers to crash the engine by sending inconsistent multimodal metadata to the disaggregated /inference/v1/generate endpoint. Attackers can replay a valid multimodal item with repeated or overlapping mm_placeholders ranges or mismatched kwargs_data and mm_hashes lists, producing a negative text span in Qwen2.5-VL position construction that terminates the engine.	5.9	More Details
CVE-2026-115329	Lotus through 0.10.2 contains an insecure direct object reference vulnerability in SubscriptionViewSet.get_object() that uses an unscoped queryset without organization filtering or permission checks. Authenticated users or API keys from any organization can supply another organization's subscription_id to read, cancel, update, switch plans, or modify addons, altering another tenant's billing.	5.9	More Details
CVE-2026-114190	Insertion of Sensitive Information Into Sent Data vulnerability in Hai Zheng / Lite Speed Cache LiteSpeed Cache litespeed-cache allows Retrieve Embedded Sensitive Data.This issue affects LiteSpeed Cache: from n/a through 7.9.	5.9	More Details
CVE-2026-114458	A flaw was found in mod_qos. An unauthenticated remote attacker can cause a Denial of Service (DoS) by maintaining numerous concurrent low-rate connections to a target server. When connection rate monitoring is active, the module fails to verify array boundaries when recording client IP addresses, allowing more connection records to be added than the allocated buffer can hold. This issue leads to heap memory corruption, which can crash the web server process.	5.9	More Details
	Use of a less trusted source for the envelope sender in the external-account download of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote unauthenticated attacker to have mail they write in the name of a domain the server hosts sent on by the server with an envelope sender at that		

CVE-2026-114474	domain, so that it passes the domain's SPF and DMARC at other receivers. A message the server downloads from an external POP3 or IMAP account has no SMTP envelope of its own, and the server took the address in its From header as its envelope sender, whatever that address was. A forward or redirect of the downloaded message to an outside address then left with that address as its envelope sender, which the Sender Rewriting Scheme leaves alone for a hosted domain, so that the receiver's SPF check passed for the forged domain, aligned even under strict alignment, with no non-default setting, host name record or DMARC policy needed. An attacker needs to know a mailbox elsewhere that an account here collects and forwards to an outside address; an account holder who adds an external account of their own can collect from a server of their own, which matters where SmtptAuthenticatedSenderCheck holds signed-in senders to their own addresses.	5.9	More Details
CVE-2026-114478	Insufficient preservation of the original sender in the forwarding of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote unauthenticated attacker to have mail they write, naming in its From an address at a domain the server hosts, delivered to other receivers as passing that domain's DMARC policy. When the server sent on a copy of mail that came from outside - an account's forward, a mail rule's forward, a Sieve redirect, a list's copy, or an alias or list expanding to an outside address - it could give the copy an SMTP identity in the organisation of the domain the copy's From names, whatever the message's origin: with the Sender Rewriting Scheme or the rewrite of the envelope sender on forwarding enabled (both off by default), an address at the forwarding domain; for a list, its bounce address; for a message with an empty envelope sender, none, under the server's own host name; and an envelope sender the message came with at a subdomain of a hosted domain was kept. Where that identity's SPF record authorises the server, other receivers' SPF check of the copy passed for a domain aligned with the From, and DMARC passed with it, although the message had not passed DMARC when the server received it. Exploitation needs a forward, redirect, alias or list that carries the hosted domain's mail to an outside address, and one of: SRS or the envelope rewrite enabled; a list whose bounce address is in the domain's organisation; an envelope sender at a subdomain whose SPF record authorises the server; or, for an empty envelope sender, an SPF record that authorises the server under its host name. The DKIM route to the same result is CVE-2026-107567.	5.9	More Details
CVE-2026-115333	Lotus through 0.10.2 contains an improper authentication vulnerability in the Go event ingestion middleware that validates API keys by prefix only, without checking the secret against the stored hashed_key. Attackers who know or obtain an API key prefix can submit usage events to POST /api/track/ as that organization, altering its metered usage and billing.	5.9	More Details
CVE-2026-114479	Not failing securely in the anti-virus attachment-stripping action of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote unauthenticated attacker to have a message that the server's virus scanner found infected delivered to the recipient's mailbox unchanged. With virus scanning enabled and its action set to remove attachments and deliver the text (the default action deletes the message), the server keeps one text part of an infected message and replaces every attachment with a note. A message whose single part is not text - a program or an archive sent as the whole message, a message carried whole, a file declared an attachment, or an S/MIME message signed as one part, whose content the recipient's mail client opens - had no part the action could pick out, and was delivered as it arrived, with the content the scanner had found and nothing but a line in the application log to show it. Any sender can write such a message.	5.8	More Details
CVE-2026-114609	fpdf2 is a PDF creation library for Python. Prior to 2.8.8, fpdf2's shared image loader automatically dereferences attacker-controlled HTTP(S) URLs and raw local paths supplied through FPDF.image(), FPDF.write_html() HTML values, SVG <image href> or xlink:href values, and template image fields. This can cause outbound requests to attacker-selected internal or external services and local file open attempts, although local content is exposed only when it is a valid supported image or SVG resource. file:/// URLs are not specially handled, so the local-path case requires a raw filesystem path. This issue is fixed in version 2.8.8.	5.8	More Details
CVE-2026-20221	Multiple Cisco products are affected by a vulnerability in the generic routing encapsulation (GRE) tunnel decapsulation feature of the Snort 3 Detection Engine that could allow an unauthenticated, remote attacker to cause an unexpected restart of the Snort 3 Detection Engine. This vulnerability is due to improper handling of nested GRE traffic. An attacker could exploit this vulnerability by sending a crafted GRE payload through an affected device. A successful exploit could allow the attacker to trigger an unexpected reload of the Snort 3 Detection Engine, resulting in either a bypass or denial of service (DoS) condition, depending on device configuration. For more information, see the section of this advisory. The Snort 3 Detection Engine will restart automatically. No manual intervention is required.	5.8	More Details
CVE-2026-76489	Multiple Cisco products are affected by a vulnerability in the memory management implementation of Snort 3 that could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart, resulting in an interruption of packet inspection. This vulnerability is due to incomplete error checking when Snort 3 processes files, which could cause unreleased memory. An attacker could exploit this vulnerability by sending a large number of files to be processed. A successful exploit could allow the attacker to cause the Snort 3 Detection Engine to restart unexpectedly, resulting in a DoS	5.8	More Details

	condition.		
CVE-2026-20218	A vulnerability in the Snort 3 Visual Basic for Applications (VBA) Decompression Engine of multiple Cisco products could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart unexpectedly, resulting in a denial of service (DoS) condition. This vulnerability is due to improper range checking when Snort 3 decompresses VBA data, which is user controlled. An attacker could exploit this vulnerability by sending crafted VBA data to the Snort 3 Detection Engine of a targeted device. A successful exploit could allow the attacker to cause an overflow of stack data, which could result in a DoS condition.	5.8	More Details
CVE-2026-20197	Multiple Cisco products are affected by a vulnerability in the Snort 3 Teredo Detection Engine that could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart, resulting in an interruption of packet inspection. This vulnerability is due to incomplete error checking when Snort 3 parses Teredo packets. An attacker could exploit this vulnerability by sending a crafted Teredo packet to a targeted device to be parsed by Snort. A successful exploit could allow the attacker to cause the Snort 3 Detection Engine to restart unexpectedly, resulting in a DoS condition.	5.8	More Details
CVE-2026-20019	Multiple Cisco products are affected by a vulnerability in the Snort 3 NetFlow Detection Engine that could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart, resulting in an interruption of packet inspection. This vulnerability is due to incomplete error checking when Snort 3 parses UDP Netflow Protocol packets. An attacker could exploit this vulnerability by sending a single crafted Netflow packet through an established connection to be parsed by Snort 3. A successful exploit could allow the attacker to cause the Snort 3 Detection Engine to restart unexpectedly, resulting in a DoS condition.	5.8	More Details
CVE-2026-20011	Multiple Cisco products are affected by a vulnerability in the Snort 3 SMTP MIME Decoder that could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart. This vulnerability is due to a lack of error checking when the MIME fields of the SMTP header are parsed. An attacker could exploit this vulnerability by sending crafted SMTP packets through an established connection to be parsed by Snort 3. A successful exploit could allow the attacker to cause a denial of service (DoS) condition when the Snort 3 Detection Engine unexpectedly restarts.	5.8	More Details
CVE-2026-115813	Server-side request forgery in the model download path in Amazon Web Services Multi Model Server 1.1.11 and earlier might allow remote actors to send requests to internal network destinations via a crafted model URL. Multi Model Server is end of life and no fixed version is available. Users should discontinue use of Multi Model Server and migrate to a supported model serving solution.	5.8	More Details
CVE-2026-114726	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a mail header injection vulnerability in the Bounce (resend) action. rcube_mime::decode_address_list() did not strip control characters from the decoded display name, so a received message whose From display-name carries CRLF sequences in an encoded word could inject arbitrary headers into the message emitted when the user bounces it. Upstream fixed this in the commit "Fix mail header injection via the From display-name in the Bounce (resend) action".	5.8	More Details
CVE-2024-7478	When a WSO2 product is configured to use an external user store, the system fails to adequately obscure the existence or validity of usernames during authentication attempts. This oversight allows an attacker to repeatedly query the authentication mechanism with different username guesses, enabling the enumeration of valid usernames present in the external user store. The discovery of valid usernames can increase the risk of targeted brute-force attacks, social engineering, and information leakage. Attackers can leverage this information to craft more convincing phishing campaigns or other social engineering tactics to trick users into revealing sensitive data. The presence of username enumeration could also negatively impact an organization's reputation, leading to loss of customer trust, regulatory non-compliance, and potential legal and financial repercussions.	5.8	More Details
CVE-2026-20179	Multiple Cisco products are affected by a vulnerability in the Snort 3 Generic Routing Encapsulation (GRE) Detection Engine that could allow an unauthenticated, remote attacker to cause the Snort 3 Detection Engine to restart, resulting in an interruption of packet inspection. This vulnerability is due to incomplete error checking when Snort 3 parses GRE packets. An attacker could exploit this vulnerability by sending a crafted GRE packet to a targeted device to be parsed by Snort. A successful exploit could allow the attacker to cause the Snort 3 Detection Engine to restart unexpectedly, resulting in a DoS condition.	5.8	More Details
CVE-2026-115391	Wasmtime 46.0.0 before 46.0.3 and 47.0.0 before 47.0.4 contains an uncontrolled resource allocation vulnerability in its WASIp3 implementation that buffers guest-specified-length writes on the host without limit. Malicious guests can offer large buffers to WASIp3 file or HTTP write streams, causing the host to copy and buffer all bytes, bounded only by guest linear memory, exhausting host memory.	5.8	More Details
	Coraza from 3.0.4 before 3.8.1 contains a request body inspection bypass because request-body		

CVE-2026-114100	processor selection compares the entire Content-Type value exactly, so URL-encoded bodies whose Content-Type carries a media-type parameter are never parsed. Unauthenticated attackers can send application/x-www-form-urlencoded; charset=UTF-8 to keep ARGS_POST and REQUEST_BODY empty, evading body rules while the backend still processes the payload.	5.8	More Details
CVE-2026-114099	Coraza from 3.0.0 before 3.8.1 has a JSON body processor that flattens nested properties into dot-separated ARGS_POST names without escaping literal dots, so distinct properties collide. Unauthenticated attackers can add a later dotted property such as account.role that overwrites a malicious nested value, hiding SQL injection payloads from OWASP CRS rules while backends still receive them.	5.8	More Details
CVE-2026-115378	Incus before 6.0.7, 6.1.0 before 7.0.2, and 7.1.0 before 7.5.0 contains an arbitrary file write vulnerability in sftpRecursivePullFile that allows a malicious incus-agent to escape the target directory during incus file pull -r. Attackers with root in a virtual machine can return duplicate entries reported first as symlinks then as files or directories, writing attacker-controlled files such as /etc/shadow on the operator's host.	5.8	More Details
CVE-2026-114657	Grav from 2.0.0 before 2.2.5 contains an information disclosure vulnerability that allows page authors to capture visitors' HTTP Basic credentials through Twig page content. Authors can print the request Uri object, which embeds PHP_AUTH_USER and PHP_AUTH_PW, inside an inline style background image to exfiltrate administrators' credentials to an external host.	5.7	More Details
CVE-2026-115390	Wasmtime 47.0.0 before 48.0.3 and 49.0.0 before 49.0.1 contains a fuel accounting flaw in Cranelift-compiled call_ref and try_table catch paths. Malicious guests can reset consumed fuel after calls to run exponentially long on linear fuel, evading timeouts and periodic yields for denial of service.	5.7	More Details
CVE-2026-115379	Wasmtime 47.0.0 before 48.0.4 and 49.0.0 before 49.0.2 contains a missing GC rooting vulnerability that leaves references live across try_call unrooted when compiling Wasm modules. Attackers can supply Wasm modules whose callee triggers garbage collection, leaving stale or dangling references that cause GC heap corruption or incorrect execution.	5.7	More Details
CVE-2026-76303	In Splunk SOAR versions below 8.7.0, a user that holds a role with the install_apps capability could store attacker-controlled content that Splunk SOAR displays in app documentation by installing an app. If another user who holds a role with permission to view apps opens the app documentation, Splunk SOAR could run that content in the browser of the user who opened the app documentation. With this access, the first user could access data and make changes that the user who opened the app documentation could make. The vulnerability is possible because Splunk SOAR does not sanitize app documentation content before it displays that content. The vulnerability requires the attacker to phish the affected user by tricking them into initiating a request within their browser. The user that holds a role with the install_apps capability should not be able to exploit the vulnerability at will. 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/7.1.0/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.	5.7	More Details
CVE-2026-114361	A vulnerability was detected in HarnessRouter up to 0.31.10. This affects an unknown function of the file gateway/app.py of the component Identity Mode. Performing a manipulation of the argument X-Harness-Org/X-Harness-Member results in improper authorization. The attack is possible to be carried out remotely. The attack is considered to have high complexity. The exploitability is reported as difficult. Upgrading to version 0.31.11 mitigates this issue. The patch is named 48a26cdf9c4544d2c88d1a07f0777f063f1d753f. The affected component should be upgraded.	5.6	More Details
CVE-2026-116120	A vulnerability was determined in NocoDB up to 2026.07.0. This vulnerability affects the function UserRefreshToken.deleteAllUserToken of the file /auth/user/signup of the component PublicApiLimiterGuard. This manipulation causes allocation of resources. The attack is possible to be carried out remotely. The attack is considered to have high complexity. It is stated that the exploitability is difficult. Code-verified absence of in-app rate limiting on the unauthenticated signup and password-forgot endpoints and per-request refresh/OAuth token invalidation, but the vendor documents the pass-through limiter as intentional design. The vendor deleted the GitHub issue for this vulnerability without any explanation. Afterwards the vendor was contacted early about this disclosure via email but did not respond in any way.	5.6	More Details
	The DHCPv4 client in subsys/net/lib/dhcpv4/dhcpv4.c maintained its transaction identifier (iface->config.dhcpv4.xid) as a counter rather than as a per-transaction random value. It was seeded once with sys_rand32_get() in dhcpv4_start_internal(), incremented by one in dhcpv4_send_discover() and		

CVE-2026-79980	dhcpv4_send_decline(), and never changed by dhcpv4_send_request() — so every RENEWING and REBINDING exchange reused, unchanged, the identifier that had already appeared on the wire in the broadcast DHCPDISCOVER at boot, and where the identifier did move it moved predictably by one. RFC 2131 4.1 requires a single identifier per transaction, chosen so that clients are unlikely to collide. This identifier is the client's only unpredictable reply-validation token. net_dhcpv4_input() accepts a BOOTP reply on op, a matching xid, and a matching chaddr (the interface MAC, which is public); the UDP callback is registered with a wildcard remote address and port, and no server-identifier or source-address check is applied. An attacker who learns one identifier — for example any host on the same segment that sees the broadcast DHCPDISCOVER, or a DHCP server or relay on a network the device previously joined — therefore retains a valid token for the device's whole uptime, and can predict the identifiers of later exchanges, without having to observe the unicast renew traffic or race the legitimate server. A forged reply that passes those checks has effect regardless of client state, because dhcpv4_parse_options() applies its side effects before the state machine gate in dhcpv4_handle_reply(): the router option calls net_if_ipv4_set_gw() and the DNS option calls dns_resolve_reconfigure(). A single accepted datagram can thus repoint the device's default gateway and resolvers at attacker-controlled addresses, redirecting its traffic; during RENEWING or REBINDING a forged DHCPACK re-binds the lease with attacker-chosen parameters and a forged DHCPNAK removes the address and restarts configuration. The fix draws a fresh sys_rand32_get() value at each point where a transaction begins — entering SELECTING, RENEWING and REBINDING, and at the INIT/INIT-REBOOT entry taken when an interface with a lease goes down — while REQUESTING and DECLINE keep the identifier of the exchange they belong to, as RFC 2131 requires. As a side effect the client also no longer discards a slow server's offer after retransmitting a discover, and the initial-discover backoff now draws its own randomness instead of deriving it from the identifier.	5.6	More Details
CVE-2026-115396	Ghidra 10.3 before 12.2 contains an improper handling of highly compressed data vulnerability in ThemeReader.processIconFile() that fails to limit decompressed size or compression ratio. Attackers can persuade a user to import a crafted theme ZIP whose icon entries expand to fill the host filesystem, causing later writes to fail.	5.5	More Details
CVE-2025-60493	A segmentation violation in the check_media_profile function (/mp4box/mp4box.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60479	A NULL pointer dereference in the StatSFVec2f function (/scene_manager/scene_stats.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MPEG-4 file.	5.5	More Details
CVE-2025-60476	A heap use-after-free in the gf_odf_ac4_cfg_clean_list function (/odf/descriptors.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted AC4 file.	5.5	More Details
CVE-2026-114064	Snipe-IT before 8.7.0 contains a server-side request forgery vulnerability in the webhook testing feature that allows super-admin users to probe internal network addresses, including loopback and RFC1918 ranges. Attackers can point webhooks at internal hosts and ports, then compare 4xx versus 5xx responses from Test Webhooks to map open ports and active hosts.	5.5	More Details
CVE-2026-111980	The GiveWP WordPress plugin before 4.19.0.1 does not correctly determine which page a page-capability request applies to, allowing users with a custom fundraising role to edit, publish and read pages owned by other users, including administrators.	5.5	More Details
CVE-2026-115397	Ghidra before 12.2 contains a resource allocation without limits vulnerability in ExtensionUtils.unzipToInstallationFolder() that extracts every extension ZIP entry without limiting file count. Attackers who persuade a user to install a crafted extension archive with many files can exhaust filesystem inodes, causing project writes and other processes to fail.	5.5	More Details
CVE-2025-60492	A stack overflow in the gf_bifs_dec_node function (/bifs/field_decode.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60469	A NULL pointer dereference in the gf_ac4_presentation_ch_mode function (/media_tools/av_parsers.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted AC4 file.	5.5	More Details
CVE-2026-115395	Ghidra 10.3 before 12.2 contains a resource allocation without limits vulnerability in ThemeReader.readZipTheme() that extracts unlimited theme ZIP entries to the images directory. Attackers can persuade users to import a crafted theme archive with many small files, exhausting inodes and disk space to deny service to the host.	5.5	More Details
CVE-	Trigger.dev CLI before 4.5.9 contains an information disclosure vulnerability that exposes resolved environment secret values in debug deployment logs by serializing unredacted envVars. Attackers with		More

CVE-2026-114073	access to terminal transcripts, CI debug logs, or support bundles can read the Starting buildWorker record to recover database connection strings and service credentials.	5.5	Details
CVE-2026-115394	Ghidra through 12.2 contains an uncontrolled memory allocation vulnerability in BinaryField.read() that allows attackers to exhaust heap memory via an unchecked 32-bit length field. Attackers can persuade users to import a crafted GZF packed database file, triggering an OutOfMemoryError in PackedDatabase.open() that aborts the import and denies service.	5.5	More Details
CVE-2026-46531	Cacti is an open source performance and fault management framework. From version 1.2.27 to before version 1.2.31, Cacti contains an authenticated stored second-order SQL injection in the automation tree rule item handler. A user holding the Automation realm (default General Administration role) can persist an arbitrary SQL fragment in the automation_tree_rule_items.field column. The fragment is later concatenated directly into a SELECT ... AS source FROM ... clause in lib/api_automation.php :: create_all_header_nodes() and executed when an administrator triggers automation via the standard graphs.php?action=actions&drp_action=8 ("Apply Automation Rules") bulk action. The result of the injected SELECT is written into graph_tree_items.title and rendered in the Cacti tree UI, providing a direct exfiltration channel. This issue has been patched in version 1.2.31.	5.5	More Details
CVE-2026-111738	The Speed Optimizer WordPress plugin before 7.8.5 does not verify network-level privilege before applying an option change to the whole multisite network, nor restrict which option may be changed, allowing a site administrator who is not a Super Admin to create or overwrite network-wide options.	5.5	More Details
CVE-2026-116119	A vulnerability was found in NocoDB up to 2026.08.0. This affects an unknown part of the file packages/nocodb/src/services/users/users.service.ts of the component Password-Reset Flow. The manipulation results in race condition. The attack can be executed remotely. The vendor deleted the GitHub issue for this vulnerability without any explanation. Afterwards the vendor was contacted early about this disclosure via email but did not respond in any way.	5.5	More Details
CVE-2026-114223	NT OneAPP developed by NewType Infortech has an Exposure of Sensitive Information vulnerability. Unauthenticated local attackers can extract plaintext usernames and passwords from backup data after obtaining ADB backup access authorization and user confirmation of the backup, resulting in sensitive information disclosure.	5.5	More Details
CVE-2026-95717	A stored cross-site scripting (XSS) vulnerability exists in the web-based management interface of AOS that could allow an authenticated remote attacker with high privileges to inject malicious scripts. Successful exploitation could allow an attacker to execute arbitrary script code in the context of another authenticated user's browser session.	5.5	More Details
CVE-2026-111947	The Pricing Table WordPress Plugin WordPress plugin before 4.1.3 does not perform authorisation checks when duplicating a legacy pricing table, allowing users with the contributor role and above to clone and publish any post or page on the site, including private and draft content belonging to other users.	5.5	More Details
CVE-2025-71438	An out-of-bounds read in gf_odf_ac4_cfg_dsi_v1 (/odf/descriptors.c) in GPAC Project/MP4Box before v26.07.0 allows attackers to cause a Denial of Service (DoS) via supplying a crafted AC4 file.	5.5	More Details
CVE-2025-60489	A NULL pointer dereference in the dasher_insert_timeline_entry function (/filters/dasher.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60478	A use-after-free in the gf_filter_packet_destroy function (/filter_core/filter_pck.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MPEG-2 file.	5.5	More Details
CVE-2026-115398	Ghidra before 12.2 contains a resource allocation without limits vulnerability in FileSystemIndexHelper used by TarFileSystem and ZipFileSystem, which store unbounded archive entries. Attackers can persuade a user to open or batch-import a crafted TAR or ZIP archive with excessive entries to exhaust the Java heap and hang the session.	5.5	More Details
CVE-2026-116280	A vulnerability was determined in Kameleo up to 5.0.1.26. Affected is the function InternalController.TestProxy of the file client-engine-application/f.cs of the component Proxy Validation. Executing a manipulation of the argument host/port can lead to server-side request forgery. The attack may be launched remotely. Upgrading to version 5.2.0 is able to address this issue. It is suggested to upgrade the affected component. The route has been removed and replaced with /general/test-proxy in 5.2.0.	5.5	More Details
	A flaw has been found in R6410418 Jackrong-llm-finetuning-guide up to ef2b17f2798cd18cd85375080c35a58397ad0d09. This affects the function validate_indexed_tensors of		

CVE-2026-113737	the file qwen-mtp-gguf/scripts/qwen_mtp_gguf_pipeline.py. Executing a manipulation of the argument weight_map can lead to path traversal. The attack may be performed from remote. The exploit has been published and may be used. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	5.5	More Details
CVE-2025-60491	A heap use-after-free in the dasher_flush_segment function (/filters/dasher.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60490	A heap use-after-free in the gf_filter_pid_send_event_downstream function (/filter_core/filter_pid.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60497	A segmentation violation in the GetESD function (/isomedia/track.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted MP4 file.	5.5	More Details
CVE-2025-60487	A segmentation violation in the gf_isom_get_esd function (/isomedia/isom_read.c) of GPAC Project/MP4Box before 2.5 allows attackers to cause a Denial of Service (DoS) via supplying a crafted XML file.	5.5	More Details
CVE-2026-76462	A vulnerability in the interprocess communication (IPC) channel of the Secure Firewall Posture module in Cisco Secure Client could allow an authenticated, local attacker to elevate privileges on an affected device. This vulnerability is due to insufficient validation of user-supplied input. An attacker could exploit this vulnerability by sending a series of crafted IPC messages to a local socket. A successful exploit could allow the attacker to perform limited file write operations on the filesystem with SYSTEM or root privileges. Although this vulnerability affects the Secure Firewall Posture module in Cisco Secure Client on Windows, macOS, and Linux operating systems, this vulnerability could be used to obtain full root privileges on macOS and Linux. For that reason Cisco has assigned this security advisory a Security Impact Rating (SIR) of High rather than Medium as the score indicates.	5.5	More Details
CVE-2026-115503	Ghidra before 12.2 contains an improper handling of highly compressed data vulnerability in ExtensionUtils.unzipToInstallationFolder() that allows attackers to exhaust disk space. Attackers can persuade users to install a crafted zip bomb extension archive that expands without size limits, causing later writes by Ghidra and other applications to fail.	5.5	More Details
CVE-2026-115424	A flaw was found in the OpenShift client oc. When extracting container image layers with oc image extract, the application fails to properly sanitize symbolic and hard link targets pointing to absolute paths or to sibling directories of the extraction destination. An attacker who supplies a malicious container image can exploit this vulnerability by embedding links, and a chained file entry, that reference locations outside the intended extraction directory. This allows the attacker to write attacker-controlled content to arbitrary locations on the host filesystem with the privileges of the user running the command, potentially leading to file modification outside the destination directory.	5.5	More Details
CVE-2026-115207	A flaw was found in GNU patch. An attacker can achieve arbitrary file creation outside the intended working directory by enticing a user or automated workflow to apply a specially crafted patch file. Due to missing validation during automatic patch reversal, absolute file paths in patch headers are accepted rather than rejected. Consequently, the utility can create files in arbitrary locations with the permissions of the executing user.	5.5	More Details
CVE-2026-116497	A NULL pointer dereference vulnerability was found in gssproxy. When a GSSX_INIT_SEC_CONTEXT request is received with the target_name field absent, the daemon dereferences the resulting NULL pointer and crashes. The request requires no authentication, allowing any local user to crash the gssproxy daemon via its Unix socket. The daemon does not automatically restart, resulting in a persistent denial of service that also affects dependent services such as Kerberized NFS and IdM.	5.5	More Details
CVE-2026-81163	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Drupal Media Library Importer allows Path Traversal. This issue affects Media Library Importer versions: from 0.0.0 to 2.1.6.	5.4	More Details
CVE-2026-116072	Infinity for Reddit through 8.3.4 contains a login cross-site request forgery vulnerability that allows attackers to log victims into an attacker-controlled account because the exported LoginChromeCustomTabActivity redirect handler checks a hardcoded OAuth state value. Attackers can authorize the app's client ID with their own Reddit account, then deliver an infinity://localhost redirect containing their code and the constant state, causing the victim's posts, messages, and browsing to occur on the attacker's account.	5.4	More Details
	The Depicter — Popup & Slider Builder WordPress plugin before 4.8.2 does not properly validate		

CVE-2026-113472	uploaded files in its editor media-upload endpoint, checking only the client-supplied MIME type and a case-sensitive block of the .php extension, which allows users with the Editor role and above to upload files with other active extensions into the public uploads directory, leading to stored script execution in the site origin and, depending on the server's handler mapping, to code execution.	5.4	More Details
CVE-2026-116212	Webasyst Framework through 4.4.0 contains a stored cross-site scripting vulnerability in the Photos app that allows authenticated backend users to inject JavaScript via the album url parameter. Attackers with album edit rights can save a url value that breaks out of the unescaped attribute in DialogAlbumSettings.html, executing JavaScript in administrators' backend sessions.	5.4	More Details
CVE-2026-84727	Use-after-free in the RTMP protocol implementation in Apache bRPC (version < 1.18.0) on all platforms allows a malicious RTMP peer to crash the process or potentially corrupt memory via a crafted Abort message. Root Cause: When processing an RTMP Abort (protocol control message type 2), the RTMP policy implementation deleted the target RtmpChunkStream object. The send path (RtmpUnsentMessage::AppendAndDestroySelf -> RtmpChunkStream::SerializeMessage) runs concurrently on sender threads and dereferences the same RtmpChunkStream object without synchronization. If a remote peer sends an Abort message naming a chunk stream that is concurrently being used to serialize an outgoing message, the stream object is freed while still in use, resulting in a use-after-free. Additionally, deleting the chunk stream discarded its previous message header, so subsequent compressed (type-3) chunk headers on that stream could no longer be decoded. Affected Scenarios: Any application using the bRPC RTMP protocol implementation (rtmp policy), either as an RTMP server or as an RTMP client, where the remote endpoint is untrusted. The attacker needs to send an Abort message targeting a chunk stream id that is concurrently used for sending, which can crash the process or lead to memory corruption. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3515	5.4	More Details
CVE-2026-113800	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-alpha.5, hid and usb permission checks made from an iframe caused setPermissionCheckHandler to receive the top-level frame's origin and details instead of the requesting frame's. An application that delegated hid or usb to a cross-origin untrusted iframe and allowed device access with origin-based logic could therefore grant that embedded frame access intended for the top-level page. This issue is fixed in versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-alpha.5.	5.4	More Details
CVE-2026-116262	A vulnerability was detected in libimobiledevice up to 1.4.0. The affected element is the function webinspector_receive_with_timeout of the file src/webinspector.c of the component Web Inspector. The manipulation results in heap-based buffer overflow. 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.	5.4	More Details
CVE-2026-14556	The HappyForms WordPress plugin before 1.26.15 does not sanitise or escape a user-controlled referral value before embedding it in the form notification email, allowing unauthenticated attackers to inject arbitrary HTML into the emails site administrators receive.	5.4	More Details
CVE-2026-115196	An Open Redirect vulnerability exists in the public form submission endpoint (/form/submit) of Mautic. After a form is submitted, the application redirects the visitor to a return address supplied with the submission without verifying that it points to the Mautic instance or another trusted host. An unauthenticated attacker can exploit this to send visitors from a legitimate Mautic domain to an arbitrary external website, for example to phishing pages that imitate the original landing page or to sites that deliver malware. Any Mautic instance with a published form is affected. Exploitation requires the victim to submit a form prepared by the attacker.	5.4	More Details
CVE-2026-87945	URL Redirection to Untrusted Site ('Open Redirect') vulnerability in Drupal SAML SSO - Service Provider allows Fake the Source of Data. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	5.4	More Details
CVE-2026-116202	YesWiki before 4.7.0 contains a stored cross-site scripting vulnerability that allows low-privilege authenticated users to upload unsanitized SVG files through the bazar ImageField class. Attackers can submit a malicious SVG via the image field of the entries API, which is served as image/svg+xml and executes script for users viewing the entry or file.	5.4	More Details
CVE-2026-114121	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-alpha.5, requests to open external protocol URLs could be attributed to a document other than the one that initiated them, causing the requestingUrl passed to session.setPermissionRequestHandler to identify trusted content for a launch requested by untrusted content. An application that grants openExternal based on the requesting origin and allows untrusted content to open or navigate windows could therefore allow that content to launch an OS-registered external application. This issue is fixed in versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-	5.4	More Details

	alpha.5.		
CVE-2026-114120	SchedulePress before 5.3.5 contains a missing authorization vulnerability that allows authenticated Contributor-level users to change any post's status and publish date via the wp_sp_el_editor_form AJAX action. Attackers can reuse the wp_sp-el-editor nonce from the Elementor modal to unpublish or reschedule administrator posts, or publish other users' pending or draft posts.	5.4	More Details
CVE-2026-116139	WWBN AVideo 26.0 through 29.3.0 contains a stored cross-site scripting vulnerability that allows authenticated commenters to inject JavaScript because markdownToHTML() preserves event-handler attributes in leading HTML blocks. Attackers can post a comment starting with an img tag carrying an onerror handler, which bypasses the inline sanitizer and executes in viewers' browsers under the application origin.	5.4	More Details
CVE-2026-115877	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, twisted.web.http.HTTPChannel.headerReceived rejects a NUL byte in an HTTP header value but accepts a bare carriage return or bare line feed. Because Twisted splits header lines only on CRLF, a hidden Content-Length header can remain folded into another value while a lenient front end treats the bare character as a line terminator. The resulting framing disagreement can smuggle a following request past front-end access controls, poison a shared connection or cache, and expose a later client to the wrong response. This issue is an incomplete-fix residual of earlier Twisted request-smuggling protections. This issue is fixed in version 26.10.0rc1.	5.4	More Details
CVE-2026-93633	A CSS sanitization vulnerability in Zimbra Collaboration Suite (ZCS) allows an authenticated user with Briefcase access to publish a document that displays attacker-controlled content under the Zimbra origin. When processing a stylesheet containing non-ASCII characters, DefangFilter appends the final ASCII segment without sanitizing it. HTML stored as application/x-zimbra-doc can therefore retain CSS URL references that the normal sanitizer removes. An attacker can share a crafted document that loads external resources and displays a credential-collection form using Zimbra's own visual assets. Exploitation requires the victim to open the document in a browser that renders the content type inline; credential theft additionally requires form submission. ZCS 10.1.16 and 10.1.20 are affected.	5.4	More Details
CVE-2026-115954	Missing Authorization vulnerability in Dylan Kuhn Geo Mashup geo-mashup allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Geo Mashup: from n/a through 1.13.22.	5.4	More Details
CVE-2026-114398	Trudesk through 1.2.11 contains an unrestricted file upload vulnerability that allows any authenticated user to place HTML or SVG files in the public webroot via the uploadlogo, uploadpagelogo and uploadfavicon endpoints. Attackers can spoof an image MIME type to upload script-bearing files served unauthenticated from /assets/, enabling stored cross-site scripting and replacement of instance branding.	5.4	More Details
CVE-2026-114074	Trigger.dev before 4.5.10 contains a server-side request forgery vulnerability that allows authenticated project members to bypass webhook SSRF protections by targeting the non-global 198.18.0.0/15 benchmarking range. Attackers can configure alert-channel webhooks pointing at internally routed services in this range, which isUnsafelPv4 omits, to perform blind internal probing and limited state-changing POST requests.	5.4	More Details
CVE-2026-59188	Part-DB is an open source inventory management system for electronic components. Prior to version 2.12.3, Part-DB-server sanitizes SVG files based solely on the user-controlled filename extension rather than actual file content type. An authenticated attacker can upload a malicious SVG with JavaScript payload using a non-SVG extension (e.g., .jpg), bypassing sanitization entirely. When any user views the attachment inline, the browser receives Content-Type: image/svg+xml and executes the embedded script in the Part-DB origin. This issue has been patched in version 2.12.3.	5.4	More Details
CVE-2026-113798	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.9, 43.7.6, 44.5.0, and 45.0.0-alpha.13, pointerLock permission requests from some cross-origin iframes caused session.setPermissionRequestHandler() to receive the top-level frame's requestingUrl and isMainFrame values instead of those of the requesting frame. An application that embeds untrusted iframe content and grants pointer lock based on the requesting origin could therefore grant the permission under the top-level page's identity. Applications that deny pointerLock for untrusted embedded content are not affected. This issue is fixed in versions 42.11.9, 43.7.6, 44.5.0, and 45.0.0-alpha.13.	5.4	More Details
CVE-2026-	Authlib is a Python library which builds OAuth and OpenID Connect servers. Prior to 1.9.0, Authlib's RFC 9101 JWTAuthenticationRequest passes the unauthenticated, client-controlled request_uri parameter to the abstract get_request_object() hook before applying a destination or registration policy; the hook itself performs no network request, while the documentation demonstrates an unrestricted HTTP fetch. In deployments whose integrator implements that hook as an unrestricted outbound fetch, a remote attacker who knows a valid public client identifier can make the authorization server request internal,	5.4	More

114620	link-local, or otherwise unintended URLs and can distinguish fetch outcomes through different error responses and timing. Deployments whose custom hook already applies a secure allowlist, destination controls, and retrieval limits are not affected. The JWTAuthenticationRequest implementation in authlib/oauth2/rfc9101/authorization_server.py forwards request_uri to get_request_object() before an application-approved retrieval policy accepts the destination. This issue is fixed in version 1.9.0.		Details
CVE-2026-115221	A flaw was found in FreeIPA. An improper authorization check allows an authenticated host to retrieve and revoke digital certificates assigned to other users or services. This issue occurs because FreeIPA verifies ownership solely by matching the host name to the certificate's subject Common Name (CN), without confirming actual ownership in the directory service. Consequently, an enrolled host can disrupt services by revoking certificates belonging to other entities.	5.4	More Details
CVE-2026-116020	BookStack before 26.09.2 contains an incorrect authorization vulnerability in AttachmentController that allows authenticated users to access other users' draft page attachments by ID. Attackers with attachment permissions can enumerate sequential attachment IDs to download, rename, relink, or replace files on another user's unpublished draft pages.	5.4	More Details
CVE-2026-87940	Use of Web Browser Cache Containing Sensitive Information vulnerability in Drupal Key auth allows Exploiting Incorrectly Configured Access Control Security Levels. This issue affects Key auth versions: from 0.0.0 to 2.2.4.	5.4	More Details
CVE-2026-114643	The MongoDB C# Driver writes GeoJSON extra members without checking whether their names collide with reserved GeoJSON members such as type and coordinates. An application that places user-supplied field names into these extra members can store documents with duplicate geometry fields, so the server and the driver resolve different values for the same record. Affected records can also become unreadable through the driver, interrupting application code that reads them.	5.4	More Details
CVE-2026-116169	SiYuan from 2.10.0 before 3.8.7 contains a stored cross-site scripting vulnerability that allows users with edit permission to inject script by placing javascript: URLs in math block content rendered with KaTeX trust enabled. Attackers can store \href, \url, or \includegraphics payloads that mathRender assigns via innerHTML, executing script in the application origin when a reader clicks the rendered link.	5.4	More Details
CVE-2025-65200	An issue was discovered in localstack 4.9.2 allowing attackers to execute arbitrary code via crafted content hosted on a web-accessable URL and when passed said URL in the the templateUrl to the _localstack/cloudformation/deploy endpoint, the hosted content can get executed in the context of the docker container.	5.4	More Details
CVE-2026-96388	Cross-Site Request Forgery (CSRF) vulnerability in Drupal Tawk.to - Live chat application allows Cross Site Request Forgery. This issue affects Tawk.to - Live chat application versions: from 0.0.0 to 3.0.4.	5.4	More Details
CVE-2026-115994	BOLA Allows Cross-Project Target Tampering in Cervantes version v1.3-beta could allow an attacker modify or delete targets belonging to other projects without authorization, compromising the integrity of penetration testing scopes, results, and associated security reports.	5.4	More Details
CVE-2026-114379	A flaw has been found in EKKOLearnAI ekko-studio up to 1.0.6. This vulnerability affects the function saveGeneratedImages of the file packages/server/src/modules/studio/controllers/media.ts of the component Media Generation API. This manipulation of the argument output_path causes file inclusion. Remote exploitation of the attack is possible. Patch name: 1d42459c292f89e412f7f6243829e7fed2e85bf5. It is recommended to apply a patch to fix this issue.	5.4	More Details
	Improper Neutralization of Special Elements used in an Expression Language Statement ('Expression Language Injection') vulnerability in Apache Camel AWS2-s3 component. This issue affects Apache Camel: from 4.8.0 before 4.18.5, from 4.19.0 before 4.22.2. For the object it uploads, reads or deletes, the camel-aws2-s3 producer takes the key from the CamelAwsS3Key message header and falls back to the keyName option configured on the endpoint only when that header is absent. Since 4.8.0, AWS2S3Utils.determineKey() also passed the resulting value to the Simple language whenever it contained \${ or \${simple{, so that a route could use a dynamic key, and the evaluation applied to a value taken from the header as well as to the configured option. The header is not route configuration, however: the aws2-s3 consumer sets CamelAwsS3Key on the exchanges it creates to the key of the object it read, that is, the name of the object in the bucket. In a route that does nothing more than consume from one aws2-s3 endpoint and produce to another, the name of each consumed object therefore reached the producer, and a name containing either marker was evaluated as an expression, although the route author never asked for any evaluation, and configuring keyName on the producer did not prevent it, because the header took precedence. A key that a route copied into the header from other message content was evaluated in the same way. Anyone able to choose the name of an object in a bucket that such a route consumes, typically a principal allowed to upload to that bucket, could therefore have Simple expressions of their choosing evaluated in the Camel process. The Simple language can read JVM system properties, environment variables and configuration properties, which		

CVE-2026-95664	commonly carry credentials, and, where camel-bean is on the classpath, invoke methods on beans in the registry. The evaluated result became the key the producer used, so reading a value back additionally required the ability to see object names in the destination bucket, while a method invocation took effect without it; together with the need to name an object in a consumed bucket, this is why the severity is assessed medium. The bucket name went through the same evaluation, but the consumer only ever places the bucket configured on its own endpoint in a bucket header, never a value derived from a remote object, so the bucket name was not reachable in this way. Releases before 4.8.0 used the key from the header as given and are not affected. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The 4.8.x to 4.17.x and 4.19.x to 4.21.x release lines reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. After upgrading, a key or bucket name supplied through a message header is used literally, so a route that put an unevaluated Simple expression in the CamelAwsS3Key or CamelAwsS3OverrideBucketName header must evaluate it in the route instead, for example with setHeader and simple(), or configure it as the keyName or bucketName option of the endpoint. For deployments that cannot upgrade immediately, do not let a key derived from an object name or from other untrusted message content reach an aws2-s3 producer unchecked: in routes that pass exchanges from an aws2-s3 consumer to an aws2-s3 producer, either remove the CamelAwsS3Key header before the producer with removeHeader("CamelAwsS3Key"), so that the keyName configured on the endpoint is used, or reject exchanges whose CamelAwsS3Key header contains \${ or \$simple{. Configuring keyName alone is not a workaround, because the header takes precedence over it. As defence in depth, allow only trusted principals to write to buckets that Camel routes consume.	5.4	More Details
CVE-2026-84924	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Email Verification / SMS Verification / OTP Verification allows Reflected XSS. This issue affects Email Verification / SMS Verification / OTP Verification versions: from 0.0.0 to 2.4.0.	5.4	More Details
CVE-2026-114698	Authorization Bypass Through User-Controlled Key vulnerability in myCred myCred mycred allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects myCred: from n/a through 4.0.0-beta.1.	5.4	More Details
CVE-2026-75524	OpenReception's appointment booking software provides an end-to-end encrypted appointment booking platform. Prior to version 2.2.2, an attacker overbooks an agent, placing multiple appointments into a single time slot that the system is designed to keep exclusive. The number of overlapping appointments equals the number of concurrent requests, and one booking token is enough because the token is consumed after the insert. This results in the integrity of the booking-calendar breaking and availability being denied to other clients. This issue has been patched in version 2.2.2.	5.4	More Details
CVE-2026-111847	Missing Authorization vulnerability in Kings Plugins MarketKing marketking-multivendor-marketplace-for-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects MarketKing: from n/a through 2.1.72.	5.4	More Details
CVE-2026-116107	A vulnerability was detected in Datafeedr Product Sets Plugin up to 1.3.24 on WordPress. Affected by this issue is the function dfrps_ajax_delete_saved_search of the file functions/ajax.php of the component AJAX Handlers. The manipulation results in missing authorization. The attack may be launched remotely. The exploit is now public and may be used. Upgrading to version 1.3.25 can resolve this issue. Upgrading the affected component is advised. The vendor was contacted early about this disclosure.	5.4	More Details
CVE-2026-113707	A security flaw has been discovered in a2aproject a2a-rs up to 0.4.4. The affected element is the function InMemoryTaskStore::update of the file a2a-server/src/task_store/inmemory.rs. The manipulation results in race condition. The attack may be performed from remote. The patch is identified as 7050c7d329238539424ee46440e23b91c4825bee. It is advisable to implement a patch to correct this issue.	5.4	More Details
	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Apache Camel Docling component. This issue affects Apache Camel: from 4.15.0 before 4.18.5, from 4.19.0 before 4.22.2. DoclingProducer runs the external docling tool in CLI mode, the default, and tells it where to write the converted document through the --output flag. When the CamelDoclingOutputFilePath header was set, addOutputDirectoryArguments handed its value to that flag verbatim: the value was not normalized, an absolute path or a path containing .. segments was accepted as given, and no endpoint option could confine it to a directory. The custom-argument validation added by the fix for CVE-2026-40047 refuses the --output and -o flags in the CamelDoclingCustomArguments header and tells callers to choose the output directory with this header instead, so the one value that decides where docling writes reached the command line without any of the checks applied to custom arguments. Docling writes the conversion result, as Markdown, HTML, JSON or text, into that directory, naming the file after the input file, and the producer afterwards removes only its own temporary		

CVE-2026-95629	directory, so the file is left in place. A party able to set the header could therefore have conversion output created in any directory the Camel process can write to. The content of such a file is the conversion of the document carried by the message, so it is shaped by whoever supplies that document without being written byte for byte, and when the document content arrives in the message body the producer stores it under a random temporary name, which the output file then carries. The header belongs to the Camel header namespace, which consumers applying Camel's default header filtering, such as the JMS and Kafka consumers, remove from incoming messages, so the issue is reached where a route forwards externally influenced values into the header or where a consumer maps incoming headers without that filtering. The docling-serve mode does not use the header. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The fix adds an outputBaseDirectory endpoint option that is unset by default, so after upgrading, set it on every docling endpoint whose CamelDoclingOutputFilePath header can carry a value influenced from outside the deployment; without it the header is normalized but can still name any directory. The 4.15.x to 4.17.x and 4.19.x to 4.21.x release lines reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. For deployments that cannot upgrade immediately, do not build the CamelDoclingOutputFilePath header from values that originate outside the deployment, remove that header from messages that arrive from outside before they reach a docling endpoint, and make sure that consumers which map incoming message headers drop the Camel header namespace, so that only the route itself chooses the output directory.	5.4	More Details
CVE-2026-95628	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Apache Camel Docling component. This issue affects Apache Camel: from 4.15.0 before 4.18.5, from 4.19.0 before 4.22.2. DoclingProducer builds the command line of the external docling tool and appends to it the custom arguments supplied in the CamelDoclingCustomArguments header, a list of strings whose flags the fix for CVE-2026-40047 restricted to an allowlist. That fix also introduced validatePathSafety to keep path-like argument values inside the intended directory, but the method only looked for relative traversal: it rejected a value containing a traversal sequence such as ../ and a value whose normalized form still held a .. element. It never checked whether a value was an absolute path, and an absolute path without traversal segments contains neither, so it passed unchanged. Every path-like value went through this check, whether it was a separate list element following its flag or was attached to the flag after an equals sign, so an absolute path was accepted in either form, for example as the directory given to the allowlisted --artifacts-path flag, which tells docling where to find its model artifacts. Such a value reached the docling command line as given, so a party able to choose the custom arguments could point the docling process at any directory the Camel process can access, outside the working directory that the check was meant to keep path values in, which is the outcome the CVE-2026-40047 fix set out to prevent. Only the CLI mode, which is the default, passes custom arguments to docling. The header belongs to the Camel header namespace, which consumers applying Camel's default header filtering, such as the JMS and Kafka consumers, remove from incoming messages, so the issue is reached where a route forwards externally influenced values into the header or where a consumer maps incoming headers without that filtering. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The 4.15.x to 4.17.x and 4.19.x to 4.21.x release lines reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. For deployments that cannot upgrade immediately, do not build the CamelDoclingCustomArguments header from values that originate outside the deployment, remove that header from messages that arrive from outside before they reach a docling endpoint, and make sure that consumers which map incoming message headers drop the Camel header namespace, so that only the route itself chooses the custom arguments.	5.4	More Details
CVE-2026-114133	Nuxt is an open-source web development framework for Vue.js. From 3.13.0 until 4.6.2, the server-island renderer serializes raw useHeadSafe entry input into the island response without preserving the safe-input marker, and NuxtIsland later inserts that input into the page head as an ordinary entry. An application that passes attacker-controlled data to useHeadSafe in a .server.vue component or page, or in a component rendered through NuxtIsland, can therefore render values normally removed by the allowlist, including event-handler attributes or executable script content, resulting in stored cross-site scripting in the application origin. This issue is fixed in version 4.6.2 .	5.4	More Details
CVE-2026-97353	The News Kit Addons For Elementor WordPress plugin through 1.4.7 does not perform capability or ownership checks before updating post metadata in one of its authenticated AJAX actions, allowing users with subscriber-level access and above to write arbitrary post meta, under attacker-chosen keys, to any post including those owned by other users. This can be used to overwrite protected core metadata and tamper with content created by higher-privileged users.	5.4	More Details
CVE-2026-	Incorrect Authorization vulnerability in Drupal Webform allows Server Side Request Forgery. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.4	More Details

96370			
CVE-2026-76287	In Splunk Enterprise versions below 10.4.4, a user that holds a role with the <code>preview_data_management_pipeline</code> capability or <code>edit_data_management_pipeline_job</code> capability could read or delete Data Management SPL2 preview jobs and results outside their ownership scope. The vulnerability is possible because the Data Management Representational State Transfer (REST) API verifies the required capability but does not verify that the requested preview job belongs to the requester. The affected Data Management REST API is available only when the Edge Processor feature is turned on. The Edge Processor feature is turned off by default. For more information see Set up a data management control plane (https://help.splunk.com/en/splunk-enterprise/process-data-at-the-edge/use-edge-processors-for-splunk-enterprise/10.4/getting-started/set-up-a-data-management-control-plane) and Getting sample data for previewing data transformations (https://help.splunk.com/en/splunk-enterprise/process-data-at-the-edge/use-edge-processors-for-splunk-enterprise/10.4/working-with-pipelines/getting-sample-data-for-previewing-data-transformations) in the Splunk documentation.	5.4	More Details
CVE-2026-114079	Trigger.dev before 4.5.6 contains an improper authorization vulnerability that allows holders of lower-trust environment API keys to cancel other environments' deployments because <code>DeploymentService.getDeployment</code> scopes lookups by <code>projectId</code> only. Attackers with a dev, preview, or CI secret key can POST to the deployment cancel endpoint to cancel in-progress production deployments within the same project.	5.4	More Details
CVE-2026-114076	Trigger.dev before 4.5.2 contains a blind server-side request forgery vulnerability that allows authenticated tenants to make the webapp server send POST requests to internal hosts by setting an unvalidated webhook alert channel URL. Attackers can point the webhook at internal services or cloud metadata addresses and trigger alerts to scan internal hosts, send state-changing POSTs, or reach IMDSv1 via redirects.	5.4	More Details
CVE-2026-114075	Trigger.dev before 4.5.12 contains an authorization bypass vulnerability that allows holders of capability-restricted user-actor tokens to create alert channels by using the legacy personal access token helper. Attackers with a read-only UAT can POST to the project <code>alertChannels</code> endpoint to create production webhook channels, routing future deployment and error alerts to attacker-controlled infrastructure.	5.4	More Details
CVE-2026-112333	GFI Archiver before 15.14 contains a stored cross-site scripting vulnerability in the MailInsights scheduled report configuration that allows authenticated attackers to inject arbitrary web script or HTML via the report name parameter to <code>/Archiver/MailInsights.aspx</code> . The injected payload is stored by <code>ReportScheduling.btnSaveReport_Click()</code> without output encoding and is executed in the browser of the user who created the scheduled report when they subsequently view the MailInsights page.	5.4	More Details
CVE-2024-4989	The callback URL parameter in WSO2 API Manager fails to validate user-supplied input, allowing it to be controlled by malicious actors. Exploitation of this vulnerability could lead to users being redirected to attacker-controlled web pages. This can be used in conjunction with social engineering techniques to trick users into interacting with malicious links, potentially resulting in phishing attacks aimed at acquiring sensitive information or causing other harm.	5.4	More Details
CVE-2026-115248	A vulnerability was detected in subhajitkhan online-clinic-management-system up to <code>e9ee77a8827a1446220fa07ee693dc4d9a29a578</code> . This issue affects some unknown processing of the file <code>patappointtime.php</code> of the component Booking Handler. The manipulation results in enforcement of behavioral workflow. It is possible to launch the attack remotely. The exploit is now public and may be used. The patch is identified as <code>e9ee77a8827a1446220fa07ee693dc4d9a29a578</code> . It is advisable to implement a patch to correct this issue. The project was informed of the problem early through an issue report but has not responded yet.	5.4	More Details
CVE-2026-107887	The Advanced Twenty Seventeen WordPress plugin through 1.3.1 does not perform any capability or nonce check on the action that creates and activates its bundled child , allowing any authenticated user such as a subscriber to switch the site's active site-wide, and allowing the same action to be triggered against any logged-in user including an administrator via CSRF.	5.4	More Details
CVE-2026-110733	Vulnerability in the SQL storage module of Apache SIS. When sorting a stream of features, the clause is copied after ORDER BY without quoting the identifier character. A crafter sort clause could inject SQL in the query. This issue affects the sql-storage module of Apache SIS from version 1.0 through 1.6 inclusive. Other modules such as sis-referencing are not impacted. Users are recommended to upgrade to version 1.7, which fixes the issue.	5.4	More Details
CVE-2026-115326	QloApps through 1.7.0 contains an insecure direct object reference vulnerability in <code>OrderDetailController::displayAjaxSubmitRefundRequest()</code> that allows authenticated customers to submit refunds for other customers' service lines. Attackers can submit foreign <code>id_service_product_order_detail</code> values with their own <code>id_order</code> to create refund records, cancel victims'	5.4	More Details

	service lines, and corrupt refund quantities.		
CVE-2026-90458	OpenStack Ironi before 39.0.1 allows DELETE and GET via vendor_passthru, e.g., to delete accounts from the BMC's Redfish service.	5.4	More Details
CVE-2026-115356	OpenBao before 2.6.4 and 2.7.x before 2.7.1 does not verify, when renewing a token issued by the Kubernetes Auth method, that the service account token originally used to authenticate is still valid. An authenticated attacker can keep renewing such tokens to persist access after the Kubernetes service account has been deleted.	5.4	More Details
CVE-2026-114175	Missing Authorization vulnerability in DigitalME eRoom eroom-zoom-meetings-webinar allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects eRoom: from n/a through 1.7.1.	5.4	More Details
CVE-2026-96371	Privilege Defined With Unsafe Actions vulnerability in Drupal Webform allows Stored XSS. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.4	More Details
CVE-2026-96372	Insufficient Granularity of Access Control vulnerability in Drupal Webform allows Forceful Browsing. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.4	More Details
CVE-2026-76435	A vulnerability in Cisco Expressway Series could allow an unauthenticated, remote attacker to obtain sensitive information from an affected device. This vulnerability is due to insufficient authorization checks of HTTP requests during Mobile and Remote Access (MRA) provisioning. An attacker could exploit this vulnerability by submitting a crafted URL request to an affected device. A successful exploit could allow the attacker to access sensitive information that they would otherwise not have authorization to access. Note: Cisco Expressway Series refers to Cisco Expressway Core (Expressway-C) devices and Cisco Expressway Edge (Expressway-E) devices.	5.3	More Details
CVE-2026-111596	The Kliken: Ads + Pixel for Meta WordPress plugin before 1.2.3 does not verify that a visitor holds an order's key before outputting that order's details on the order confirmation page, allowing unauthenticated users to retrieve the billing location, purchased products and totals of other customers' orders by enumerating order IDs.	5.3	More Details
CVE-2026-84920	Incorrect Authorization vulnerability in Drupal Unpublished Node Permissions allows Forceful Browsing. This issue affects Unpublished Node Permissions versions: from 0.0.0 to 1.8.0.	5.3	More Details
CVE-2026-53986	Girafit camera app (Android and iOS) version 1.9.1 and earlier, and Girafit camera devices, contain a hardcoded credentials vulnerability that allows attackers to obtain the shared TURN relay username by decompiling the app or intercepting unencrypted network traffic during live view registration. Attackers can extract the static TURN username and session token, which are identical across all installs and not tied to any user or device, potentially enabling tracking, impersonation, or further exploitation if combined with the TURN relay password.	5.3	More Details
CVE-2026-107789	The SmartyParcel WordPress plugin before 1.23.4 does not perform authorisation checks when serving shipping label downloads, allowing unauthenticated users to obtain the shipping labels of arbitrary orders, which contain customer names, addresses and phone numbers.	5.3	More Details
CVE-2026-109934	The Otter Blocks WordPress plugin before 3.2.8 does not correctly determine whether a post is publicly readable before returning its data through a REST API endpoint, treating any post with the published status as public, allowing unauthenticated users to read manual excerpts of password-protected posts, the content and titles of non-public post types, and, on WooCommerce sites, coupon codes and their descriptions.	5.3	More Details
CVE-2026-112826	Missing Authorization vulnerability in eLightUp Falcon - WordPress Optimizations & Tweaks falcon allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Falcon - WordPress Optimizations & Tweaks: from n/a through 2.12.0.	5.3	More Details
CVE-2026-111957	The SUMIT Payment Gateway for WooCommerce WordPress plugin before 4.0.0 does not bind the payment it confirms with the provider to the order being completed, nor check the payment amount against the order total, allowing unauthenticated users to mark a pending order as paid by replaying a payment identifier obtained from any single genuine payment.	5.3	More Details
CVE-2026-84914	Missing Authorization vulnerability in Drupal Calculate Working Days allows Forceful Browsing. This issue affects Calculate Working Days versions: from 0.0.0 to 2.0.3.	5.3	More Details

CVE-2026-108073	The SureCart WordPress plugin before 4.9.5 does not restrict which related objects an anonymous visitor may include when reading a single upsell funnel through its REST API, nor strip private fields from those objects, allowing unauthenticated users to read merchant-private catalog data such as pricing, commission, stock, SKU and tax internals, including unpublished products.	5.3	More Details
CVE-2026-114728	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a remote content blocking bypass in the handling of FuncIRI attributes. rcube_washtml detected the CSS url() token with a pattern that enumerated only some of its escaped spellings, so values using other CSS escape sequences, for example \000075r(or u\72 l, were not recognized and remote content was loaded from an HTML message without the user's consent. Upstream fixed this in the commit "Fix remote content blocking bypass via CSS escapes in FuncIRI attributes (again)", which resolves CSS escapes before matching and also inspects the `to`, `from` and `by` attributes.	5.3	More Details
CVE-2026-15261	The TI WooCommerce Wishlist WordPress plugin before 2.13.0 does not enforce the same owner-only access control on its REST API that it applies on the front-end, allowing unauthenticated users to read the metadata and item contents of private wishlists.	5.3	More Details
CVE-2026-109935	The Otter Blocks WordPress plugin before 3.2.8 does not enforce its block visibility conditions when a published post is rendered through the WordPress REST API, allowing unauthenticated users to read block content that the site owner restricted to specific roles, logged-in users, the post author, categories or tags, or (on the Pro tier) Stripe purchasers.	5.3	More Details
CVE-2026-115800	A vulnerability was identified in UbiquitousLearning mllm up to 2.0.2. The affected element is the function load_model of the file pymllm/mobile/convertor/_init_.py of the component Convertor. The manipulation of the argument shard_file leads to path traversal. The attack needs to be performed locally. The exploit is publicly available and might be used. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-76445	Multiple Cisco products are affected by a vulnerability in the Snort 3 Detection Engine that could allow an unauthenticated, remote attacker to cause the detection engine to unexpectedly restart, resulting in a denial of service (DoS) condition. The vulnerability is due to a flaw in logic handling when General Packet Radio Service (GPRS) Tunneling Protocol (GTP) extension headers are processed within the codec. An attacker could exploit this vulnerability by sending a crafted, fragmented UDP datagram that contains these headers to the Snort 3 Detection Engine of an affected device. A successful exploit could allow the attacker to cause the detection engine to enter an infinite loop, resulting in a DoS condition.	5.3	More Details
CVE-2026-105999	The Booking Package WordPress plugin before 1.7.31 does not perform authorization checks before returning calendar-account configuration in a response, allowing unauthenticated users to disclose sensitive settings including the administrator's notification email addresses.	5.3	More Details
CVE-2026-114086	Trigger.dev before 4.5.5 contains a missing authentication vulnerability in the realtime streams POST handler that allows unauthenticated attackers to inject data into run streams. Attackers can send POST requests to /realtime/v1/streams/:runId/:streamId with a predictable run friendlyId to inject arbitrary data into any environment's realtime stream, including production.	5.3	More Details
CVE-2026-115678	Arm Keil MDK-Middleware 8.0.0-8.3.0 may allow pre-authentication out-of-bounds read in SNMP processing	5.3	More Details
CVE-2026-114348	webpack-dev-server is a development server that provides live reloading for webpack projects. In versions from 5.1.0 before 5.2.7 and in version 6.0.0, the unauthenticated open-editor endpoint validates the Host header with a regular expression meant for browser Origin schemes, so a Host value beginning with file: or with an extension scheme short-circuits the hostname and allowedHosts checks. A network peer that can reach the dev server can send a request with a malformed Host that passes validation and then causes the endpoint's URL construction to throw, and in the default non-production mode the server returns an HTTP 500 body containing the full exception stack trace, which discloses absolute filesystem paths, installed module locations, and internal source line numbers. The server keeps running, so there is no denial of service. The issue is fixed in webpack-dev-server 5.2.7 and 6.1.0, and users should upgrade to one of those versions or later.	5.3	More Details
CVE-2026-116154	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak in its XML parsing code, where allocated memory is not released when parsing reaches the nesting depth limit. Attackers can supply crafted, deeply nested XML data to trigger the leak, causing memory exhaustion and denial of service.	5.3	More Details
CVE-2026-114062	SvelteKit before 3.0.0 contains a server-side request forgery vulnerability that allows remote attackers to reach internal services because self-fetch fallbacks follow HTTP redirects automatically. Attackers can manipulate the Host or forwarded headers to point fallback requests at their server, which redirects them to internal services.	5.3	More Details

CVE-2026-116155	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak in the handling of XMP profiles by various decoders. Attackers can supply images with a crafted XMP profile that leak memory each time they are processed, leading to resource exhaustion and denial of service.	5.3	More Details
CVE-2026-116394	Improper Neutralization of Script-Related HTML Tags in a Web Page (Basic XSS) vulnerability in Magepeople inc. WpEvently mage-eventpress allows Code Injection.This issue affects WpEvently: from n/a through 5.7.9.	5.3	More Details
CVE-2026-110893	Missing Authorization vulnerability in Brainstorm Force SureDash suredash allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects SureDash: from n/a through 1.12.2.	5.3	More Details
CVE-2026-114044	FreeRDP before 3.32.1 contains a heap out-of-bounds read vulnerability in the Android client that allows malicious RDP servers to read beyond surface buffers by supplying unvalidated mappedWidth and mappedHeight values. Attackers can send a crafted RDPGFX_MAP_SURFACE_TO_WINDOW_PDU with oversized dimensions, causing android_UpdateWindowFromSurface to copy heap memory beyond the allocation, leading to crashes or memory disclosure.	5.3	More Details
CVE-2026-114714	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a path traversal vulnerability in the squirrelmail_usercopy plugin. With the file based backend, the source path was built from the login name without rejecting path separators, so a login name containing a slash could cause files outside the configured squirrelmail_data_dir to be read and imported. Upstream fixed this in the commit “Fix unsafe use of a login name in SquirrelMail import file path”, which rejects login names containing a slash.	5.3	More Details
CVE-2026-110894	Missing Authorization vulnerability in Brainstorm Force SureDash suredash allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects SureDash: from n/a through 1.12.2.	5.3	More Details
CVE-2026-114491	pdfmake 0.3.8 before 0.3.12 fails to apply the setLocalAccessPolicy() restriction to images referenced inside SVG content, allowing attackers to read blocked local files. Attackers supplying a document definition can add an SVG image href pointing to a local path, embedding restricted file contents in the generated PDF.	5.3	More Details
CVE-2026-111842	Missing Authorization vulnerability in WP Sunshine Sunshine Photo Cart sunshine-photo-cart allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Sunshine Photo Cart: from n/a through 3.7.5.	5.3	More Details
CVE-2026-112218	Missing Authorization vulnerability in Ahmad Majestic Support majestic-support allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Majestic Support: from n/a through 1.2.1.	5.3	More Details
CVE-2026-114699	Missing Authorization vulnerability in yoomoney ЮKassa для WooCommerce yookassa allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects ЮKassa для WooCommerce: from n/a through 2.17.3.	5.3	More Details
CVE-2026-114696	Missing Authorization vulnerability in miniOrange Page Restriction WordPress (WP) page-and-post-restriction allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Page Restriction WordPress (WP): from n/a through 1.6.0.	5.3	More Details
CVE-2026-116542	Netty netty-tcnative before 2.0.84.Final contains an SNI hostname spoofing vulnerability that allows remote unauthenticated clients to alter hostnames via JNI NewStringUTF decoding. Attackers can send overlong encodings like 0xC0 0x80 in the ClientHello server_name to confuse SNI-based certificate selection, access decisions, and logging.	5.3	More Details
CVE-2026-114195	Insertion of Sensitive Information into Externally-Accessible File or Directory vulnerability in WP Desk ShopMagic shopmagic-for-woocommerce allows Retrieve Embedded Sensitive Data.This issue affects ShopMagic: from n/a through 4.9.1.	5.3	More Details
CVE-2013-10077	A weakness has been identified in OctoPrint 1.0.0. This impacts the function getSettings of the file src/octoprint/server/ajax/settings.py of the component API Key Handler. This manipulation causes information disclosure. The attack is possible to be carried out remotely. The exploit has been made available to the public and could be used for attacks. Upgrading to version 1.1.0-rc1 will fix this issue. Patch name: eccc9d6fbde0c46cdd6e255fd5ce9a30077ba954. You should upgrade the affected component. Affected is a 2013 release where access control was not yet implemented and the endpoint was part of first-run setup). The API key has been admin-gated since 1.2.0 and GET /api/settings is protected by Permissions.SETTINGS_READ plus reauthentication in all current releases.	5.3	More Details
CVE-	rclone serve s3 versions 1.75.0 before 1.75.2 contain an integer overflow vulnerability in multipart reorder-memory accounting that bypasses the configured buffer limit. Attackers can open two		More

2026-116188	concurrent current-part PUT requests declaring huge lengths to wrap the signed int64 counter and retain out-of-order parts in memory, exhausting the host.	5.3	Details
CVE-2026-114091	Trigger.dev before 4.5.0 looks up runs by friendlyld alone in the api.v1.runs.\$runParam.replay route without scoping the lookup to the caller's runtime environment. An attacker holding any tenant's API key who knows a victim's run friendlyld can replay that run into their own environment and read the victim's task payload.	5.3	More Details
CVE-2026-114460	A flaw was found in mod_qos. An off-by-one pointer calculation error occurs when merging request body parameters with an existing URL query, which truncates the query string during security evaluation while leaving the actual request unchanged. An unauthenticated remote attacker could exploit this vulnerability by submitting a crafted request that conceals prohibited values from inspection, resulting in a bypass of configured security rules.	5.3	More Details
CVE-2026-103306	The Online Payments with iK Pay Gateway WordPress plugin before 3.0.4 does not properly verify the authenticity of its payment gateway callback on a default installation, allowing unauthenticated attackers to forge a successful payment notification and mark arbitrary orders as paid without any payment being made, and to mark arbitrary unprocessed orders as failed.	5.3	More Details
CVE-2026-105998	The Booking Package WordPress plugin before 1.7.31 does not perform any authorization or validity check before resetting its premium subscription state, allowing unauthenticated attackers to disable the Booking Package WordPress plugin before 1.7.31's paid features.	5.3	More Details
CVE-2026-116053	A weakness has been identified in dbcli mycli up to 2.29.4/2.30.1. The impacted element is the function source_special_command_is_safe of the file mycli/packages/special/source.py of the component Whitelist. Executing a manipulation can lead to os command injection. It is possible to launch the attack on the local host. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2026-105992	The Ettrick Escape Room Booking WordPress plugin before 1.3.1 does not verify a booking's payment status against the payment provider before marking it paid and confirmed, allowing unauthenticated attackers to mark bookings as paid without paying.	5.3	More Details
CVE-2026-103302	The Payment Gateway for Idram WordPress plugin before 2.1.6 does not verify the authenticity of incoming payment notifications, because the merchant credentials used to compute their checksum are empty by default, allowing unauthenticated attackers to forge a valid payment confirmation and mark arbitrary orders as paid, and to force arbitrary orders into a failed state regardless of the gateway's configuration.	5.3	More Details
CVE-2026-116141	ImageMagick before 7.1.2-33 contains a policy bypass vulnerability in the UHDR decoder that allows attackers to evade image width and height limits because extent checks occur after decoding. Attackers can supply crafted UHDR images with oversized dimensions that are fully decoded before policy limits are enforced, causing excessive resource consumption.	5.3	More Details
CVE-2026-76487	A vulnerability in Cisco Unified Communications Manager (Unified CM) and Cisco Unified Communications Manager Session Management Edition (Unified CM SME) could allow an unauthenticated, remote attacker to conduct access control bypass attacks on an affected device. This vulnerability is due to improper input validation for specific HTTP requests. An attacker could exploit this vulnerability by sending a crafted HTTP request to an affected device. A successful exploit could allow the attacker to access parts of the network that are beyond the scope of their intended access.	5.3	More Details
CVE-2026-114180	Missing Authorization vulnerability in Wpmet Metform metform allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Metform: from n/a through 4.3.2.	5.3	More Details
CVE-2026-103319	The OPay Payment Gateway for WooCommerce WordPress plugin through 2.5.5 does not properly verify that payment notifications originate from the payment provider, as the secret used to validate their signature is empty by default, allowing unauthenticated attackers to forge a valid notification and mark arbitrary orders as paid, cancelled or failed.	5.3	More Details
CVE-2026-114116	Vyper 0.1.0b1 before 0.1.0b17 and 0.4.1 through 0.5.0b1 builds its source-line table with str.splitlines(), which treats characters like form feed and U+2028 as line breaks the Python parser ignores, so numeric literals are classified and validated from shifted source text. An attacker contributing source code can hide these characters in comments, strings or indentation to silently change address, integer or decimal constants.	5.3	More Details
CVE-2026-116152	ImageMagick before 7.1.2-33 contains a memory leak vulnerability in the MAT decoder that fails to release allocated memory in a specific error path. Attackers can supply crafted MAT files that trigger this error path repeatedly, causing small memory leaks that gradually consume process memory.	5.3	More Details

CVE-2026-103300	The Eurobank WooCommerce Payment Gateway WordPress plugin through 2.1.1 does not verify that a payment-confirmation request genuinely originates from the payment provider, and ships with the secret that would authenticate it empty by default, allowing unauthenticated attackers to mark any pending order as paid, and to force any order into a failed status without any secret at all.	5.3	More Details
CVE-2026-114114	Vyper 0.1.0b1 before 0.1.0b12, 0.1.0b14, and 0.4.1 through 0.5.0b1 mix character-based line offsets with Python's UTF-8 byte col_offset values in vyper/ast/parse.py, so multi-byte characters shift the source slices used to reconstruct later numeric literals. An attacker contributing source with non-ASCII identifiers that NFKC-normalize to ASCII, or an invisible byte-order mark, can silently substitute addresses, integers or other constants in compiled code.	5.3	More Details
CVE-2026-103298	The CCAvenue Payment Gateway for WooCommerce WordPress plugin through 3.1 does not verify that a payment-confirmation request genuinely originates from the payment provider, and derives its verification key from a value that is empty on a default install, allowing unauthenticated attackers to forge a confirmation and mark any WooCommerce order as paid without payment.	5.3	More Details
CVE-2026-114110	Astral uv before 0.11.6 contains a path traversal vulnerability that allows attackers to delete arbitrary files by supplying wheels with relative traversal paths in RECORD entries. Attackers can craft a malicious wheel whose RECORD references files outside the installation prefix, causing uv to delete those files when the user uninstalls the package.	5.3	More Details
CVE-2026-114501	Pillow is a Python imaging library. From 11.2.1 until 13.0.0, Arrow export assumes that an image's rows are tightly contiguous and exports from im->blocks[0].ptr even when PILLOW_ALIGNMENT makes aligned_linesize include row padding or an aligned offset before the first row. The resulting Arrow buffer can include uninitialized padding, bytes preceding the first row, or stale heap contents from previously freed images. Applications that configure alignment and export images to Arrow can therefore disclose one user's pixels or other process memory to another user. This issue is fixed in version 13.0.0.	5.3	More Details
CVE-2026-115479	baserCMS is a website development framework. In versions 5.2.5 before 5.3.2 and 5.4.0 before 5.4.2, an unauthenticated visitor can supply a comma-separated address list in a public mail form's email field and cause a single auto-reply to be addressed to multiple third parties. MailMessageMailer::sendFormToUser() passes the visitor-controlled userMail value to parseAndNormalizeMailAddresses(), which splits it with explode(',') and supplies every resulting address to setTo(). Newly created forms with an email-type field are affected when the separate, optional VALID_EMAIL server-side validation is not enabled; browser type=email validation does not prevent crafted requests. The bundled contact form enables VALID_EMAIL and is not affected. The auto-reply uses the site's sender identity and includes partially attacker-controlled content, allowing abuse of the site to deliver unsolicited mail. This issue has been patched in versions 5.3.2 and 5.4.2.	5.3	More Details
CVE-2026-114103	adm-zip 0.6.0 contains a decompression bomb protection bypass that allows attackers to evade the maxLength guard by targeting the asynchronous getDataAsync() path in methods/inflater.js. Attackers can craft ZIP entries with falsified small declared sizes and high-ratio DEFLATE payloads to exhaust memory in applications reading entries asynchronously.	5.3	More Details
CVE-2026-116193	Rclone from 1.64.0 before 1.75.2 contains an authentication bypass vulnerability in serve webdav that lets unauthenticated attackers skip authentication by sending OPTIONS requests. Attackers can probe arbitrary paths and compare Allow header values to determine whether files or collections exist, enabling namespace enumeration.	5.3	More Details
CVE-2026-87837	The Wava Payment Plugin WordPress plugin before 0.3.10 does not verify the authenticity of incoming payment-provider notifications, nor restrict who may send them, allowing unauthenticated attackers to change the status of any WooCommerce order in the store, including marking an unpaid order as paid and completed, and cancelling other customers' orders. In versions up to 0.3.7 the response additionally returns the affected order's billing details, disclosing customer names, addresses, email addresses and phone numbers to the same unauthenticated caller.	5.3	More Details
CVE-2026-97255	Authorization Bypass Through User-Controlled Key vulnerability in WProyal Royal Elementor Addons royal-elementor-addons allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Royal Elementor Addons: from n/a through 1.7.1069.	5.3	More Details
CVE-2026-115480	baserCMS is a website development framework. In versions 5.0.0 before 5.3.2 and 5.4.0 before 5.4.2, the public POST /baser/api/bc-mail/mail_messages/add/{mailContentId}.json endpoint accepts submissions without enforcing the form's auth_captcha setting. BcMail\Controller\Api\MailMessagesController::add() and MailMessagesService::create() omit the BcCaptchaService::check() verification of captcha_id and auth_captcha performed by MailController::submit() in the web UI. An unauthenticated attacker who can reach the API can therefore automate submissions without solving the enabled image challenge. The API disables FormProtection and has no rate limit; submissions can generate administrator-inbox spam, send attacker-influenced auto-replies to third parties when an email field is present, and accumulate mail_messages database	5.3	More Details

	records when save_info is enabled. The API already enforces the isAccepting() check added for CVE-2026-30878; the missing CAPTCHA check is a separate control failure. This issue has been patched in versions 5.3.2 and 5.4.2.		
CVE-2026-115959	Missing Authorization vulnerability in WPExperts APIExperts Square for WooCommerce woosquare allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects APIExperts Square for WooCommerce: from n/a through 4.7.5.	5.3	More Details
CVE-2026-113492	The GeoDirectory WordPress plugin before 2.8.194 does not perform status, ownership or password checks when loading listings for one of its unauthenticated AJAX handlers, allowing unauthenticated users to disclose the content of non-public (draft, pending and private) listings by supplying their IDs.	5.3	More Details
CVE-2026-116095	A vulnerability was identified in smg-project smg 1.7.0/1.8.0/1.9.0/1.10.0/1.10.1. Affected is the function jpeg_turbo::decode_jpeg_rgb of the file crates/multimodal/src/jpeg_turbo.rs of the component llm-multimodal. The manipulation leads to uncontrolled memory allocation. The attack can be initiated remotely. The exploit is publicly available and might be used. Upgrading to version 1.11.0 is able to address this issue. The identifier of the patch is 8f51b57e219a914089e345b704c3c8de37a3bb8e. It is advisable to upgrade the affected component.	5.3	More Details
CVE-2026-103634	Out-of-bounds read and write in the Count-Min sketch deserialization of Apache DataSketches C++ (repo: datasketches-cpp). count_min_sketch::deserialize() did not include the preamble in its input size check, so a truncated sketch could cause a read of up to 16 bytes past the end of the input. In addition, the table size was computed from the serialized number of buckets and number of hash functions in 32-bit arithmetic. A crafted sketch could make it wrap to zero, so that the sketch deserialized with an empty table, and later updates and estimate queries read and wrote outside the heap allocation. This can corrupt heap memory, causing a crash and potentially enabling further exploitation. This issue affects Apache DataSketches C++: from 4.1.0 before 5.3.0. Only applications that deserialize Count-Min sketches from untrusted sources are affected. Users are recommended to upgrade to version 5.3.0, which fixes this issue.	5.3	More Details
CVE-2026-114401	Trudesk through 1.2.11 contains a regular expression injection vulnerability that allows unauthenticated attackers to enumerate usernames by submitting regex patterns to POST /api/v1/login. Attackers can send alternations or prefixes like (admin root) or a.* and distinguish 'Invalid User' from 'Invalid Password' responses to discover valid accounts.	5.3	More Details
CVE-2026-116096	A weakness has been identified in ggml-org llama.cpp. Affected by this issue is the function list_pockettts_state_slots of the file tools/mtmd/models/pockettts-gen.cpp of the component Pocket-TTS. This manipulation of the argument mmproj causes out-of-bounds write. The attack requires local access. The exploit has been made available to the public and could be used for attacks. Upgrading to version b11796 can resolve this issue. Patch name: 0a1ef309cddb9f2024f58a38fb4c42352c5db569. The affected component should be upgraded. The reported heap out-of-bounds write (gen-audio / pocket-tts state_out path) is no longer a silent exploitable heap write on current master and the latest release b11871 - a malicious mmproj that previously triggered it now trips a release-active GGML_ASSERT and aborts. But the submitter's recommended by-construction fix (PR #29360 "prevent OOB accesses from malformed mmproj files") is still open. So the root cause is detected-and-aborted, not prevented-by-design.	5.3	More Details
CVE-2026-115595	Missing Authorization vulnerability in Supsysic Contact Form by Supsysic contact-form-by-supsysic allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Contact Form by Supsysic: from n/a through 1.10.7.	5.3	More Details
CVE-2026-114392	Trudesk through 1.2.11 contains an authorization bypass vulnerability that allows authenticated users to read tickets outside their groups via GET /api/v1/tickets/:uid. Attackers can enumerate sequential ticket uids to read every ticket's subject, issue, comments, owner and assignee, because the apiTickets.single handler skips group membership checks.	5.3	More Details
CVE-	The device multiplexer pseudo-driver in drivers/misc/devmux/devmux.c exposes two system calls, devmux_select_get() and devmux_select_set(). When CONFIG_USERSPACE is enabled, their user-mode verifiers z_vrfy_devmux_select_get() and z_vrfy_devmux_select_set() forwarded directly to the kernel implementations without performing any kernel-object validation — no k_object_access_check() and no driver object check. Compounding this, DEVICE_DT_INST_DEFINE() registered devmux instances with a NULL driver API pointer, which causes scripts/build/gen_kobject_list.py to skip them entirely, so a devmux device had no struct k_object and no permission on it could be granted or withheld in the first place. Any thread running in user mode can reach the defect: device pointers such as DEVICE_DT_GET(DT_NODELABEL(devmux0)) are compile-time constants available to application code, and the syscall accepted them from a thread that had never been granted access to the device. The kernel implementations do validate the pointer against the static devmux instance table and bounds-check the selection index, so this is not a memory-safety issue; the failure is purely one of		More

2026-19947	authorization. Exposure requires CONFIG_USERSPACE=y together with the experimental CONFIG_DEVMUX=y (which depends on CONFIG_DEVICE_MUTABLE) and a zephyr,devmux devicetree node. An unprivileged user-mode thread can read and change which downstream device a multiplexer forwards to, bypassing the userspace object-permission model. Because a devmux forwards the data, configuration and API of the selected device — its documented use being zephyr,console and zephyr,shell-uart switching — an attacker can suppress console, log and shell output, redirect that output to a downstream device it is able to observe, or flip the selection underneath another thread that is legitimately using the multiplexer. The impact is bounded to the devicetree-fixed set of downstream devices; no out-of-bounds access or arbitrary pointer dereference is reachable. The fix declares __subsystem struct devmux_driver_api and attaches a DEVICE_API(devmux, devmux_api) instance to each devmux device so that it becomes a tracked K_OBJ_DRIVER_DEVMUX kernel object, and adds K_OOPS(devmux_inst_get(dev) == SIZE_MAX) plus K_OOPS(k_object_access_check(dev) != 0) to both verifiers.	5.3	Details
CVE-2026-114390	Chamilo LMS from 2.0.0-alpha.1 before 2.0.0-RC.3 contains a missing authorization vulnerability in legacy AJAX handlers under public/main/inc/ajax/ that exposes user names and certificates. Unauthenticated or low-privileged users can call actions like get_users_by_group_course, user_by_all_roles, and export_all_certificates to harvest user names and other learners' certificate PDFs.	5.3	More Details
CVE-2026-114388	Chamilo LMS 2.0.0 pre-releases from 2.0.0-alpha.1 before 2.0.0-RC.3 contain a missing access control vulnerability in GradebookController that allows unauthenticated users to modify gradebook certificates. Attackers can request the set_default_certificate endpoint with any course id and document id to tamper with certificate configuration for any course.	5.3	More Details
CVE-2026-115801	A security flaw has been discovered in Kakao khaiii up to 0.4. The impacted element is the function khaiii::nn::Linear::open of the file src/main/cpp/khaiii/nn/Linear.cpp of the component JSON Config File Handler. The manipulation of the argument hidden_dim/class_num results in out-of-bounds read. The attack requires a local approach. The exploit has been released to the public and may be used for attacks. The vendor explains: "khaiii treats the resource directory (--rsc-dir) as trusted input, like a binary or a shared library. To use this defect, an attacker must already be able to write to that directory. With that access, an attacker can change the analysis results without this defect." This vulnerability only affects products that are no longer supported by the maintainer.	5.3	More Details
CVE-2026-114248	A security vulnerability has been detected in dhikshitha29 Insurance-Management-System 91de5edda5baa109ad8e9baf03ce6a3ceda84183/125362cc21dca20d74091a4913a5096bb58dd837. This affects an unknown part of the file /Admin/includes/AddNewPlan-inc.php of the component Plan Management Form. The manipulation of the argument submit/Plan_no/Name leads to missing authentication. 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 continuous delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-111974	The GiveWP WordPress plugin before 4.19.0.1 does not restrict the currency of a donation to the one the site or form accepts, while comparing its amount limits as bare numbers, allowing unauthenticated users to make donations below a form's configured minimum or fixed amount.	5.3	More Details
CVE-2026-110285	The WP Stripe Checkout WordPress plugin before 1.2.2.62 does not verify the Stripe webhook signature when no signing secret is configured, which is its default state, allowing unauthenticated attackers to forge a payment completion notification and create a completed order without paying.	5.3	More Details
CVE-2026-96356	Incorrect Authorization vulnerability in Drupal Webform allows Forceful Browsing. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.3	More Details
CVE-2026-115972	A vulnerability was determined in FaveoSuite Faveo Helpdesk up to 2.0.3. Affected by this vulnerability is the function UnAuthController.changeStatus/UnAuthController.close/UnAuthController.open/UnAuthController.resolve of the file app/Http/Controllers/Client/helpdesk/UnAuthController.php. Executing a manipulation of the argument id/status can lead to missing authorization. The attack can be launched remotely. The vendor was contacted early about this disclosure.	5.3	More Details
CVE-2026-100142	The FluentCart A New Era of eCommerce WordPress plugin before 2.2.0 does not verify that a stored customer address selected during checkout belongs to the requester, allowing unauthenticated attackers to attach any customer's stored address to an order they place and read it back, disclosing customers' postal addresses, phone numbers and company details.	5.3	More Details
CVE-2026-96364	Improper Control of Interaction Frequency vulnerability in Drupal Webform allows Functionality Misuse. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.3	More Details

CVE-2026-115614	Missing Authorization vulnerability in Magepeople inc. Booking and Rental Manager booking-and-rental-manager-for-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Booking and Rental Manager: from n/a through 2.8.2.	5.3	More Details
CVE-2026-55465	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 22.04 to before version 26.05, the unauthenticated REST endpoint GET /api/v1/index/{solrField}/spatial/search concatenates the user-supplied region query parameter directly into a Solr query string with no escaping, no local-param stripping, and no runtime schema enforcement (the OpenAPI @Schema(pattern=...) annotation is documentation-only at the JAX-RS layer). An attacker can inject Solr local-param syntax ({!parent ...}, {!join ...}) and boolean operators to make access-restricted georeferenced records match the query, bypassing the access-condition filter. This issue has been patched in version 26.05.	5.3	More Details
CVE-2026-111817	Apache Traffic Server can stall a chunked transfer that it passes through with its trailers removed (proxy.config.http.drop_chunked_trailers) until the connection times out, because it waits for the trailer bytes it has already dropped. Trailer dropping is on by default in 10.x. In 8.1.11 and 9.2.x it is off by default, and those versions are affected only if it is enabled. This issue affects Apache Traffic Server: 8.1.11, from 9.2.5 through 9.2.15, from 10.0.0 through 10.1.4, 10.2.0. Users are recommended to upgrade to version 10.1.5 or 10.2.1, which fix the issue.	5.3	More Details
CVE-2026-55444	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. Prior to version 26.05, the legacy RSS servlet (RssResolver, GET /rss/) accepts the q and filterQuery parameters and concatenates them into Solr queries without validation. The servlet predates the REST API layer and lacks the input validation present in GET /api/v1/records/rss. q forms the main query body (wrapped only in literal parentheses); filterQuery is passed directly as a Solr fq. An unauthenticated attacker can inject Solr boolean operators to break out of the wrapping and render the trailing access-condition filter-query suffix optional, so that access-restricted records are included in the generated RSS feed. The q and filterQuery parameters are the same underlying vulnerability in the same servlet; q (the main query body) has the broader scope. This issue has been patched in version 26.05.	5.3	More Details
CVE-2026-55442	The Goobi viewer is a web application that allows digitised material to be displayed in a web browser. From version 23.03 to before version 26.05, SolrTools.cleanUpQuery() strips Solr local-parameter syntax ({!...}) from user-supplied query strings before they reach Solr. Its regex can be bypassed by wrapping the local parameter in a second pair of curly braces ({!{!...}}): the outer braces are stripped, the inner {!...} survives intact, and a valid Solr local parameter is forwarded. The flaw is in the shared sanitizer, so it affects any endpoint that routes user input through cleanUpQuery. The double-brace bypass applies from v23.10 (the greedy \{(.+)\} regex); a related bare-!join rewrap path is present. This issue has been patched in version 26.05.	5.3	More Details
CVE-2026-49442	Cacti is an open source performance and fault management framework. Prior to version 1.2.31, a User Enumeration vulnerability allows an attacker to identify valid usernames on the Cacti service. When a user attempts to log in, the application returns slightly different error messages depending on whether the provided username exists in the system. When username exists on Cacti, in the response, they see Access Denied! Login failed. When it does not exist on Cacti, they see Access Denied! Login Failed. The only difference between the responses is the capitalization of the letter "f" in the word failed. This issue has been patched in version 1.2.31.	5.3	More Details
CVE-2026-115971	A vulnerability was found in Travisun Opptrix 0.4.0. Affected is the function Buffer.concat of the file /api/channels/webhook/ of the component preParsing handler. Performing a manipulation results in denial of service. The attack can be initiated remotely. The exploit has been made public and could be used. The project has been completely taken offline and deprecated. The project maintainer explains: "The mentioned security risk was a design flaw that manifested primarily in local environments during earlier stages of version development. We previously addressed and resolved this specific risk; however, because the project was restricted to a small learning scope and the main versions were intended to run exclusively on localhost, the security fix was only pushed directly to our GitHub repository before it went private." This vulnerability only affects products that are no longer supported by the maintainer.	5.3	More Details
CVE-2026-96366	Incorrect Authorization vulnerability in Drupal Webform allows Forceful Browsing. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	5.3	More Details
CVE-2026-116255	A vulnerability has been found in PGYER codefever up to 0c35e60f006f9d558cc33a159baf110659351990. Affected is the function _detail of the file application/controllers/doc.php of the component URI Sanitizer. The manipulation of the argument uri_string leads to path traversal. The attack may be initiated remotely. The exploit has been disclosed to the public and may 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 e4433eee96883a222403854e8ac38b3146774a07. Applying a patch is the recommended	5.3	More Details

	action to fix this issue.		
CVE-2026-115958	Missing Authorization vulnerability in WP Discussion Board WP Discussion Board wp-discussion-board allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Discussion Board: from n/a through 2.6.0.	5.3	More Details
CVE-2026-116205	YesWiki before 4.7.0 contains an unauthenticated server-side request forgery vulnerability that allows remote attackers to fetch arbitrary URLs by supplying [[[URL]]] wiki syntax to the render handler. Attackers can pass the syntax in the content parameter of the page render handler to reach internal hosts and ports and read their responses.	5.3	More Details
CVE-2026-110377	Missing Authorization vulnerability in VillaTheme Bopo – WooCommerce Product Bundle Builder bopo-woo-product-bundle-builder allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Bopo – WooCommerce Product Bundle Builder: from n/a through 1.2.3.	5.3	More Details
CVE-2026-113799	Slack Morphism is a client library for Slack Web/Events API/Socket Mode and Block Kit. Versions 1.0.0 through 2.27.0 expose the correct Slack request HMAC when signature verification fails. For an invalid request, `slack-morphism` computes the correct signature for the attacker-controlled timestamp and request body and stores it in `SlackEventWrongSignatureError.generated_hash`. The error's `Debug` and `Display` representations expose that value, and the default Axum listener error handler logs it at error level. An attacker who can both reach the application's Slack callback endpoint and read its application/service logs can therefore use the verifier as a chosen-message HMAC oracle. The attacker can send an arbitrary request with an invalid signature, recover the correct signature from the resulting log entry, and resend the same request within the five-minute freshness window. The second request passes Slack Morphism's signature verification and reaches the protected application callback without the attacker knowing the Slack signing secret. Version 2.28.0 patches the issue.	5.3	More Details
CVE-2026-113490	The GeoDirectory WordPress plugin before 2.8.194 does not perform any status or ownership check when rendering a listing's map-popup through one of its single-item REST routes, allowing unauthenticated users to read the title, permalink and card fields of non-public (draft, pending, private and trashed) listings by supplying the listing ID.	5.3	More Details
CVE-2026-82953	NULL pointer dereference in the RTMP protocol implementation in Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to crash the server (denial of service) via a crafted createStream command. Root Cause: When an RTMP service rejects a createStream request by returning a null stream from RtmpService::NewStream(), the server sends an error response to the client. If writing this rejection response fails (for example, because the client closes the connection immediately after sending the createStream command), the error-handling path in RtmpChunkStream::OnCreateStream() dereferences the null stream pointer to access stream->_onfail_id, causing the process to crash. Affected Scenarios: Run an Apache bRPC RTMP server whose RtmpService::NewStream() may return null (i.e., the service rejects createStream requests). A remote client that sends a createStream command and then closes the connection before the rejection response is written can crash the server. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3508	5.3	More Details
CVE-2026-111976	The GiveWP WordPress plugin before 4.19.0.1 does not perform any status or capability check when rendering a donation form on a public front-end route, allowing unauthenticated visitors to view the content and settings of unpublished (draft, pending, private and scheduled) donation forms that WordPress core otherwise keeps non-public.	5.3	More Details
CVE-2026-111849	Missing Authorization vulnerability in Magepeople inc. WpTravelly tour-booking-manager allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WpTravelly: from n/a through 2.3.5.	5.3	More Details
CVE-2026-115327	libyang through 6.4.13 contains an uninitialized pointer free vulnerability in lydxml_env_netconf_rpc() in src/parser_xml.c when parsing NETCONF rpc action envelopes. Peers with a NETCONF session can send an action element containing non-whitespace text via lyd_parse_op() to crash servers like netopeer2 or corrupt the heap.	5.3	More Details
CVE-2026-107447	The Contest Gallery WordPress plugin before 34.0.1 does not reliably validate the visitor identity used to enforce its one-vote-per-entry restriction, trusting a client-supplied request header and an unbound client-generated token, which allows unauthenticated users to bypass the restriction and cast multiple votes for the same contest entry, inflating its vote total.	5.3	More Details
CVE-2026-114731	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by an uncontrolled resource consumption vulnerability in the CSS proxy (the utils/modcss action). The fetched stylesheet was read into memory in full, with no size limit and no read timeout, so a stylesheet URL referencing an attacker-controlled host could be used to exhaust server memory. Upstream fixed this in the commit “Add a size limit to the CSS proxy content processing”, which caps the response at 5 MB, reads it as a stream and	5.3	More Details

	applies a read timeout.		
CVE-2026-94295	The Chikkili- Google Pay India for Woocommerce WordPress plugin through 1.0.6 does not verify that a payment notification actually originated from the payment provider before marking the associated order as paid, allowing unauthenticated attackers to mark arbitrary orders as paid without any payment being made.	5.3	More Details
CVE-2026-116259	A security flaw has been discovered in libimobiledevice up to 1.4.0. This vulnerability affects the function <code>afc_get_device_info_key</code> of the file <code>src/afc.c</code> of the component Device Info Handler. Performing a manipulation of the argument Value results in memory corruption. The attack requires a local approach. The exploit has been released to the public and may be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2026-105265	A vulnerability exists in the the time zone parameter and does not properly validate user input, which may allow an authenticated user to cause <code>/etc/localtime</code> to point to a file outside the intended <code>zoneinfo</code> directory .	5.3	More Details
CVE-2026-111100	The ONVIF device service on the E3 responds to certain informational requests without requiring authentication. A remote, unauthenticated party can retrieve detailed device identifiers along with a list of configured user accounts and their privilege levels. This aids attackers in fingerprinting vulnerable devices and preparing targeted credential attacks.	5.3	More Details
CVE-2026-109125	In NLnet Labs Unbound up to and including version 1.26.1, an observable cache-level side-channel can be exposed with 'unwanted-reply-threshold > 0' configurations (not a default configuration). If a response-like UDP datagram reaches the upstream source port currently used by a recursive query, but its TXID does not match any pending query, Unbound counts the packet as an unwanted reply. When the accumulated count reaches the configured threshold, Unbound clears the worker's RRset cache and message cache. This cache clear action does not destroy the current upstream UDP pending query. Therefore, an external observer can distinguish whether a candidate port matches the current upstream source port by checking whether an existing cached item has been cleared. This is a cache side-channel vulnerability that can be used to probe the upstream source port. Although the attack is not feasible on default configurations and the suggested value (10 million) makes it an unreliable attack on a public resolver, in scenarios where the attacker can keep a concurrent query window alive, the attacker can first filter the source port and then perform birthday-style guessing against the TXID, making cache poisoning a feasible attack path.	5.3	More Details
CVE-2026-105256	The LatePoint WordPress plugin before 5.7.5 does not validate the client-supplied value that determines how much is charged for a booking against the amounts the service and site actually allow, allowing an unauthenticated visitor to make the amount due resolve to zero and obtain a confirmed appointment and a recorded order without completing any payment.	5.3	More Details
CVE-2026-93628	An attendee authentication vulnerability in the email calendar processing of Zimbra Collaboration Suite (ZCS) allows an unauthenticated sender to forge another attendee's response to an existing meeting. Zimbra passes a null sender identity when processing iCalendar REPLY components, causing the attendee authorization filter to be skipped. An attacker who knows the meeting UID, organizer address and attendee address can send a non-outdated reply that changes the attendee's participation status in the organizer's calendar. The change occurs automatically during delivery without organizer or attendee interaction. ZCS 10.1.16 and 10.1.21 are affected.	5.3	More Details
CVE-2026-104641	The Simple Shopping Cart WordPress plugin before 5.3.3 does not verify that a PayPal payment notification is for a payment received by the site owner's PayPal account in its default configuration, allowing unauthenticated attackers to have their orders marked as paid using a payment made to an account they control.	5.3	More Details
CVE-2026-116046	A flaw has been found in eventcatalog sdk up to 2.29.0. This vulnerability affects the function <code>getCustomDoc/writeCustomDoc/rmCustomDoc</code> of the file <code>packages/sdk/src/custom-docs.ts</code> of the component Custom Documentation Helpers. This manipulation of the argument <code>options.path</code> causes path traversal. It is possible to launch the attack on the local host. Upgrading to version 2.30.0 is able to resolve this issue. Patch name: <code>a4f072a03f440d58230092e88d8b4893f4c6d6ca</code> . The affected component should be upgraded. The vendor was contacted early, responded in a very professional manner and quickly released a fixed version of the affected product.	5.3	More Details
CVE-2026-103301	The Payment Gateway for Telcell WordPress plugin before 2.0.5 does not verify the authenticity of incoming payment notifications, because the secret used to sign them is empty by default, allowing unauthenticated attackers to forge a valid notification and mark arbitrary WooCommerce orders as paid. The same handler also marks an attacker-chosen order as failed when given an invalid signature.	5.3	More Details
CVE-2026-101138	The CMP - Coming Soon & Maintenance WordPress plugin before 4.1.21 does not correctly apply maintenance mode to pages added to its page whitelist by manually entered URL pattern, allowing unauthenticated visitors to view the content of those pages while maintenance mode is active.	5.3	More Details

CVE-2026-92999	The Click n Chat (Chat Widget Integration) WordPress plugin before 1.1.3 does not perform any authorization check on a lead-export handler, allowing unauthenticated users to download every captured contact record (names, email addresses and phone numbers).	5.3	More Details
CVE-2026-84921	Missing Authorization vulnerability in Drupal Webform Submissions Delete allows Forceful Browsing. This issue affects Webform Submissions Delete versions: from 0.0.0 to 1.2.0.	5.3	More Details
CVE-2026-93627	An organizer authentication vulnerability in the email calendar processing of Zimbra Collaboration Suite (ZCS) allows an unauthenticated sender to cancel or modify an existing meeting in a recipient's calendar. Zimbra accepts iCalendar CANCEL and REQUEST components without verifying that the email sender is authorized to act as the supplied organizer. An attacker who knows the meeting UID and organizer address can supply a non-outdated sequence number to move the recipient's meeting to Trash or replace its details. Processing occurs automatically during email delivery without recipient interaction. ZCS 10.1.16 and 10.1.21 are affected.	5.3	More Details
CVE-2026-116216	MacCMS v10 through 2026.1000.4056 contains a mass assignment vulnerability in index/Comment::saveData() that allows unauthenticated attackers to set arbitrary mac_comment columns including user_id. Attackers can submit anonymous comments with a crafted user_id parameter to attribute them to any registered member, displaying that member's avatar and impersonating them.	5.3	More Details
CVE-2026-93629	An organizer authentication vulnerability in Zimbra Collaboration Suite (ZCS) allows an unauthenticated email sender to insert calendar invitations attributed to another user. The email delivery path permits external invitations but does not verify that the message sender is authorized to use the iCalendar ORGANIZER identity. When automatic invitation insertion is enabled, a crafted REQUEST component creates an event containing attacker-selected meeting details and reminders under a forged organizer without recipient interaction. No corresponding event is created in the claimed organizer's calendar. ZCS 10.1.16 and 10.1.21 are affected.	5.3	More Details
CVE-2026-95718	An authentication bypass vulnerability in the PAPI protocol of AOS-8 could allow an unauthenticated remote attacker to circumvent existing authentication controls. Successful exploitation could allow an attacker to make limited unauthorized modifications to certain settings of the affected system.	5.3	More Details
CVE-2026-114729	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a remote content blocking bypass. rcube_washtml inspected the SVG SMIL animation attributes `to`, `from` and values for URIs, but not `by`; so markup such as in an HTML message caused remote content to be loaded without the user's consent, disclosing the client IP address and confirming that the message was opened. Upstream fixed this in the commit "Fix remote-content blocker bypass via SVG SMIL `by` attribute".	5.3	More Details
CVE-2026-87950	Use of Hard-coded Credentials vulnerability in Drupal SAML SSO - Service Provider allows Retrieve Embedded Sensitive Data. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	5.3	More Details
CVE-2026-116196	Rclone before v1.75.2 contains a credential exposure vulnerability in the b2 backend that leaks the live account Authorization token and SSE-C keys on HTTP redirects from --b2-download-url hosts. Attackers can observe a same-host HTTPS-to-HTTP downgrade redirect to capture the cleartext B2 token, or receive cross-host redirects disclosing SSE-C customer encryption keys.	5.3	More Details
CVE-2026-114470	Cross-site request forgery in the REST API of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote attacker to make an authenticated user's browser perform a state-changing request. The REST API accepts HTTP Basic authentication and answers a browser with a Basic sign-in box; a browser caches the credential once the box is answered and sends it with later requests to that origin, including a top-level form post from another site, and the server reads a request body whatever its declared content type. The cross-site protection covered only the cookie session (a header a cross-site form cannot set, and a SameSite=Strict cookie); HTTP Basic was not held to it. A page under an attacker's control, visited by a user whose browser holds a cached Basic credential for the REST API - typically an administrator who signed in through the box - could therefore cause a state-changing request such as creating a domain, an account or an API key, or changing a setting, with the victim's credential. The attacker cannot read the cross-origin response, so the forgery is blind.	5.3	More Details
CVE-2026-75525	OpenReception's appointment booking software provides an end-to-end encrypted appointment booking platform. In versions 2.1.0 and prior, an unauthenticated attacker can store an arbitrary ML-KEM-768 public key under any existing user's identity in any tenant's staff-crypto recipient directory. The product's design guarantees that appointment data is end-to-end encrypted to staff members and that the staff recipient directory is not attacker-writable, so that even administrators and the operator cannot read patient data. This endpoint breaks that guarantee for an unauthenticated remote actor: it allows writing an arbitrary encryption recipient into the recipient directory of any tenant. The injected	5.3	More Details

	key receives the tunnel key of every appointment booked afterward in that tenant, because booking clients encrypt to all active staff keys. The attacker holds the corresponding secret key and is a persistent, unauthorized recipient of that key material. At time of publication, there are no publicly available patches.		
CVE-2026-96377	Allocation of Resources Without Limits or Throttling vulnerability in Drupal Combined image style allows Excessive Allocation. This issue affects Combined image style versions: from 0.0.0 to 1.0.7.	5.3	More Details
CVE-2026-101259	Generation of Error Message Containing Sensitive Information vulnerability in Apache Camel CoAP component. This issue affects Apache Camel: from 4.0.0 before 4.18.5, from 4.19.0 before 4.22.2. When the CoAP consumer could not build its response once the route had completed, CamelCoapResource caught the resulting exception and sent its message back to the CoAP client as the payload of a 5.00 Internal Server Error response. The typical case is a type conversion failure: the consumer converts the final message body to the byte[] payload, and the CamelCoapMaxAge and CamelCoapETag headers to their types, after the route has returned, and therefore outside the route's error handler. On the affected releases the message of such a failure names the class of the value being converted, carries the value's string form, and ends with the class name and message of the underlying cause, so a remote CoAP client could read internal detail about the deployment: class and package names, the string form of the value, such as the path of a file body, and whatever the cause carries, which in practice can include paths, host names and connection detail of a backend. The consumer had no option to suppress this, unlike the HTTP consumers, which offer a muteException option for that purpose. The exception raised by a failing route was not sent: the consumer did not check whether the exchange had failed, and answered a route failure with 2.05 Content and whatever body the message held at that point. Reaching the disclosure therefore requires a route whose completed response a remote client can make unconvertible, which depends on how the route is written. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The 4.14.x releases stream and earlier, and the 4.15.x to 4.17.x and 4.19.x to 4.21.x release lines, reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. The upgrade is enough on its own: the new muteException consumer option defaults to true on all three fixed release lines, so the error response carries no payload without any configuration change. For deployments that cannot upgrade immediately, make sure the response can always be built once the route has completed: convert the body to byte[] or String as the last step of each route that consumes from a coap endpoint, for example with convertBodyTo, give the CamelCoapMaxAge and CamelCoapETag headers only values of their own types, and add an onException handler that marks the failure as handled and sets the response body explicitly, so that a conversion failure is raised and replaced inside the route rather than in the consumer.	5.3	More Details
CVE-2026-114603	OpenTelemetry Python provides the Python API and SDK for OpenTelemetry. Prior to 1.45.1, the opentelemetry-api function TraceState.from_header in opentelemetry/trace/span.py, reached through TraceContextTextMapPropagator.extract, parses inbound W3C tracestate headers with an unbounded optional-whitespace regular expression. A malformed header containing a long run of spaces or tabs without a following comma causes quadratic backtracking before the member-count limit is applied. Services using the default W3C Trace Context propagator can therefore spend excessive CPU time on one unauthenticated request when the HTTP server, parser, gateway, or custom carrier accepts a sufficiently large header. The processing can occupy a synchronous worker or delay unrelated requests on the same asynchronous event loop. The impact is limited to availability and does not expose or modify telemetry data or allow code execution. This issue is fixed in version 1.45.1.	5.3	More Details
CVE-2026-96405	NLnet Labs Unbound up to and including version 1.26.1 limits compression-pointer traversal for each individual domain-name operation but has no aggregate inbound decompression-work budget. Many chained compression pointers in a message results in a lot of CPU usage following them. A malicious actor with recursive-query access and control of an ordinarily delegated authority that serves specially crafted replies can exploit the vulnerability and degrade Unbound's general performance.	5.3	More Details
CVE-2026-114645	A 32-bit integer wraparound in the MongoDB C Driver's Windows-only Kerberos/SSPI authentication code can cause an undersized heap buffer to be allocated when an excessively long credential value is supplied in the client's connection configuration. A party able to control that client-side configuration, in conjunction with the affected authentication mechanism, can cause the driver to write outside the bounds of the allocated buffer. This may result in corruption of client process memory and denial of service of the client application. The issue affects only client builds compiled with the Windows-native authentication backend.	5.2	More Details
CVE-2026-102485	Out-of-bounds read in the Motorola Oncore refclock driver in NTPsec ntpd before 1.2.6 allows an attacker who controls the GPS receiver's serial data to crash ntpd or disclose memory via a PPS Control (@@Gc) reply with an out-of-range mode byte.	5.2	More Details

CVE-2026-78116	The Broadcom AFBR-S50 time-of-flight sensor driver double-completes an RTIO submission when sensor data evaluation fails. In <code>data_ready_work_handler()</code> (<code>drivers/sensor/broadcom/afbr_s50/afbr_s50.c</code>), a non-<code>STATUS_OK</code> return from <code>Argus_EvaluateData()</code> calls <code>handle_error_on_result()</code>, which finalises the request with <code>rtio_iodev_sqe_err()</code>, but execution then fell through to the tail of the function and finalised the very same struct <code>rtio_iodev_sqe</code> a second time with <code>rtio_iodev_sqe_ok()</code>. Unlike the other error branches in the same function, this one did not return. Both completion helpers reach <code>rtio_executor_done()</code> in <code>subsys/rtio/rtio_executor.c</code>, which returns the submission to the pool via <code>rtio_sqe_pool_free()</code> and, for streaming (multishot) submissions, also releases the RX buffer with <code>rtio_release_buffer()</code>. The second completion therefore reads a freed <code>rtio_iodev_sqe</code>, pushes the same <code>mpsc</code> node onto the submission queue while it is already on the pool free list — splicing the two lists together — and re-submits the freed request so the next measurement is written through an already-released buffer. For one-shot submissions it emits a spurious success completion carrying stale userdata and frees the same pool entry twice. The failure branch is driven by sensor conditions rather than by host software. The vendor API documents statuses such as <code>ERROR_ARGUS_BIAS_VOLTAGE_REINIT</code> and <code>ERROR_ARGUS_LASER_MONITOR_INACTIVE</code> as transient results of harsh ambient light, so a party in front of the optical sensor can provoke the invalid-data condition without any physical contact with the board or any software access; SPI transfer errors reach the same branch. In releases before the payload.Status check was dropped, ordinary sharp distance transitions in the sensor's field of view were sufficient. Impact is corruption of the RTIO submission pool and mempool bookkeeping: aliased submissions, writes through released buffers, and duplicate completions delivered to the application. The attacker does not control the corrupted content, so the practical consequence is memory corruption and loss of availability on devices that instantiate <code>brcm,afbr-s50</code>, rather than directed code execution. Only builds that enable this driver are affected.	5.2	More Details
CVE-2026-114038	A flaw was found in <code>device-mapper-multipath</code>. A local user can trigger an out-of-bounds read by sending a command with a format string ending in a trailing percent character to the <code>multipath</code> daemon. This issue can allow the user to expose unintended daemon memory, leading to information disclosure, or cause the daemon to crash, resulting in a Denial of Service (DoS).	5.1	More Details
CVE-2026-55222	Isar is an integration system for automated root filesystem generation. In versions 0.11 and newer, <code>/etc/apt/trusted.gpg.d</code> is created manually during bootstrap setup to be able to deploy local keys. By that, the directory is owned by the calling user (the one that executes Isar) instead of root. If the calling user's id is identical to one of an unprivileged user inside the image, this user is able to alter existing keys and deploy new ones, silently breaking the apt repo integrity protection. The issue is has no impact on those images that use an immutable rootfs, or if apt is never invoked. Changing only the keys does not trigger an exploit, it's just one step in a chain. An attacker would also have to manipulate a remote repo or intercept non-secure communication with it. This issue has been patched via commit 29e84ed.	5.1	More Details
CVE-2026-89419	OpenStack Ironi before 39.0.1 sometimes allows a user to change the maintenance setting on a node via PATCH to <code>/v1/nodes/{ident}</code> - in other words, not using the specific maintenance endpoint.	5.0	More Details
CVE-2026-107128	The Uncanny Toolkit for LearnDash WordPress plugin before 3.8.2 does not verify that the acting user is permitted to manage the targeted accounts before changing their verification state, allowing a subsite administrator on a multisite network to mark arbitrary network users outside their own site as unverified and prevent those users from logging in.	5.0	More Details
CVE-2026-114080	Trigger.dev before 4.5.2 contains a blind server-side request forgery vulnerability that allows authenticated project members to make the server send requests by setting unvalidated alert-channel webhook URLs. Attackers can create alert channels pointing to loopback, internal, or cloud metadata addresses, then trigger alerts to probe internal service reachability and send signed POST requests.	5.0	More Details
CVE-2026-116187	Rclone from v1.68.0 before v1.75.2 contains an authorization bypass vulnerability in <code>serve s3 --auth-proxy</code> that allows authenticated tenants to manipulate other tenants' multipart uploads by referencing their upload IDs. Attackers who learn a victim's bucket, object, and upload ID can inject parts, complete, or abort in-flight uploads and view upload and object metadata across tenants.	5.0	More Details
CVE-2026-115720	Frappe Builder before 1.36.1 contains a server-side request forgery vulnerability in <code>builder/api.py</code> where the <code>assert_not_private_url()</code> guard validates only the initial hostname before callers invoke <code>requests.get(url)</code> with default redirect-following and no IP pinning, allowing an attacker-supplied public URL to HTTP-redirect to internal or link-local addresses the guard never validated. An authenticated attacker can call the <code>convert_to_webp</code> endpoint, which lacks a permission decorator, with a URL that redirects to internal targets such as cloud metadata services to perform SSRF, internal port scanning, or trigger a request-hang denial of service due to the missing timeout.	5.0	More Details
CVE-	phpMyFAQ before 4.1.10 contains a server-side request forgery vulnerability in the PDF export feature that fetches external images referenced in FAQ answers. Attackers who can influence FAQ content can		More

2026-116182	embed image URLs pointing to loopback or internal services, causing the server to request them and embed responses in generated PDFs.	5.0	Details
CVE-2026-116185	OpenRemote core before 1.31.0 contains an authorization bypass vulnerability that allows authenticated users to read other realms' gateway tunnel information via getGatewayActiveTunnellInfos. Attackers can request the gateway tunnel endpoint with their own realm and a victim gateway ID to disclose internal target hosts, ports, and routing data.	5.0	More Details
CVE-2026-115411	crawlseo before version 0.2.4 contains a server-side request forgery vulnerability in its crawler engine where isPrivatelp() fails to block reserved IP ranges such as RFC 6598 Carrier-Grade NAT, RFC 6890, TEST-NET ranges and certain IPv6 addresses, and additionally fails to pin the validated address for the actual outbound connection, allowing a TOCTOU DNS rebinding bypass. An authenticated attacker can add a malicious domain that resolves to an internal or restricted address, trigger a crawl to reach internal services or cloud metadata endpoints like 169.254.169.254, and retrieve the fetched page's title, description, headings, status code and links through the crawl results API, enabling internal network reconnaissance.	5.0	More Details
CVE-2026-114725	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a mail header injection vulnerability. rmail_sendmail stripped newline characters only from the Subject and Organization headers, leaving other header values taken from request parameters unsanitized, so a remote authenticated attacker could inject arbitrary headers into an outgoing message via the References, In-Reply-To, Message-ID and From parameters. Upstream fixed this in the commit "Fix mail header injection via References/In-Reply-To/Message-ID/From parameters", which normalizes whitespace in every string header value.	5.0	More Details
CVE-2026-114473	Incorrect authorization in the Sender Rewriting Scheme (SRS) of Progressive Robot hMailServer 6.2.1 through 6.3.6 allows a remote attacker to relay mail through the server to outside addresses. With SRS enabled, the server rewrites the envelope sender of mail it forwards to a signed SRS address at the forwarding domain, and passes a bounce to that address back to the original sender; but it took every message to a valid SRS address, whatever its sender, as delivery to a local recipient, so the relay rules and the SMTP authentication they require were not applied, and the message was delivered to the outside address the SRS address encodes. Anyone holding a valid SRS address of the server, such as the recipient of a forwarded message, could have the server relay mail to the address it encodes for 21 days, and an account holder whose account forwards to an outside mailbox could obtain one for any address. SRS is off by default.	5.0	More Details
CVE-2026-115774	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, twisted.mail.imap4.IMAP4Server reaches twisted.mail.imap4.parseNestedParens with a negative literal size in an authenticated SEARCH command after the client enters the selected mailbox state. The negative size moves the parser cursor below zero, and Python's negative indexing prevents the synchronous scanning loop from terminating. The loop freezes the single-threaded reactor, stopping other IMAP connections and any SMTP or POP3 services hosted in the same process until restart. This issue is fixed in version 26.10.0rc1.	4.9	More Details
CVE-2026-116170	SiYuan before v3.8.7 contains an sql injection vulnerability that allows admin-token holders to read history.db by injecting unescaped op and type=3 query values. Attackers can send crafted UNION SELECT payloads to searchHistory or getHistoryItems endpoints to exfiltrate deleted document content, titles, and paths.	4.9	More Details
CVE-2026-114570	WordPress Core is vulnerable to Stored Cross-Site Scripting via Quick Edit in all versions up to, and including, 7.0.2 due to insufficient escaping of user display names. 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 higher privileged user accesses the affected administrative screens. This issue only affects sites with a large number of users.	4.9	More Details
CVE-2026-110982	The AI Engine WordPress plugin before 3.8.6 does not check that the caller is allowed to edit the targeted object before updating its title and excerpt, allowing users with the Editor role and above to modify arbitrary posts of any type, including content they cannot otherwise edit such as navigation menus, block templates, global styles and WooCommerce coupons.	4.9	More Details
CVE-2026-111972	The Photo Engine (Media Organizer & Lightroom) WordPress plugin before 6.5.8 does not verify that a user is authorised to act on the targeted media attachment, allowing authenticated users with the author role and above to delete arbitrary media attachments, including those uploaded by other users and administrators.	4.9	More Details
CVE-2026-110514	The WPZOOM Portfolio Lite WordPress plugin before 1.4.34 does not check that a user is allowed to edit a post before changing its type and order, which could allow users with the contributor role and above to convert any other user's posts and pages into portfolio items, removing them from their original listings and URLs.	4.9	More Details

CVE-2026-116213	Webasyst Framework through 4.4.0 contains an arbitrary file read vulnerability in <code>waArchiveTar::writeltem()</code> that creates symlinks from theme archives without validating link targets. Authenticated users with theme management rights can upload themes containing symlinks to <code>wa-config/db.php</code> or <code>/etc/passwd</code> and read them through the public theme path.	4.9	More Details
CVE-2026-95719	An arbitrary file deletion vulnerability exists in the command line interface of AOS-10 that could allow an authenticated attacker with administrative privileges to cause a denial of service. Successful exploitation could allow an attacker to crash the system, preventing it from rebooting without manual intervention and disrupting normal operation of the affected device.	4.9	More Details
CVE-2026-112219	The Download Monitor WordPress plugin before 5.3.5 does not correctly restrict management of the list of filesystem locations it may serve files from to network administrators on multisite, allowing a subsite administrator who is not trusted at the network level to approve arbitrary locations and have the Download Monitor WordPress plugin before 5.3.5 serve readable files from outside the uploads directory, including files belonging to other subsites, to unauthenticated visitors; single-site installations are unaffected.	4.9	More Details
CVE-2026-114740	The WPVR WordPress plugin before 9.1.3 does not correctly check user capabilities when creating a tour through its setup wizard, allowing users with the Contributor role and above to create and publish tour content and set arbitrary post metadata they could not otherwise create or publish.	4.9	More Details
CVE-2026-110010	The MB Custom Post Types & Custom Taxonomies WordPress plugin before 2.12.3 does not perform any capability or authentication check before allowing an unauthenticated request to re-run a legacy data migration, letting anyone revert a custom post type or taxonomy definition on sites upgraded from a pre-2.0 release back to its pre-upgrade state. This discards configuration changes an administrator made after upgrading and, where those changes had restricted a definition's visibility, silently re-exposes the definition and its content to anonymous visitors.	4.8	More Details
CVE-2026-114386	Pydantic AI Harness before 0.59.0 fails to enforce <code>UsageLimits.tool_calls_limit</code> on nested tool calls dispatched from <code>run_code</code> in the <code>CodeMode</code> toolset. Users or prompt-injected content can make the sandboxed program invoke tools via <code>run_nested_call()</code> without <code>check_before_tool_call</code> , exceeding operator limits on costly or side-effecting tools.	4.8	More Details
CVE-2026-95663	Insufficient Session Expiration vulnerability in Apache Camel Keycloak component. This issue affects Apache Camel: from 4.16.0 before 4.18.5, from 4.19.0 before 4.22.2. When token introspection is enabled on the camel-keycloak KeycloakSecurityPolicy (<code>useTokenIntrospection=true</code>), the security processor validates the access token of a request by sending it to Keycloak's OAuth 2.0 token introspection endpoint, and accepts the token while the response reports it as active. The introspection endpoint reports a token as inactive once its exp (expiration) time has passed, so on this path the active flag is what enforces token expiry. Introspection results are cached by default (<code>introspectionCacheEnabled=true</code> , <code>introspectionCacheTtl=60</code> seconds), and both cache implementations shipped with the component, <code>ConcurrentMapTokenCache</code> and <code>CaffeineTokenCache</code> , expired an entry a fixed time after it was inserted, without bounding that lifetime by the exp claim that the cached response already carried. KeycloakSecurityProcessor in turn checked only the active flag of the result, on every path that uses introspection, and never compared exp with the current time. A response cached shortly before the token expired was therefore served to later requests carrying the same token for up to the configured time-to-live after the token's expiry, and those requests were admitted with the roles and permissions the expired token had carried, although RFC 7662, which defines token introspection, states that a response containing exp must not be cached beyond that time. The access token is supplied by the caller with each request, typically as a bearer token in the Authorization header, so a caller holding a token that was valid moments earlier could keep using it after it expired, for as long as the cached entry lived. The window is bounded by <code>introspectionCacheTtl</code> , which is 60 seconds by default and 300 to 600 seconds where the documentation's recommendation for internal services was followed, and it opens only for a token that the same Camel process introspected and cached while the token was still active. The issue grants nothing beyond what the token already carried: it allows no token forgery and no privilege escalation, and token introspection is not enabled by default, which is why the severity is assessed medium. The local JWT verification path, used when introspection is disabled, is not affected by this issue. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. If users are on the 4.18.x releases stream, then they are suggested to upgrade to 4.18.5. The 4.16.x, 4.17.x and 4.19.x to 4.21.x release lines reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. For deployments that cannot upgrade immediately, set <code>introspectionCacheEnabled=false</code> on the KeycloakSecurityPolicy, so that every request is checked against the introspection endpoint, which reports an expired token as inactive, at the cost of one introspection call per request. Lowering <code>introspectionCacheTtl</code> only narrows the window and does not close it.	4.8	More Details
CVE-			

2026-87946	Observable Timing Discrepancy vulnerability in Drupal SAML SSO - Service Provider allows Brute Force. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	4.8	More Details
CVE-2021-48019	Vyper before 0.2.9 deploys forwarder proxies through the built-in create_forwarder_to function that report a fixed 4096-byte RETURNDATASIZE rather than the actual return data size. Calls through these proxies can yield corrupted data when a function returns more than 4096 bytes without return data sanitation, and fail strict return-size checks such as SafeERC20.safeTransfer.	4.8	More Details
CVE-2026-87947	Improper Verification of Cryptographic Signature vulnerability in Drupal SAML SSO - Service Provider allows Signature Spoof. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	4.8	More Details
CVE-2026-105996	The Booking Package WordPress plugin before 1.7.31 does not generate booking cancellation tokens using a secure source of randomness, deriving them instead from values it exposes publicly, so an unauthenticated attacker can reconstruct an arbitrary booking's token, disclose that booking's personal data and, where visitor cancellation is enabled, cancel the booking.	4.8	More Details
CVE-2026-114118	Vyper 0.3.4 through 0.5.0b1 computes the remaining blueprint size in create_from_blueprint by uint256 subtraction without first checking that code_offset is below the blueprint's code size, so a huge offset wraps to a small positive value. An attacker controlling code_offset can make a factory deploy contracts with no runtime code, potentially trapping ETH and tokens.	4.8	More Details
CVE-2026-96375	Improper Certificate Validation vulnerability in Drupal Cloud allows Adversary in the Middle (AiTM). This issue affects Cloud versions: from 0.0.0 to 7.0.1.	4.8	More Details
CVE-2026-94386	Use of Cache Containing Sensitive Information vulnerability in Apache Camel OAuth component. This issue affects Apache Camel: from 4.19.0 before 4.22.2. OAuthClientCredentialsTokenResolver obtains OAuth 2.0 bearer tokens through the client credentials grant on behalf of the components that authenticate their outbound calls with a named OAuth profile, such as camel-openai, camel-langchain4j-agent, camel-huggingface, camel-docling, camel-a2a, camel-spring-ai-chat and the IBM watsonx components, and of applications that call the OAuthClientAuthenticationFactory SPI directly. It kept the tokens it obtained in a static cache shared by every resolver instance and every CamelContext in the JVM, and keyed that cache on the token endpoint and the client ID only, although the token request also carries the client secret and the requested scope. Two profiles that share a token endpoint and a client ID but request different scopes, or authenticate with different client secrets, therefore mapped to the same cache entry: the profile that resolved a token first populated the entry, and the other profile was then handed that token, with no token request made under its own configuration, for as long as the token remained valid. Token caching is enabled by default, and pointing several profiles at one identity provider client is an ordinary use of the profile support rather than an opt-in to an insecure setting. The outcome is an outbound request that carries a token issued for another profile's scope, which can grant broader authority than the calling profile was configured for, or reach a service the token was not issued for; with identity providers that derive a token's audience from the requested scope, a different scope can mean a token for a different service altogether. A profile configured with a different client secret was likewise served the token obtained with the other secret, without authenticating with its own. The profiles are configured by the deployment, which is trusted, so no untrusted input reaches the cache: the party that gains is the service that receives a token it was not meant to receive, which, if hostile or compromised, can replay that token within its scope. Reaching the issue requires a deployment that defines more than one profile sharing a token endpoint and client ID while differing in scope or client secret. Users are recommended to upgrade to version 4.23.0, which fixes the issue. If users are on the 4.22.x releases stream, then they are suggested to upgrade to 4.22.2. The 4.19.x to 4.21.x release lines reached end of life before this issue was resolved and will not receive a fix, so deployments on them must move to a supported release. For deployments that cannot upgrade immediately, disable token caching on every OAuth profile that shares a token endpoint and client ID with another profile, by setting camel.oauth.NAME.cache-tokens=false for a named profile or camel.oauth.cache-tokens=false for the default one, so that each resolution requests a token with that profile's own scope and client secret, at the cost of a token request on every resolution. Alternatively, register a separate client at the identity provider for each profile that needs a different scope or client secret, so that no two such profiles share a client ID.	4.8	More Details
CVE-2026-114437	In the Golang Chi web framework, the ClientIPFromXFF middleware skips any X-Forwarded-For entry in a configured trusted CIDR but never verifies that the direct TCP peer in http.Request.RemoteAddr belongs to that CIDR. A client that can reach the Go server directly can therefore append a trusted-looking proxy address and choose the client IP returned by middleware.GetClientIP. This issue affects chi before 5.4.0.	4.8	More Details
CVE-2026-	Wasmtime 47.0.0 before 48.0.4 and 49.0.0 before 49.0.2 contains an improper validation vulnerability that accepts components whose core modules import WebAssembly tags with a subtype of the original	4.8	More

115383	type. Malicious guests can mis-type exception object payloads to corrupt the GC heap, causing Wasmtime to panic when the exceptions proposal is enabled.		Details
CVE-2026-115382	Wasmtime 46.0.0 before 48.0.4 and 49.0.0 before 49.0.2 contains a denial-of-service vulnerability in its wasi:http implementation that allows guests to panic the host. Guests can specify a zero between-bytes-timeout request option, which is forwarded directly to a Tokio API that panics, crashing the host embedding.	4.8	More Details
CVE-2026-94263	The Advanced Custom Fields: Extended and Advanced Custom Fields: Extended PRO WordPress plugins from 0.9.2.4 to 0.9.2.7 do not consistently disable shortcode processing in the confirmation message of their front-end forms, allowing unauthenticated visitors to have arbitrary shortcodes executed server-side when they submit a form whose stored configuration does not carry the relevant setting.	4.8	More Details
CVE-2026-103234	The WP Ultimate Review WordPress plugin before 2.4.5 does not neutralise shortcodes in all of the review fields submitted through its public review form, which is available to unauthenticated visitors, allowing them to execute arbitrary shortcodes when the affected field is displayed, provided user reviews and that field's display are enabled.	4.8	More Details
CVE-2026-115381	Wasmtime before 36.0.17, 37.0.0 before 48.0.4, and 49.0.0 before 49.0.2 contains an unbounded memory allocation in the wasmtime-wasi WASIp2 write-zeroes method of SinkOutputStream, the default stream used when stdout/stderr are not configured. A malicious WebAssembly guest can request huge zero-filled writes, exhausting host memory and causing denial of service.	4.8	More Details
CVE-2024-5563	Due to an incomplete implementation, reCaptcha is not applied to certain authentication flows even when the reCaptcha option is enabled. This oversight permits attackers to bypass the intended protection mechanism. Brute force attacks, if successful, can lead to unauthorized access to user accounts and sensitive information by systematically attempting various credential combinations.	4.8	More Details
CVE-2026-87944	Improper Certificate Validation vulnerability in Drupal SAML SSO - Service Provider allows Adversary in the Middle (AiTM). This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	4.8	More Details
CVE-2026-115754	serialize-javascript from 7.0.1 before 7.1.3 contains a cross-site scripting vulnerability that allows attackers to break out of HTML script elements by ending serialized function source text with an unescaped </script. Attackers controlling a function's source via a spoofed toString() can combine surrounding output delimiters and later values to inject live markup such as img onerror handlers.	4.7	More Details
CVE-2026-116538	HtmlSanitizer before 9.2.1039 and 10.0 prereleases before 10.0.1037-beta contains a cross-site scripting vulnerability because SanitizeDom(string) never sanitizes attributes on the html and body wrapper elements. Attackers can supply input with a second body or html tag carrying event handlers like onload, executing script when applications render serialized document or body outer HTML.	4.7	More Details
CVE-2026-116063	A weakness has been identified in DedeCMS up to 5.7.118. Impacted is the function require_once of the file dede/update_guide.php. This manipulation of the argument vtime causes code injection. Remote exploitation of the attack is possible. The exploit has been made available to the public and could be used for attacks.	4.7	More Details
CVE-2026-109681	The FTP client trusts the server's PASV response IP:port without validation against the control connection target. A rogue FTP server can specify any IP:port in the 227 Passive Mode response, causing the operator's machine to open a TCP connection to an arbitrary host.	4.7	More Details
CVE-2026-114060	SvelteKit (@sveltejs/kit) before 3.0.0 contains an open redirect vulnerability in enhanced form actions (use:enhance) on routes accepting paths beginning with a double slash. Attackers can craft action paths starting with // that become protocol-relative URLs, causing non-redirect form results to navigate victims to attacker-controlled sites.	4.7	More Details
CVE-2026-114061	SvelteKit (@sveltejs/kit) before 3.0.0 contains a cross-site scripting vulnerability in prerendered pages because a cache-control header value set via setHeaders is inserted without escaping into a <meta http-equiv="cache-control"> tag. An attacker who can influence that untrusted value at prerender time can break out of the content attribute and inject HTML or JavaScript that executes when users visit the page.	4.7	More Details
CVE-2026-115351	Mistune from 3.2.1 before 3.3.0 contains a CSS injection vulnerability that allows attackers controlling markdown input to inject arbitrary style declarations via the unvalidated figwidth option of the figure directive. Attackers can append declarations like background:url() to the rendered figure style attribute, forcing victim browsers to send requests to attacker servers and leak attribute values.	4.7	More Details
CVE-2026-115360	OpenBao before 2.6.3 contains an open redirect vulnerability in the UI's OIDC identity provider that affects users not authenticated to OpenBao. Attackers can craft OIDC links that redirect unauthenticated users to attacker-controlled destinations, enabling phishing or credential theft.	4.7	More Details

CVE-2026-114718	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a cross-site scripting vulnerability in the password plugin. Error text returned by a password driver was concatenated into the message presented to the user without HTML escaping, so a backend able to influence that text, such as the DirectAdmin API, could inject script into the Roundcube interface. Upstream fixed this in the commit "Fix XSS issue in handling of password driver error responses", which escapes the driver message before display.	4.7	More Details
CVE-2026-114264	A security flaw has been discovered in ProjectWorlds Advanced Library Management System 1.0. The impacted element is an unknown function of the file edit_member.php. Performing a manipulation of the argument firstname/middlename/lastname/address/branch/type results in sql injection. Remote exploitation of the attack is possible. The exploit has been released to the public and may be used for attacks.	4.7	More Details
CVE-2026-109682	Metasploit Framework versions before 6.5.7 allow HTTP requests made through send_request.cgi! to follow redirects to a different host without restricting the redirect destination. A malicious HTTP server can cause the system running Metasploit to make requests to arbitrary HTTP services reachable from the operator's network, including internal services that are not otherwise accessible to the attacker.	4.7	More Details
CVE-2026-114471	Insertion of sensitive information into a log file in the SMTP server of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a local user who can read the server's log, or anyone given a copy of it, to obtain credentials. With SMTP logging enabled, the server hid the credential of an AUTH PLAIN command only when it was sent after EHLO in one exact form, and wrote the rest of the SASL exchange as the client sent it: a PLAIN password sent before EHLO or with extra spacing, the OAuth 2.0 bearer token of AUTH XOAUTH2 and AUTH OAUTHBEARER (from 6.2.0), the SCRAM-SHA-256 client proof (from 6.2.0), which with the logged salt and iteration count allows offline password guessing, and the GSSAPI authenticator (from 6.3.2). SMTP logging is off by default; on Windows before 6.3.6 the log folder was readable by every local user.	4.7	More Details
CVE-2026-96361	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Cross-Site Scripting (XSS). This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
CVE-2026-96367	Privilege Defined With Unsafe Actions vulnerability in Drupal Webform allows Stored XSS. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
CVE-2026-96368	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Stored XSS. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
CVE-2026-96382	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Diba carousel slider allows Stored XSS. This issue affects Diba carousel slider versions: from 0.0.0 to 3.0.2.	4.6	More Details
CVE-2026-55422	conda-forge is a GitHub organization containing repositories of conda recipes. In commits 85748ad2368a and prior, two React components render externally-sourced strings directly as HTML using dangerouslySetInnerHTML, with no sanitization (the project has no DOMPurify or equivalent dependency). JSON files are produced automatically by the conda-forge autotick bot (regro/cf-scripts) while it processes every conda-forge feedstock and its upstream source artifacts. The error/status strings embed raw output captured from those processes (build logs, patch output, tracebacks). The website renders them verbatim as HTML. At time of publication, there are no publicly available patches.	4.6	More Details
CVE-2026-86020	Part-DB is an open source inventory management system for electronic components. Prior to version 2.13.2, Part-DB contains a spreadsheet formula injection vulnerability in the part import/export workflow. An authenticated user with permission to import parts can import a CSV file containing formula-prefixed values, such as =1+1, into user-controlled part fields. When the affected part is later exported as an XLSX spreadsheet, Part-DB writes the attacker-controlled value using PhpSpreadsheet's setCellValue() method without neutralising spreadsheet formula characters. As a result, the exported XLSX file contains an executable spreadsheet formula cell rather than a plain text value. If another user exports and opens the spreadsheet, the formula may be evaluated by their spreadsheet software. This issue has been patched in version 2.13.2.	4.6	More Details
CVE-2026-96358	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Cross-Site Scripting (XSS). This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
	Authlib is a Python library which builds OAuth and OpenID Connect servers. Prior to 1.9.0, OpenIDImplicitGrant.validate_authorization_request() checks the client-controlled raw		

CVE-2026-114619	request.payload.scope for openid after the parent grant has filtered allowed scopes into request.scope. A registered client permitted an ID token response type but not the openid scope can add openid to the request, obtain user consent, and receive a signed ID token containing the user's subject and allowed profile claims. Authorization code flow and deployments that allow openid for every client permitted an ID token response are not affected. In authlib/oidc/core/grants/implicit.py, OpenIDImplicitGrant.validate_authorization_request() checks request.payload.scope instead of the filtered request.scope. This issue is fixed in version 1.9.0.	4.6	More Details
CVE-2026-115249	A flaw has been found in subhajitkhan online-clinic-management-system up to e9ee77a8827a1446220fa07ee693dc4d9a29a578. Impacted is an unknown function of the file patappointtime.php of the component Appointment Booking. This manipulation of the argument day/time/appdate causes improper input validation. 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.	4.6	More Details
CVE-2026-96359	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Cross-Site Scripting (XSS). This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
CVE-2026-116263	A flaw has been found in HaschekSolutions PictShare up to 3.7.1. The impacted element is the function architect of the file src/inc/core.php of the component File Upload. This manipulation causes race condition. The attack can be initiated remotely. The attack's complexity is rated as high. The exploitability is regarded as difficult. The exploit has been published and may be used. Upgrading to version 3.7.2 is sufficient to resolve this issue. Patch name: 6306e33da630c04ddff60703922c2ec7b7539452. It is advisable to upgrade the affected component. The vendor was contacted early about this disclosure.	4.6	More Details
CVE-2026-96357	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Stored XSS. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.6	More Details
CVE-2026-116487	NXP MCUXpresso SDK USB middleware (mcuxsdk-middleware-usb) before 2.13.0 contains an out-of-bounds read vulnerability in USB_DeviceCdcRndisSetCommand() that trusts a host-supplied informationBufferOffset. A malicious USB host can send RNDIS SET messages for OID_GEN_CURRENT_PACKET_FILTER or OID_802_3_MULTICAST_LIST with an out-of-range offset to read and disclose device memory.	4.6	More Details
CVE-2026-84915	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Component blocks allows Stored XSS. This issue affects Component blocks versions: from 0.0.0 to 1.2.7.	4.6	More Details
CVE-2026-84913	Missing Authorization vulnerability in Drupal AI Translate allows Forceful Browsing. This issue affects AI Translate versions: from 0.0.0 to 1.3.2, from 1.4.0 to 1.4.1.	4.6	More Details
CVE-2026-84912	Missing Authorization vulnerability in Drupal AI (Artificial Intelligence) allows Forceful Browsing. This issue affects AI (Artificial Intelligence) versions: from 0.0.0 to 1.3.13, from 1.4.0 to 1.4.8.	4.6	More Details
CVE-2026-84919	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal PhotoSwipe - Responsive JavaScript Modal Image Gallery allows Stored XSS. This issue affects PhotoSwipe - Responsive JavaScript Modal Image Gallery versions: from 0.0.0 to 5.0.9.	4.6	More Details
CVE-2026-87942	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal SafeDelete allows Stored XSS. This issue affects SafeDelete versions: from 0.0.0 to 1.0.88.	4.6	More Details
CVE-2026-116486	Objenesis before 3.6 contains a type confusion vulnerability that allows code in one class loader to substitute objects into another by exploiting instantiator caching keyed on class name. When a shared ObjenesisStd instance caches instantiators by clazz.getName() in ObjenesisBase.getInstantiatorOf(), same-named classes collide, causing newInstance() to return wrong-class objects and ClassCastException failures.	4.5	More Details
CVE-2026-112250	A vulnerability has been identified in BRAIN2 versions up to and including version 3.09. Under specific installation conditions, the password of a Windows service account may be written in cleartext to a setup log file. The issue occurs when Microsoft SQL Server access is configured using Integrated Security and the Windows account intended to run the BRAIN2 MainService is specified during execution of the BRAIN2 installation wizard. The setup log file is stored in the temporary directory of the Windows administrator performing the installation. This directory is not accessible to standard Windows users and can generally only be accessed by users with administrative privileges. Successful	4.4	More Details

	exploitation could result in the disclosure of credentials belonging to the configured BRAIN2 service account.		
CVE-2026-115400	crawlseo before version 0.2.5 contains a cleartext storage vulnerability that allows attackers with database read access to obtain Google OAuth access and refresh tokens stored unencrypted in the User.googleTokens column. Attackers can read the plaintext long-lived refresh tokens from a compromised or exposed Postgres database to gain indefinite, unauthenticated access to all registered users' Google Search Console data.	4.4	More Details
CVE-2026-95720	A buffer overflow vulnerability exists in AOS-8 that could allow an authenticated attacker with administrative access to cause denial of service. Successful exploitation could allow an attacker to crash the system, preventing it from rebooting without manual intervention and disrupting network operations.	4.4	More Details
CVE-2026-115536	A security flaw has been discovered in gedelumbung HospitalManagement up to c2d45543789a3887067d3915f69d44cfc2cf76a8. The impacted element is the function simpan of the file application/modules/admin/controllers/routing_pages.php of the component Admin Routing Page. The manipulation of the argument content results in cross site scripting. The attack may be performed from remote. The exploit has been released to the public and may be used for attacks. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-114176	Missing Authorization vulnerability in DigitalME eRoom eroom-zoom-meetings-webinar allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects eRoom: from n/a through 1.7.1.	4.3	More Details
CVE-2026-93636	An LDAP filter injection vulnerability in Zimbra Collaboration Suite (ZCS) allows an authenticated user to query otherwise unreturned attributes of calendar resources through SearchCalendarResourcesRequest. Zimbra escapes search values but inserts the supplied attribute names directly into LDAP filter expressions without validation. An attacker can inject additional conditions and use repeated matching queries to recover readable resource attributes such as internal mailbox hostnames and mail transport routes. The search remains restricted to calendar-resource objects and the LDAP service account's access rights. ZCS 10.1.16 and 10.1.21 are affected.	4.3	More Details
CVE-2026-93634	A cross-account authorization vulnerability in Zimbra Collaboration Suite (ZCS) allows an authenticated user to insert forged shared-file records into another user's mailbox through FileSharedWithMeRequest. The request bypasses delegated-account authorization and accepts attacker-selected file metadata, owner identifiers and declared sizes without validating a corresponding document share. An attacker can plant entries attributed to another real local account without that account sharing a document. No victim interaction is required. ZCS 10.1.16 and 10.1.20 are affected.	4.3	More Details
CVE-2026-93630	A hostname canonicalization vulnerability in the email HTML sanitizers of Zimbra Collaboration Suite (ZCS) allows an unauthenticated sender to bypass the restriction on forms that target the Zimbra host. DefangFilter and the OWASP ActionAttributePolicy compare the form-action hostname with the request hostname before browser URL canonicalization. A percent-encoded hostname can therefore survive sanitization while the browser submits the form to the normal Zimbra origin with the recipient's authentication cookie. This permits authenticated form requests from a crafted email after the recipient opens it in Classic and submits the form. The issue is present in ZCS 10.1.16 and 10.1.21. It does not bypass an endpoint's CSRF-token checks; a state-changing attack additionally requires an endpoint that accepts the resulting request without that protection.	4.3	More Details
CVE-2026-115718	Frappe Builder versions 1.33.0 before1.36.0 contain a missing authorization vulnerability in the improve_prompt endpoint that allows any authenticated non-Guest session to invoke LLM completions without proper permission checks. Attackers can call the improve_prompt API to leverage the site owner's stored OpenRouter API key via resolve_api_key, consuming paid LLM credits and effectively turning the site into a free LLM proxy.	4.3	More Details
CVE-2026-114647	The MongoDB C# Driver does not reject replacement documents whose top-level field names begin with a dollar sign in FindOneAndReplace, unlike other replace operations. An application that builds such a document from user-supplied field names can have the server apply update operators instead of a full replacement. The effect is limited to the document matched by the application's own filter, whose contents can be modified in ways the application did not intend.	4.3	More Details
CVE-2026-115785	An SMTP command injection vulnerability in the data-source testing functionality of Zimbra Collaboration Suite (ZCS) allows an ordinary authenticated user with account options enabled to inject CR/LF into the emailAddress inserted into an SMTP envelope command by a JSON SOAP request. After a successful preliminary POP3 or IMAP login, Zimbra sends the injected commands to the selected SMTP service. If that service accepts the pipelined transaction, the attacker can submit mail with attacker-	4.3	More Details

	selected recipients, headers and body content without victim interaction. The vulnerable command-construction path is present in ZCS 10.1.16 and 10.1.20.		
CVE-2026-115370	Incus 6.0.5 before 6.0.7, 6.12.0 before 7.0.2, and 7.1.0 before 7.5.0 contains a restriction bypass vulnerability that allows restricted project users to use pull mode migration via clustered refresh requests. Attackers with instance creation rights can send a refresh migration request for an existing instance to make the server open outbound TLS websocket connections and probe reachable services.	4.3	More Details
CVE-2026-115537	A weakness has been identified in gedelumbung HospitalManagement up to c2d45543789a3887067d3915f69d44cfc2cf76a8. This affects the function simpan of the file application/modules/admin/controllers/routing_pages.php of the component URL Validation. This manipulation of the argument url_route causes open redirect. It is possible to initiate the attack remotely. The exploit has been made available to the public and could be used for attacks. This product is using a rolling release to provide continuous delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-114402	Trudesk through 1.2.11 contains a path traversal and authorization bypass vulnerability in accountsController.uploadImage that allows authenticated users to overwrite other users' avatars via form-supplied _id and username fields. Attackers can supply a username like ../.././x to bypass the incomplete CVE-2022-1752 sanitization and write png, jpg, jpeg, or gif files outside public/uploads/users.	4.3	More Details
CVE-2026-114393	Trudesk through 1.2.11 contains an authorization bypass vulnerability that allows any authenticated user to modify tickets outside their groups via PUT /api/v1/tickets/:id and PUT /api/v2/tickets/batch. Attackers can supply arbitrary ticket ObjectIds to change status, subject, group, priority or assignee, or mass-close tickets, since only the default tickets:update role grant is checked.	4.3	More Details
CVE-2026-97295	Cross-Site Request Forgery (CSRF) vulnerability in David Lingren Media Library Assistant media-library-assistant allows Cross Site Request Forgery. This issue affects Media Library Assistant: from n/a through 3.42.	4.3	More Details
CVE-2026-87849	The Discovery naming-service reply parser in Apache bRPC (versions < 1.18.0) calls RapidJSON object-only accessors (FindMember) on JSON values without first validating that those values are objects. A malformed Discovery response whose data/service members or instances list contain non-object entries (e.g. a bare integer) triggers the RapidJSON object-access precondition on the client, reading memory out of bounds and potentially crashing the bRPC client. Root Cause: In src/brpc/policy/discovery_naming_service.cpp, DiscoveryNamingService::GetServers() and DiscoveryNames::ListDiscoveryNodes() access JSON values with member functions on non-object values. RapidJSON only defines these operations for object types; invoking them on a non-object violates the class invariant and results in undefined behavior (out-of-bounds read), commonly crashing the process. Affected Scenarios: An operator-configured (or compromised) Discovery backend returns a malformed /discovery/nodes or /discovery/fetchs response whose data/service/instance entries are not JSON objects, crashing any bRPC client that refreshes its server list. How to Fix: 1. Upgrade bRPC to version 1.18.0 or later, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3533 The fix validates that the Discovery response data, per-service data, and instance entries are JSON objects (and skips non-object node/instance entries and rejects non-object data/service values) before calling member accessors.	4.3	More Details
CVE-2026-114395	Trudesk through 1.2.11 contains an authorization bypass vulnerability in apiMessages.send that allows authenticated users to post messages as other users due to an incomplete CVE-2022-1931 fix. Attackers can supply arbitrary owner and cld values to POST /api/v1/messages/send to inject spoofed messages into conversations they do not belong to.	4.3	More Details
CVE-2026-114394	Trudesk through 1.2.11 contains an authorization bypass vulnerability that allows authenticated users to post content attributed to other users by supplying user-controlled owner IDs. Attackers can set ownerId or owner in postComment, postInternalNote, and ticket create requests to forge administrator-authored comments, internal notes, and tickets on any ticket.	4.3	More Details
CVE-2026-63243	An authentication vulnerability in the Bluetooth Low Energy (BLE) implementation of Creative Sound Blaster Katana could allow an attacker within Bluetooth range to connect to an affected device without prior pairing or authentication.	4.3	More Details
CVE-2026-116108	A flaw has been found in openbq-org OpenBB up to 4.7.2/5.0.0. This affects an unknown part of the file openbb_platform/core/openbb_core/app/model/api_settings.py of the component Local API Server. This manipulation causes improper access controls. Remote exploitation of the attack is possible. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details

CVE-2026-112828	Missing Authorization vulnerability in Meow Apps AI Engine ai-engine allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects AI Engine: from n/a through 3.8.6.	4.3	More Details
CVE-2026-115638	OpenStack Ironi through 39.0.1 allows a denial of service via a configdrive that results in a very large amount of uncompressed data, aka a gzip bomb.	4.3	More Details
CVE-2026-110376	Missing Authorization vulnerability in VillaTheme DEPART depart-deposit-and-part-payment-for-woo allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects DEPART: from n/a through 1.0.10.	4.3	More Details
CVE-2026-114389	Chamilo LMS from 2.0.0-alpha.4 before 2.0.0-RC.3 contains a missing authorization vulnerability in AiController::openAnswerGrade that allows any authenticated user to obtain AI grading of other students' exam attempts. Attackers can enumerate exeld and questionId values sent to the open_answer_grade endpoint to expose other students' answers through feedback and consume AI quota.	4.3	More Details
CVE-2026-115960	Missing Authorization vulnerability in WPExperts APIExperts Square for WooCommerce woosquare allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects APIExperts Square for WooCommerce: from n/a through 4.7.5.	4.3	More Details
CVE-2026-116224	A vulnerability has been found in fr0tt benotes up to 2.8.1. This affects the function getUrlInfo of the file app/Http/Controllers/PostController.php of the component Bookmark Metadata Endpoint. The manipulation of the argument url leads to server-side request forgery. The attack may be initiated remotely. The exploit has been disclosed to the public and may be used. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-114477	Missing authorization in the REST metrics route of Progressive Robot hMailServer 6.2.25 through 6.3.6 allows a remote authenticated user whose credential is confined to named domains to read the server-wide traffic history. GET /api/v1/metrics/history returns the server's traffic counters over the last 30 days; those counters are server-wide and carry no domain, which is why the report routes already refuse the same counters to a credential confined to named domains. The metrics route was not among the routes the server's authorization gate refuses a confined credential and its handler was not given the caller, so a domain administrator, an account given another role that holds view_configuration in named domains (the auditor role given a domain list among them), or an API key issued for named domains could read the whole server's traffic totals rather than its own domains'. The disclosure is of aggregate operational counters, not of any message or account.	4.3	More Details
CVE-2026-112868	Authorization Bypass Through User-Controlled Key vulnerability in H5P H5P h5p allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects H5P: from n/a through 1.17.9.	4.3	More Details
CVE-2026-113876	SmileCDR HAPI FHIR does not correctly check authorization for conditional create actions. A user with CREATE but not READ permissions could submit a conditional create whose predicate matches an existing resource and receive the target resource's contents.	4.3	More Details
CVE-2026-116018	ONEARMY community-platform before 2.165.2 contains an insecure direct object reference vulnerability that allows authenticated users to mark other users' notifications as read via POST /api/notifications/:id/read. Attackers can iterate sequential notification IDs to suppress unread indicators of users in the same tenant, causing them to miss replies, comments and new-content notices.	4.3	More Details
CVE-2026-116097	A security vulnerability has been detected in CloudClassroom-PHP-Project. This affects an unknown part of the file addnewstudent.php. Such manipulation of the argument fname leads to cross site scripting. The attack may be launched remotely. The exploit has been disclosed publicly and may be used. This product operates on a rolling release basis, ensuring continuous delivery. Consequently, there are no version details for either affected or updated releases. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-87955	Missing Authorization vulnerability in Drupal Ultimate Table Field allows Forceful Browsing. This issue affects Ultimate Table Field versions: from 0.0.0 to 1.1.1, from 2.0.0 to 2.0.1.	4.3	More Details
CVE-2026-116010	RosarioSIS before 12.9.1 contains a missing authorization vulnerability that allows authenticated teachers to write records for students outside their course period by supplying arbitrary student IDs. Attackers can modify the st_arr, values, or attendance POST arrays in MakeReferral.php, EnterEligibility.php, InputFinalGrades.php, and TakeAttendance.php to create referrals, eligibility codes, attendance and final grades for any student in the school.	4.3	More Details

CVE-2026-96390	Privilege Defined With Unsafe Actions vulnerability in Drupal Editoria11y Accessibility Checker allows Privilege Escalation. This issue affects Editoria11y Accessibility Checker versions: from 0.0.0 to 2.2.23, from 3.0.0 to 3.0.9.	4.3	More Details
CVE-2026-116015	yudao-cloud through 2026.09 contains an insecure direct object reference vulnerability in the mall trade module that allows authenticated members to read other members' after-sale operation logs by supplying arbitrary afterSaleId values. Attackers can enumerate sequential after-sale IDs against GET /app-api/trade/after-sale-log/list to expose refund amounts, status transitions and timestamps belonging to other members.	4.3	More Details
CVE-2026-116017	Xiaozhi-esp32-server before 0.9.8 contains an insecure direct object reference vulnerability in the manager-api DeviceController.getAddressBook handler that allows authenticated users to read other users' device address books. Attackers with a normal or self-registered account can request /device/address-book/{macAddress} with another user's device MAC to obtain contact MACs, aliases, call permissions, and owner user IDs.	4.3	More Details
CVE-2026-96378	Improperly Controlled Modification of Dynamically-Determined Object Attributes vulnerability in Drupal Commerce Decoupled Checkout allows Object Injection. This issue affects Commerce Decoupled Checkout versions: from 1.0.0 to 1.8.0.	4.3	More Details
CVE-2026-115982	A vulnerability was detected in phpIPAM up to 1.8.3. Affected by this issue is the function create Rack Link of the file app/tools/racks/show-rack-popup.php of the component Rack Link Creation. Performing a manipulation of the argument deviceid results in cross site scripting. Remote exploitation of the attack is possible. The exploit is now public and may be used. The vendor was contacted early about this disclosure.	4.3	More Details
CVE-2026-105777	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 17.6.5 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, under certain conditions, an unauthenticated user could have retrieved sensitive internal server credentials due to improper filtering of internal response headers.	4.3	More Details
CVE-2026-87954	Missing Authorization vulnerability in Drupal Taxonomy Term Glossary allows Forceful Browsing. This issue affects Taxonomy Term Glossary versions: from 0.0.0 to 4.6.0.	4.3	More Details
CVE-2026-115998	IDOR in Cervantes version v1.3-beta could allow an authenticated attacker ca access sensitive findings, proof-of-concept details, credentials, and client information from projects they do not belong to. Additionally, they can modify finding statuses and transfer findings to their own projects, compromising data confidentiality and integrity.	4.3	More Details
CVE-2026-116045	A vulnerability was detected in 1999AZZAR designer-mcp up to 2.1.0. This affects the function get_reference of the file skills/motion-designer/references/motion-dev-guidelines.md of the component MCP Interface. The manipulation of the argument Name results in path traversal. It is possible to launch the attack remotely. The exploit is now public and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The patch is identified as c8c2951dcbcd9b1710a745993f028317f10a3f10. Applying a patch is advised to resolve this issue.	4.3	More Details
CVE-2026-112378	GitLab has remediated an issue in GitLab EE affecting all versions from 17.8 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that under certain conditions could have allowed an authenticated user to access paid AI features without a valid subscription due to insufficient server-side validation in billing eligibility checks.	4.3	More Details
CVE-2026-116047	A vulnerability has been found in Privoce Vocechat up to 0.9.121. This issue affects some unknown processing of the file widget.html of the component Query Parameter Handler. Such manipulation of the argument welcome leads to cross site scripting. The attack can be launched remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-114574	A security vulnerability has been detected in massCodeIO massCode 5.10. Affected is an unknown function of the file src/main/api/index.ts of the component Local HTTP Service. Such manipulation leads to information disclosure. It is possible to launch the attack remotely. The exploit has been disclosed publicly and may be used. The name of the patch is d25d8fa3f5a1e34285c8634f18729650c0c05cda. It is advisable to implement a patch to correct this issue.	4.3	More Details
CVE-2026-116052	A security flaw has been discovered in LinkStackOrg LinkStack up to 4.8.8. The affected element is an unknown function. Performing a manipulation results in cross-site request forgery. Remote exploitation of the attack is possible. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details

CVE-2026-114634	WWBN AVideo before 29.4.0 contains a blind server-side request forgery vulnerability in <code>getHeaderContentTypeFromURL()</code> because <code>get_headers()</code> follows redirects without re-validating targets with <code>isSSRFSafeURL()</code> . Authenticated uploaders can store a <code>videoLink</code> pointing to a redirector, causing internal GET requests when visitors view the video, enabling internal port scanning via the embed-versus-video render oracle.	4.3	More Details
CVE-2026-13319	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 15.10 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, could have allowed an authenticated to read milestone metadata, including titles, descriptions, and full project paths, from arbitrary private projects across the instance without any project membership, due to missing authorization checks on a user-supplied project filter parameter in the group milestones API.	4.3	More Details
CVE-2026-116008	DaybydayCRM through 2.2.1 contains a missing authorization vulnerability in <code>ProjectsController::updateAssign()</code> that allows any authenticated user to reassign projects without the <code>can-assign-new-user-to-project</code> permission. Low-privileged employees can send PATCH or POST requests to <code>/projects/updateassign/{external_id}</code> with a chosen <code>user_assigned_id</code> to redirect project ownership and notifications.	4.3	More Details
CVE-2026-96398	Incorrect Authorization vulnerability in Drupal Webform allows Forceful Browsing. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.3	More Details
CVE-2026-115796	A security flaw has been discovered in Tenda F6 12.02.01.73. This impacts an unknown function of the component Web Management Interface. Performing a manipulation results in cross-site request forgery. The attack may be initiated remotely. The exploit has been released to the public and may be used for attacks.	4.3	More Details
CVE-2026-115322	Bytebase before 3.16.0 contains an authorization bypass vulnerability in <code>GetTaskRunLog</code> , <code>GetTaskRunSession</code> and <code>PreviewTaskRunRollback</code> that allows authenticated users to access other projects' task runs by UID. Attackers with permission on one project can supply foreign task or task run UIDs to read execution logs, database session details, and rollback SQL from other projects.	4.3	More Details
CVE-2026-114374	A vulnerability was identified in leomoon-studios wiki-go up to 1.9.1. This impacts an unknown function of the file <code>internal/handlers/metadata_api.go</code> of the component <code>fetch-metadata</code> Endpoint. Such manipulation leads to server-side request forgery. It is possible to launch the attack remotely. The exploit is publicly available and might be used. Upgrading to version 1.9.2 will fix this issue. The name of the patch is 25498ebf5c012b870967e38eec2a1e5133fa1a83. Upgrading the affected component is recommended.	4.3	More Details
CVE-2026-78801	An issue in SKYWORTH DIGITAL CO., LTD SK-WR6640X Wi-Fi 6 Dual-Band Fiber Router (AX1800M) V1.0.3 allows a remote attacker to obtain sensitive information via the <code>hidden_factory_switch.gch</code>	4.3	More Details
CVE-2026-116198	Snipe-IT before 8.7.0 contains a missing encryption vulnerability in <code>Api\AssetsController::store</code> that writes plaintext into encrypted-flagged custom-field columns when callers lack <code>assets.view.encrypted_custom_fields</code> . Authenticated API users without that permission can create assets whose sensitive custom-field values persist unencrypted in the database, exposing them through backups, replicas, or direct database access.	4.3	More Details
CVE-2026-115247	A security vulnerability has been detected in subhajitkhan online-clinic-management-system up to e9ee77a8827a1446220fa07ee693dc4d9a29a578. This vulnerability affects unknown code of the file <code>patientregistration.php</code> . The manipulation of the argument <code>pat_pswd/doc_pswd</code> leads to missing encryption of sensitive data. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-113489	The GeoDirectory WordPress plugin before 2.8.194 does not verify that a referenced attachment belongs to the listing being saved, allowing users with Subscriber-level access and above to modify the title, caption, moderation status and featured flag of images belonging to other users' listings, and to claim another user's image as their own featured image.	4.3	More Details
CVE-2026-116274	A vulnerability has been found in JIZHICMS up to 2.5.6. This affects the function <code>Pay</code> of the file <code>app/home/c/OrderController.php</code> of the component Front-end Order Handler. Such manipulation of the argument <code>orderno</code> leads to improper authorization. The attack can be launched remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-	Mattermost versions 11.7.x <= 11.7.11 fail to validate null elements when decoding playbook import conditions which allows an authenticated attacker to crash the Playbooks plugin instance-wide via a	4.3	More Details

13411	crafted import request containing a null condition. Mattermost Advisory ID: MMSA-2026-00709		
CVE-2026-115321	Bytebase before 3.16.0 contains an authorization bypass vulnerability in the UpdateIssueComment handler that allows low-privileged project users to modify comments in other projects by supplying a foreign comment UID. Attackers with any project role, including Project Viewer, can set their own project's issue as parent and reference another project's comment UID to overwrite comments they cannot access.	4.3	More Details
CVE-2026-112843	Missing Authorization vulnerability in WP Chill Strong Testimonials strong-testimonials allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Strong Testimonials: from n/a through 3.3.14.	4.3	More Details
CVE-2026-97352	The News Kit Addons For Elementor WordPress plugin before 1.4.9 does not perform a capability check on one of its AJAX actions, allowing authenticated users with subscriber-level access and above to install and activate a specific News Kit Addons For Elementor WordPress plugin before 1.4.9 from the WordPress.org repository.	4.3	More Details
CVE-2026-76295	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, and Splunk Secure Gateway versions below 3.10.12, 3.9.26, and 3.8.73, a user who does not hold the "admin" or "power" Splunk roles could view dashboard and report listings beyond the requested filter scope while remaining within objects that the user can access. The vulnerability is possible because Splunk Secure Gateway does not neutralize special characters in values before it creates filters for dashboard and report requests. For more information see About configuring role-based user access (https://help.splunk.com/en/splunk-enterprise/administer/manage-users-and-security/10.2/manage-splunk-platform-users-and-roles/about-configuring-role-based-user-access) in the Splunk documentation.	4.3	More Details
CVE-2026-115802	A weakness has been identified in InternScience InternAgent 1.5.0. This affects the function train of the file tasks/AutoChem/code/experiment.py of the component LoRA Loading. This manipulation of the argument base_model_name_or_path causes inclusion of functionality from untrusted control sphere. The attack may be initiated remotely. The exploit has been made available to the public and could be used for attacks.	4.3	More Details
CVE-2026-115803	A flaw has been found in Scada-LTS 2.8.0. Affected by this vulnerability is the function HttpServletRequest.getParameter of the file AmChartValuesAPI.java of the component AmCharts Data API. Executing a manipulation of the argument ids can lead to authorization bypass. The attack can be executed remotely. The exploit has been published and may be used. The vendor announces: "At this moment we are working on new release with security enhancements which will be introduced to the community soon."	4.3	More Details
CVE-2026-76296	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could create, move, edit, or delete sourcetype renaming configurations that they own in apps that they do not have permission to access. The vulnerability is possible when Splunk Web is turned on because Splunk Web does not verify the required app permissions before it performs sourcetype renaming operations. For more information see Rename source types at search time (https://help.splunk.com/en/splunk-enterprise/get-started/get-data-in/10.6/configure-source-types/rename-source-types-at-search-time) in the Splunk documentation.	4.3	More Details
CVE-2026-76304	In Splunk SOAR versions below 8.7.0, an authenticated user who does not hold a role with permission to edit Automation Brokers could download the Splunk SOAR Automation Broker Docker Compose configuration file. The file contains the Uniform Resource Locator (URL) of the Splunk SOAR instance, the container image location, and the Splunk SOAR Automation Broker version. The vulnerability is possible because Splunk SOAR returns the configuration file without enforcing the Automation Broker edit permission. For more information see Install Splunk Automation Broker on a Docker host (https://help.splunk.com/en/splunk-soar/splunk-automation-broker/install-and-create-classic-splunk-soar-automation-brokers/install-splunk-automation-broker-on-a-docker-host) 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.	4.3	More Details
CVE-2026-114182	Missing Authorization vulnerability in IdeaBox Creations PowerPack Addons for Elementor powerpack-lite-for-elementor allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects PowerPack Addons for Elementor: from n/a through 3.1.0.	4.3	More Details
CVE-2026-	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, a user who does not hold the "admin" or "power" Splunk roles could move field extractions or field transformations to apps for which the user does not have read access. The vulnerability is possible because the Representational State Transfer (REST) API does not validate destination app access before it processes field extraction and field transformation move operations. For more information see Use the Field extractions page (https://help.splunk.com/en/splunk-enterprise/manage-knowledge-objects/knowledge-management-	4.3	More Details

76294	manual/10.6/use-the-settings-pages-for-field-extractions-in-splunk-web/use-the-field-extractions-page) and 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) in the Splunk documentation.		
CVE-2026-115534	A vulnerability was determined in gedelumbung HospitalManagement up to c2d45543789a3887067d3915f69d44cfc2cf76a8. Impacted is the function kirim of the file application/modules/web/controllers/contact.php of the component Contact Form. Executing a manipulation of the argument nama/email/alamat/telepon/pesan can lead to cross site scripting. The attack can be executed remotely. The exploit has been publicly disclosed and may be utilized. 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-2026-114058	ArcadeDB from 26.2.1 before 26.10.1 contains a missing authorization vulnerability in the Bolt protocol handler that lets authenticated users read other databases' schema metadata. Attackers with credentials for one database can set the db field in RUN or BEGIN messages and call db.labels() or SHOW INDEXES to enumerate unauthorized schemas.	4.3	More Details
CVE-2026-114093	Trigger.dev before 4.5.6 contains a cross-site request forgery vulnerability that allows remote attackers to log out authenticated users by sending a link to the GET-based /logout endpoint. Attackers can lure victims into clicking an obfuscated, URL-encoded link that terminates their session, causing loss of unsaved work.	4.3	More Details
CVE-2026-116282	A security flaw has been discovered in Kameleo up to 5.3.1. Affected by this issue is the function CookieController.ListCookies of the file client-engine-application/w.cs of the component Cookie Controller. The manipulation of the argument guid results in information disclosure. The attack can be executed remotely. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-116283	A weakness has been identified in Kameleo up to 5.0.1.26. This affects an unknown part of the file /profiles of the component Local API. This manipulation causes information disclosure. The attack is possible to be carried out remotely. Upgrading to version 5.2.0 is able to mitigate this issue. Upgrading the affected component is advised.	4.3	More Details
CVE-2026-113741	A vulnerability has been found in itsourcecode Online Cake Ordering System 1.0. Affected by this issue is some unknown functionality of the file /404.php. Such manipulation of the argument userid leads to cross site scripting. The attack may be performed from remote. The exploit has been disclosed to the public and may be used.	4.3	More Details
CVE-2026-93856	In OpenStack Cinder through 29.0.0, cross-project information disclosure can occur in the Messages API. A non-admin user with the project-scoped 'member' or 'reader' role making the Block Storage API v3.3 "List messages" call finds the non-expired user messages for other projects included in the response, resulting in the leakage of UUIDs of resources owned by other projects, in addition to the user message about the other projects' resources (which, by design, are sanitized generic messages). All Cinder deployments are affected.	4.3	More Details
CVE-2026-86611	The Download Manager WordPress plugin before 3.3.71, Download Manager WordPress plugin before 7.5.9 does not have CSRF checks in some places, allowing attackers to make a logged in admin update its privacy and storage settings via a crafted link, including the credential that gates its unauthenticated maintenance endpoint.	4.3	More Details
CVE-2026-109926	The Productive Commerce WordPress plugin before 1.1.49 does not perform a capability check in one of its AJAX actions and exposes the nonce that guards it on every admin-area page, allowing any authenticated user such as a subscriber to activate arbitrary Productive Commerce WordPress plugin before 1.1.49 already installed on the site.	4.3	More Details
CVE-2026-114056	ArcadeDB from 26.8.1 before 26.10.1 contains an authorization bypass vulnerability that allows read-only users to modify edge type lightweight and unique schema flags without UPDATE_SCHEMA permission. Attackers can send a bare ALTER TYPE edgeType WITH lightweight = true command over the HTTP, gRPC, or Bolt command API to persistently alter edge storage declarations.	4.3	More Details
CVE-2026-114057	ArcadeDB from 26.2.1 before 26.10.1 contains a missing authorization vulnerability in the Bolt protocol handler that allows authenticated users to enumerate all database names. Attackers with credentials for any single database can issue a SHOW DATABASES query over the Bolt port to reveal unauthorized database names, exposing tenant or project identifiers.	4.3	More Details
CVE-2026-	The SmartyParcel WordPress plugin before 1.23.4 does not perform authorisation checks when deleting shipping label records, allowing any authenticated user such as a subscriber to delete the shipping label	4.3	More

107788	and tracking records of arbitrary orders store-wide.		Details
CVE-2026-114072	The Trigger.dev webapp before 4.5.12 contains a cross-site request forgery vulnerability in its organization-level Slack OAuth flow, which uses the organization ID as the OAuth state without binding it to the initiating session. Attackers who know a target organization ID can trick an authenticated member into opening an unredeemed /integrations/slack/callback URL, connecting an attacker-controlled Slack workspace that may receive the organization's operational alerts.	4.3	More Details
CVE-2026-116284	A security vulnerability has been detected in Kameleo up to 5.3.1. This vulnerability affects unknown code of the file /general/user-info of the component Local API. Such manipulation leads to information disclosure. The attack may be performed from remote. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-116340	A vulnerability was identified in iFlytek astron-agent up to 1.1.2. This impacts the function FileInfoV2Service.listKnowledgeByPage of the file FileInfoV2Service.java of the component Console Knowledge-List Endpoint. Such manipulation of the argument space-id leads to authorization bypass. The attack may be launched remotely. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2026-111965	The Profile Builder WordPress plugin from 4.0.0 to 4.0.4 does not enforce its reCAPTCHA verification on all comment submission paths, allowing authenticated users to bypass the anti-spam protection an administrator configured for the WordPress comment form.	4.3	More Details
CVE-2026-115752	Signal K Server from 2.28.0-beta.2 before 2.34.0 contains a server-side request forgery vulnerability in the App Store icon fetcher that trusts absolute signalk.appicon URLs. Attackers publishing npm packages with Signal K keywords can make servers request loopback, vessel network, or cloud metadata endpoints, storing image-typed responses served back via the admin UI.	4.2	More Details
CVE-2026-84910	Missing Authorization vulnerability in Drupal Advanced Search allows Forceful Browsing. This issue affects Advanced Search versions: from 0.0.0 to 2.4.5.	4.2	More Details
CVE-2026-76297	In Splunk Enterprise versions below 10.4.4, 10.2.8, 10.0.11, and 9.4.16, an unauthenticated user could trick a user that holds a role with the edit_federated_providers capability into turning off all federated search providers. Successful exploitation could temporarily prevent federated searches from retrieving data from remote datasets. The vulnerability is possible because the federated provider Representational State Transfer (REST) API accepts GET requests for actions that change provider state. 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 Run federated searches over remote Splunk platform deployments (https://help.splunk.com/en/splunk-enterprise/search/federated-search/10.6/run-federated-searches-across-other-splunk-deployments/run-federated-searches-over-remote-splunk-platform-deployments) 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) in the Splunk documentation.	4.2	More Details
CVE-2026-114438	A flaw was found in the Update User Locale required action of the Keycloak identity and access management solution. When this action is ordered before the Configure OTP action in a realm, the system may fail to enforce mandatory multi-factor authentication enrollment during the login process. This issue allows an attacker with a valid password to access an account without setting up the required one-time password (OTP) credential.	4.2	More Details
CVE-2026-78117	The Infineon PSOC 6 BLESS HCI driver in drivers/bluetooth/hci/hci_infineon_psoc6_bless.c stored the outgoing HCI packet descriptor in a file-scope variable, static cy_stc_ble_hci_tx_packet_info_t hci_tx_pkt. psoc6_bless_send() cleared and populated that descriptor's packetType, dataLength and data fields before acquiring psoc6_bless_operation_sem, so the driver's serialising semaphore protected only the call into Cy_BLE_SoftHciSendAppPkt() and not the shared state that call reads. The Bluetooth host does not guarantee a single HCI sender. Most transmissions are funnelled through the cooperative TX-processor work queue, but bt_hci_cmd_send() deliberately bypasses bt_dev.cmd_tx_queue for BT_HCI_OP_HOST_NUM_COMPLETED_PACKETS and calls bt_send() directly from the caller's thread; that caller is bt_hci_host_num_completed_packets(), the acl_in_pool destroy callback, which runs on whichever thread drops the last reference to a received ACL buffer. CONFIG_BT_HCI_ACL_FLOW_CONTROL defaults to y on this driver's boards, and an application that returns -EINPROGRESS from its L2CAP recv callback and later calls bt_l2cap_chan_rcv_complete() performs that unref from its own preemptible thread. A remote peer that generates Bluetooth traffic concurrently with the device's own transmissions raises the collision rate. When two senders interleave, the semaphore holder can hand the controller a descriptor whose dataLength comes from one net_buf and whose data pointer comes from another — or one whose buffer has already been returned to its pool. The controller then copies and transmits up to a full HCI/ACL payload starting at a shorter or	4.2	More Details

	recycled buffer, disclosing adjacent RAM contents to the connected peer; other interleavings duplicate one HCI packet and silently drop another, disrupting the link. The driver only reads through the shared descriptor, so no memory corruption or write primitive results, and the disclosed bytes are not attacker-chosen. The fix makes the descriptor a local variable in psoc6_bless_send(), removing the shared state; Cy_BLE_SoftHciSendAppPkt() copies the packet before returning, so it need not outlive the call.		
CVE-2026-96374	Cross-Site Request Forgery (CSRF) vulnerability in Drupal Project Browser allows Cross Site Request Forgery. This issue affects Project Browser versions: from 0.0.0 to 2.0.3, from 2.1.0 to 2.1.5.	4.2	More Details
CVE-2026-96362	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Cross-Site Scripting (XSS). This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	4.2	More Details
CVE-2026-113796	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-alpha.6, Electron treated pages opened into a new window from links as application-initiated navigations instead of preserving the page containing the link as the initiator. The resulting top-level request could be classified as same-site, causing SameSite=Strict cookies to be sent and Sec-Fetch-Site not to identify the initiating page. An application that loads untrusted content and permits it to open new windows could therefore allow that content to make a credentialed request to another site that relies on SameSite cookies for CSRF protection. This issue is fixed in versions 42.11.3, 43.7.0, 44.3.0, and 45.0.0-alpha.6.	4.2	More Details
CVE-2026-115725	Chatwoot before 4.17.0 contains an authorization bypass vulnerability that allows authenticated agent bots to access other accounts by supplying a foreign contact inbox source_id. Attackers can call the conversations create endpoint with another account's source_id to create or retrieve conversations and post messages to that account's contacts.	4.2	More Details
CVE-2026-115509	Jitsu console before 2.14.0-beta.41 contains an incorrect authorization vulnerability in sources/state.ts handlers that call findFirst without await, so the syncId ownership check never fails. Authenticated users knowing another workspace's sync ID can read, overwrite or delete its stream cursors, causing data re-ingestion or skipped data.	4.2	More Details
CVE-2026-87939	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Feed Block allows Stored XSS. This issue affects Feed Block versions: from 0.0.0 to 2.0.2, from 3.0.0 to 3.0.2.	4.2	More Details
CVE-2026-114617	Authlib is a Python library which builds OAuth and OpenID Connect servers. Prior to 1.9.0, OpenIDMixin.parse_id_token() reads the nonce_supported decision from the untrusted ID token being validated and sets the expected nonce to null when that claim is the JSON boolean false. In implicit and hybrid relying parties, or applications that directly accept externally supplied ID tokens, an attacker with a captured validly signed token can bypass the nonce-to-session binding and replay the victim's token into an attacker-controlled session. Plain authorization code flow is not affected, and the token must still pass signature, issuer, and audience validation. OpenIDMixin.parse_id_token() trusts nonce_supported in both authlib/integrations/base_client/sync_openid.py and authlib/integrations/base_client/async_openid.py to disable the expected nonce. This issue is fixed in version 1.9.0.	4.2	More Details
CVE-2026-115361	OpenBao before 2.6.3 contains an ACL bypass vulnerability that silently drops explicit deny rules based on forbidden identity metadata templates within wildcard grants. Attackers whose identity metadata contains wildcard or path segments can access paths that should be denied within the broader wildcard-granted scope.	4.1	More Details
CVE-2026-114101	Coraza before 3.8.1 extracts multipart upload filenames with Go's mime.ParseMediaType, which silently ignores RFC 5987 filename* parameters whose charset is not utf-8 or us-ascii. Unauthenticated attackers can send a decoy filename with filename*=iso-8859-1"shell.php, or filename* alone, so FILES-based rules never see the real name, letting dangerous uploads reach backends that honor filename*.	4.0	More Details
CVE-2026-115380	Wasmtime before 36.0.17, 37.0.0 before 48.0.4, and 49.0.0 before 49.0.2 contains a denial of service vulnerability in the wasmtime-wasi WASI preview 0 poll_oneoff adapter, which performs O(n) host work per subscription without consuming or checking fuel. Malicious guest modules can pass many subscriptions to force long-running host computation that ignores configured fuel limits.	4.0	More Details
CVE-2026-114717	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a mail header injection vulnerability when a message is forwarded as an attachment. The subject of the forwarded message was used verbatim to generate the attachment name, so a received message whose subject contains CRLF sequences could inject arbitrary MIME part headers into the outgoing message. Upstream fixed this in the commit "Fix mail header injection via subject of a message forwarded as attachment", which collapses whitespace in the generated name.	4.0	More Details

CVE-2026-114651	A memory-handling weakness in the BSON deserialization component of the MongoDB C# Driver may cause data left over from previously processed content to be returned when specially crafted serialized input is parsed. An unauthenticated party able to supply serialized data to an application that maps certain primitive collection types can cause residual bytes from unrelated data processed earlier on the same execution thread to be surfaced as legitimate values. This can result in disclosure of a limited amount of information belonging to other requests or principals of the consuming application. No authentication or special privileges within the application are required.	4.0	More Details
CVE-2024-6423	The product failed to enforce validation constraints on user-supplied input in URL parameters, allowing sensitive data to be passed to flows where it was not intended to be processed. Although the product did not utilize this data, its acceptance could lead to logging. If a user mistakenly sends sensitive data via URL parameters to any flow, this data could be inadvertently logged in access logs. This has the potential to expose confidential information, compromising the security and privacy of affected systems if the product is used incorrectly.	4.0	More Details
CVE-2026-114122	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.6, 43.7.3, 44.4.3, and 45.0.0-alpha.9, frame data received from the GPU process for webContents.beginFrameSubscription was not fully validated by the main process. An attacker who had separately gained code execution in the GPU process could supply inconsistent frame geometry and cause the main process to read out-of-bounds memory while producing captured frame images, disclosing memory or crashing the application. This issue is fixed in versions 42.11.6, 43.7.3, 44.4.3, and 45.0.0-alpha.9.	3.9	More Details
CVE-2026-116145	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-32 contains a security policy bypass caused by an incomplete fix for the policy bypass reported in GHSA-jjp4-3fwf-393j. An attacker able to supply input to ImageMagick can evade configured security policy restrictions, allowing operations the policy was intended to block.	3.9	More Details
CVE-2026-114124	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.6, 43.7.3, 44.4.3, and 45.0.0-alpha.9, frame data received from the GPU process for macOS offscreen rendering was not fully validated by the main process. An attacker who had separately gained code execution in the GPU process could provide inconsistent frame geometry and cause the main process to read out-of-bounds memory while producing paint event images, disclosing memory or crashing the application. This issue is fixed in versions 42.11.6, 43.7.3, 44.4.3, and 45.0.0-alpha.9.	3.9	More Details
CVE-2026-115393	Wasmtime before 24.0.12, 25.0.0 before 36.0.13, 37.0.0 before 46.0.2, and 47.0.0 before 47.0.3 contains a type confusion vulnerability because several Store APIs fail to validate that objects belong to the Store's Engine. Embedders mixing objects between Engines via Instance::new, StructRefPre::new, or Tag::new can trigger trampoline signature mismatches causing out-of-bounds stack writes, GC heap corruption, or panics.	3.8	More Details
CVE-2026-96387	Authentication Bypass Using an Alternate Path or Channel vulnerability in Drupal Stop administrator login allows Authentication Bypass. This issue affects Stop administrator login versions: from 1.0 to 1.6.	3.8	More Details
CVE-2026-114741	The WPVR WordPress plugin before 9.0.3 does not perform object-level authorization checks when reading and updating virtual tours, allowing users with contributor-level access and above to read and modify tours created by other users.	3.8	More Details
CVE-2026-87951	Insecure Storage of Sensitive Information vulnerability in Drupal SAML SSO - Service Provider allows Retrieve Embedded Sensitive Data. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	3.8	More Details
CVE-2026-110537	The Essential Addons for Elementor WordPress plugin before 6.8.7 does not verify that an administrator may manage the specific user accounts targeted by its login-approval feature, checking only a site-level capability, which on WordPress Multisite could allow a subsite administrator to change the approval and login state of accounts belonging to other sites on the network.	3.8	More Details
CVE-2026-21811	HCL AppScan Azure Boards Plugin is potentially affected by an Information Disclosure vulnerability involving secrets accessible in plaintext during pipeline execution. Depending on the pipeline configuration, these secrets could be exposed through task outputs, debug logs, environment variable dumps, or build agent scripts, despite being masked in the user interface.	3.8	More Details
CVE-2026-93846	Mattermost versions 11.7.x <= 11.7.10 fail to restrict which built-in roles can be assigned to a team member via the API, which allows a team admin to grant a team member permissions beyond those permitted by the configured permission scheme via the {{update-team-member-roles}} REST API.. Mattermost Advisory ID: MMSA-2026-00734	3.8	More Details
CVE-			

2026-16648	Insertion of Sensitive Information into Log File vulnerability in Drupal Mailer Plus Log allows Remote Services with Stolen Credentials. This issue affects Mailer Plus Log versions: from 0.0.0 to 1.2.7.	3.8	More Details
CVE-2026-114377	A security vulnerability has been detected in leomoon-studios wiki-go up to 1.9.1. Affected by this issue is the function DeleteComment of the file internal/comments/comments.go of the component Comment Deletion Handler. The manipulation of the argument documentPath leads to path traversal. The attack may be initiated remotely. The exploit has been disclosed publicly and may be used. Upgrading to version 1.9.2 can resolve this issue. The identifier of the patch is 224ef2a11facddef6b3a3c1f0a69156d1e574a33. It is advisable to upgrade the affected component.	3.8	More Details
CVE-2026-114576	A vulnerability has been found in Linksys RE6250 1.0.04.001. This affects the function strcmp of the component Web Management Interface. The manipulation of the argument session_id leads to insufficiently random values. The attack may be initiated remotely. The complexity of an attack is rather high. It is indicated that the exploitability is difficult. The exploit has been disclosed to the public and may be used.	3.7	More Details
CVE-2026-114641	Traefik 2.11.x before 2.11.58 and 3.x before 3.7.14 contains a race condition in NTLM and Kerberos (Negotiate) backend authentication, where the per-connection sticky round tripper slot is read and written without synchronization. Remote attackers can send concurrent HTTP/2 or HTTP/3 requests over one connection to orphan backend connections, break connection-bound authentication, or crash Traefik.	3.7	More Details
CVE-2026-114459	A flaw was found in mod_qos. When inspecting HTTP response bodies, improper boundary checking can cause the module to read memory outside allocated buffer boundaries. An unauthenticated remote attacker can exploit this vulnerability by triggering specifically crafted server responses, causing an Apache worker process to crash and resulting in a Denial of Service (DoS).	3.7	More Details
CVE-2026-114457	A flaw was found in mod_qos. An unauthenticated remote attacker could cause a Denial of Service (DoS) by sending a specially crafted JSON request when JSON body filtering is enabled. Due to improper string termination during the processing of percent-encoded input, the module reads past allocated memory boundaries, which can cause an Apache worker process to crash.	3.7	More Details
CVE-2026-116151	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak in the JNG encoder that fails to release allocated memory when processing a crafted image. Attackers can supply crafted images that are encoded to JNG to leak memory, potentially causing resource exhaustion and a limited denial of service.	3.7	More Details
CVE-2026-115750	mcp-from-openapi from 2.5.0 before 2.9.0 contains a credential exposure vulnerability that allows attackers to capture Authorization and Cookie headers because safeFetch() forwards them on cross-origin redirects. Attackers controlling a spec host or abusing an open redirect can redirect fromURL() spec requests to their own origin to steal bearer tokens or cookies.	3.7	More Details
CVE-2026-109931	The Infility Global WordPress plugin before 2.16.25 does not perform any authorisation or nonce check on an AJAX action available to unauthenticated users, allowing any visitor to overwrite a machine-translation configuration value belonging to a companion Infility Global WordPress plugin before 2.16.25.	3.7	More Details
CVE-2026-114723	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 allows the ImageMagick resource limits to be bypassed during thumbnail generation. When rcube_image falls back to GD, the available-memory check was applied only on Windows, so on other platforms a crafted image with large dimensions was processed by GD without any resource limit and could exhaust server memory. Upstream fixed this in the commit "Fix ImageMagick resource limits bypass by GD fallback during thumbnail generation", which applies the memory check on all platforms.	3.7	More Details
CVE-2026-35501	When NLnet Labs Unbound 1.0.1 up to and including version 1.26.1 is configured with both a root trust anchor and an additional 'trust-anchor:' for an island-of-security zone whose parent is provably insecure under the root chain, a malicious actor with a forged NSEC reply with the SOA bit cleared can trick Unbound to treat the configured trust island as insecure. This happens because the DNSSEC validator in Unbound accepts and caches the NSEC insecure response bypassing the RRSIG validation that the validator was supposed to do. The acceptance of the NSEC record at the trust anchor apex was allowed due to and old fix for type DS and NSEC queries from clients when trust anchor for that name existed.	3.7	More Details
CVE-2026-112230	The Popup Builder WordPress plugin before 4.4.6 does not perform authorization checks before trashing popup content in a maintenance routine that runs on unauthenticated requests, allowing unauthenticated attackers to move affected popups into the Trash and overwrite their stored preview data.	3.7	More Details
CVE-2026-	The GeoDirectory WordPress plugin before 2.8.194 does not correctly verify ownership of listings submitted by logged-out users before allowing them to be edited, allowing unauthenticated attackers to	3.7	More

113674	modify, unpublish, attach images to, or reassign any guest-submitted listing when guest posting is enabled.		Details
CVE-2026-116190	rcldone from v1.56.0 before v1.75.2 contains an information disclosure vulnerability in the WebDAV backend that forwards operator-configured custom credential headers to different hosts on HTTP 3xx redirects. Attackers who control or influence a cross-host redirect can capture headers such as X-Auth-Token set via --webdav-headers and replay them against the original WebDAV service.	3.7	More Details
CVE-2026-115797	A weakness has been identified in Tenda F6 12.02.01.73. Affected is an unknown function of the component Web Management Interface. Executing a manipulation of the argument ecos_pw can lead to cleartext transmission of sensitive information. The attack may be launched remotely. This attack is characterized by high complexity. The exploitability is told to be difficult. The exploit has been made available to the public and could be used for attacks.	3.7	More Details
CVE-2026-114370	A flaw was found in the PathResourceManager component of the Undertow web server. This component is responsible for serving static files from the file system. The issue occurs because the mechanism used to verify if a symbolic link points to a safe location only checks if the target path starts with a specific string, without ensuring it matches a full directory name. This could allow an attacker to access files in a sibling directory if that directory's name starts with the same characters as an authorized safe directory.	3.7	More Details
CVE-2026-116105	A weakness has been identified in HeyCyan App 1.0.103 on Android. Affected is an unknown function of the file AndroidManifest.xml of the component Network Security Configuration. Executing a manipulation can lead to cleartext transmission of sensitive information. The attack can be launched remotely. The attack requires a high level of complexity. The exploitability is told to be difficult. The vendor was contacted early about this disclosure but did not respond in any way.	3.7	More Details
CVE-2026-116114	A weakness has been identified in Yoosee GM8135S Camera 21.00.01.36. This affects an unknown function of the file /npc/npc of the component Live UDP H.264 Video Stream Handler. This manipulation causes cleartext transmission of sensitive information. The attack can be initiated remotely. The attack's complexity is rated as high. The exploitability is reported as difficult. The exploit has been made available to the public and could be used for attacks. The vendor was contacted early about this disclosure but did not respond in any way.	3.7	More Details
CVE-2026-114115	Vyper 0.3.8 through 0.5.0b1 contains a compiler miscompilation vulnerability in parse_IfExp that reads nested arrays with narrower inner bounds using the wider result type's memory layout. When a ternary selects the narrower array, inner arrays after the first are read at wrong offsets, causing contracts to silently store corrupted values or revert.	3.7	More Details
CVE-2026-114472	Insufficient resource pool in the ACME HTTP-01 challenge responder of Progressive Robot hMailServer 6.0.0 through 6.3.6 allows a remote unauthenticated attacker to cause a certificate issuance or renewal to fail. When the server obtains or renews a TLS certificate over ACME it runs an HTTP-01 challenge responder on port 80 for the certificate authority to read the challenge token; the responder served one connection at a time, accepting a socket, waiting up to five seconds to receive a request, answering it, and only then accepting the next. A client that connected and sent nothing held the single worker for the whole timeout, so enough idle connections during a validation keep the certificate authority's own request from being answered in time and the validation fails. Sustained across a certificate's renewals this leads to the certificate expiring, after which clients refuse the server's TLS; the mail service itself continues on its current certificate until then. No authentication is required to connect to the responder. Version 6.3.7 serves up to 64 connections at once and closes a connection past that, so that where a handful of idle connections, one after another, sufficed before, a client must keep 64 idle connections open at once throughout a validation to have the certificate authority's connection closed.	3.7	More Details
CVE-2026-115359	OpenBao before 2.6.4 and 2.7.x before 2.7.1 contains a flaw in its rate limit quota subsystem where role resolution fails for requests with form-encoded bodies or forwarded certificates, so quota rules scoped to backend-specific roles are not enforced. Valid callers can repeatedly issue tokens without consuming the configured role-based rate limit budget.	3.7	More Details
CVE-2026-116121	A vulnerability was identified in NocoDB up to 2026.08.0. This issue affects some unknown processing of the file users.service.ts of the component Registration. Such manipulation leads to race condition. The attack may be performed from remote. The attack requires a high level of complexity. The exploitability is assessed as difficult. The vendor deleted the GitHub issue for this vulnerability without any explanation. Afterwards the vendor was contacted early about this disclosure via email but did not respond in any way.	3.7	More Details
CVE-2026-116122	A security flaw has been discovered in NocoDB up to 2026.07.0. Impacted is the function AuthService.validateUser of the component Sign-in Endpoint. Performing a manipulation of the argument salt results in information disclosure. It is possible to initiate the attack remotely. The attack's complexity is rated as high. The exploitability is considered difficult. The vendor deleted the GitHub	3.7	More Details

	issue for this vulnerability without any explanation. Afterwards the vendor was contacted early about this disclosure via email but did not respond in any way.		
CVE-2026-116123	A weakness has been identified in NocoDB up to 2026.07.0. The affected element is the function validatePassword of the file /auth/user/signup. Executing a manipulation can lead to improper restriction of excessive authentication attempts. It is possible to launch the attack remotely. This attack is characterized by high complexity. The exploitability is described as difficult. The vendor documents this as intentional. The vendor deleted the GitHub issue for this vulnerability without any explanation. Afterwards the vendor was contacted early about this disclosure via email but did not respond in any way.	3.7	More Details
CVE-2026-116013	OpenResty lua-resty-dns before 0.24 contains an improper input validation vulnerability in parse_section() and _decode_name() that fails to check resource record lengths against the received message. Attackers controlling or spoofing a queried nameserver can send truncated record headers or RDATA, raising uncaught Lua exceptions in resolver:query() that abort requests and break name resolution.	3.7	More Details
CVE-2026-115751	Agentfront mcp-from-openapi from 2.2.0 before 2.9.0 contains a server-side request forgery vulnerability that allows attackers to bypass the ssrf.blockedHosts denylist using trailing-dot hostnames. Attackers who control an OpenAPI spec URL or redirect target can append a root dot to a blocked hostname to fetch content from operator-blocked hosts.	3.7	More Details
CVE-2026-96901	The Cart Abandonment Recovery for WooCommerce WordPress plugin before 2.2.0 does not verify the integrity of its cart-recovery token before acting on it and still accepts an unsigned legacy token format, allowing unauthenticated attackers who know a victim's cart session identifier to modify that victim's abandoned-cart record, for example unsubscribing them from recovery emails or applying an arbitrary coupon to their restored cart.	3.7	More Details
CVE-2026-114049	vLLM before 0.31.0 contains a cross-request mask leak in InputBatch.condense() that allows remote attackers to constrain other clients' outputs by leaving stale allowed_token_ids rows. Attackers can submit concurrent constrained generate requests, including via unauthenticated /invocations, so recycled batch slots apply their token allowlist to victims' unrestricted requests.	3.7	More Details
CVE-2026-116137	adm-zip through 0.6.1 contains a duplicate-entry validation bypass that allows attackers to evade the readEntries() check by using canonically equivalent entry names like './hello.txt' and 'hello.txt'. Attackers can craft ZIP archives where getEntry() or readAsText() validates benign content while extractAllTo() writes malicious content to the same path, bypassing integrity checks.	3.7	More Details
CVE-2026-84916	Missing Authorization vulnerability in Drupal Islandora allows Forceful Browsing. This issue affects Islandora versions: from 0.0.0 to 2.19.0.	3.7	More Details
CVE-2026-84918	Improper Verification of Cryptographic Signature vulnerability in Drupal Monobank payment API allows Signature Spoof. This issue affects Monobank payment API versions: from 0.0.0 to 1.0.3.	3.7	More Details
CVE-2026-114719	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 is affected by a CSRF validation bypass in the mark-as-junk action of the markasjunk plugin. The plugin read the _uid parameter from POST data but passed it to rcmail::get_uids() without constraining the input mode, so values re-read inside that helper could be taken from the query string and the action carried out without a valid CSRF token. Upstream fixed this in the commit "Fix CSRF validation bypass in the mark-as-junk action", which passes rcube_utils::INPUT_POST explicitly.	3.7	More Details
CVE-2026-116544	Open5GS 2.6.1 through 2.8.0 contains an out-of-bounds read vulnerability in parse_scheme_output() in lib/sbi/conv.c during SUCI de-concealment due to incomplete length validation. Remote unauthenticated UEs or SBI peers can submit a crafted SUCI with a 42-byte scheme output to a UDM, triggering a heap over-read past the buffer.	3.7	More Details
CVE-2026-113735	A vulnerability was determined in D-Link DCS-6501LH 3.03.02. Affected by this vulnerability is the function mydlink_set_root_pwd of the file /mnt/mtd/aoni_ipc. Executing a manipulation can lead to use of weak credentials. The attack may be launched remotely. The attack requires a high level of complexity. The exploitation appears to be difficult. The exploit has been publicly disclosed and may be utilized.	3.7	More Details
CVE-2026-116007	A flaw has been found in YunaiV/zhijiantianya ruoyi-vue-pro up to 2026.09. The impacted element is the function SmsCodeServiceImpl.validateSmsCode0 of the file /service/sms/SmsCodeServiceImpl.java of the component SMS Code Verification. This manipulation causes improper restriction of excessive authentication attempts. The attack may be initiated remotely. The complexity of an attack is rather high. The exploitability is regarded as difficult. The exploit has been published and may be used. The	3.7	More Details

	vendor was contacted early about this disclosure but did not respond in any way.		
CVE-2020-37288	Vyper from 0.1.0b10 before 0.2.3 generates incorrect code for a return statement inside a for loop in an internal function called two or more internal calls deep, producing an invalid jump destination (stack underflow). Affected contracts revert such calls unnecessarily, which in limited circumstances could enable denial of service against public contracts.	3.7	More Details
CVE-2026-116147	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a null pointer dereference in the DJVU decoder because the result of a memory allocation is not checked before use. Attackers can supply a DJVU image that, when processed under conditions where the allocation fails, crashes the process and causes a denial of service.	3.7	More Details
CVE-2026-116146	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak in the MSL decoder that fails to release allocated memory when processing a crafted MSL image. Attackers can supply crafted MSL files for processing to repeatedly leak memory, causing resource exhaustion and denial of service.	3.7	More Details
CVE-2020-37289	Vyper before 0.3.1 lacks a uint8 type, allowing interface return values declared as uint8, such as ERC20.decimals(), to be stored in uint256 variables without bounds checking. Malicious contracts can return values greater than 255 to manipulate or bypass contract logic that relies on these values.	3.7	More Details
CVE-2026-115996	A vulnerability was detected in Unitree Robotics Unitree Go2 up to 1.1.13/1.1.14.1/1.1.15.1. This affects the function WebSocketClient.connect of the file websocket.py of the component chat_go. Performing a manipulation results in improper certificate validation. The attack may be initiated remotely. The attack's complexity is rated as high. The exploitability is reported as difficult. The vendor confirmed early: "These issues were previously submitted by the researcher through our Security Response Center platform and are tracked internally as follows: Submission 990504 corresponds to our internal ticket #616 (Go2 chat_go WSS client disabled TLS server certificate validation, CWE-295)."	3.7	More Details
CVE-2026-79978	The RTIO streaming completion callback adxl345_process_fifo_samples_cb() in drivers/sensor/adi/adxl345/adxl345_stream.c clamps the amount of FIFO data it will report (read_len = MIN(fifo_bytes, buf_avail)) to the space actually returned by rtio_sqe_rx_buf(), but then queues one 6-byte bus read per hardware FIFO entry at consecutive offsets from read_buf, bounded by the device-reported entry count rather than by the clamp. When the allocated buffer is smaller than a full batch, the driver writes up to 186 bytes past the end of the allocation; a companion defect left data->fifo_total_bytes un-reset at batch start, so an aborted or errored chain could leave a stale write offset that pushes the next batch's writes out of bounds as well. A short buffer is an expected, first-class outcome of the RTIO allocator: rtio_block_pool_alloc() walks the requested size down one block at a time until a contiguous run is available, and the driver only demands sizeof(struct adxl345_fifo_data) + 6 bytes as its minimum. Any streaming configuration whose mempool block granularity is smaller than a full FIFO batch (204 bytes) — for example the sensor shell's RTIO_DEFINE_WITH_MEMPOOL(sensor_read_rtio, 8, 8, 32, 64, 4) used by sensor stream <dev> on fifo_wm incl, or the 16-byte-block streaming example in the documentation — takes this path whenever the pool is fragmented or the consuming thread falls behind. The out-of-bounds writes are performed by the I2C/SPI RTIO backend (DMA on many SoCs) into the shared RTIO block pool, corrupting neighbouring live buffers or, for a pool smaller than one batch, memory beyond the pool itself. The written bytes are accelerometer samples and the destination follows the allocation, so there is no controllable corruption primitive: the realistic impact is silent corruption of other RTIO consumers' data or a crash. Reaching it requires a board with an interrupt-wired ADXL345 built with CONFIG_ADXL345_STREAM, plus buffer-pool starvation, which a co-resident local workload can induce but which also occurs spontaneously under load; there is no network, USB or radio path to this code.	3.6	More Details
CVE-2026-116012	Ghidra before 12.2 contains a numeric truncation vulnerability in PackedDecode::skipAttribute and skipAttributeRemaining in marshal.cc that stores 64-bit string lengths in 32-bit variables. Attackers can supply a crafted packed stream, such as a malicious .sla file, to inject attributes into the decompiler's input or crash the decompiler process.	3.6	More Details
CVE-2026-116148	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains an injection vulnerability in the MVG encoder, which does not properly neutralize user-specified drawing options before writing them into generated MVG output. Attackers who can supply drawing options can inject arbitrary MVG drawing commands into the resulting MVG file.	3.6	More Details
CVE-2026-114048	WWBN AVideo before 24.0 contains a cross-site request forgery vulnerability in legacy endpoints like playlistStatus.php that bypass the .json.php-only autoCSRFGuard() check. Attackers can lure logged-in users to a page with an auto-submitting form to make private playlists public, remove playlist videos, or forge video suggestions.	3.5	More Details
CVE-2026-96392	Improper Neutralization of Special Elements Used in a Template Engine vulnerability in Drupal AI CKEditor allows Code Injection. This issue affects AI CKEditor versions: from 0.0.0 to 1.4.3.	3.5	More Details

CVE-2026-115449	A flaw was found in libslirp. When handling ICMP echo requests, the reply length is not updated to match the actual data received from the host. Under specific conditions where the normal ICMP socket is unavailable, this length mismatch causes the host's real IP address to be leaked to the guest in the echo reply, breaking the NAT isolation that user-mode networking is expected to provide.	3.5	More Details
CVE-2026-115507	Vikunja 2.2.0 before 2.7.1 contains an HTML injection vulnerability in CreateConversationalHeader in pkg/notifications/mail.go that interpolates display names and task and project titles raw into email anchors. Authenticated users can craft titles or display names that close the anchor and inject phishing links into task-comment and task-mention notification emails sent to project members.	3.5	More Details
CVE-2026-113706	A vulnerability was identified in a2aproject a2a-rs up to 0.4.3. Impacted is an unknown function of the file a2a-server/src/task_store/inmemory.rs. The manipulation of the argument historyLength leads to information disclosure. The attack is possible to be carried out remotely. The identifier of the patch is 68f1b7e772ae5d6c9e581ac9b48bcddd10fe70eb. To fix this issue, it is recommended to deploy a patch.	3.5	More Details
CVE-2026-87948	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal SAML SSO - Service Provider allows Cross-Site Scripting (XSS). This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	3.5	More Details
CVE-2026-116261	A security vulnerability has been detected in libimobiledevice up to 1.4.0. Impacted is the function afc_file_open of the file src/afc.c of the component AFC File Handler. The manipulation of the argument bytes leads to out-of-bounds read. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.5	More Details
CVE-2026-116229	A vulnerability was determined in ffxiv-teamcraft FFXIV Teamcraft up to 11.4.31. Impacted is the function clipboard.readText of the file apps/electron/src/preload.ts of the component Renderer-Process IPC Bridge. This manipulation causes cross site scripting. The attack is possible to be carried out remotely. The exploit has been publicly disclosed and may be utilized. Upgrading to version 11.4.32 is recommended to address this issue. Patch name: def118fdf36f8f3b615966893a132dd81d9cc3dc. Upgrading the affected component is recommended.	3.5	More Details
CVE-2026-93847	Mattermost versions 11.7.x <= 11.7.10 fail to validate reserved system post types on the scheduled post creation and publication paths which allows an authenticated channel member to publish messages rendering as trusted system events via the scheduled post endpoint.. Mattermost Advisory ID: MMSA-2026-00732	3.5	More Details
CVE-2026-114040	A security vulnerability has been detected in DT42 BerryNet up to 3.10.2. The affected element is the function DLModelManager.get_model_meta of the file berrynet/dlmodelmgr.py of the component Model Package Metadata Handler. The manipulation of the argument label leads to path traversal. The attack is possible to be carried out remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15705	The Contact Form by Supsystic WordPress plugin before 1.8.0 does not sanitise and escape a form name, allowing users with access to the form editor to inject arbitrary web scripts that execute when another user views the form. On multisite installations, where administrators do not hold the unfiltered_html capability, this crosses a privilege boundary.	3.5	More Details
CVE-2026-114070	Trigger.dev 4.5.0 before 4.6.0 contains an information disclosure vulnerability in the Dashboard Agent input proxy that stores a delegated user-actor token in browser-readable session input records. Attackers holding a chat's public session token can read the in records endpoint to recover the token from payload.metadata.userActorToken and replay it for environment JWT exchange.	3.5	More Details
CVE-2026-18748	The Bluetooth Classic A2DP host stores a codec Information Element (IE) in struct bt_a2dp_codec_ie, a one-byte len field plus a fixed 8-byte codec_ie[] array. When an AVDTP peer sent a codec IE longer than 8 bytes, a2dp_process_config_ind() in subsys/bluetooth/host/classic/a2dp.c copied only the first 8 bytes but assigned the full, attacker-chosen length to len, and then persisted that inconsistent structure into the application's struct bt_a2dp_stream via stream->codec_config. bt_a2dp_get_config_cb() and bt_a2dp_set_config_cb() had the same truncate-the-copy-but-keep-the-length pattern. Every later consumer trusts len. In particular the responder for an incoming AVDTP GET_CONFIGURATION, a2dp_get_config_ind(), executes net_buf_add_mem(rsp_buf, &ep->stream->codec_config.codec_ie[0], ep->stream->codec_config.len), so a paired BR/EDR peer that first sends a SET_CONFIGURATION carrying an oversized codec IE and then requests the configuration back causes the stack to read up to roughly 245 bytes past the end of the 8-byte array and transmit them in the response. codec_config is the last member of struct bt_a2dp_stream, so the read runs into memory adjacent to the application's stream object. The same inconsistency is also handed to the config_req(), reconfig_req() and get_config_rsp() application callbacks, where a natural memcpy() of len bytes performs a further out-of-bounds read. Reaching the on-air disclosure requires CONFIG_BT_CLASSIC with CONFIG_BT_A2DP, an	3.5	More Details

	AVDTP peer that has completed pairing (the AVDTP L2CAP channel requires BT_SECURITY_L2), and a local endpoint registered for a non-SBC codec — BT_A2DP_MPEG1/BT_A2DP_MPEG2/BT_A2DP_ATRAC/BT_A2DP_VENDOR via BT_A2DP_EP_INIT() — because the SBC branch already pins the IE length to 4 bytes. The bt_a2dp_get_config_cb() response path has no such gate and hands the lying length to any application. There is no out-of-bounds write: the net_buf tailroom check bounds the destination, so the impact is limited to disclosure of a bounded region of adjacent RAM to the peer. The fix replaces the hardcoded BT_A2DP_MAX_IE_LENGTH with the configurable CONFIG_BT_A2DP_CODEC_MAX_IE_LEN (default 8, range 4–253), sizes codec_ie[] from it, and rejects any received IE that exceeds it instead of silently truncating, so len can no longer exceed the number of valid bytes.		
CVE-2026-96384	Incorrect Authorization vulnerability in Drupal Mermaid Diagram Field allows Forceful Browsing. This issue affects Mermaid Diagram Field versions: from 1.0.4 to 1.0.9.	3.5	More Details
CVE-2026-111737	The Speed Optimizer WordPress plugin before 7.8.5 does not sanitise or escape some of its settings before outputting them in the head of front-end pages, which could allow high-privilege users such as admins to perform Stored Cross-Site Scripting attacks even when the unfiltered_html capability is disallowed (for example in a multisite setup).	3.5	More Details
CVE-2026-116004	A security vulnerability has been detected in onyx-dot-app Onyx up to 5.0.0. Impacted is an unknown function of the file web/src/app/craft/components/output-panel/PreviewTab.tsx. The manipulation leads to cross site scripting. The attack can be initiated remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.5	More Details
CVE-2026-110510	The Essential Addons for Elementor WordPress plugin before 6.8.7 does not validate the HTML tag setting of some of its widgets before using it to build markup, which could allow users with the author role and above to inject iframe and form elements into published pages.	3.5	More Details
CVE-2026-105711	GitLab has remediated an issue in GitLab CE/EE affecting all versions from 15.0 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, when a merge request in a public project referenced a confidential issue from a private project, an authenticated user with no project membership could have read the confidential issue's title, due to a viewer-independent fragment cache key that incorrectly shared per-user-redacted rendered content across all users.	3.5	More Details
CVE-2026-115246	A weakness has been identified in subhajitkhan online-clinic-management-system up to e9ee77a8827a1446220fa07ee693dc4d9a29a578. This affects an unknown part of the file patdocmsg.php. Executing a manipulation of the argument Query can lead to cross site scripting. The attack may be performed from remote. The exploit has been made available to the public and could be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2026-96360	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal Webform allows Cross-Site Scripting (XSS). This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	3.5	More Details
CVE-2026-87953	Server-Side Request Forgery (SSRF) vulnerability in Drupal SAML SSO - Service Provider allows Server Side Request Forgery. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	3.3	More Details
CVE-2026-116143	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a policy bypass in the PS, PDF and XPS decoders that allows the security policy's list length limit on scene selection to be circumvented. Attackers supplying a specific argument can cause ImageMagick to process more scenes than permitted, potentially exhausting resources.	3.3	More Details
CVE-2026-116142	ImageMagick before 7.1.2-33 contains a heap buffer over-read vulnerability that allows attackers to read beyond allocated memory when printing a minimum-bounding-box. Attackers can supply a crafted image that triggers the over-read during minimum-bounding-box output, causing crashes or exposing adjacent heap memory.	3.3	More Details
CVE-2026-116281	A vulnerability was identified in Kameleo up to 5.0.1.26. Affected by this vulnerability is the function ListProfiles of the file client-engine-application/t.cs of the component Profile Controller. The manipulation leads to information disclosure. The attack can only be performed from a local environment. Upgrading to version 5.2.0 addresses this issue. Upgrading the affected component is recommended.	3.3	More Details
CVE-2026-95721	A memory corruption vulnerability exists in the web-based management interface of AOS-8. Successful exploitation could allow an authenticated low privilege attacker with local access to view some system files and interrupt the normal operation of the affected service.	3.3	More Details

CVE-2026-116156	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak caused by a missing cleanup when the -set profile option is used. An attacker who induces a user to process images with the -set profile option can leak memory, potentially exhausting resources and causing a denial of service.	3.3	More Details
CVE-2026-115363	OpenBao 2.2.0 before 2.6.4 and 2.7.0 before 2.7.1 contains an information disclosure vulnerability that exposes the tls_acme_eab_mac_key value through the sys/config/state/sanitized endpoint. Callers of the sanitized configuration endpoint can retrieve the ACME External Account Binding MAC key used for TLS listener certificates, letting them bind to the operator's ACME account.	3.3	More Details
CVE-2026-87949	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal SAML SSO - Service Provider allows Cross-Site Scripting (XSS). This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	3.3	More Details
CVE-2026-116160	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a heap-based buffer overflow in the X11 animate operation when processing crafted images. Attackers can supply a malicious image that, when a user runs an X11 animate operation on it, overwrites heap memory and can crash the application.	3.3	More Details
CVE-2026-116157	ImageMagick before 7.1.2-33 contains a memory leak vulnerability in PerlMagick that fails to free values returned when reading image properties. Attackers can repeatedly trigger property reads through five affected PerlMagick code locations to exhaust memory and cause denial of service.	3.3	More Details
CVE-2026-114236	storage/pkg/archive.ApplyLayer and ApplyUncompressedLayer may, when processing corrupt or malicious inputs, leak a file descriptor; unbounded repeated calls can exhaust the process' file descriptor limit	3.3	More Details
CVE-2026-115749	hydra-ray-launcher 1.1.0.dev1 before 1.2.2 and 1.4.0.dev4 before 1.4.0rc2 contains an arbitrary module import vulnerability via the ray_pkg_version OmegaConf resolver passing input to importlib.import_module(). Attackers controlling configuration values or command-line overrides can supply a ray_pkg_version interpolation to import any installed module, triggering its import-time side effects in the application process.	3.3	More Details
CVE-2026-116159	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains a memory leak in the handling of the -cdl (color decision list) option, where allocated memory is not cleaned up after use. Attackers can supply input that causes the -cdl option to be processed, repeatedly leaking memory and degrading availability through resource exhaustion.	3.3	More Details
CVE-2026-96373	Incorrect Authorization vulnerability in Drupal Webform allows Forceful Browsing. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	3.3	More Details
CVE-2026-116140	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 incorrectly parses delegate invocations, which can allow an executable blocked in policy.xml to still run. An attacker who gets ImageMagick to process input that triggers a delegate invocation can bypass the administrator-defined policy and execute the blocked program.	3.3	More Details
CVE-2026-115569	Fallow is a codebase intelligence tool for TypeScript and JavaScript. From 2.68.0 until 3.34.0, Fallow follows attacker-controlled Jest projects paths outside the project root through absolute paths, parent segments, dot-prefixed wildcard globs, or symlinks. In version 3.33.1, the same behavior also affects Vitest test.projects paths, Vitest inline-project extends paths, and symlinked script --config paths. Fallow parses the outside files as test-runner configuration and uses them to credit dependencies, setup files, and JSX import sources, which can change unused dependency and unused file findings in CI or agent analysis. The outside file contents are not printed or sent over the network. This issue is fixed in version 3.34.0.	3.3	More Details
CVE-2026-116228	A vulnerability was found in 777genius agent-teams-ai up to 2.17.10. This issue affects the function doStart of the file src/main/services/infrastructure/HttpServer.ts of the component Local HTTP Service. The manipulation results in missing authentication. The attack is only possible with local access. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.3	More Details
CVE-2026-116144	ImageMagick before 6.9.13-58 and 7.x before 7.1.2-33 contains uncontrolled recursion in the MVG decoder that can lead to a stack overflow. Attackers can supply a crafted MVG file that, when read with specific arguments, exhausts the stack and crashes the process, causing a denial of service.	3.3	More Details
CVE-2026-112326	GitLab has remediated an issue in GitLab EE affecting all versions from 17.8 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, under certain conditions, could have allowed an authenticated user with read-only permissions to permanently delete existing vulnerability records due to improper cleanup logic that failed to distinguish between pre-existing records from newly created ones during error handling.	3.1	More Details

CVE-2026-113875	On a SmileCDR HAPI FHIR server with REST-hook subscriptions enabled, if a user is granted subscription-creation permissions, they may be able to query and read resources that they are not otherwise permitted to access.	3.1	More Details
CVE-2026-84911	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal AI (Artificial Intelligence) allows Stored XSS. This issue affects AI (Artificial Intelligence) versions: from 0.0.0 to 1.3.13, from 1.4.0 to 1.4.8.	3.1	More Details
CVE-2026-116150	ImageMagick before 7.1.2-33 contains a memory leak vulnerability in the FPX encoder that allows attackers to exhaust memory by supplying a crafted image. Attackers can submit crafted images for FPX encoding to leak memory repeatedly, causing resource exhaustion and denial of service.	3.1	More Details
CVE-2026-54238	Supabase Realtime provides Broadcast, Presence, and Postgres Changes via WebSockets. Prior to version 2.103.1, the REST broadcast endpoints accept an expired JWT when the exp claim is encoded as a JSON string instead of a JSON number. The same route rejects an expired integer exp token signed with the same tenant secret. The accepted request returns 202 and emits the broadcast to subscribers. This could allow an unauthenticated attacker can broadcast realtime messages to subscribers. The attack can be performed via the Internet/only from within the supabase project. The impact of broadcasting messages to subscribers is application specific. This issue has been patched in version 2.103.1.	3.1	More Details
CVE-2026-114721	Roundcube Webmail before 1.6.20 and 1.7.x before 1.7.5 allows the disabled_actions configuration setting to be bypassed. The check compared only the action name as requested, so a remote authenticated attacker could reach an action listed in disabled_actions by requesting it through its deprecated alias instead. Upstream fixed this in the commit "Fix bug where deprecated action aliases bypassed `disabled_actions` setting", which moves the check into rcmail::action_handler() and resolves aliases before comparing.	3.1	More Details
CVE-2026-96369	Use of Web Browser Cache Containing Sensitive Information vulnerability in Drupal Webform allows Exploiting Incorrectly Configured Access Control Security Levels. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	3.1	More Details
CVE-2026-115730	i18next-http-backend 4.0.2 contains an incomplete fix in interpolateUrl() that allows attackers to redirect translation requests by supplying an ns value starting with a slash. Attackers controlling the namespace value can form a protocol-relative URL in loadPath or addPath templates to load attacker-hosted translations, enabling UI content injection and cross-site scripting.	3.1	More Details
CVE-2026-116115	A security vulnerability has been detected in Yoosee GM8135S Camera 21.00.01.36. This impacts an unknown function of the file wpa_supplicantX.conf of the component WiFi Configuration. Such manipulation leads to improper initialization. The attack can only be initiated within the local network. This attack is characterized by high complexity. The exploitability is said to be difficult. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2026-114743	The InstaWP Connect WordPress plugin before 0.1.4.2 does not perform a capability check on its InstaWP Connect WordPress plugin before 0.1.4.2-settings update action, verifying only a nonce, allowing an authenticated user holding a low-privileged role that an administrator has granted access to the InstaWP Connect WordPress plugin before 0.1.4.2's sync feature to overwrite the InstaWP Connect WordPress plugin before 0.1.4.2's settings, including the setting controlling which roles may access that feature. The role grant is not enabled by default.	3.1	More Details
CVE-2026-87952	Authentication Bypass by Capture-replay vulnerability in Drupal SAML SSO - Service Provider allows Remote Services with Stolen Credentials. This issue affects SAML SSO - Service Provider versions: from 0.0.0 to 3.2.0.	3.1	More Details
CVE-2026-114071	Trigger.dev before 4.5.12 contains an authorization bypass vulnerability that allows authenticated project members to complete another member's DEVELOPMENT environment waitpoint by skipping the per-user environment ownership check. Attackers who know a victim's waitpoint friendly ID can submit type=DATETIME or type=MANUAL with isTimeout to prematurely mark the pending waitpoint completed.	3.1	More Details
CVE-2026-83813	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, twisted.names.dns.DNSDatagramProtocol accepts a UDP DNS response for a pending query when its 16-bit transaction ID matches but does not verify that the response came from the queried host and port. An attacker able to deliver a packet to the resolver's listening port with the correct transaction ID can inject a forged answer. The standard twisted.names.client.Resolver uses a new random source port for each query, increasing off-path attack complexity, while applications that reuse a DNSDatagramProtocol instance expose only the transaction-ID search space after the first query. A	3.1	More Details

	successful forgery can redirect hostname resolution and enable mail interception, server-side request forgery, or man-in-the-middle positioning. This issue is fixed in version 26.10.0rc1.		
CVE-2026-14288	GitLab has remediated an issue in GitLab EE affecting all versions from 17.4 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that under certain conditions could have allowed an authenticated to bypass a required merge request approval policy on a protected branch due to improper access control during approval rule synchronization.	3.1	More Details
CVE-2026-113797	Electron is a framework for writing cross-platform desktop applications using JavaScript, HTML and CSS. Prior to versions 42.11.1, 43.5.1, 44.1.1, and 45.0.0-alpha.4, chrome.tabs.query() returned each tab's url and title to extensions that had neither the tabs permission nor matching host permissions. An application that loaded an extension it did not fully trust through session.extensions.loadExtension could therefore allow that extension to read the URL and title of every webContents in its session. This issue is fixed in versions 42.11.1, 43.5.1, 44.1.1, and 45.0.0-alpha.4.	3.1	More Details
CVE-2026-113501	The LiteSpeed Cache WordPress plugin before 7.9.2 does not perform a capability check on its admin-notice dismissal handler, and the nonce it relies on is not bound to the specific notice being dismissed, so on sites in certain configurations a low-privileged authenticated user (Subscriber) can dismiss administrative notices the LiteSpeed Cache WordPress plugin before 7.9.2 marks as non-dismissible, including a warning that sensitive files may be publicly accessible.	3.1	More Details
CVE-2026-116215	MacCMS v10 through 2026.1000.4056 contains a race condition vulnerability in Card::useData() that allows authenticated members to redeem one recharge card multiple times via the user/buy endpoint. Attackers can send many concurrent flag=card requests that pass the non-locking card read before it is marked used, inflating user_points to buy VIP upgrades and paid content.	3.1	More Details
CVE-2026-115724	VLC for Android before 3.7.2 contains a confused-deputy file overwrite vulnerability in FileUtils.getUri() that trusts content URIs based only on the com.fsck.k9.attachmentprovider or gmail-ls authority. A co-installed app without permissions can declare these authorities and send a crafted VIEW intent to the exported StartActivity, making VLC write attacker-chosen content under an attacker-chosen name in the shared Download directory, overwriting other apps' files when VLC holds All files access.	2.9	More Details
CVE-2026-88849	The MasterStudy LMS WordPress Plugin WordPress plugin before 3.7.55 does not verify that a user targeted through its student-enrolment feature has any relationship to the requesting instructor, allowing users with the Instructor role to overwrite the first name, last name and public display name of any registered user on the site, including administrators, enabling impersonation and phishing.	2.7	More Details
CVE-2026-111981	The GiveWP WordPress plugin before 4.19.0.1 does not perform a capability check when handling donor address changes, allowing users who have been granted donor-report (view) access but not donor-edit access to add, modify, or delete the stored address of any donor on the site.	2.7	More Details
CVE-2026-114742	The WPVR WordPress plugin before 9.0.3 does not perform object-level ownership checks when updating virtual tours, allowing users with contributor-level access and above to modify, reconfigure and publish tours created by other users.	2.7	More Details
CVE-2026-103515	GitLab has remediated an issue in GitLab EE affecting all versions from 19.2 before 19.3.6, 19.4 before 19.4.3, and 19.5 before 19.5.1 that, when an instance administrator locked a group's Duo availability, an authenticated user with group Owner permissions could have re-enabled AI-backed data egress to a third-party service for that group and its descendant projects, due to insufficient granularity of the admin-lock enforcement controls.	2.7	More Details
CVE-2026-105257	The Appointment Booking Plugin WordPress plugin before 5.7.5 does not verify that an authenticated agent-level user is assigned to a customer before returning or searching that customer's record, allowing any agent to read and enumerate the personal information of every customer on the site, including those managed by other agents.	2.7	More Details
CVE-2026-110011	The MB Custom Post Types & Custom Taxonomies WordPress plugin before 2.12.3 does not perform a capability check in its export feature, allowing users with a role as low as Contributor to read the configuration of admin-only custom post types and taxonomies, as well as the titles, dates, statuses and some stored content of other users' non-public posts and pages.	2.7	More Details
CVE-2026-93848	Mattermost versions 11.7.x <= 11.7.10 fail to apply outbound internal-connection filtering to the OpenID discovery endpoint request which allows a System Administrator to determine whether internal addresses and ports are reachable from the server via an unauthenticated login request that triggers the discovery fetch.. Mattermost Advisory ID: MMSA-2026-00774	2.7	More Details
CVE-2026-104642	The Simple Shopping Cart WordPress plugin before 5.3.3 does not verify submitted product prices against a value only the site owner can set, allowing users with the Contributor role and above to register attacker-chosen prices for existing products, which any visitor can then use to add those products to the cart at that price.	2.7	More Details

CVE-2026-115366	OpenBao from 2.3.0 before 2.6.3 contains an improper authorization vulnerability that allows privileged child-namespace tokens to modify plugin registrations because <code>/v1/sys/plugins/catalog/*</code> is not restricted to the root namespace. Attackers can remove existing registrations, create invalid ones, or register installed plugins through a child namespace's <code>sys/</code> mount, causing denial of service.	2.6	More Details
CVE-2026-116197	Edgeless Systems Contrast from 1.21.0 before 1.25.0 contains a non-deterministic registry selection flaw in imagepuller that can prefer wildcard configuration entries over concrete domain matches. When a pull configuration lists both a concrete domain and its parent wildcard, registry credentials, CA bundles, or mirror settings intended for the wildcard may be applied to the concrete registry, leaking credentials or pull metadata.	2.6	More Details
CVE-2026-114614	Authlib is a Python library which builds OAuth and OpenID Connect servers. From 0.13 until 1.8.0, <code>authlib.common.security.is_secure_transport()</code> uses a textual prefix comparison for loopback HTTP exemptions, so a URL whose userinfo begins with a loopback host can be accepted even when its parsed host is an external non-loopback system. An external party able to supply a URL to OIDC RP-Initiated Logout registration, authorization-server metadata validation, or another call site can bypass the TLS-required guard and cause cleartext redirection or acceptance of cleartext endpoints. The issue does not bypass redirect URI exact matching and does not by itself forge tokens or authenticate users. The vulnerable <code>authlib.common.security.is_secure_transport()</code> implementation in <code>authlib/common/security.py</code> misinterprets the userinfo component before OIDC RP-Initiated Logout dynamic client registration and other consumers enforce transport security. This issue is fixed in version 1.8.0.	2.6	More Details
CVE-2026-115386	Wasmtime before 36.0.17, 37.0.0 before 48.0.4, and 49.0.0 before 49.0.2 contains an information disclosure vulnerability in the WASIp1 <code>fd_readdir</code> implementation that copies uninitialized struct padding into guest memory. Guests with a preopened directory can list entries to receive three bytes of leftover host stack or heap memory per directory entry.	2.6	More Details
CVE-2026-116153	ImageMagick before 7.1.2-33 contains a heap buffer overflow vulnerability on 32-bit systems that allows attackers to overwrite heap memory when convex-hull information is requested. Attackers can supply crafted images processed with convex-hull calculations to corrupt heap memory, causing crashes or potential code execution.	2.5	More Details
CVE-2026-115319	fzf 0.36.0 through 0.74.5 contains an inefficient algorithmic complexity vulnerability in <code>maskActionContents</code> that allows local attackers to exhaust the <code>--listen</code> HTTP server through quadratic string concatenation. Attackers can repeatedly send POST bodies of about 1 MB to the unauthenticated localhost port, blocking the single-threaded server for about 30 seconds each time.	2.5	More Details
CVE-2026-116158	ImageMagick before 7.1.2-33 contains a memory leak vulnerability in the <code>magick compare</code> command caused by missing cleanup when an operation fails. Attackers can supply inputs that trigger repeated failed compare operations to exhaust memory and cause denial of service.	2.5	More Details
CVE-2026-115981	A security vulnerability has been detected in phpIPAM up to 1.8.3. Affected by this vulnerability is an unknown functionality of the file <code>upgrade/index.php</code> . Such manipulation of the argument <code>REQUEST_URI</code> leads to cross site scripting. The attack may be launched remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure.	2.4	More Details
CVE-2026-115535	A vulnerability was identified in gedelumbung HospitalManagement up to <code>c2d45543789a3887067d3915f69d44cfc2cf76a8</code> . The affected element is the function <code>sistem.php::simpan</code> of the file <code>application/modules/admin/controllers/sistem.php</code> of the component Settings Handler. The manipulation of the argument <code>Title</code> leads to cross site scripting. The attack is possible to be carried out remotely. The exploit is publicly available and might be used. This product adopts a rolling release strategy to maintain continuous delivery. Therefore, version details for affected or updated releases cannot be specified. The project was informed of the problem early through an issue report but has not responded yet.	2.4	More Details
CVE-2026-115183	The Qi Addons For Elementor WordPress plugin before 1.11.2 does not properly sanitise a widget's heading text after decoding HTML entities in it, allowing users with the Contributor role and above to embed live HTML forms, including password fields that submit to an external host, into posts where WordPress would otherwise strip them.	2.4	More Details
CVE-2026-116227	A vulnerability has been found in Edimax BR-6208AC 1.32. This vulnerability affects unknown code of the component Virtual Server Page. The manipulation leads to cross site scripting. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	2.4	More Details
CVE-2026-	A vulnerability has been found in gougucms up to 4.08.18. Impacted is the function <code>Add</code> of the file <code>app/admin/controller/Level.php</code> of the component User Level Handler. The manipulation of the argument <code>Title</code> leads to cross site scripting. The attack is possible to be carried out remotely. The	2.4	More

116100	exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.		Details
CVE-2026-115389	Wasmtime before 36.0.16, 37.0.0 before 48.0.3, and 49.0.0 before 49.0.1 does not adequately account for host allocations when lifting guest values through the dynamically-typed Val API, bypassing the hostcall fuel limit. A malicious WebAssembly guest can return crafted values that allocate at least ~100x more host memory than tracked, causing denial of service.	2.3	More Details
CVE-2026-115369	OpenBao from 2.5.0 before 2.6.2 fails to enforce the PKI role allowed_ip_sans_cidr restriction on IP SANS taken from a CSR when use_csr_sans is enabled. Authenticated callers with sign permission can submit a CSR with an out-of-range IP SAN to pki/sign/<role> and obtain CA-signed certificates for unauthorized IPs.	2.2	More Details
CVE-2026-116054	A security vulnerability has been detected in SH3. This affects an unknown function of the component Módulo de Gestão/Módulo de Arrecadação. The manipulation leads to password hash with insufficient computational effort. The attack is possible to be carried out remotely. A high degree of complexity is needed for the attack. The exploitability is reported as difficult. The exploit has been disclosed publicly and may be used. Upgrading to version 2026.0.2.0 mitigates this issue. Upgrading the affected component is recommended.	2.2	More Details
CVE-2026-115364	OpenBao before 2.6.3 contains a missing authorization check that allows attackers to stop Agent and Proxy services by omitting the X-Vault-Header. When require_request_header is enabled, attackers can send requests via server-side request forgery to /agent/v1/quit or /proxy/v1/quit to shut down these services.	2.2	More Details
CVE-2026-115357	OpenBao before 2.6.4 and 2.7.x before 2.7.1 fails to write response audit entries for the unauthenticated sys/generate-root, sys/rekey, and sys/rekey-recovery-key endpoints when request audit logging succeeds but response audit logging fails. An attacker using these endpoints can complete root token generation or rekey operations without a corresponding response audit record, hindering forensic detection.	2.2	More Details
CVE-2026-115362	OpenBao before 2.6.4 and 2.7.x before 2.7.1 does not check Secret ID expiration at use in the AppRole auth method, leaving expired Secret IDs valid until the periodic tidy removes them. An attacker holding an expired, untidied Secret ID can log in and gain access beyond its TTL.	2.2	More Details
CVE-2026-21812	HCL AppScan Azure Boards Plugin is affected by a Cleartext Transmission of Sensitive Information vulnerability when configured to communicate with the AppScan server over HTTP. This could allow an attacker to intercept credentials or authentication tokens during transmission. The plugin has been updated to display a warning when HTTP is used and recommend HTTPS for secure communication.	2.2	More Details
CVE-2026-115358	OpenBao before 2.6.3 contains an information disclosure vulnerability that allows attackers to leak sensitive data into audit logs by sending malformed plugin parameters. Attackers can submit requests with malformed TypeKVPair or TypeHeader field values, causing error messages containing sensitive input to be recorded in audit logs.	2.0	More Details
CVE-2026-78066	Resource exhaustion in Apache bRPC (version < 1.18.0) on all platforms: multiple parsers that accept wire-controlled input bound no limit on the sizes declared by the client or on decompressed output, which attackers can use to drive disproportionate memory use and crash the process. Root Cause: 1) The gzip/zlib/snappy decompressors enforce no output limit: -max_body_size is checked against the compressed bytes only, so a cap-compliant body (e.g. 64 MiB) can decompress to tens of GiB. 2) The AMF string readers allocate the declared length before the bytes are actually available. 3) RTMP chunk headers may re-declare a message length mid-stream; the declared length was never bounded and the partial reassembly buffer kept growing. 4) The sofa-pbrpc parser does not bound meta_size. 5) mcpack2pb trusts wire-declared item counts for Reserve() and a string value_size of 0 underflows the buffer size, aborting the process. Affected Scenarios: Any bRPC service accepting gzip/zlib/snappy-compressed bodies or rtmp, sofa-pbrpc, or mcpack2pb messages from untrusted clients. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3476 Note: After the fix, requests whose decompressed size, RTMP message length, or mcpack declared count exceed the new limits are rejected. Operators who legitimately need larger values can raise the corresponding flag (-max_decompressed_body_size, max_body_size).	N/A	More Details
	An out-of-bounds read vulnerability exists in the streaming RPC client of Apache bRPC (versions 1.11.0 through 1.17.0) on all platforms. A client that creates multiple streams in a single RPC expects the server response to carry exactly one extra_stream_id for every additional stream. The response handling path indexed the returned identifiers using the locally created stream count without first validating the size of the returned list, so a malicious or compromised server (or an on-path attacker) can trigger an out-of-bounds read in the client by returning a malformed StreamSettings response with		

CVE-2026-79704	fewer extra_stream_ids than expected. This may crash the client process (denial of service) or leak unintended memory contents into subsequent stream handling; surplus identifiers have no valid client-stream mapping and are processed inconsistently. Root Cause: Since bRPC 1.11.0, a client can create multiple streams in one RPC via batch StreamCreate. The expected behavior is that the server response must return exactly (N-1) extra_stream_ids, where N is the number of locally created streams. In Controller::HandleStreamConnection, the response path reads extra_stream_ids[i-1] for each additional local stream without first comparing the returned list size with (N-1). If the server returns fewer identifiers, the read runs past the end of the returned list (out-of-bounds read); if it returns more, the surplus identifiers have no corresponding client-side stream. Affected Scenarios: A bRPC client using multi-stream RPC (creating more than one stream in a single RPC) that connects to an untrusted, malicious, or compromised server or to a server reached through an untrusted network path. The server sends a malformed StreamSettings response to trigger the issue. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3485 Note: The fix makes the client reject responses whose extra_stream_ids count does not exactly match the number of additional client streams. Such mismatched responses now fail the RPC with ERESPONSE and go through the existing stream cleanup path, instead of being partially processed.	N/A	More Details
CVE-2026-78334	Missing authentication in the streaming RPC path of Apache bRPC (version < 1.18.0) on all platforms allows a remote attacker to bypass server authentication and inject or deliver streaming frames, when a server is configured with an Authenticator. Root Cause: When authentication is configured, streaming RPC frames are consumed directly by the streaming parser (ParseStreamingMessage) and never pass through InputMessenger's authentication hook, so an unauthenticated peer can exchange frames without being authenticated. In addition, a server-accepted stream is bound lazily to the first frame that happens to carry its id, so a stream could accept a frame carrying its id from a socket other than the socket that created the stream. Affected Scenarios: A bRPC server configured with authentication that serves streaming RPC from an untrusted network, or a deployment where a foreign connection can impersonate a stream by id, can be made to accept unauthorized streaming frames. How to Fix: Upgrade bRPC to 1.18.0, which fixes this issue, or apply the patch in PR https://github.com/apache/bRPC/pull/3481 . Note: The fix rejects streaming frames on server sockets when configured authentication has not completed, and keeps server-accepted streams bound to the socket of their creating RPC and rejects frames arriving from another socket. Authenticated streaming RPC behavior and the wire format are unchanged; only unauthenticated or foreign-socket frames are now rejected.	N/A	More Details
CVE-2026-76529	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76529 are related to improper protection of alternate path and are grouped under the Common Weakness Enumeration (CWE) CWE-424.	N/A	More Details
CVE-2026-80463	Unbounded memory accumulation in the stream message reassembly path of Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to exhaust the memory of a server (denial of service) by sending an endless sequence of continuation data frames. Root Cause: The streaming RPC protocol bounds every single data frame by max_body_size, but a fragmented stream message accumulates continuation frames into one buffer (_pending_buf) without any limit on the accumulated size. A malicious peer that can open a stream can therefore keep sending continuation frames to make the receiver accumulate an arbitrarily large in-memory message. Since the default flow-control settings (socket_max_streams_unconsumed_bytes=0) do not throttle the sending rate, the receiver's memory grows until the process runs out of memory and crashes or becomes unavailable. Affected Scenarios: Services that accept streams from untrusted clients (for example using brpc::StreamAccept or brpc::StreamCreate with the built-in streaming protocol) are affected. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3489 Note: The fix limits the accumulated payload of a fragmented stream message to max_body_size (default 64MB). When the accumulated payload exceeds the limit, the pending fragments are released and the stream is closed with error EMSGSIZE. After applying the fix, legitimate fragmented stream messages larger than max_body_size will be rejected; you can adjust the gflag max_body_size to change the limit.	N/A	More Details
CVE-2026-89052	A hardcoded cryptographic key management vulnerability exists in Archer C6 due to the use of a universal cryptographic key shared across devices. An attacker who obtains an encrypted configuration file or other protected credential data can recover plaintext credentials. Successful exploitation may allow unauthorized disclosure of sensitive information, including administrative credentials.	N/A	More Details
CVE-2026-	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76301	N/A	More Details

76301	are related to protection mechanism failure and are grouped under the Common Weakness Enumeration (CWE) CWE-693.		
CVE-2026-78058	Uncontrolled resource consumption in the HTTP/2 (h2) receive path of Apache bRPC (version < 1.18.0) on all platforms allows a remote unauthenticated attacker to crash the server or grow its memory without bound by sending crafted HTTP/2 frames. Root Cause: The h2 parser, enabled by default on every bRPC server port and used by h2/gRPC clients, trusts the peer to respect receive-side limits that bRPC advertises but never enforces: 1. SETTINGS_MAX_CONCURRENT_STREAMS defaulted to unlimited and was never checked on the server, letting a peer pin an unbounded number of pending streams per connection (~100x memory amplification). 2. SETTINGS_MAX_HEADER_LIST_SIZE defaulted to unlimited and was never enforced; 1-byte HPACK indexed references amplify a small HEADERS frame ~4000x into an unbounded decoded header list ("HPACK bomb"), and a never-completed HEADERS/CONTINUATION block accumulates memory without limit. 3. Control-frame acks (PING/SETTINGS acks, RST_STREAM, GOAWAY, WINDOW_UPDATE) unconditionally bypassed the - socket_max_unwritten_bytes write cap, so a PING flood combined with withheld TCP reads grew the connection buffer without the documented bound. 4. The HPACK dynamic table size update was validated against the protocol default (4096) instead of the advertised SETTINGS_HEADER_TABLE_SIZE, which together with the dynamic-table queue/accounting issue can desynchronize the decoder and dereference NULL (crash). Affected scenarios: Any bRPC server serving HTTP/h2/gRPC to an untrusted network, and any bRPC h2/gRPC client connecting to an untrusted backend. The h2 parser runs before any Authenticator check on the shared port. How to fix: Upgrade bRPC, or apply the patch in https://github.com/apache/bRPC/pull/3474. Note: After applying the fix the following limits are enforced by default: - max_concurrent_streams: 1024 per connection (4294967295 restores unlimited). - max_header_list_size: 1 MiB decoded size per stream (4294967295 restores unlimited; - h2_client_max_header_list_size for clients). - h2 control-frame acks respect - socket_max_unwritten_bytes (-h2_ack_ignore_eovercrowded restores the old behavior). Requests exceeding these limits are rejected (RST_STREAM(REFUSED_STREAM) or a connection error).	N/A	More Details
CVE-2026-97200	Allocation of Resources Without Limits or Throttling vulnerability in hackney_conn in benoitc hackney allows a malicious HTTP server to exhaust client memory via an oversized or highly compressed response body. hackney:request/5 and hackney:body/1 return the whole response body, and nothing bounds how much the connection process buffers. Any server the client requests can trigger it. hackney_conn:read_full_body/2 and the HTTP/2 and HTTP/3 accumulators append every chunk until the body ends. With {auto_decompress, true} the buffered body is then inflated in one step with zlib:gunzip/1 or zlib:inflate/2, which return the whole output at once. This issue affects hackney: from 2.0.0-beta.1 before 4.9.0.	N/A	More Details
CVE-2026-103377	A reflected Cross-Site Scripting (XSS) vulnerability has been identified in the web application of the Castile-La Mancha Regional Parliament's video library. The application embeds user-supplied content into the page without properly sanitising it before it is rendered by the browser. Exploitation of this vulnerability could allow a remote attacker to execute arbitrary JavaScript code in the browser of a victim who interacts with the content crafted for the attack.	N/A	More Details
CVE-2026-78048	A reachable-assertion flaw in the HPACK header-compression decoder of the HTTP/2 (h2/gRPC) stack in Apache bRPC (version < 1.18.0) on all platforms allows a remote attacker to crash the server process by sending a crafted sequence of HPACK headers, resulting in a denial of service. Root Cause: In src/bRPC/details/hpack.cpp, IndexTable::Init sizes the dynamic-table ring buffer as num_headers = max_size / (32 + 2), i.e. it assumes a minimum table entry of 34 bytes (name >= 1 byte + value >= 1 byte + 32 bytes of overhead). However, RFC 7541 section 4.1 specifies that the minimum entry size is 33 bytes, because the value may be empty (name >= 1 byte + value == 0 + 32 bytes). For the default 4096-byte table this allocates only 120 slots (4096/34). Eviction in AddHeader is driven by byte usage (_max_size), not by slot count, so a peer can add 121 minimal 33-byte entries (121*33 = 3993 <= 4096) without triggering eviction. The 121st AddHeader then hits CHECK(!_header_queue.full()), which aborts the process. Affected Scenarios: Use a bRPC server that serves HTTP/2 or gRPC requests (e.g. via the h2 protocol) from an untrusted network. The malicious header block is processed during HPACK decoding of incoming requests without authentication. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3462 Note: The fix sizes the dynamic table for the 33-byte minimum entry (max_size / (32 + 1)) and ensures at least one ring-buffer slot is kept when max_size is smaller than the minimum entry, so that initialization cannot fail on malloc(0).	N/A	More Details
CVE-2026-74758	A heap use-after-free in the RTMP protocol component of Apache bRPC (versions < 1.18.0) allows a remote attacker to crash a bRPC RTMP server by sending a crafted Abort control message. No authentication or application-level request is required. Root Cause: When an RTMP Abort control message names the chunk stream that is currently being parsed (the chunk stream carrying the message itself), RtmpChunkStream::OnAbortMessage calls RtmpContext::ClearChunkStream, which deletes the RtmpChunkStream while its Feed() is still running. After the message handler returns, Feed() continues to access members of the deleted object, resulting in a heap-use-after-free. Affected	N/A	More Details

	Scenarios: bRPC servers that enable RTMP support (ServerOptions.rtmp_service set) and accept connections from untrusted networks. How to Fix: 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3452 The fix skips clearing the chunk stream when the Abort message targets the chunk stream that is currently being parsed.		
CVE-2026-111943	On most platforms, Go uses the given system root certificates. However, in some narrower cases, the system trust stores may not be updated without further action. To defend against the narrower set of cases that may remain affected, a denylist was added in crypto/x509 to return an error when calling Verify with any of the known intermediate certificates that are no longer considered to be trustworthy.	N/A	More Details
CVE-2026-74755	An integer underflow in the mcpack2pb component of Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to crash a bRPC server by sending a request containing a string field with a zero value_size. The UnparsedValue::as_string() function computes `_size - 1` without validating `_size`; when value_size is 0 the expression underflows to SIZE_MAX, resize() throws std::length_error, and because the exception is not caught on the request parsing path the process aborts (SIGABRT). Root Cause: In src/mcpack2pb/parser.cpp, UnparsedValue::as_string() resizes the output string to `_size - 1`, where `_size` is the value_size of a string field taken directly from the network data. A value_size of 0 underflows the unsigned size_t, and the subsequent resize() throws an exception that is not caught. Affected Scenarios: bRPC servers that accept nshead frames wired to the ubrpc mcpack2 adaptor (for example the NsheadService + UbrpcMcpack2Adaptor deployment pattern) can be crashed by a single unauthenticated request. Stock open-source bRPC registers the ubrpc/nshead_mcpack protocols as client-only, so default server configurations are not affected. How to Fix: 1. Upgrade bRPC to a version that includes this fix. 2. Apply the patch in PR https://github.com/apache/bRPC/pull/3450	N/A	More Details
CVE-2026-66392	Apache bRPC versions before 1.18.0 validate the server certificate chain when client TLS peer verification is enabled, but do not bind the certificate to an expected DNS name or IP address. A network attacker who can redirect or intercept a connection and who possesses a certificate issued by a CA trusted by the client for another name may impersonate the requested TLS endpoint. Affected Scenarios: Applications use a bRPC client Channel to connect to an untrusted network with certificate verification enabled and trust a CA that can issue certificates for identities other than the intended endpoint. How to Fix: Upgrade to Apache bRPC 1.18.0 and enable VERIFY_PEER with a positive verify_depth. HTTPS Channels automatically verify the certificate against the host in the Channel URL. Callers may set VerifyOptions::expected_peer_name explicitly when the certificate identity differs from the URL host. Alternatively, apply the patch from https://github.com/apache/bRPC/pull/3408 . Note: The new option performs DNS SAN or IP SAN verification during the TLS handshake. Configurations that set expected_peer_name while peer verification is disabled are rejected.	N/A	More Details
CVE-2026-105044	Improper Neutralization of Special Elements used in an OS Command in Security Toolbox versions before 20260929 in Qingteng Host Adaptive Security Platform (Host Monitoring) V3.0 allows an authenticated console administrator to execute arbitrary commands on attached clients via OS command injection	N/A	More Details
CVE-2026-61916	Out-of-bounds reads in URL parsing and case-insensitive header lookup in Apache bRPC before version 1.18.0 on platforms where plain char is unsigned, such as ARM64, allow remote attackers to cause a process crash and denial of service by sending crafted input containing bytes with the high bit set. Root Cause: Two 256-entry lookup tables are accessed through pointers biased by 128 so that valid indices are in the range -128 to 127. Input bytes are promoted from plain char to int without first normalizing them to the signed-char range. On affected platforms, bytes in the range 0x80 to 0xFF therefore become indices 128 to 255 and read up to 127 bytes beyond the end of the underlying tables. Affected Scenarios: A remotely accessible bRPC HTTP server on an affected platform can reach the vulnerable URL lookup through an unauthenticated request URI containing a high-bit byte. The case-insensitive lookup is also reachable when untrusted header names are processed through APIs such as the HTTP/2 HPACK path. HTTP/1 header names containing high-bit bytes are rejected by the HTTP/1 parser before reaching the vulnerable lookup. Builds where plain char is signed, which commonly include x86_64 builds, are not affected by this issue. How to Fix: Choose one of the following options: 1. Upgrade Apache bRPC to version 1.18.0 or later. 2. Apply the patch from https://github.com/apache/bRPC/pull/3376 . The fix normalizes lookup-table indices to the signed-char range before accessing the biased tables.	N/A	More Details
CVE-2026-	Uncontrolled resource consumption in the streaming RPC component of Apache bRPC (version < 1.18.0) allows a remote attacker to exhaust server memory and cause denial of service by sending a single RPC request that declares an excessive number of streams. Root Cause: In the streaming RPC feature, a client can declare multiple streams in one RPC request by putting extra stream identifiers into the repeated field `stream_settings.extra_stream_ids` of the RPC meta. On the server side, `StreamAccept` creates one response stream for every stream identifier declared by the request, without limiting the number of streams accepted in a single RPC. The request body size is only bounded by `max_body_size` (default 64MB), which is enough to carry millions of stream identifiers. Each accepted stream allocates its own bthread execution queue and associated resources, so a single crafted request	N/A	More

81755	can make the server create a huge number of streams and exhaust memory, leading to denial of service. Affected Scenarios: Use a bRPC server that serves streaming RPC and calls `StreamAccept` to accept streams from untrusted clients. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to the 1.18.0 version. 2. Apply this patch: https://github.com/apache/bRPC/pull/3493 Note: The fix adds a new gflag `stream_max_streams_per_request` with a default value of 64. Requests declaring more streams than the limit are rejected before any response stream is created. Applications that legitimately need to accept more than 64 streams in one RPC must raise this flag; otherwise such requests will fail after applying the fix.		Details
CVE-2026-82231	A Thrift client in Apache bRPC (version < 1.18.0, built with WITH_THRIFT which is off by default) can be crashed by the server peer it talks to. When the server returns a T_EXCEPTION reply whose exception struct is malformed, the client decodes it without catching what Apache Thrift throws (TProtocolException, TTransportException or std::length_error). The exception unwinds out of the bthread task function and the process calls std::terminate(), terminating the whole process along with every other in-flight RPC. This is the client response path, so the adversary is the server peer, which THREAT_MODEL.md treats as untrusted. It requires a WITH_THRIFT build (off by default). How to Fix: 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3500	N/A	More Details
CVE-2026-76528	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76528 are related to protection mechanism failure and are grouped under the Common Weakness Enumeration (CWE) CWE-693.	N/A	More Details
CVE-2026-76408	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76408 are related to improper control of a resource through its lifetime and are grouped under the Common Weakness Enumeration (CWE) CWE-664.	N/A	More Details
CVE-2026-82819	A vulnerability in Cerberus FTP Server allows an authenticated remote attacker to perform UNC path injection via a manipulated parameter. By supplying a malicious SMB share path, an attacker can execute authentication coercion against the underlying server, resulting in the leakage of the Cerberus service account's Net-NTLMv2 hash. This issue has been resolved in Cerberus FTP Server: 2026.3.0.	N/A	More Details
CVE-2026-104948	WarehousePG (WHPG) 6.27.1 through 6.27.6 and 7.2.1 through 7.6.0 are affected by a missing authorization vulnerability (CWE-862) in the built-in cluster expansion functions gp_expand_lock_catalog() and gp_expand_bump_version(). Any authenticated database role can run them, and neither checks privileges. A low-privileged user can call gp_expand_lock_catalog() inside an open transaction to block all catalog changes across the cluster for every user, superusers included, until that transaction ends, which in practice requires a superuser to terminate the session. Blocked changes include DDL, temporary table creation, GRANT and CREATE ROLE. The same user can call gp_expand_bump_version() repeatedly. Each call disconnects sessions that hold temporary tables, with a FATAL error on their next catalog change, and makes every other session rebuild its segment connections.	N/A	More Details
CVE-2026-104947	WarehousePG (WHPG) 7.2.1 through 7.6.0 is affected by a missing authorization vulnerability (CWE-862) in the built-in function gp_log_backend_memory_contexts(). Any authenticated database role can run this function without a GRANT. It also gets around the superuser restriction that PostgreSQL places on the underlying pg_log_backend_memory_contexts() function. A low-privileged user can force the segment backends of any other session, superuser sessions included, to write their memory context details to the segment server logs, which the attacker cannot read. Those details can include query text from other users' cached plans, and the call can be repeated to generate large amounts of log output. WarehousePG 6.x is not affected.	N/A	More Details
CVE-2025-0326	The system fails to invalidate existing access tokens associated with sub-organization users during sensitive account management operations such as password reset, role change, or sub-organization disablement. This oversight allows previously issued tokens to remain active even after the user's access context has been modified or revoked through administrative actions. The identified issue could lead to unauthorized access by a malicious actor, but only when all of the following conditions are met: previously issued access tokens remain active, and these tokens have been exposed to malicious actors.	N/A	More Details
CVE-2026-76407	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76407 are related to improper access control and are grouped under the Common Weakness Enumeration	N/A	More Details

(CWE) CWE-284.

CVE-2026-86249

On a `brpc::Server` that sets both `ServerOptions.auth` and `ServerOptions.internal_port`, the HTTP built-in service exemption on `internal_port` can be abused to bypass authentication entirely. `VerifyHttpRequest()` returns `true` without checking any credential when a built-in service is requested on `internal_port`, and this verdict is not per - request: verification runs only for the first message of a connection and the result latches the connection. Sending a request such as `GET /status` to `internal_port` first authenticates the whole connection, after which every subsequent request dispatched on the same connection is handled without ever calling `VerifyHttpRequest()`. This lets a remote attacker reach ordinary (non-built-in, non-tabbed) authenticated services without supplying valid credentials whenever the server uses `ServerOptions.internal_port` together with an Authenticator. The fix restricts `internal_port` so that it carries only built-in and Tabbed services; requests for ordinary services sent to `internal_port` are rejected with `EPREM / HTTP 403`. Mitigation / Fix: upgrade to 1.18.0 or apply PR <https://github.com/apache/brpc/pull/3525>. Users who rely on serving normal services on `internal_port` must move those requests to the port passed to `Server::Start()`.

N/A

[More Details](#)

CVE-2026-86147

External control of file name or path and uncontrolled resource consumption in Apache bRPC (version < 1.18.0) on all platforms allows remote attackers who can reach the builtin service port to create directories and truncate/overwrite files with the privileges of the server process, or to hold the process-wide CPU/contention profiler for an arbitrarily long time. Root Cause: 1. The builtin `/flags` service allows `gflags` carrying a validator to be modified at runtime via `?setvalue=`. The `bvar/mbvar` dump path `gflags` (`bvar_dump_file`, `bvar_dump_tabs`, `mbvar_dump_file`) were registered with an accept-anything validator whose only purpose was to wake the dumping thread. A remote attacker could therefore repoint the dump path (e.g. `GET /flags/bvar_dump_file?setvalue=/victim/dir/target` followed by `GET /flags/bvar_dump?setvalue=true`, or the equivalent for `mbvar_dump_file`, which is used verbatim without a forced suffix), making the dumping thread recursively create directories and open/truncate files at attacker-chosen locations. Tab names in `bvar_dump_tabs` were appended to the path without sanitization, allowing `'..'` to escape the configured directory. 2. The `/pprof` builtin service did not clamp the `?seconds=` parameter to `FLAGS_max_profiling_seconds` (a limit that `/hotspots` enforced). A single request such as `?seconds=2147483647` made the handler sleep for decades while holding the process-wide `gperftools` CPU profiler or `bthread` contention profiler, so `/hotspots/cpu` and `/hotspots/contention` returned 503 for everyone else during that time. Affected Scenarios: Any Apache bRPC server exposing the builtin services (enabled by default on the same port as normal RPC traffic) to untrusted networks. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: <https://github.com/apache/brpc/pull/3523> Note: After the fix, `bvar_dump_file`, `bvar_dump_tabs` and `mbvar_dump_file` can no longer be modified at runtime through the `/flags` builtin service. They can still be changed programmatically via `google::SetCommandLineOption()`, taking effect on the next dump cycle. The profiling duration of `/pprof` is now capped by the `-max_profiling_seconds` `gflag` (default 300 seconds), which can be adjusted by the server owner.

N/A

[More Details](#)

CVE-2026-86146

Inconsistent interpretation of HTTP requests (HTTP request/response smuggling) in the bundled HTTP/1.1 parser of Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to smuggle requests and inject headers when bRPC is deployed behind a proxy or load balancer that parses the same bytes differently. Root Cause: The bRPC fork of `http-parser` accepts and acts on several HTTP/1.1 header forms that a conforming front-end either rejects or interprets differently, so the framing or the header set of a message can differ between the front-end and bRPC: 1. A space before the colon of framing headers (e.g. `Content-Length : 5`) is accepted; the value still drives bRPC's framing while the reported field name contains the trailing space. 2. Interior spaces inside a `Content-Length` value (e.g. `Content-Length: 1 3`) are skipped, so bRPC reads 13 where a front-end reads 1 or rejects the message (same class as CVE-2022-32213 in Node.js). 3. Obsolete line folding (obs-fold, RFC 7230 3.2.4) of `Content-Length` and `Transfer-Encoding` values is accepted, so a folded framing value drives bRPC's framing while a front-end that does not unfold sees a different message. 4. Chunked trailer fields are fed to the same header callbacks as the header section and are merged into the header map like real headers. A front-end only inspects the header section, so any chunked request can smuggle fields such as `Authorization` or `X-Forwarded-For` past it. 5. `Transfer-Encoding` lists with optional whitespace after a comma (e.g. `Transfer-Encoding: gzip, chunked`) are not recognized as chunked: such requests are rejected, and responses are read until connection close instead of as chunked, causing response desync. Affected Scenarios: bRPC server or client communicating over HTTP/1.1 with a front proxy, load balancer, or any peer that parses the message independently. An attacker able to send requests through the front-end can desynchronize the two parsers to smuggle requests, forge headers (e.g. `Authorization`, `X-Forwarded-For`) via chunked trailers, or desynchronize responses. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: <https://github.com/apache/brpc/pull/3521> Note: After applying the fix, the four header forms above are rejected by default and chunked trailers are dropped instead of merged into the header map. Two `gflags` are provided for compatibility: `http_allow_obs_fold=true` restores acceptance of folded `Content-Length/Transfer-Encoding` values, and `http_strict_header_token=true` additionally rejects spaces inside any header field name (not only the ones that decide framing). Legitimate clients that do

N/A

[More Details](#)

	not use these obsolete forms are unaffected.		
CVE-2026-84748	Improper authorization in Apache bRPC (versions < 1.18.0) on all platforms allows remote attackers to access builtin services through the public port. Root Cause: ServerOptions.internal_port is documented to make builtin services (and Tabbed services) reachable only from the internal port, and the HTTP/h2 dispatch path enforced this. However, all protobuf-based protocol dispatch paths (baidu_std, hulu_pbrpc, sofa_pbrpc, and the nshead-based protocols public_pbrpc, nshead_mcpack, nova_pbrpc and ubrpc2pb) computed the security_mode flag but never enforced it between method resolution and method invocation. Since builtin services share the method map with user services, a client able to reach the public port could invoke builtin services such as brpc.vars or hotspots by name over these pb protocols and disclose sensitive runtime information (metrics, flags, profiler data) that the operator intended to expose only via the internal port. Affected Scenarios: Use a bRPC server with ServerOptions.internal_port configured (or with builtin/Tabbed services otherwise expected to be gated by it) while the public port is reachable from an untrusted network and any pb-based protocol is enabled. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3511	N/A	More Details
CVE-2026-84747	Interpretation conflict in the HTTP protocol implementation in Apache bRPC (version < 1.18.0) on all platforms allows remote attackers to bypass access control rules enforced by a front proxy. Root Cause: The method dispatch function FindMethodPropertyByURI() in bRPC's HTTP protocol splits the request path with a string splitter that skips empty fields, so paths containing empty path segments (consecutive slashes), such as //flags, /flags// and /flags//port, are dispatched to the same method as their collapsed form /flags. Per RFC 3986 these are distinct paths. A front proxy that enforces an ACL (e.g. blocking access to builtin services such as /flags) matches the collapsed form only, and therefore forwards the padded literal to the bRPC server, where it is dispatched as if it were the collapsed form. This differential allows //flags?setvalue= to reach a builtin service that /flags is not allowed to reach. Additionally, the HTTP/2 implementation never validated the :path pseudo-header, so requests without a :path, or with a :path that neither begins with '/' nor is the asterisk-form '*', contrary to RFC 9113 section 8.3.1, were not rejected. Affected Scenarios: Deployments of Apache bRPC behind a front proxy or gateway that performs access control (ACL) decisions on the literal request path, where the bRPC server accepts HTTP or HTTP/2 traffic from that proxy. In such deployments an attacker can send a path containing empty path segments to reach a resource that the front proxy intended to block. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3512 Note: After applying the fix, HTTP and HTTP/2 request paths containing empty path segments (consecutive slashes) are rejected by default, and HTTP/2 :path values that are empty or neither begin with '/' nor equal '*' are rejected. Deployments that depend on the previous lenient behavior can restore it by turning on the gflag - http_allow_empty_path_segments, which is not recommended for servers exposed to untrusted traffic.	N/A	More Details
CVE-2026-106781	An issue in WCH CH340/CH341 USB-to-Serial Windows Driver setup package CH341SER.EXE version 4.0 allows a local attacker to escalate privileges.	N/A	More Details
CVE-2026-82942	A data race in the RDMA transport of Apache bRPC (versions < 1.18.0) allows a remote client to crash the server process (denial of service) when RDMA is enabled at build time and runtime. Root Cause: When a bRPC server negotiates RDMA with a client, the handshake completes on the TCP fd, but the server kept parsing that fd afterwards as if it were still an RPC channel. The TCP-fd input stream and the RDMA QP input stream (driven by CQ events) shared the same Socket::read_buf, preferred_index and parsing_context without synchronization. In addition, CQ event processing could start while the handshake parsing was still in progress, and the handshake ACK phase could consume bytes belonging to the next protocol. A remote, unauthenticated client that has completed the RDMA handshake can send data on the TCP fd at the same time RDMA messages are being processed, corrupting the shared input buffer and parser state and crashing the server process (CWE-362, availability impact). Affected Scenarios: Use a bRPC server with RDMA support enabled (built with BRPC_WITH_RDMA and serving rdma:// / SOCKET_MODE_RDMA sockets) and accept connections from untrusted clients. Deployments without RDMA are not affected. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3505	N/A	More Details
CVE-2026-4316	Reflected Cross-Site Scripting (XSS) in Gym Management Software IsMyGym by Zuing Studio affecting the reservation endpoint where the 'ret' parameter is reflected in the HTTP response without any proper sanitization. Successfull exploitation of this vulnerability could allow an attacker to inject arbitrary JavaScript code executed in the context of the victim's browser.	N/A	More Details
CVE-2026-61902	Apache bRPC versions before 1.18.0 do not consistently enforce the configured max_body_size limit in non-progressive HTTP/1 and HTTP/2 parsing and in memcache and couchbase response parsing. A remote unauthenticated peer can cause an affected process to buffer message data beyond the configured limit, which may lead to memory exhaustion and denial of service. Users should upgrade to 1.18.0 or apply the patch from PR https://github.com/apache/brpc/pull/3382 . The fix rejects oversized	N/A	More Details

	buffered HTTP bodies, returns HTTP 413 for HTTP/1 requests, resets oversized HTTP/2 streams, and rejects oversized memcache and couchbase responses from their declared lengths. Progressive HTTP reading remains supported because it uses bounded buffering.		
CVE-2026-82425	A reachable assertion in Apache bRPC (version < 1.18.0) on all platforms allows a malicious remote peer to crash the process, resulting in denial of service. Root Cause: When a Socket with health checking enabled fails, Socket::WaitAndReset closes its file descriptor while the event dispatcher may still be delivering an input event (EPOLLIN/EVFILT_READ) that was fetched for the just-closed fd. If the health check then succeeds and revives the Socket, Socket::Address succeeds again while the fd is still -1 (the next write reconnects on demand). In that window, Socket::OnInputEvent hits CHECK(!(events & EPOLLIN)) (or the kqueue equivalent) because the Socket is addressable but has no fd. CHECK failures terminate the process when bRPC is built with glog or when the gflag -crash_on_fatal_log is enabled (CHECK always terminates in OFFICIAL_BUILD/NDEBUG builds as well). Affected Scenarios: Applications using bRPC clients with health check enabled (-health_check_interval). A malicious remote peer that the client connects to can deliberately abort connections and immediately accept the health check probes, racing connection teardown against health check revival to trigger the assertion and abort the process. How to Fix: (Choose one of the following options) 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3510 Note: After applying the fix, stale input events on a Socket revived without an fd are silently dropped instead of triggering the assertion. No user-visible interface or wire protocol is changed.	N/A	More Details
CVE-2026-104949	WarehousePG (WHPG) 6.27.1 through 6.27.6 and 7.2.1 through 7.6.0 are affected by a missing authorization vulnerability (CWE-862) in gp_add_coordinator_standby(). The function is named gp_add_master_standby() in WHPG 6, and WHPG 7 keeps that name as an alias. Unlike the other segment add and remove functions, it does not require superuser privileges. A low-privileged user who can open a utility-mode connection can use it to add a fake standby coordinator entry to gp_segment_configuration on a cluster that has no standby configured. Administrators then cannot add a real standby until a superuser removes the fake entry. User data is not exposed or modified.	N/A	More Details
CVE-2026-110748	This issue affects Apache SIS from version 0.7 through 1.6 inclusive. When a data file contains links to other files, Apache SIS opened these files regardless their paths. It was possible, for example, to specify a path to a local file. However, these files were always expected to be in some specific formats and the reading stops at the first parsing error, which makes this vulnerability unlikely to leak sensitive information. The impacted formats are: * Geographic Markup Language (GML) with `xlink:href` attributes for other files in GML format. * Well Known Text (WKT) with parameter elements for datum shift grids in NADCON, NTV2 or RGF93 formats. * Landsat scene with paths to files for each band in the GeoTIFF format. Users are recommended to upgrade to version 1.7, which fixes the issue by restricting the paths to the same directory as the file containing the references to other files, or to a predetermined directory for datum shift grids.	N/A	More Details
CVE-2026-82210	Enforceable assertion in the streaming feature of Apache bRPC (version < 1.18.0) on all platforms allows a remote attacker to crash the server process by sending a malformed stream feedback frame. Root Cause: When handling FRAME_TYPE_FEEDBACK frames, bRPC calls Stream::SetRemoteConsumed() which contains CHECK(_cur_buf_size > 0) and unconditionally asserts CHECK(buf.IsEmpty()) on the frame payload. A feedback frame that carries an unexpected payload, or a feedback frame received on a stream where flow control is disabled (cur_buf_size == 0), reaches these assertions. On failure, CHECK failures abort the process, allowing an unauthenticated single client (once a stream is established) to terminate the server and cause a denial of service. Affected Scenarios: Use bRPC server with the streaming control feature (Stream API) served from untrusted network. How to Fix: 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3494 The fix rejects malformed feedback frames with EPROTO and closes only the affected stream instead of crashing the whole process, and it bounds buffer shrinking so cur_buf_size never drops to 0 while flow control is enabled.	N/A	More Details
CVE-2026-114405	Improper validation of array index vulnerability in Velneo vServer. This issue affects vServer: from 36.0 before 36.6 and from 37.0 before 37.2.	N/A	More Details
CVE-2026-76527	As part of Cisco's ongoing commitment to proactive security and product quality, Cisco has conducted a comprehensive internal security review. This review resulted in a software hardening release that addresses multiple internally discovered vulnerabilities. The vulnerabilities tracked by CVE-2026-76527 are related to improper check or handling of exceptional conditions and are grouped under the Common Weakness Enumeration (CWE) CWE-703.	N/A	More Details
	MISP Guard, a proxy-based security filter for MISP feed distribution, contains a case-sensitivity flaw in multiple policy enforcement checks. The functions check_required_taxonomy_exists, check_required_taxonomy_tags_allowed, check_blocked_attribute_categories, and check_blocked_analyst_data_sharing_groups_uuids performed case-sensitive string comparisons when matching event tags, attribute categories, and sharing-group UUIDs against configured allow/block lists.		

CVE-2026-116469	An attacker with the ability to create or modify events in a MISP instance (or a connected instance that pushes events through the guard) could craft events using different letter case in taxonomy tag names, attribute categories, or sharing-group UUIDs to bypass the configured security policies. For example, if the policy requires the 'tlp' taxonomy with only 'tlp:clear' allowed, an event tagged 'TLP:AMBER' would not be recognized as belonging to the 'tlp' taxonomy and would therefore not be checked against the allowed-tags list, resulting in the event being forwarded to feed consumers. Similarly, a blocked attribute category such as 'Person' would not match a configured block entry of 'person', allowing the attribute to pass through the filter. The security impact is unauthorized disclosure of threat-intelligence data that the operator intended to restrict. Events carrying non-allowed TLP (Traffic Light Protocol) markings, blocked attribute categories, or blocked sharing-group associations are forwarded to downstream feed consumers in violation of the configured distribution policy. This undermines the confidentiality controls that MISP Guard is designed to enforce. The vulnerability is trivially exploitable: the attacker only needs to alter the letter case of a tag, category, or UUID string. No authentication beyond normal MISP event-creation privileges is required, and no user interaction is needed.	N/A	More Details
CVE-2026-107157	Specially crafted CBOR input passed to QCborValue::fromCbor() in Qt Core can cause an out-of-bounds read, which may crash the application and could expose information.	N/A	More Details
CVE-2012-20006	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-111103	Improper Cleanup on Thrown Exception vulnerability in the HTTP/2 handler of mtrudel bandit allows an unauthenticated remote attacker to bypass per-connection resource limits via a single TCP write that carries several frames, one of which raises a stream error. When a frame raises StreamError partway through a chunk, Bandit.HTTP2.Handler.handle_data/3 sends RST_STREAM and continues with the connection state from before the chunk, dropping the rest of it. Streams opened by earlier frames in the chunk keep running their Plug calls, but the stream counts, last_stream_id, the rapid-reset timestamps, the HPACK decoder state and the buffer roll back, so max_concurrent_streams, max_requests and the rapid-reset limiter never see those streams. A client that sends one HEADERS frame more than max_concurrent_streams in each write has the extra refused and the others forgotten: with a limit of 10, five such writes run 50 Plug calls at once. This issue affects bandit: from 1.9.0 before 1.14.0.	N/A	More Details
CVE-2026-112880	Improper Neutralization of CRLF Sequences in HTTP Headers ('HTTP Request/Response Splitting') vulnerability in hackney_http_connect in benoitc hackney allows HTTP Request Splitting via the value a proxy_auth_fun callback returns. After a 407 response from an HTTP CONNECT proxy, hackney_http_connect:handle_proxy_auth/11 calls the caller's proxy_auth_fun with the status code and the response headers and writes the returned binary into the Proxy-Authorization line of the CONNECT retry without checking it. A CR, LF or NUL byte in that value adds header lines to the request sent to the proxy. The value is attacker-influenced when the callback derives it from untrusted data, for example from the proxy's challenge headers. Other caller-supplied values on the CONNECT path, such as the target host and the request headers, were already checked for control bytes. This issue affects hackney: from 2.0.0 before 4.9.0.	N/A	More Details
CVE-2026-65226	Rejected reason: This CVE ID has been rejected by runZero as unused.	N/A	More Details
CVE-2026-113436	Inefficient Algorithmic Complexity vulnerability in the header integer parser of mtrudel bandit allows an unauthenticated remote attacker to consume server CPU via a content-length value or a Host or :authority port made of thousands of digits. Bandit.Headers.parse_integer/2 in lib/bandit/headers.ex builds the number one digit at a time with no limit on the digit count, allocating a new bignum at each step, so one value costs time quadratic in its length. It runs on raw request input over HTTP/1 and HTTP/2, before any range check. Over HTTP/2, a 4,000-digit value can be stored once in the HPACK dynamic table and re-sent on each new stream for about one byte: 124 KB of traffic kept the server busy for 27.9 s. This issue affects bandit: from 1.0.0-pre.7 before 1.14.0.	N/A	More Details
CVE-2026-113440	Insufficiently Protected Credentials vulnerability in hackney in benoitc hackney allows a redirecting server to obtain the basic_auth and cookie credentials of a request via a same-host redirect to another scheme or port. With follow_redirect enabled, hackney:follow_redirect/10 in src/hackney.erl reissues the request to the Location target after hackney:maybe_strip_auth_on_redirect/3, which compared only the hostname of the current and the redirect URL. A redirect from https to http on the same host resent the Authorization and Cookie headers in cleartext, where an on-path attacker can read them, and a redirect to another port on the same host handed them to a different service. location_trusted forwards credentials to any target by design and is unchanged. This issue affects hackney: from 2.0.0 before 4.9.0.	N/A	More Details

CVE-2026-114213	Allocation of Resources Without Limits or Throttling vulnerability in the WebSocket fragment reassembly of mtrudel bandit allows an unauthenticated remote attacker to exhaust server memory via a fragmented message made of many 1-byte continuation frames. Bandit.WebSocket.Connection checks max_fragmented_message_size against the payload bytes received, but keeps every non-final continuation frame as a new list cell around its own binary, and nothing limits the number of fragments. A 1-byte fragment costs 7 bytes on the wire but about 150 to 170 bytes of server memory, so one connection can hold about 150 times the configured limit, roughly 1.2 to 1.4 GB at the default 8 MB. Ping frames between fragments keep the connection alive past the idle timeout. This issue affects bandit: from 1.11.0 before 1.14.0.	N/A	More Details
CVE-2026-85297	Bucket Squatting in Google Cloud Run MCP versions prior to 1.11.0 on Google Cloud Platform allows an attacker to intercept and access proprietary source code, execute arbitrary code and exfiltrate data. Users must upgrade to Cloud Run MCP version 1.11.0 or later.	N/A	More Details
CVE-2026-114216	Inconsistent Interpretation of HTTP Requests ('HTTP Request/Response Smuggling') vulnerability in the HTTP/1.1 chunked body reader of mtrudel bandit allows an unauthenticated remote attacker to smuggle a second request past a front-end proxy via a chunk extension that contains a bare LF. The chunked reader in lib/bandit/http1/socket.ex takes the chunk-size line to be everything up to the first CRLF and discards everything after ; without checking it against the RFC 9112 chunk-ext grammar, so bare LF, bare CR and other control bytes pass. A proxy that ends the size line at a bare LF reads 2;\nxx\r\n40\r\n... as one request whose second chunk holds a GET /admin, while Bandit ends the body early and dispatches GET /admin as a separate pipelined request. Deployments without such a proxy in front of Bandit are not affected. This issue affects bandit: from 1.12.4 before 1.14.0.	N/A	More Details
CVE-2026-54426	Uncontrolled recursion in the Redis reply parser in Apache bRPC < 1.18.0 allows a remote Redis server peer to crash a bRPC Redis client process by returning a deeply nested RESP array reply. The parser recursed once for each nested array level without enforcing a depth limit. The fix adds a configurable Redis reply depth limit with default value 128 and rejects replies exceeding that limit. Affected Scenarios: Applications using bRPC as a Redis client with ChannelOptions.protocol="redis" and receiving replies from an untrusted or faulty Redis server peer. How to Fix: 1. Upgrade Apache bRPC to 1.18.0. 2. Apply the patch from PR https://github.com/apache/bRPC/pull/3344 . Note: After applying the fix, Redis replies with array nesting deeper than redis_max_reply_depth are rejected.	N/A	More Details
CVE-2026-54425	Uncontrolled recursion in the HTTP/2 HPACK decoder in Apache bRPC before 1.18.0 allows remote attackers to cause denial of service by sending a crafted HTTP/2 header block containing many consecutive Dynamic Table Size Update entries. Each update was decoded by recursively invoking the HPACK decoder, which could exhaust the thread stack before a header field or end of input was reached. Root Cause: The bRPC HPACK decoder recursively decoded the next entry after processing a Dynamic Table Size Update. The decoder did not bound the recursion depth, so a long sequence of update entries could consume excessive stack space. Affected Scenarios: Use bRPC with the built-in HTTP/2 implementation to accept HTTP/2 traffic from an untrusted network. How to Fix: 1. Upgrade bRPC to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/bRPC/pull/3343 Note: The fix changes HPACK Dynamic Table Size Update decoding to use iterative parsing instead of recursive parsing.	N/A	More Details
CVE-2026-48830	Improper authentication in Apache bRPC through 1.17.0 allows a remote unauthenticated attacker to bypass authentication for protected non-builtin HTTP endpoints by first sending a builtin HTTP request on the same HTTP/1.1 keep-alive connection. When ServerOptions.auth is enabled and builtin HTTP services are exposed on the public listener, a builtin request received as the first message on a socket may cause the connection to be treated as authenticated for later requests on that socket. Users are recommended to upgrade to version 1.18.0, which fixes this issue, or apply this PR separately: https://github.com/apache/bRPC/pull/3308	N/A	More Details
CVE-2026-87096	Sensitive information disclosure and manipulation due to improper authorization checks. The following products are affected: Acronis Cyber Infrastructure (ACI) before build 6.0.1-112, Acronis Cyber Infrastructure (ACI) before build 6.1.1-64, Acronis Cyber Infrastructure (ACI) before build 6.2.1-71.	N/A	More Details
CVE-2026-87095	Sensitive information disclosure and manipulation due to improper authorization checks. The following products are affected: Acronis Cyber Infrastructure (ACI) before build 6.0.1-112, Acronis Cyber Infrastructure (ACI) before build 6.1.1-64, Acronis Cyber Infrastructure (ACI) before build 6.2.1-71.	N/A	More Details
CVE-2026-58663	An authentication bypass vulnerability exists in Apache IoTDB's Thrift RPC service This issue affects Apache IoTDB: from 2.0.0 before 2.0.11. Users are recommended to upgrade to version 2.0.11, which fixes the issue.	N/A	More Details
CVE-2026-41114	Uncontrolled Resource Consumption vulnerability in Apache IoTDB JDBC driver. This issue affects Apache IoTDB JDBC driver: from 1.3.3 before 2.0.11. Users are recommended to upgrade to version 2.0.11, which fixes the issue.	N/A	More Details

CVE-2026-105923	A malicious SVG can trigger an integer overflow which in turn triggers a null-pointer dereference in VectorImage QML Type.	N/A	More Details
CVE-2026-115269	Allocation of Resources Without Limits or Throttling vulnerability in the HTTP/2 request body reader of mtrudel bandit allows an unauthenticated remote attacker to hold stream processes open indefinitely via WINDOW_UPDATE or empty SETTINGS frames sent while a Plug reads the request body. Bandit.HTTP2.Stream.do_read_data/4 in lib/bandit/http2/stream.ex waits for each frame with a fresh read_timeout. A stream WINDOW_UPDATE, or the zero-delta window update that every empty SETTINGS frame delivers to all streams, ends the wait without delivering any body data, and the loop starts a new full countdown. A client that sends one of these 9-byte frames before each countdown expires keeps read_body blocked, together with the stream process and whatever the application holds while reading, for as long as it likes. This issue affects bandit: from 1.4.0 before 1.14.0.	N/A	More Details
CVE-2026-115693	Improper Certificate Validation vulnerability in hackney_ssl in benoitc hackney allows an on-path attacker to impersonate a server pinned to a private certificate authority via a certificate chain that ends in a publicly trusted root. hackney adds {partial_chain, fun hackney_ssl:partial_chain/1} to every verify_peer configuration. The callback returns any certificate in the peer chain whose public key matches a root in the bundled certifi store as the trust anchor, and OTP honours that callback over the configured cacerts. hackney_ssl:merge_ssl_opts/3 kept the callback when the caller replaced the trust store with cacerts or cacertfile in ssl_options, so a connection pinned to a private CA still accepted any chain that carried a publicly trusted root. The hostname check still applies, so the attacker needs a publicly trusted certificate for the pinned hostname. This issue affects hackney: from 1.24.0 before 4.9.0.	N/A	More Details
CVE-2026-115700	Interpretation Conflict vulnerability in hackney_url in benoitc hackney allows Server Side Request Forgery via a percent-encoded bracketed IPv6 literal or an IP literal with a trailing dot. A host such as %5B%3A%3Affff%3A127%2E0%2E0%2E1%5D (decoding to [::ffff:127.0.0.1]) or 127.0.0.1. is not an IP address to a caller's allowlist built on uri_string:parse/1 and inet:parse_address/1, while hackney connects to the address it denotes. The same applies to a Location target when redirects are followed. The IP-literal gate that hackney_url:normalize/2 applies to a percent-decoded host, added to close the percent-encoded IP bypass of CVE-2026-47076, passes the decoded value to inet_parse:address/1 as is. The connect path in hackney_happy:parse_address/1 drops a surrounding bracket pair, and the resolver ignores a trailing dot, so both forms pass the gate and still reach the literal address. Domain names with a trailing dot and unencoded IPv6 literals are not affected. This issue affects hackney: from 4.0.1 before 4.9.0.	N/A	More Details
CVE-2026-70387	Exposed dangerous method or function and Incorrect Authorization vulnerability in OnePlus OxygenOS on Android. The olc2 logging HAL exposes a binder method that fork()'s and calls execl("/vendor/bin/sh", "sh", "-c", cmd, NULL) with a caller-supplied command string. The only authorization check is getCallingUid() == 0. Because SELinux policy grants binder access to this HAL to a wide client attribute, any uid-0 process in any of those domains (including dumpstate, which is reachable from an unprivileged app via Vulnerability 1) obtains arbitrary command execution. The spawned shell lands in u:r:vendor_qti_init_shell:s0 SELinux context via a policy type_transition, a domain holding all Linux capabilities. This issue affects OxygenOS: from 16.1.0 before 16.1.10.500, from 16.0.0 before 16.0.5.1200, and from 15.0.0 before 15.0.0.2000.	N/A	More Details
CVE-2026-70386	Improper neutralization of special elements used in an OS command ('OS command injection') and Missing Authorization vulnerability in OnePlus OxygenOS on Android. Any locally installed Android application, with no permissions declared and no user interaction beyond launching it, can execute arbitrary shell commands as uid 0 in the u:r:dumpstate:s0 SELinux domain. Two defects combine: (a) OnePlus's atlasservice binder interface accepts setEvent(String8 name, String8 value) from any caller without authorization, and (b) one event handler (OplusAtlasLogWriter::handleEvent, matching event name atlas_event_multimedia_audio_dumpsys) copies the caller-supplied value into the system property oplus.audio.dumpinfo.type and starts the root init service audiodumpinfo, which passes that property value unsanitized to system(). This issue affects OxygenOS: from 16.1.0 before 16.1.10.500, from 16.0.0 before 16.0.5.1200, and from 15.0.0 before 15.0.0.2000.	N/A	More Details
CVE-2026-110938	In Eclipse OpenJ9 versions up to 0.63.0, the JITServer component is vulnerable to denial of service due to improper boundary checks during protocol message and reply handling. An unauthenticated remote attacker or malicious JITServer can transmit malformed network payloads that cause out-of-bounds heap reads, immediately crashing the JITServer daemon or client JVM process via a segmentation fault (SIGSEGV).	N/A	More Details
CVE-2026-111075	In Eclipse OpenJ9 versions prior to 0.63.0, the JITServer component fails to perform cross-field validation on attacker-controlled `_size`, `_paddingSize`, and `_dataOffset` fields during `DataDescriptor` message deserialization. An unauthenticated remote attacker can send a crafted network payload that triggers an unsigned integer underflow in `DataDescriptor::getPayloadSize()` and `DataDescriptor::getNextDescriptor()`. This results in an out-of-bounds heap read and an immediate	N/A	More Details

	process crash (SIGSEGV), causing a denial of service.		
CVE-2026-116165	Parse Server before 8.6.101 and 9.0.0 before 9.10.4-alpha.4 contains a missing authorization vulnerability on MongoDB deployments that lets clients without the master key delete the stored schema of any class. Unauthenticated attackers can send REST API delete requests targeting the internal _SCHEMA collection to remove field definitions, class-level permissions, and protected fields, exposing data to unauthorized reads and modifications.	N/A	More Details
CVE-2026-111590	A wide range of Harbor versions allow SSRF via its webhook delivery component which does not properly validate web hook target addresses, allowing usage of private and link-local IP ranges. An authenticated attacker with ProjectAdmin permissions can configure a webhook policy pointing to a private internal address. When a project event fires, Harbor jobservice makes unauthenticated outbound HTTP POST requests to the target. If the target answers with an error status, the response body is written to the webhook execution log, which project members can read, potentially exposing cloud metadata or other network-restricted resources.	N/A	More Details
CVE-2026-111101	Allocation of Resources Without Limits or Throttling vulnerability in the HTTP/1.1 chunked body reader of mtrudel bandit allows an unauthenticated remote attacker to exhaust server memory and CPU via a request body made of many 1-byte chunks. The chunked reader in lib/bandit/http1/socket.ex keeps every chunk as its own list cell and binary term until IO.iodata_to_binary/1 joins them, and nothing limits the number of chunks. The :length option of read_body/2 counts only data bytes, so tiny chunks cost far more memory than their payload: 48 MB of 1-byte chunks grew one handler process to about 500 MB and took about 6 s of CPU. The route does not need to read the body, because on keep-alive connections ensure_completed/1 drains unread bodies through the same reader. This issue affects bandit: from 0.1.0 before 1.14.0.	N/A	More Details
CVE-2026-59562	Missing authentication for critical function vulnerability in Apache IoTDB. This issue affects Apache IoTDB: from 2.0.0 before 2.0.11. Users are recommended to upgrade to version 2.0.11, which fixes the issue.	N/A	More Details
CVE-2026-110499	Allocation of Resources Without Limits or Throttling vulnerability in the HTTP/2 implementation of mtrudel bandit allows an unauthenticated remote attacker to grow per-connection memory and CPU cost without bound via a large SETTINGS_HEADER_TABLE_SIZE. Bandit resizes its HPACK encoder table to whatever size the client advertises, up to $2^{32}-1$ bytes, and indexes response headers into it. Because responses carry distinct values such as date or content-length, nothing is evicted and the table grows with every response for the life of the connection, while each lookup scans it linearly. The SETTINGS handling in Bandit.HTTP2.Connection (lib/bandit/http2/connection.ex) passes remote_settings.header_table_size to HPAX.resize/2 with no server-side cap, and encode_headers encodes every response header as :store. Serving N responses on one connection therefore costs $O(N^2)$ CPU in the connection process. This issue affects bandit: from 0.3.6 before 1.14.0.	N/A	More Details
CVE-2026-59229	The Legacy Pipe receiver protocol (handshake / sendPipeData) uses a hardcoded SUPER_USER (root) identity when executing data loading and deletion operations. A user with only SYSTEM-level privileges (such as USE_PIPE or MAINTAIN) who does NOT have WRITE_DATA or DELETE_DATA permissions on target databases can send specially crafted PipeData messages through the Legacy Pipe API, resulting in data loading/deletion operations executed as root — effectively bypassing IoTDB's permission model entirely. This issue affects Apache IoTDB: from 2.0.0 before 2.0.11. Users are recommended to upgrade to version 2.0.11, which fixes the issue.	N/A	More Details
CVE-2026-58064	The Redis protocol implementation in Apache bRPC allowed inline command parsing to allocate and copy the full connection buffer before a complete CRLF-terminated command was available. A remote attacker with access to a Redis-protocol bRPC service could send long inline inputs without a terminating CRLF to increase memory usage and CPU copy overhead, potentially causing denial of service. Root Cause: The Redis inline command parser allocated memory based on the current input buffer size and copied the entire buffer before determining the current inline command boundary. The existing redis_max_allocation_size limit was applied to RESP bulk strings and arrays, but not effectively to the inline command path before allocation. Affected Scenarios: Use Apache bRPC server with Redis protocol enabled and reachable by untrusted clients. How to Fix: 1. Upgrade Apache bRPC to version 1.18.0, which fixes this issue. 2. Apply the patch from PR https://github.com/apache/bRPC/pull/3363 . Note: The fix applies redis_max_allocation_size to Redis inline command lines and scans for CRLF before allocating parser arena memory. Inline command lines exceeding this limit are rejected.	N/A	More Details
CVE-	The Redis server command parser in Apache bRPC accepts an empty RESP array as a command header. A remote unauthenticated attacker who can reach a bRPC server with Redis service enabled can send an empty array header followed by a bulk string, causing an out-of-bounds write in the parser and crashing the process. Root Cause: RedisCommandParser::ConsumelImpl rejected negative RESP array counts but accepted zero. It resized the command argument vector to zero elements and continued parsing. When a following bulk string was parsed, the only index guard was a CHECK. In configurations		More

2026-56388	where fatal logs do not abort immediately, execution could continue and write to <code>_args[0]</code> even though the vector was empty. Affected Scenarios: Deployments that enable <code>ServerOptions.redis_service</code> and accept Redis protocol traffic from untrusted clients. The issue is reached during parsing before command handling or authentication. How to Fix: 1. Upgrade <code>brpc</code> to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3354 Note: The fix rejects empty RESP arrays as invalid Redis commands and replaces the internal index CHECK with an explicit parse error so malformed command streams cannot fall through to an indexed vector write.	N/A	Details
CVE-2026-65202	Rejected reason: The reported issue was assessed as intended functionality. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-56387	The RTMP AMF parser in Apache <code>brpc</code> accepts AMF string lengths and array counts directly from untrusted input. A remote unauthenticated attacker who can reach a <code>brpc</code> server with RTMP enabled can send a small AMF payload with an excessive declared long string length or array count, causing excessive memory allocation or resource consumption before the parser rejects the incomplete payload. Root Cause: The AMF long string parser resized the destination string to the declared 32-bit length before validating that the stream contained that many bytes. ECMA array and strict array parsers also accepted 32-bit element counts without a configurable upper bound. Affected Scenarios: Deployments that enable the <code>brpc</code> RTMP server component and accept RTMP traffic from untrusted clients. Direct uses of the AMF parser on untrusted data are also affected. How to Fix: 1. Upgrade <code>brpc</code> to version 1.18.0, which fixes this issue. 2. Apply this patch: https://github.com/apache/brpc/pull/3353 Note: The fix introduces configurable AMF parsing limits. The default maximum string size is 64 MiB and the default maximum array element count is 1,048,576. Payloads above these limits are rejected. The limits can be adjusted with <code>amf_max_string_size</code> and <code>amf_max_array_size</code> .	N/A	More Details
CVE-2026-54427	Uncontrolled recursion in the AMF parser in Apache <code>brpc</code> < 1.18.0 allows a remote RTMP peer to crash a <code>brpc</code> process by sending deeply nested AMF objects or arrays. The parser recursively processed nested AMF containers without enforcing a depth limit, which could exhaust the native stack. Affected Scenarios: Applications using <code>brpc</code> RTMP support and parsing AMF data from an untrusted RTMP peer. This includes RTMP servers parsing client command messages, RTMP clients parsing server responses or status messages, and either side parsing inbound AMF script data such as metadata or cue points. Direct callers of <code>ReadAMFObject</code> or <code>ReadAMFArray</code> on untrusted input are also affected. How to Fix: 1. Upgrade Apache <code>brpc</code> to 1.18.0. 2. Apply the patch from PR https://github.com/apache/brpc/pull/3344 . Note: The fix introduces a configurable AMF nesting depth limit with default value 128. AMF objects or arrays deeper than <code>amf_max_depth</code> are rejected after applying the fix.	N/A	More Details
CVE-2026-96432	A malicious or compromised HTTP/2 server withholding <code>SETTINGS ACK</code> can send <code>RST_STREAM</code> on a completed stream, planting a null pointer, and then send a <code>PUSH_PROMISE</code> for that stream, causing a null pointer dereference in the <code>QNetworkAccessManager</code> class of the Qt Network module that crashes the application.	N/A	More Details
CVE-2026-107090	An issue in Maltego Technologies GmbH Maltego Windows Installer 4.13.0 allows a local attacker to execute arbitrary code via the Maltego Windows installer executable-resolution logic	N/A	More Details
CVE-2026-108525	Allocation of Resources Without Limits or Throttling vulnerability in the HTTP/2 implementation of <code>mtrudel bandit</code> allows an unauthenticated remote attacker to exhaust server memory via <code>DATA</code> frames sent past the advertised receive window. The attacker requests a response larger than the 65,535-byte initial window and withholds <code>WINDOW_UPDATE</code> , so the stream process blocks while sending and stops reading, then floods small <code>DATA</code> frames on that stream. Every frame is queued in the stream's mailbox until the <code>BEAM</code> node runs out of memory. <code>Bandit.HTTP2.Connection.handle_frame/3</code> in <code>lib/bandit/http2/connection.ex</code> forwards each <code>DATA</code> frame to the stream process without consulting the receive window, and <code>Bandit.HTTP2.FlowControl.compute_recv_window/2</code> replenishes the window as data arrives rather than as it is consumed, and never answers a violation with <code>FLOW_CONTROL_ERROR</code> . This issue affects <code>bandit</code> : from 0.3.1 before 1.14.0.	N/A	More Details
CVE-2026-108529	Inefficient Algorithmic Complexity vulnerability in the HTTP/2 implementation of <code>mtrudel bandit</code> allows an unauthenticated remote attacker to exhaust server CPU via small <code>SETTINGS</code> or connection-level <code>WINDOW_UPDATE</code> frames. The attacker opens thousands of streams whose responses block on the connection send window, then sends cheap frames that each make the server walk the whole queue of blocked sends. In the reporter's run, with 3,000 blocked streams a single 9-byte <code>SETTINGS</code> frame took 133 ms to process. <code>do_pending_sends/2</code> in <code>lib/bandit/http2/connection.ex</code> retries every queued send and, before each retry, removes that entry from the queue with <code>List.delete/2</code> , so one frame costs $O(n^2)$ list operations for n blocked streams. <code>max_concurrent_streams</code> defaults to <code>:infinity</code> , so the attacker chooses n . This issue affects <code>bandit</code> : from 0.3.4 before 1.14.0.	N/A	More Details
	Improper Handling of Highly Compressed Data (Data Amplification) vulnerability in the HTTP/2		

CVE-2026-109584	implementation of mtrudel bandit allows an unauthenticated remote attacker to exhaust server memory via a small HPACK-compressed HEADERS block. The block stores one large cookie value in the HPACK dynamic table and then references it tens of thousands of times with 1-byte indexed fields. In the reporter's run, a 49 KB block made a server with default settings allocate a 183 MB cookie binary before any application code ran, and every concurrent stream can repeat the cost. Bandit.HTTP2.Connection caps only the compressed block (max_header_block_size), and its check of the decoded list against max_header_list_size never fires because that setting defaults to :infinity. combine_cookie_crumbs/1 in lib/bandit/http2/stream.ex then joins every crumb into one binary. This issue affects bandit: from 0.3.3 before 1.14.0.	N/A	More Details
CVE-2026-65203	Rejected reason: The reported issue could not be reproduced. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-65212	Rejected reason: The reported issue was assessed as intended functionality. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-65214	Rejected reason: The reported issue is tracked under CVE-2026-18449. This CVE ID is rejected as a duplicate.	N/A	More Details
CVE-2026-65215	Rejected reason: The reported issue could not be reproduced. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-110314	Client-Side Enforcement of Server-Side Security vulnerability in the WebSocket connection handling of mtrudel bandit allows a remote client to keep running WebSock handler code after the server has stopped its connection, via data frames sent after ignoring the server's Close frame. When a WebSock callback returns {:stop, ...}, Bandit.WebSocket.Connection calls terminate/2, sends Close and shuts down the write side, but keeps the connection process running in a :closing state. handle_frame/3 and handle_info/3 never check that state, so later frames still reach handle_in/2 and later messages reach handle_info/2, with the state from the stop tuple, until the client sends Close, TCP drops or the idle timeout fires. A client refused in init/1 after a failed token check could still run an admin-only action. This issue affects bandit: from 0.5.1 before 1.14.0.	N/A	More Details
CVE-2026-65216	Rejected reason: The reported issue is tracked under CVE-2026-18450. This CVE ID is rejected as a duplicate.	N/A	More Details
CVE-2026-65217	Rejected reason: The reported issue could not be reproduced. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-110474	This vulnerability exists in the CreditNirvana Agentic AI Call Centre application due to hardcoded static API Key in the client side JavaScript bundle. An unauthenticated remote attacker could exploit this vulnerability by downloading client side JavaScript files and extracting API key, which could lead to make arbitrary API calls to retrieve call transcripts and recordings through the application API of the targeted system.	N/A	More Details
CVE-2026-110475	This vulnerability exists in the CreditNirvana Collections Platform due to improper authentication and authorization controls at the platform login service. An unauthenticated remote attacker could exploit this vulnerability by manipulating authentication requests and accessing APIs without a valid authenticated session on the targeted system. Successful exploitation of this vulnerability could allow the attacker to enumerate platform users, establish sessions as arbitrary users including administrative users, access and modify sensitive information with the privileges of the impersonated user.	N/A	More Details
CVE-2026-65218	Rejected reason: The reported issue is tracked under CVE-2026-18451. This CVE ID is rejected as a duplicate.	N/A	More Details
CVE-2026-65221	Rejected reason: The reported issue was assessed as intended functionality. This CVE ID is unused by the CNA.	N/A	More Details
CVE-2026-65222	Rejected reason: The reported issue could not be reproduced, and the affected products are end-of-life. This CVE ID is unused by the CNA.	N/A	More Details

CVE-2026-113786	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.1, an attacker can craft a PDF containing image data labeled with the /DCTDecode filter that causes Pillow to select and invoke an unintended image parser when the image is loaded. Environments without the optional Pillow dependency are not affected. This issue is fixed in version 6.21.1.	N/A	More Details
CVE-2026-77054	boneIO Black stores passwords, including the MQTT broker and web panel passwords, in cleartext. The GET /api/config endpoint returns the full configuration with sub-services passwords exposed. Any authenticated user can call the endpoint and access the administrative passwords in cleartext form. Critically, due to CVE-2026-77053 (authentication disabled by default), an unauthenticated attacker on the same local network can retrieve all passwords. This issue was fixed in version 1.6.0: web panel passwords are stored as salted scrypt hashes in a file readable only by the service account, and the API and logs no longer return secret values. The MQTT broker password has to stay recoverable because the application presents it to the broker as a client; it is kept in a separate file readable only by the service account and is no longer exposed through the API.	N/A	More Details
CVE-2026-107080	OpenJPEG 2.4.0 through 2.5.4, and git master through commit 8314119b, contains an out-of-bounds read and write in opj_dwt_encode_and_deinterleave_v() in src/lib/openjp2/dwt.c. This allows an attacker to cause a denial of service.	N/A	More Details
CVE-2026-92208	NoMachine mDNS Heap-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows network-adjacent attackers to execute arbitrary code on affected installations of NoMachine. Authentication is not required to exploit this vulnerability. The specific flaw exists within the parsing of mDNS resource records. 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 code in the context of the service account. Was ZDI-CAN-30982.	N/A	More Details
CVE-2026-92183	GIMP APNG File Parsing Stack-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of GIMP. 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 APNG files. The issue results from the lack of proper validation of the length of user-supplied data prior to copying it to a stack-based buffer. An attacker can leverage this vulnerability to execute code in the context of the current process. Was ZDI-CAN-29400.	N/A	More Details
CVE-2026-93328	Path Traversal and Missing Authorization in cloud-storage source in Google Cloud MCP Toolbox for Databases (googleapis/mcp-toolbox) versions 1.2.0 through 1.4.0 allows an unauthenticated remote attacker to exfiltrate host files to external Google Cloud Storage buckets and overwrite arbitrary local files via crafted tool invocation parameters.	N/A	More Details
CVE-2026-92202	Luxion KeyShot BIP File Parsing Uncontrolled Search Path Element Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of Luxion KeyShot. 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 BIP files. The issue results from loading a library from an unsecured location. An attacker can leverage this vulnerability to execute code in the context of the current process. Was ZDI-CAN-29268.	N/A	More Details
CVE-2026-96581	Authorization Bypass Through User-Controlled Key vulnerability in ash-project ash allows a user of one tenant to read, overwrite and take over another tenant's records via an upserting create action. Under attribute multitenancy, Ash.Actions.Create (lib/ash/actions/create/create.ex) and Ash.Actions.Create.Bulk (lib/ash/actions/create/bulk.ex) upserted on the primary key, or on an all_tenants? identity, without scoping the conflicting record to the request's tenant. A caller of a create action with upsert? true that accepts the primary key, supplying a record id from another tenant, updated that record's upsert fields and received all of its attributes in the result. With the default upsert fields the tenant attribute was written too, moving the record into the caller's tenant. Read and update policies do not run for upserts, and tenant-scoped create policies pass because the tenant is the caller's own. This issue affects ash: from 1.22.0 before 3.37.2.	N/A	More Details
CVE-2026-92203	Airbyte SharePoint Connector _get_shared_drive_object Server-Side Request Forgery Information Disclosure Vulnerability. This vulnerability allows remote attackers to initiate arbitrary server-side requests on affected installations of Airbyte. Authentication is required to exploit this vulnerability. The specific flaw exists within the implementation of the _get_shared_drive_object method. The issue results from the lack of proper validation of a URI prior to accessing resources. An attacker can leverage this vulnerability to disclose information in the context of the service account. Was ZDI-CAN-28215.	N/A	More Details
CVE-2026-92204	Airbyte OneDrive Connector _get_shared_drive_object Server-Side Request Forgery Information Disclosure Vulnerability. This vulnerability allows remote attackers to initiate arbitrary server-side requests on affected installations of Airbyte. Authentication is required to exploit this vulnerability. The specific flaw exists within the _get_shared_drive_object method. The issue results from the lack of	N/A	More Details

	proper validation of a URI prior to accessing resources. An attacker can leverage this vulnerability to disclose information in the context of the service account. Was ZDI-CAN-28216.		
CVE-2026-92205	BusyBox libarchive Symlink Directory Traversal Arbitrary File Creation Vulnerability. This vulnerability allows remote attackers to create arbitrary files on affected installations of BusyBox. 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 libarchive component. 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 create files in the context of the current user. Was ZDI-CAN-27556.	N/A	More Details
CVE-2026-111074	In Eclipse OpenJ9 versions prior to 0.63.0, the JITServer component is vulnerable to a pre-authentication remote denial of service. When processing the initial 4-byte length field of an incoming TCP connection, an integer overflow in `MessageBuffer::computeRequiredCapacity()` causes the receiving compilation worker thread to enter an infinite loop consuming ~100% CPU. An unauthenticated remote attacker can exploit this by opening a small number of crafted TCP connections to exhaust the entire worker thread pool, silently halting compilation services for all connected client JVMs.	N/A	More Details
CVE-2026-92206	CrewAI crewAI Framework Agent Loading Unsafe Reflection Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of CrewAI crewAI. User interaction is required to exploit this vulnerability in that the target must load a malicious agent configuration from the repository. The specific flaw exists within the load_agent_from_repository function. The process does not properly restrict a user-supplied argument before using it to import a module. An attacker can leverage this vulnerability to execute code in the context of the service account. Was ZDI-CAN-28387.	N/A	More Details
CVE-2026-92207	MindsDB OpenBBtable Code Injection Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of MindsDB. Authentication is required to exploit this vulnerability. The specific flaw exists within the OpenBBtable class. The issue results from the lack of proper validation of a user-supplied string before using it to execute Python code. An attacker can leverage this vulnerability to execute code in the context of the service account. Was ZDI-CAN-28388.	N/A	More Details
CVE-2026-114463	The webhook platform gateway in Nous Research Hermes Agent versions builds the session chat_id as webhook:{route_name}:{delivery_id} and re-derives the authorizing route by splitting that string, so a route name containing a ":" resolves toolset authorization to a different, higher-privileged prefix route than the one authenticated, which allows a remote attacker with the target HMAC to gain that higher-privileged route's toolsets.	N/A	More Details
CVE-2026-92209	NoMachine Redis Improper Authentication Local Privilege Escalation Vulnerability. This vulnerability allows local attackers to escalate privileges on affected installations of NoMachine. 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 authentication mechanism controlling access to the Redis database. The issue results from using insufficient entropy in the generation of credentials. An attacker can leverage this vulnerability to escalate privileges and execute arbitrary code in the context of the service account. Was ZDI-CAN-30984.	N/A	More Details
CVE-2026-31158	SQL Injection vulnerability in CRM SOFTTRADE Infotech (aka softtradeindia) 1.0 allows a remote attacker to execute arbitrary code via the Login Page	N/A	More Details
CVE-2026-114462	Improper link resolution before file access ('link following') vulnerability in Velneo vServer. This issue affects vServer: from 32.3 before 36.6 and from 37.1 before 37.2.	N/A	More Details
CVE-2026-92210	NoMachine nxhtd Server-Side Request Forgery Information Disclosure Vulnerability. This vulnerability allows remote attackers to initiate arbitrary server-side requests on affected installations of NoMachine. Authentication is not required to exploit this vulnerability. The specific flaw exists within the web service, which listens on UDP port 4443 by default. The issue results from the lack of proper validation of a URI prior to accessing resources. An attacker can leverage this vulnerability to initiate requests in the context of the service account. Was ZDI-CAN-31619.	N/A	More Details
	The Apache Pulsar broker exposes a transaction administration REST API under /admin/v3/transactions. When authentication and authorization are enabled, most of the read endpoints of this API verify only that the request is authenticated and that the broker serves the requested topic or transaction coordinator. They do not check whether the caller is permitted to access the tenant, namespace or topic whose data is returned. The endpoints that modify state, which abort a transaction or scale the transaction coordinators, correctly require super-user access and are not affected. As a result, any authenticated user, including one that has been granted no permissions, can read transaction coordinator and transaction information belonging to other tenants. This includes the list of transaction		

CVE-2026-105931	coordinators with the broker service URLs that serve them, transaction coordinator statistics and internal statistics, transaction metadata including the names of the topics and subscriptions a transaction involves and the principal that opened it, the list of long-running transactions on a transaction coordinator, transaction buffer and pending acknowledgement statistics and internal statistics for any topic and subscription, and whether a given message position of a subscription is currently held in a pending transactional acknowledgement. The disclosure of transaction metadata and statistics can only occur when transaction support is enabled in the broker (transactionCoordinatorEnabled, which is disabled by default). When transaction support is disabled, these endpoints return no transaction data; the only residual effect is that one endpoint's error response can reveal whether a named topic exists in another tenant. The impact is limited to disclosure of transaction-related metadata and statistics across tenant boundaries. The issue does not allow reading message contents, modifying or aborting transactions, or changing any configuration. This issue affects Apache Pulsar versions from 2.8.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users who cannot upgrade immediately and do not use Pulsar transactions should keep transaction support disabled. Where transactions are required, restricting network access to the broker's admin HTTP endpoints to trusted administrative clients reduces the exposure.	N/A	More Details
CVE-2026-113438	Pleasanter provided fails to properly verify passwords in LDAP authentication when processing specially crafted requests under certain configurations. If exploited, an attacker may bypass LDAP authentication by providing only a valid Login ID, without a password. As a result, an attacker may view, alter or delete information stored in the product with the privilege of the account used to log in.	N/A	More Details
CVE-2026-105932	The Apache Pulsar broker admin REST API provides endpoints to read, set and remove the topic-level entry filter policy of a topic. For persistent topics these endpoints check that the caller is authorized for the entry filters topic policy operation. The corresponding endpoints for non-persistent topics (/admin/v2/non-persistent/{tenant}/{namespace}/{topic}/entryFilters) do not perform this authorization check. The only authorization check these endpoints perform is an indirect lookup permission check, made while the topic's partition metadata is retrieved; it is present in all affected versions. As a result, when authentication and authorization are enabled, an authenticated user who has been granted produce or consume permission on a non-persistent topic, or on its namespace, but not permission to manage the topic's policies can read, replace or remove the topic-level entry filter policy of that topic. These operations should be restricted to users with the tenant admin role or the superuser role. This impact analysis assumes that Pulsar has been configured with the default authorization provider. For custom authorization providers, the impact could be slightly different. The practical impact is limited by two broker configuration requirements, neither of which is met by default: - Entry filters must have been installed by the operator. Entry filter implementations are loaded only at broker startup from the broker's entry filters directory, which is empty by default. A request can only select among entry filters the operator has already deployed; setting any other filter name is rejected. The issue does not allow a caller to install or load entry filter code or to run arbitrary code in the broker. - Topic-level entry filter policies must be allowed to override the broker-level configuration (allowOverrideEntryFilters, which is disabled by default). When this setting is disabled, the topic-level entry filter policy has no effect on message dispatch. In deployments where both conditions are met, an unauthorized user can remove the topic-level entry filter policy of a non-persistent topic, for example one used to restrict which messages are delivered to consumers, or apply a different installed entry filter that withholds messages from subscriptions of that topic. Removing it makes the topic fall back to the namespace-level or broker-level entry filters, which are not removed. A changed policy affects only subscription dispatchers created after the change, for example for new subscriptions or when a subscription is re-created after all of its consumers have disconnected; dispatchers that already exist keep the entry filters they were created with. Because non-persistent topics do not store messages, messages withheld by an entry filter are not delivered. Reading the policy discloses only the names of the configured entry filters. The issue does not allow reading message contents, and entry filters cannot modify message payloads. Persistent topics and namespace-level entry filter policies are not affected. Apache Pulsar versions before 3.3.0 are not affected, because in those versions the entry filters topic policy check is made in code shared by the persistent and non-persistent endpoints. This issue affects Apache Pulsar versions from 3.3.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 3.3.x and 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users who cannot upgrade immediately and have deployed entry filters should keep allowOverrideEntryFilters at its default value of false, so that topic-level entry filter policies do not affect message dispatch.	N/A	More Details

CVE-2026-105933	Apache Pulsar allows an administrator to require end-to-end message encryption for all topics in a namespace by setting the encryption_required namespace policy. When this policy is set, the broker rejects messages that are published without encryption. When the broker loads a persistent topic, it reads the namespace policies and the topic-level policies as part of topic initialization. If a step of this initialization fails, for example because of an error reading from the metadata store, a failure loading the topic-level policies or a failure in a later initialization step, the broker logs a warning, disables the encryption requirement for that topic, and completes the topic load as if it had succeeded. The topic is then served with encryption_required not enforced, even though the namespace policy requires it. Topic-level policies that could not be loaded are also not applied. While the encryption requirement is not enforced, an authenticated client with permission to produce to the affected topic can publish messages that are not encrypted, and the broker accepts and stores them. This applies both to producers created after the topic was loaded and to producers that were already connected. Messages published with encryption remain encrypted; the issue does not expose the content of encrypted messages or the encryption keys. The condition is not limited to transient failures. If the cause of the initialization failure persists, the encryption requirement is disabled again each time the topic is loaded, for example after a topic unload, a namespace bundle reassignment or a broker restart. An update to the namespace policies restores enforcement on a loaded topic, but only until the topic is loaded again. Enforcement is reliably restored only once the underlying cause of the failure has been resolved and the topic has been loaded again. The broker logs a warning when the encryption requirement is disabled, but the topic load is reported as successful and the state of the encryption requirement is not shown in the topic statistics. In fixed versions, a failure during topic initialization causes the topic load to fail so that it can be retried, instead of loading the topic with the encryption requirement disabled. This issue affects Apache Pulsar versions from 1.22.0-incubating before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users who rely on encryption_required and cannot upgrade immediately can monitor the broker logs for warnings containing "isEncryptionRequired will be set to false" (the exact wording of the message differs between releases), resolve the underlying cause of the failure, and then reload the affected topic and confirm that the warning does not recur.	N/A	More Details
CVE-2026-115425	Multiple cross site scripting (XSS) vulnerabilities affect OPNsense firewall version 26.7.1_1 and earlier allows low privileged attacker to execute Javascript in the context of another user visiting a specific page or clicking on a specifically crafted link.	N/A	More Details
CVE-2026-114113	Sharp before 0.35.0 bundles a version of libvips affected by multiple vulnerabilities (CVE-2026-33327, CVE-2026-33328, CVE-2026-35590, CVE-2026-35591 and CVE-2026-69242) in its GIF, TIFF and VIPS image decoders. Attackers can supply crafted GIF, TIFF or VIPS files to applications using sharp to process untrusted input, triggering these libvips flaws.	N/A	More Details
CVE-2026-105934	The Apache Pulsar broker admin REST API provides a truncate operation for persistent topics (DELETE /admin/v2/persistent/{tenant}/{namespace}/{topic}/truncate). Truncating a topic moves all subscription cursors of the topic to the end of the topic and deletes the topic's inactive ledgers. Truncating a non-partitioned topic, or a single partition addressed by its partition name, correctly requires the tenant admin role or the superuser role. When the operation is requested for the name of a partitioned topic, the broker does not apply this tenant admin check to the caller. Instead, it truncates each partition using the broker's own internal admin client, which authenticates with the broker's configured client credentials rather than the caller's identity. Each per-partition truncate request is still checked for the tenant admin role or the superuser role, but against the broker's internal identity instead of the caller. The partitions are therefore truncated only when the broker's configured client role is a superuser or an admin of the target tenant, which is the usual configuration when authorization is enabled. As a result, when that condition holds, an authenticated user who is not a tenant admin can truncate any partitioned topic on which the user has lookup permission. Caller authentication and lookup permission alone are not sufficient: if the broker's internal identity lacks tenant admin or superuser authorization for the target tenant, the per-partition requests are rejected. With the default authorization provider, lookup permission is granted by produce or consume permission on the topic or its namespace. A user with no permissions on the topic cannot trigger the issue. This impact analysis assumes that Pulsar has been configured with the default authorization provider. For custom authorization providers, the impact could be different. Truncating a partitioned topic this way moves the cursors of all subscriptions on every partition to the end of the topic, so any unacknowledged backlog is skipped and is no longer delivered to consumers. It also deletes eligible older ledgers of each partition regardless of the namespace's configured retention policy, so messages stored in those ledgers can be permanently lost. The ledger currently being written is not deleted by the truncate operation. The issue does not allow reading message contents or changing permissions or configuration, and it does not give access to topics on which the user has no lookup permission. This issue affects Apache Pulsar versions from 2.8.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at	N/A	More Details

	least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users who cannot upgrade immediately can reduce exposure by restricting network access to the broker's admin REST API (/admin/v2) to trusted administrative clients.		
CVE-2026-105935	The Pulsar Functions Worker provides a REST API endpoint for triggering a function, which publishes a caller-supplied value to one of the function's configured input topics. The Functions Worker performs this publish with its own broker client, which authenticates with the Worker's configured broker-client credentials rather than the credentials of the API caller. When authorization is enabled, in Apache Pulsar 2.4.0 and later versions, the trigger endpoint checks only that the caller is authorized to perform function operations in the function's namespace. It does not check that the caller is authorized to produce to the input topic that the message is published to. Versions before 2.4.0 do not perform a caller authorization check for the trigger endpoint at all. Function input topics are not restricted to the function's own tenant or namespace, and registering a function does not require produce permission on its input topics. As a result, an authenticated user who is authorized to manage functions in a namespace, for example a role granted the functions permission on that namespace, can cause the Functions Worker to publish messages with caller-controlled content to a topic that the user is not authorized to produce to, including topics in other tenants. These messages are published under the identity of the Functions Worker. The issue does not require any function instance to be running or any function code to be executed. It requires that the Functions Worker's broker-client identity is itself authorized to produce to the target topic and, in Apache Pulsar 2.3.0 and later versions, to list the subscriptions of the target topic, as is the case when that identity is a superuser. The impact is on the integrity of the data in the affected topics: consumers of those topics may receive messages from a user who is not permitted to publish them. This concerns the Functions Worker REST API as called by an external, authenticated client; it does not change the security model for function code, which must be fully trusted. In fixed versions, when authorization is enabled, triggering a function additionally requires the caller to be authorized to produce to the selected input topic before any message is published. This issue affects Apache Pulsar versions from 2.0.0-rc1-incubating before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users of Apache Pulsar 2.4.0 and later versions with authorization enabled who cannot upgrade immediately can limit the permission to manage functions to trusted roles. Users of versions before 2.4.0 should restrict access to the Functions Worker REST API to trusted callers, because in these versions the trigger endpoint does not check the caller's permission to manage functions.	N/A	More Details
CVE-2026-105936	The Pulsar Functions worker REST API provides two endpoints that return the catalog of Pulsar IO connectors installed on the functions worker: the worker connectors endpoint (/admin/v2/worker/connectors) and the deprecated functions connectors endpoint (/admin/v2/functions/connectors and /admin/v3/functions/connectors). When authorization is enabled, the worker connectors endpoint requires super-user access. The functions connectors endpoint, served both by the functions worker and by a broker running the functions worker, returns the same catalog without this check. As a result, when authentication and authorization are enabled, any authenticated client can read the connector catalog through the functions connectors endpoint, regardless of the permissions granted to its role and without super-user access. Unauthenticated requests are rejected unless anonymous access is configured with an anonymous user role. The catalog contains the connector type definitions installed in the functions worker's connectors directory: the connector name, its description, and the source, sink and configuration class names. It does not contain configured source or sink instances, their configuration values or credentials, and an installed connector type does not indicate that a connector is configured or running. The impact is limited to disclosure of this metadata, which can include the names and descriptions of custom connectors installed by the operator. The issue does not allow modifying the catalog or managing connectors. Connector definitions that have a source or sink class are also returned by the built-in sources and sinks endpoints to authenticated clients. In fixed versions, the functions connectors endpoint applies the same checks as the worker connectors endpoint, so that it requires super-user access when authorization is enabled. Clients that are not super-users can use the built-in sources and sinks endpoints to list the available connector types. This issue affects Apache Pulsar versions from 2.4.1 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/.	N/A	More Details
CVE-	Deserialization of untrusted data in Mandiant FLARE floss releases before 3.2.0 allows arbitrary code		More

2026-114198	execution when a crafted .viv or .exe.viv workspace is passed to the floss utility	N/A	Details
CVE-2026-114197	Deserialization of untrusted data in Mandiant FLARE capa releases before 10.0.0 allows arbitrary code execution when a crafted .viv or .exe.viv workspace is passed to the capa utility	N/A	More Details
CVE-2026-99099	In the Linux kernel, the following vulnerability has been resolved: x86/CPU/AMD: Fix AMD TLBI Erratum #1718 The processor may fail to invalidate a TLB entry while executing an INVLPGB/TLBSYNC instruction pair. Detect the erratum and apply a software workaround when a microcode fix is not present.	N/A	More Details
CVE-2026-111184	NoMachine Named Pipe Incorrect Permission Assignment Local Privilege Escalation Vulnerability. This vulnerability allows local attackers to escalate privileges on affected installations of NoMachine. 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 NoMachine Device Service. The issue results from incorrect permissions on a named pipe. An attacker can leverage this vulnerability to escalate privileges and execute arbitrary code in the context of SYSTEM. Was ZDI-CAN-31213.	N/A	More Details
CVE-2026-77055	boneIO black_debian_images gives the boneio service account excessive privileges, including full passwordless sudo access and membership in the docker group. A local attacker who compromises the boneio user can use these privileges to escalate to root. This issue was fixed in version 1.6.0: the image removes the sudoers rule and the docker group membership when it is sealed, and devices already in the field have both removed by migrations shipped with boneio 1.6.0 and applied on first start after the upgrade. Privileged operations now go through helpers that accept only a fixed set of commands, with migrations verified against Ed25519 signatures.	N/A	More Details
CVE-2026-103299	The Payment Gateway WordPress plugin through 2.2.0 does not ensure a shared secret is configured before using it to authenticate the bank's payment-notification callback, allowing unauthenticated attackers to forge a valid notification and mark any WooCommerce order as paid on installs where the secret has not been set, and the same callback lets them move any pending order to a failed state whether or not the secret is configured.	N/A	More Details
CVE-2026-114126	Storybook is a frontend workshop for building user interface components and pages in isolation. From 8.0.0 until 8.6.19, 9.1.21, 10.6.2, and 11.0.0-alpha.6, the Storybook development server can expose its WebSocket authentication token to clients able to load its HTML pages over a local network. A LAN attacker can obtain the token and invoke WebSocket-backed file-writing or testing functionality. Weak story-write validation can allow code to be placed in story files. If the Storybook Vitest addon is installed with an old version of Vitest, the payload can use a Vitest API to write additional files or execute commands. Depending on the payload and local tooling, the attack can lead to code execution with the developer's privileges. This issue is fixed in 8.6.19, 9.1.21, 10.6.2, and 11.0.0-alpha.6.	N/A	More Details
CVE-2026-77053	boneIO Black firmware ships with authentication disabled by default. Its management web API listens on all network interfaces and exposes all endpoints without requiring credentials. An unauthenticated remote attacker can fully control the device, causing denial of service, modifying configuration, and operating connected physical circuits. This issue was fixed in version 1.6.0: a device without an administrator account refuses API requests until the owner creates one in the first-run wizard, and every API route is restricted to an authenticated role. Versions before 1.2.0 are based on Debian 10, are no longer supported and cannot be upgraded in place; affected devices should be re-flashed with the 1.6.0 system image.	N/A	More Details
CVE-2026-77052	boneIO Black exposes a Node-RED instance without authentication at the /nodered/ path through a Caddy reverse proxy. The Node-RED Admin API allows deploying flows that can execute arbitrary system commands. An unauthenticated attacker on the local network can deploy a malicious flow to achieve remote code execution and gain full control over the home automation system. This issue was fixed in version 1.6.0.	N/A	More Details
CVE-2026-114125	Storybook is a frontend workshop for building user interface components and pages in isolation. From 8.0.0 until 8.6.19, 9.1.21, 10.6.2, and 11.0.0-alpha.6, a malicious website can interact with the preview of a running Storybook development server. Insufficient iframe embedding and message-origin restrictions can allow attacker-controlled events to reach the running preview. A developer must visit the malicious website and interact with the injected interface. The attack can cause unauthorized changes to story files and, if a changed file is later executed by a local runtime, arbitrary code execution with the developer's privileges. This issue is fixed in 8.6.19, 9.1.21, 10.6.2, and 11.0.0-alpha.6.	N/A	More Details
	Astro is a web framework for content-driven websites. From 7.1.0 until 7.3.3, Astro's server-side View Transition CSS generator does not keep string animation properties within their generated declarations. An application that passes attacker-controlled data into properties such as duration, delay, easing,		

CVE-2026-114127	direction, fillMode, or animation names can allow CSS delimiters to terminate the declaration and add attacker-selected rules to the server-rendered page. The generated stylesheet still escapes literal less-than characters, so the injected content cannot terminate the style element or directly inject HTML or JavaScript. The injected CSS can conceal or reposition content, add generated text, and otherwise misrepresent the interface; some applications may also expose data already present in selectable DOM attributes through conditional resource loads. This issue is fixed in 7.3.3.	N/A	More Details
CVE-2026-3322	Due to improper permission checks in Lobster_pro, attackers are able to interact with endpoints that should require administrative privileges. This issue affects Lobster_pro version 4.12.12-LTS and potentially other versions prior to 25.1.3.	N/A	More Details
CVE-2026-77056	boneIO black_debian_images ships with a Mosquitto MQTT broker that uses the same default credentials on all devices. These credentials are publicly documented on github, and there is no indication in the instructions to change them. A remote attacker can use them to gain full read and write access to the broker, allowing control of connected relays, blinds, and gates, and injection of fake entities into a linked Home Assistant installation. This issue was fixed in version 1.6.0 of the image: MQTT passwords are generated per device at first boot. Devices upgraded from earlier versions keep their existing broker credentials; from boneio 1.6.0 the web panel reports credentials that are still the published defaults, and the owner has to change them.	N/A	More Details
CVE-2026-36698	Xovis AG APS-R-POE, APS-R-POE-PP, APS-R, APS-R-POE-USB, and APS-B 2.10.1.1 and earlier contains an OS command injection vulnerability in /sbin/ImportParamSet.sh. The script processes symlinks from uploaded parameter archive files and embeds the symlink target path into a shell command without quoting or sanitization. The command is executed as root via APC_SendMsg IPC.	N/A	More Details
CVE-2026-77057	boneIO Black does not limit authentication attempts on its web interface (POST /api/login) or SSH service, and does not enforce password complexity. An unauthenticated remote attacker can make unlimited brute-force login attempts, potentially compromising user accounts and gaining unauthorized access to the device. This issue was fixed in version 1.6.0: web login attempts are rate-limited per source address and per account, panel passwords must be at least 8 characters long, and SSH allows three authentication attempts per connection with a shortened login grace time.	N/A	More Details
CVE-2026-111183	GStreamer rtpL16depay Heap-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of GStreamer. Interaction with this library is required to exploit this vulnerability but attack vectors may vary depending on the implementation. The specific flaw exists within the processing of RTP payload elements. 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 code in the context of the current process. Was ZDI-CAN-31143.	N/A	More Details
CVE-2026-105869	Apache Pulsar's MultiRolesTokenAuthorizationProvider authorizes a client with all roles listed in a claim of its JSON Web Token (JWT). The provider is intended for use with token authentication, which verifies the token before authorization. The provider did not verify the token itself: it parsed the token's claims with the signature removed and trusted that authentication had already checked it. When authentication and authorization are enabled with a non-empty anonymousUserRole setting, some HTTP requests whose credentials fail authentication are not rejected, but are handled as requests of the anonymous user role. For the broker's HTTP endpoints, such as the admin REST API, this fallback applies to requests that do not name an authentication method and requires strictAuthMethod to be false (the default) where that setting is available. For the WebSocket API, failed credentials can also fall back to the anonymous user role; this path does not enforce strictAuthMethod. The authentication data kept for such a request still contained the request's Authorization header. MultiRolesTokenAuthorizationProvider then read the bearer token from that header without verifying its signature and authorized the request with the roles in the token's role claim. As a result, in deployments that use MultiRolesTokenAuthorizationProvider together with a non-empty anonymousUserRole and that permit this fallback, an unauthenticated client can send an HTTP request, for example to the admin REST API or the WebSocket API, with a forged or unsigned token that lists arbitrary roles, and be authorized with those roles. The impact can include super-user access when the token lists a role configured in superUserRoles, allowing the client to read and change tenants, namespaces, topics, permissions and broker configuration. No valid credentials are required. Deployments are not affected if they use the default authorization provider (PulsarAuthorizationProvider) or another authorization provider, or if anonymousUserRole is not set (the default). Requests that authenticate successfully with a valid token are not affected, because the token that MultiRolesTokenAuthorizationProvider reads is the token that authentication verified. Connections that use the Pulsar binary protocol are not affected by this issue, because the broker does not keep client authentication data for connections that it assigns to the anonymous user role. In fixed versions, MultiRolesTokenAuthorizationProvider obtains roles only from tokens that the configured token authentication provider has verified, using its configured keys and validation settings. Broker HTTP requests that are assigned to the anonymous user role are authorized only as that role, and invalid	N/A	More Details

	tokens cannot supply additional roles to WebSocket requests. The broker fails to start if MultiRolesTokenAuthorizationProvider is configured with authorization enabled but without a token authentication provider. This issue affects Apache Pulsar versions from 2.8.1 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/ . Users who cannot upgrade immediately can remove the anonymousUserRole setting from the broker and WebSocket service configurations, or restrict network access to the HTTP endpoints of these services to trusted clients.		
CVE-2026-77058	boneIO Black is vulnerable to Cross-Site Request Forgery (CSRF) on all state-changing API endpoints because it lacks CSRF tokens and does not validate Origin or Referer headers. A remote attacker can craft a malicious webpage that, when visited by a user on the internal network, sends cross-origin requests to the device. This allows the attacker to reboot the controller, toggle relay circuits, control gates and blinds, or shut down the device, causing loss of integrity and availability. Critically, due to CVE-2026-77053 (authentication disabled by default), any user on the same internal network can trigger attacker-controlled actions by visiting the malicious website. This issue was fixed in version 1.6.0.	N/A	More Details
CVE-2026-77059	boneIO Black contains a TOCTOU race condition in the create_timedatectl_sudoers_file() function in boneio/webui/routes/timezone_sudoers.py. The function writes sudoers rules to a predictable temporary file, validates them, then installs them with root privileges. A local attacker can exploit the race window between validation and installation to substitute malicious content, or use the predictable filename for symlink attacks, resulting in arbitrary sudoers rule injection and local privilege escalation to root. This issue was fixed in version 1.6.0.	N/A	More Details
CVE-2026-114130	Nuxt is an open-source web development framework for Vue.js. Prior to 4.6.0, an incomplete fix for CVE-2026-53722 allows an application with pages disabled to render a script-capable URL through NuxtLink when attacker-controlled input bound to to is prefixed with an ASCII C0 control character. The no-pages router classifies the value as an internal route before the earlier protocol check and passes it to the rendered anchor href; browsers remove the leading control character and interpret the remaining javascript: URL. The default anchor and NuxtLink custom output are affected, the unsafe href is present in server-rendered HTML, and a victim must click the link for script to execute in the application's origin. This issue is fixed in version 4.6.0.	N/A	More Details
CVE-2026-114132	Nuxt is an open-source web development framework for Vue.js. Prior to 4.6.2, Nuxt resolves pages case-insensitively by default while Nitro matches server-side route-rule keys case-sensitively, allowing a request whose path differs only by letter case to render without the server rule intended for that page. The mismatch can skip redirect, headers, proxy, ssr: false, and noScripts, and can cause getRouteRules(event) and matchRouteRules to return no rule for the case variant. Applications that set router.options.sensitive: true in nuxt.config are not affected; setting sensitive only in app/router.options.ts does not establish this exclusion because the server follows the nuxt.config value. Access checks implemented in server middleware, page route middleware, or page data fetching remain independent of this matcher. This issue is fixed in version 4.6.2.	N/A	More Details
CVE-2026-105930	When package management is enabled, Pulsar Functions, Sources and Sinks can be registered or updated from a package stored in the Pulsar Packages Management Service by supplying a package URL. The Pulsar Functions worker authorized these requests only against the namespace in which the component was being registered or updated. It did not check whether the requesting user was permitted to access packages in the namespace that the package URL refers to, and it retrieved the package from the broker using the worker's own credentials. The permission check of the Packages Management Service was therefore applied to the worker's identity instead of the requesting user's. Package URLs of transform functions supplied when registering or updating a sink were handled in the same way. An authenticated user who is permitted to register or update functions, sources or sinks in any namespace of their own tenant could use this to obtain the contents of a package stored in a namespace of another tenant, without having the packages permission on that namespace. The impact is limited to disclosure of package contents, which may include application code, business logic or embedded credentials. The issue does not allow packages to be modified or deleted and does not affect availability. This issue only affects deployments that have authentication and authorization enabled and in which package management has been enabled: the broker setting enablePackagesManagement and, in current releases, the functions worker setting functionsWorkerEnablePackageManagement. Both settings are disabled by default. This issue affects Apache Pulsar versions from 2.8.0 before 4.0.14, from 4.1.0 through 4.1.3, and from 4.2.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/ .	N/A	More Details

CVE-2026-77060	boneIO black_debian_images stores the Mosquitto MQTT broker password file at /etc/mosquitto/passwd with world-readable permissions. Any local user can read the password hashes for all MQTT accounts and perform offline dictionary attacks. This issue was fixed in version 1.6.0.	N/A	More Details
CVE-2026-105927	The Apache Pulsar WebSocket service allows a WebSocket consumer to configure a dead letter policy through query parameters of the consumer endpoint, including the deadLetterTopic parameter that names the topic to which messages are written after the maximum number of redeliveries is exceeded. When authorization is enabled, the WebSocket service checks only that the connecting user holds the CONSUME permission on the subscribed topic. It does not check whether the user holds the PRODUCE permission on the topic named in deadLetterTopic. The dead letter messages are then written by the WebSocket service using its own broker client identity rather than the identity of the connecting user. As a result, an authenticated user who holds the CONSUME permission on at least one topic can cause messages to be written to a topic of their choosing in any existing namespace, including namespaces belonging to other tenants, provided the role used by the WebSocket service to connect to the brokers is authorized to produce to that topic and the topic accepts a producer that registers no schema. The dead letter producer created by the WebSocket service registers no schema, so it is rejected when the target topic has a schema and schema validation is enforced for that topic or its namespace; this setting is not enabled by default. The same user is refused when attempting to produce to that topic directly. Where automatic topic creation is enabled, which is the default, the named topic is created if it does not already exist. The issue can only be triggered when the WebSocket service is enabled, either in the broker (websocketServiceEnabled, which is disabled by default) or as a standalone WebSocket proxy. The impact is unauthorized injection of messages into topics the user is not permitted to produce to. The user can determine the content of the injected messages when they also control the messages on the subscribed topic. The issue does not allow reading, modifying or deleting existing messages in the target topic. Injected messages carry the dead letter message properties and a dead letter producer name, so they can be distinguished from messages produced normally; however, the values of these properties are not authenticated and should not be relied upon to establish the origin of a message. Whether an injected message has further consequences depends on how the applications consuming the target topic process it. Repeated injection also consumes storage and backlog capacity in the target namespace. This issue affects Apache Pulsar versions from 2.3.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/ . Users who cannot upgrade immediately and do not require the WebSocket service should disable it. Where the WebSocket service is required, restricting network access to it and limiting the permissions granted to the role the WebSocket service uses to connect to the brokers to the namespaces WebSocket clients are intended to access reduces the exposure.	N/A	More Details
CVE-2026-77061	boneIO Black directly excludes several API endpoints from the authentication. An unauthenticated remote attacker on the local network can retrieve system logs, firmware version, serial number, MQTT username, configuration file listings, and API documentation. This data aids reconnaissance and enables targeted attacks against the device. This issue was fixed in version 1.6.0: every API route requires an authenticated role, the interactive API documentation is disabled, and the serial number is withheld from anonymous callers.	N/A	More Details
CVE-2026-113345	Electronic delivery item Inspection Support System provided by Japan's Ministry of Land, Infrastructure, Transport and Tourism contains an improper file access permission vulnerability. Files in the installation folder of the application may be tampered with, resulting in arbitrary code execution.	N/A	More Details
CVE-2026-77062	The boneIO Black web interface lacks proper clickjacking protections. A remote attacker can embed the control panel in an iframe on a malicious web page and trick an authenticated user on the internal network into clicking buttons that activate relays, open gates, or control blinds. This can lead to unauthorized state changes with potential physical safety consequences. Critically, due to CVE-2026-77053 (authentication disabled by default), any user on the same internal network can trigger attacker-controlled actions by visiting the malicious website. This issue was fixed in version 1.6.0.	N/A	More Details
	The Pulsar broker serves two schema commands on its binary protocol: GetSchema, which returns the schema registered for a topic, and GetOrCreateSchema, which registers a new schema version for an existing topic. The broker handled both commands after checking only that the client connection was authenticated. Unlike the broker's other topic commands, such as those for creating producers and consumers or for looking up topics, these two commands did not check whether the authenticated client was authorized for the target topic. The Pulsar REST API endpoints for schemas were not affected. An authenticated client without any permission on the target topic, namespace or tenant could use these commands to: - read the schema registered for a topic in another namespace or tenant, disclosing the schema definition such as record field names and types. Message payloads are not disclosed - register a new schema version on an existing topic in another namespace or tenant. The		

CVE-2026-105926	topic's schema compatibility and schema auto-update policies still apply. The issue does not allow message data to be read or modified, but a registered schema version can cause legitimate producers and consumers of that topic to fail with schema incompatibility errors until the schema is corrected by an administrator. The standard Pulsar client performs an authorized topic lookup before it sends these commands, so the missing check could not be observed when testing with the standard client. A client that implements the binary protocol directly is not required to perform that lookup first. This issue is relevant only to deployments that have authentication and authorization enabled. The GetSchema command is affected from Apache Pulsar 2.2.0, and the GetOrCreateSchema command from Apache Pulsar 2.5.0. After the fix, the broker requires the lookup permission on the topic for GetSchema and the produce permission on the topic for GetOrCreateSchema. This impact analysis assumes that Pulsar has been configured with the default authorization provider. For custom authorization providers, the impact could be slightly different. This issue affects Apache Pulsar versions from 2.2.0 before 4.0.14, from 4.1.0 through 4.1.3, and from 4.2.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/ .	N/A	More Details
CVE-2026-115635	Incorrect Authorization vulnerability in ash-project ash_postgres allows a user to read another tenant's record through a skipped upsert. When Ash.bulk_create/4 or a single create runs as an upsert with return_skipped_upsert?: true and the upsert is skipped (for example because its upsert_condition does not hold), AshPostgres.DataLayer reads the existing row back by the upsert identity in do_bulk_create/3 and upsert/4. That lookup did not apply the resource's attribute multitenancy, so on a resource using strategy :attribute it could match and return the row of another tenant with the same identity values, exposing its fields to the caller. This issue affects ash_postgres: from 2.6.22 before 2.15.2.	N/A	More Details
CVE-2026-115627	Incorrect Authorization vulnerability in ash-project ash_postgres allows a user to load destination records linked to a source record only through another tenant's join rows. When a many_to_many relationship, or a through relationship, is loaded with a lateral join (for example with a per-record limit or sort), AshPostgres.DataLayer builds the query for the join resource, and for each resource along a through path, in lateral_join_query/3 and through_step_queries/4 without running that resource's read action, which is where attribute multitenancy is applied. When those resources use strategy :attribute multitenancy, rows belonging to other tenants were joined, so the load returned destination records that the caller's tenant is not related to. This issue affects ash_postgres: from 0.38.0 before 2.15.2.	N/A	More Details
CVE-2026-114214	Incorrect Authorization vulnerability in ash-project ash allows a client to read destination records unrelated to the source record through a through relationship. When a has_many relationship declared with through: was loaded for a single record with a limit of 1, or an unsorted has_one through relationship was loaded with a limit of 1, lateral_join?/4 in lib/ash/actions/read/relationships.ex chose the path that reads the destination without any filter linking it to the source. The load returned the first destination record by sort, or any other record with an offset. The read still runs the destination's policies but carries the relationship as accessing_from, so a client paging such a field (for example teachers(limit: 1, offset: n) through AshGraphql) can enumerate destination records that policies allow only when reached through that relationship. This issue affects ash: from 3.24.0 before 3.37.1.	N/A	More Details
CVE-2026-99091	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: MGMT: Fix status of pending commands flushed on power off cmd_complete_rsp() receives a struct cmd_lookup and, for a pending command that has no cmd_complete callback, falls through to cmd_status_rsp(cmd, data). cmd_status_rsp() reads data as a u8 *, so the status sent comes from the first byte of match->sk's representation, not from match->mgmt_status. commit f53e1c9c726d ("Bluetooth: MGMT: Fix possible crash on mgmt_index_removed") changed the callers' data from &status to &match and updated the cmd_complete branch to match->mgmt_status, but left the fall-through passing data unchanged. In mgmt_index_removed() match->sk is NULL, so every such command is answered with status 0x00 (Success) instead of MGMT_STATUS_INVALID_INDEX. In __mgmt_power_off() it is NULL too, unless a Set Powered command was pending: then settings_rsp() has stored that command's socket in match->sk, and the byte sent comes from the socket pointer. Either way the status the caller set, MGMT_STATUS_NOT_POWERED (or MGMT_STATUS_INVALID_INDEX when the device is being unregistered), is lost. Pass the status the callers set. This was seen on a QCA9377 laptop: bluetoothd's Start Discovery, pending when the adapter was powered off, was answered with Command Status 0x00 and no parameters, and bluetoothd 5.72 crashed on the zero-length success. BlueZ master has since added the length check on that path ("adapter: Fix crash on short start discovery reply", BlueZ a734b06059cb), but userspace should still be told Not Powered, not Success. Tested on Ubuntu 7.0.0-31 with a virtual controller (hci_vhci): power off is held in HCI Write Scan Enable while a Start Discovery is submitted, and the flush answers it with status 0x00 before this change and 0x0f (Not Powered) after.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ASoC: samsung: i2s: enable op_clk in probe to balance runtime PM The probe function puts the device into the runtime PM active state with only the "iis" bus clock enabled, but it also sets priv->op_clk to the parent of the RCLK_SRC mux		

CVE-2026-99089	without enabling it. The runtime PM suspend callback disables both clocks, so the first runtime suspend after probe disables op_clk once more than it was enabled. On Exynos5250 Arndale both clocks resolve to the same EXYNOS_I2S_BUS gate, so this triggers the following warning during boot: WARNING: drivers/clk/clk.c:1270 at clk_core_disable+0xe4/0x23c, CPU#0: kworker/u8:3/45 i2s_bus already disabled ... Workqueue: pm pm_runtime_work Call trace: ... clk_core_disable from clk_disable+0x28/0x34 clk_disable from i2s_runtime_suspend+0x58/0x68 [snd_soc_i2s] i2s_runtime_suspend [snd_soc_i2s] from genpd_runtime_suspend+0xc4/0x298 genpd_runtime_suspend from __rpm_callback+0x48/0x18c __rpm_callback from rpm_callback+0x4c/0x50 rpm_callback from rpm_suspend+0xf4/0x718 rpm_suspend from pm_runtime_work+0x9c/0xa0 ... Fix this by enabling op_clk in probe right after getting it, so the clock state matches the runtime PM active state, and by disabling it in remove.	N/A	More Details
CVE-2026-113790	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.0, an attacker can craft a PDF containing arrays or dictionaries with many elements, causing parsing to consume large amounts of memory. This issue is fixed in version 6.21.0.	N/A	More Details
CVE-2026-82987	There exists a screen capture endpoint in vCast media streaming service in ViewSonic ViewBoard unknown allows a remote, unauthenticated attacker to capture and exfiltrate the current screen content via GET requests to an unauthenticated /snapshot or /screen endpoint returning JPEGs	N/A	More Details
CVE-2026-112921	Robots::Validate versions from 0.4.3 before 0.4.10 for Perl report clients as imposter robots when User-Agent matcher state carries over from a previous request. Each Robots::Validate object holds one pure-Perl Aho-Corasick search machine and feeds every User-Agent into it without resetting its state. A User-Agent that ends part-way through a robot name leaves the machine mid-match, so the next User-Agent checked by the same object can complete that name. That client's address is then checked against the robot's DNS rule, which it fails. A remote client can use this to make a site that blocks bad robots block other visitors where the prefix User-Agent string is a suffix of a robot being validated. Note that deployments with Algorithm::AhoCorasick::XS are also affected, since it was unused due to a different bug.	N/A	More Details
CVE-2026-114509	Laravel is a web application framework. Prior to 12.68.0 and 13.29.0, TokenGuard configured as a standard token-based API guard does not reject a crafted non-string credential. Under affected MySQL or MariaDB configurations, the authentication lookup can match an unintended existing user, allowing an unauthenticated attacker to bypass authentication and take over an account on routes protected by TokenGuard. This issue is fixed in versions 12.68.0 and 13.29.0.	N/A	More Details
CVE-2026-94140	Rex versions through 1.16.1 for Perl accept an unknown SSH host key for remote rsync transfers because sync hardcodes StrictHostKeyChecking=no. Every remote branch of sync builds its ssh command with that option, for password authentication, for private-key authentication and for the fallback. The ssh options a caller sets through set_openssh_opt reach the OpenSSH transport but not this command, so there is no configuration that restores host key checking for the transfer. An attacker who can redirect the connection before the real host key is known presents their own and the transfer proceeds: under password authentication sync sends the connection password to that server, an upload writes the selected local files to it, and a download writes what it returns to the local destination.	N/A	More Details
CVE-2026-94423	Net::LibSSH versions before 0.004 for Perl authenticate to an unverified server because connect never checks the host key against known_hosts. connect reports success as soon as ssh_connect returns SSH_OK. libssh leaves the known-hosts check to the application, and this binding never makes it, so the strict_hostkeycheck and knownhosts options are set on the session and then decide nothing. Password, agent and public-key authentication, channel and SFTP operations all run on a session whose peer was never identified, so an attacker who can redirect the connection receives the password when password authentication is used, and commands and data sent over the session.	N/A	More Details
CVE-2026-113789	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.0, an attacker can craft a PDF that causes large memory usage or long runtimes when an application retrieves the document's form fields. This issue is fixed in version 6.21.0.	N/A	More Details
CVE-2026-100420	A sensitive information disclosure vulnerability in the cloud-client daemon of Deco BE63 allows authenticated attackers to retrieve the plaintext TP-Link ID account password. Successful exploitation may allow attackers to compromise the associated cloud account and perform unauthorized actions on behalf of the affected user.	N/A	More Details
CVE-2026-103090	Missing Authentication for Critical Function (CWE-306) and Insufficiently Protected Credentials (CWE-522) in the control plane enrollment status handler (HandleEnrollStatus at GET /enroll/status in internal/controlplane/server.go) in Google OSS Agent Mesh (SAM) versions 0.1.0-alpha.1 through 0.1.0-alpha.8 allow an unauthenticated remote attacker to retrieve an approved bootstrap-enrolled peer's signed Biscuit identity token by supplying a peer_id obtainable from the unauthenticated GET /info endpoint. A related flaw in POST /enroll allowed any bootstrap-token holder to re-fetch another peer's credential. An attacker can replay the disclosed Biscuit token as a bearer credential on /policies to read	N/A	More Details

	mesh authorization policies and on /routers/lease to overwrite advertised router multiaddrs. Users should upgrade to version 0.1.0-alpha.9 or later, which requires cryptographic proof of possession of the enrollment private key on credential-releasing endpoints, enforces rate limiting, and validates router lease addresses.		
CVE-2026-99090	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: serialize vblank counter reads against GPU reset The vblank counter callback reads display registers and scanout position without holding reset_domain->sem. During GPU recovery, display commit work can call the callback while reset owns the semaphore for writing. The hardware access then triggers the lockdep_assert_held() in amdgpu_device_skip_hw_access() and sets TAINT_WARN when lockdep is enabled. Try to acquire the reset semaphore for reading before accessing the hardware, and hold it across the counter and scanout-position reads. If reset is already in progress, return DRM's cached software vblank count instead. This avoids blocking behind reset and avoids accessing display hardware while it is being reset. The cached count is less accurate than the hardware count, but is used only when reset prevents a safe hardware read. This addresses the reset-time warning from this vblank path; it does not suppress or fix unrelated lockdep warnings. (cherry picked from commit 51a4359f7431a4be5e42803c8bdaf7e0e5be2934)	N/A	More Details
CVE-2026-99092	In the Linux kernel, the following vulnerability has been resolved: riscv: patch: fix handling of bpf-jit execmem addresses Commit 8718e5a3090b ("riscv: patch: skip fixmap mapping when kernel text is already writable") gated the core_kernel_text() branch of patch_map() on CONFIG_STRICT_KERNEL_RWX, and explicitly left the vmalloc branch unchanged. That branch has a separate problem. That branch only creates a temporary writable fixmap alias when CONFIG_STRICT_MODULE_RWX is enabled, and otherwise assumes the target is directly writable and returns the address unchanged. That assumption no longer holds: bpf_prog_pack_alloc calls set_memory_rox() on every pack unconditionally in alloc_new_pack(), independently of any CONFIG_STRICT_*_RWX option, so BPF text in the vmalloc area is read-only regardless of what CONFIG_STRICT_MODULE_RWX says. With CONFIG_STRICT_MODULE_RWX=n, patch_map() returns the read-only address directly, the subsequent copy_to_kernel_nofault() takes a page fault and returns -EFAULT. bpf_arch_text_copy() turns that into -EINVAL, which propagates through bpf_jit_binary_pack_finalize() to the WARN_ON() in bpf_int_jit_compile() and leaves the program un-JITed. Note that BPF cannot be fixed the way kprobes was in commit bdc46e507b59 ("riscv: mm: make EXECMEM_KPROBES writable without CONFIG_STRICT_MODULE_RWX"). EXECMEM_BPF is already PAGE_KERNEL, i.e. writable at allocation time; the read-only mapping is established afterwards by generic code in alloc_new_pack(), which arch code cannot influence. The only place this can be handled is patch_map(). This is the same problem that was fixed on arm64 by commit b1480ed230ac ("arm64: patching: fix handling of execmem addresses"), and the fix is the same: CONFIG_EXECMEM is what actually tracks whether the vmalloc area can hold executable memory that needs a temporary alias to be written. CONFIG_BPF_JIT, CONFIG_KPROBES and CONFIG_MODULES all select it. Note that this is not limited to CONFIG_MODULES=n. Unlike arm64, riscv selects CONFIG_ARCH_OPTIONAL_KERNEL_RWX, so CONFIG_STRICT_MODULE_RWX can be disabled with CONFIG_MODULES=y as well, and the bug is reachable in that configuration too. The !CONFIG_MMU build is unaffected: it has its own __patch_insn_set() and __patch_insn_write() implementations and never calls patch_map().	N/A	More Details
CVE-2026-114112	Sharp before 0.35.4 bundles a vulnerable libheif dependency that allows attackers to achieve possible remote code execution on glibc-based Linux by supplying crafted images. Attackers can submit malicious AVIF/HEIF input to applications processing untrusted images with sharp, triggering libheif decoding flaws to execute code.	N/A	More Details
CVE-2026-106053	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-99093	In the Linux kernel, the following vulnerability has been resolved: mm: don't schedule deferred kernel page table freeing while booting Booting with a boot-time function tracer and a filter, for example ftrace=function ftrace_filter=pud_free_pmd_page panics on 7.3-rc4 as soon as the tracer starts: [23.531178] Starting tracer 'function' [23.675800] Oops: general protection fault, probably for non-canonical address 0xdffffc0000000038: 0000 [#1] SMP KASAN NOPTI [23.819917] KASAN: null-ptr-deref in range [0x00000000000001c0-0x00000000000001c7] [23.964025] CPU: 0 UID: 0 PID: 0 Comm: swapper Not tainted 7.3.0-rc4-fe2ec83746e5-with-fixes-v2+ #195 PREEMPT(undef) [24.252248] RIP: 0010: __queue_work+0xab/0xf0 [25.981629] Call Trace: [26.125727] <TASK> [26.413912] ? pagetable_free_kernel+0x20/0x120 [26.990283] queue_work_on+0x97/0xf0 [27.134382] __cpa_collapse_large_pages+0x501/0x6f0 [27.566662] cpa_flush+0x394/0x620 [27.998953] change_page_attr_set_clr+0x321/0x4a0 [29.151729] set_memory_rox+0xa2/0xf0 [29.584018] create_trampoline+0x431/0x6f0 ... [44.343347] Kernel panic - not syncing: Attempted to kill the idle task! The boot-time tracer is started from early_trace_init(), which runs before workqueue_init_early(). Making its trampoline read-only splits a large page, and CPA collapses it again right away. The split table has been a kernel page table since commit 9e4a3ec3411b ("x86/mm/pat: Allocate split page	N/A	More Details

	tables as kernel page tables"), so the collapse frees it through pagetable_free_kernel(), which queues work on system_percpu_wq - still NULL at that point. That commit is correct in itself; it only lets CPA reach pagetable_free_kernel() before the workqueue that function relies on exists. Keep putting the table on the list, but don't schedule the work while the system is still booting. A core_initcall schedules it once to free whatever was queued by then.		
CVE-2026-116186	OpenRemote core (openremote-manager) before 1.31.0 does not perform realm or gateway-descendant authorization on inbound edge gateway events and maps event asset IDs only through the reversible GatewayConnector.mapAssetId transform. An attacker controlling a registered gateway can craft pre-image IDs in AttributeEvent or DELETE AssetEvent messages to overwrite attributes on, or delete, assets in any realm, breaking tenant isolation.	N/A	More Details
CVE-2026-113788	Lagoon is an application delivery platform for Kubernetes. Versions prior to 2.33.1 have a vulnerability in the `ssh` service in Lagoon core that would allow an attacker to execute arbitrary code which could lead to, among other things, elevated administrative permissions. It is mitigated by the fact that an attacker must have a Lagoon user account. Exploitation may be detected by checking the container logs of the `ssh` service for: `too many arguments`. The vulnerability was fixed in Lagoon v2.33.2, which is included in the `lagoon/lagoon-core` helm chart v1.62.2. As a workaround, the `ssh` service in Lagoon core can be disabled. The `ssh-portal` service and components are not effected and can continue to be used.	N/A	More Details
CVE-2026-99094	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btintel_pcie: fix TX descriptor bounds check btintel_pcie_prepare_tx() uses tfd_index to access the TFD and data buffer arrays before the index is advanced modulo txq->count. An index equal to txq->count is one past the end of both arrays, so reject it in the bounds check.	N/A	More Details
CVE-2026-116164	Parse Server before 8.6.100 and 9.0.0 before 9.10.4-alpha.3 does not reject reserved internal field names in client object writes, allowing clients to store or clear internal fields. Attackers can write a reserved field that makes reads of an object fail on MongoDB, disrupting queries on that class for all users, or clear their own account lockout state and password history to bypass lockout and password policies.	N/A	More Details
CVE-2026-99096	In the Linux kernel, the following vulnerability has been resolved: io_uring/register: don't exempt disabled rings from task restrictions A ring created with IORING_SETUP_R_DISABLED is exempt from its register restrictions until it is enabled. This lets the creator register the restrictions and then enable the ring with IORING_REGISTER_ENABLE_RINGS, whether or not that opcode is in the allowed set. Since commit ed82f35b926b ("io_uring: allow registration of per-task restrictions"), a ring also inherits restrictions from the task that creates it, and the same exemption applies to those. A task whose register opcodes are restricted can therefore create a disabled ring, issue any register opcode on it, and then enable it. Mark rings whose restrictions were inherited from the task, and keep checking register opcodes for them while the ring is disabled. Enabling such a ring now requires IORING_REGISTER_ENABLE_RINGS to be in the task's allowed set. Per-ring restrictions registered on a disabled ring work as before. [axboe: amend with inheritance flag fixup]	N/A	More Details
CVE-2026-102536	Rejected reason: This CVE has been rejected because it was issued to a development-only issue with no affected distributed version identified.	N/A	More Details
CVE-2026-99097	In the Linux kernel, the following vulnerability has been resolved: tee: shm: reject zero-sized allocations in tee_dyn_shm_alloc_helper() tee_dyn_shm_alloc_helper() derives nr_pages from a caller-supplied size and passes it to alloc_pages_exact() without checking it. For size == 0 nr_pages is 0, and alloc_pages_exact(0) calls get_order(0), which is documented as undefined and returns BITS_PER_LONG - PAGE_SHIFT. The page allocator then trips its order > MAX_PAGE_ORDER warning and fails the allocation; on a panic_on_warn kernel that ends the boot. This can be triggered by TEE_IOC_SHM_ALLOC with struct tee_ioctl_shm_alloc_data where size is 0. Reject a zero page count, as register_shm_helper() already does for the register path.	N/A	More Details
CVE-2026-99098	In the Linux kernel, the following vulnerability has been resolved: dm-integrity: fix buffer overflow in inline mode with large tag size In inline mode, dm_integrity_check and dm_integrity_inline_recheck compute the checksum of each block into an on-stack buffer of HASH_MAX_DIGESTSIZE bytes. integrity_sector_checksum pads its result with zeroes up to the tag size, so if the tag size is larger than HASH_MAX_DIGESTSIZE, the padding runs past the end of the buffer. With CONFIG_STACKPROTECTOR the kernel panics in dm_integrity_check on the first read. Enlarge both buffers to hold MAX_TAG_SIZE bytes, as commit b93b6643e9b5 ("dm integrity: fix a crash with unusually large tag size") did for integrity_metadata. Note that this is not a security bug because the buffer overflow can only be triggered by root.	N/A	More Details
CVE-2026-	OpenRemote before 1.31.0 contains an authorization bypass vulnerability that allows realm admins to impersonate gateways in other realms because GatewayService only checks the gateway ID, not the	N/A	More

116184	realm. Attackers with WRITE:ADMIN in one realm can create a service user named after another realm's gateway ID to overwrite its descendant assets and open SSH tunnels into its local network.		Details
CVE-2026-114143	toml-node is a TOML parser for Node.js and the browser. Prior to 5.0.1, toml.parse() repeatedly slices and rejoins every path prefix in lib/compiler.js while processing long dotted keys, table headers, array-of-tables headers, and inline-table keys in assign(), deepRef(), and setNestedKey(). The resulting quadratic work can block the Node.js event loop for several seconds, and sufficiently long paths can terminate the process through memory exhaustion. The maxDepth option does not limit key-segment count and therefore does not mitigate this path. This issue is fixed in version 5.0.1.	N/A	More Details
CVE-2026-114142	toml-node is a TOML parser for Node.js and the browser. Prior to 5.0.1, toml.parse() can write attacker-controlled properties to Object.prototype because quoteDottedString() in lib/compiler.js produces the same bookkeeping string for distinct quoted key paths, allowing addTableArray() to clear another key's guard. The assign() path then descends through an unowned Date value without applying the owned-container check and can reach the prototype chain. Objects created elsewhere in the process can inherit the injected property. This issue is fixed in version 5.0.1.	N/A	More Details
CVE-2026-114141	toml-node is a TOML parser for Node.js and the browser. Prior to 5.0.1, the toml.parse() compiler's addTableArray() function in lib/compiler.js rescans the complete assignedPaths and valueAssignments bookkeeping sets for every array-of-tables header. Distinct [[name]] paths are not pruned, so a document with many such headers causes quadratic synchronous work and can block the Node.js event loop, denying service to other operations in the process. This issue is fixed in version 5.0.1.	N/A	More Details
CVE-2026-114140	The nbconvert tool, jupyter nbconvert, converts Jupyter notebooks to various other formats via Jinja templates. From 5.4.0 until 7.17.2, ExtractOutputPreprocessor in nbconvert/preprocessors/extractoutput.py joins the attacker-controlled cell output metadata filename to output_files_dir without limiting the result to the intended directory. When a victim converts a malicious notebook with the markdown, LaTeX, reStructuredText, or AsciiDoc exporter, path traversal in the output filename can write attacker-controlled output data to arbitrary locations accessible to the victim. Running the notebook cell again replaces the malicious output metadata, but merely converting the supplied notebook does not. This issue is fixed in version 7.17.2.	N/A	More Details
CVE-2026-116372	MISP mispguard is a security filtering proxy that inspects and blocks MISP feed responses based on configured rules. The process_feed_response method in src/mispguard.py was responsible for evaluating feed responses before they reached the MISP instance. The original logic only explicitly handled the case where the HTTP status code was not 200 by passing the response through without filtering, on the assumption that non-200 responses would not carry ingestible event data. However, HTTP 2xx status codes other than 200 (such as 201 Created, 203 Non-Authoritative Information, and 206 Partial Content) were also passed through unfiltered. Because MISP may still ingest content from these responses, a malicious or compromised feed host could return a 203 or 206 status code with arbitrary content, effectively bypassing all of mispguard's filtering rules. This allowed unfiltered or malicious feed data to be ingested into the MISP instance, undermining the security controls that mispguard is designed to enforce. The vulnerability requires the attacker to control or compromise a feed endpoint configured in the MISP instance. No authentication or user interaction is required beyond the feed polling mechanism already in place.	N/A	More Details
CVE-2026-105937	Apache Pulsar shadow topics (PIP-180) are topics that read the stored messages of a source topic from that topic's managed ledger. The source topic is selected by setting the reserved topic property PULSAR.SHADOW_SOURCE. The admin REST API accepts this property when creating a non-partitioned or partitioned topic and when updating the properties of an existing topic. On these paths the broker authorizes the caller only against the topic being created or updated. It does not authorize the caller against the source topic named in the property. The existing authorization for shadow topic configuration (the SHADOW_TOPIC topic policy on the source topic) is applied only to the source-side shadow topics admin API, not to these paths. As a result, an authenticated user who is allowed to create topics or update topic properties in one tenant can bind a topic in that tenant to a persistent topic in any other tenant. With the default authorization provider, these operations require the tenant admin role for the user's own tenant. The user does not need any permission on the source topic, its namespace or its tenant. This impact analysis assumes that Pulsar has been configured with the default authorization provider. For custom authorization providers, the impact could be different. The user can then consume from the shadow topic, using only permissions on their own topic, and receive retained messages that were published to the source topic in the other tenant. Exposure is limited to messages that are still retained and readable in the source topic's ledgers; messages protected with client-side end-to-end encryption cannot be decrypted without the encryption keys. A producer on the shadow topic can also register a new schema version under the source topic's schema, subject to the schema compatibility settings; existing schema versions are not replaced. In addition, the shadow topic's managed ledger applies its own retention, trimming, offloading and deletion to the source topic's ledgers. This can permanently delete the source topic's ledger data from BookKeeper and from tiered storage, regardless of the source topic's own subscriptions and retention policy, including messages	N/A	More Details

	that have not yet been acknowledged. Deleting the shadow topic can also delete the source topic's ledgers. This data loss can also affect shadow topics that were configured intentionally by administrators. This issue affects Apache Pulsar versions from 3.0.0 before 4.0.14 and from 4.1.0 before 4.2.5. Apache Pulsar 2.x releases do not contain the shadow topic implementation. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. In the fixed versions, shadow topics are disabled by default and are controlled by the new broker configuration setting enableShadowTopics (default false). While it is disabled, the broker rejects creating shadow topics, setting the PULSAR.SHADOW_SOURCE property, and loading or replicating to existing shadow topics. Deployments that use shadow topics must set enableShadowTopics=true in the broker configuration of all brokers after upgrading; otherwise existing shadow topics will not be loaded. When shadow topics are enabled, binding a shadow topic to a source topic in a different tenant requires the superuser role. The fixed versions also prevent a shadow topic from deleting or trimming the ledgers and offloaded data of its source topic. Affected versions provide no configuration option to disable shadow topics. Users who cannot upgrade immediately can reduce exposure by granting the tenant admin role only to trusted principals.		
CVE-2026-105938	Apache Pulsar namespaces can restrict which roles may use a subscription, either with the Prefix subscription authentication mode, which requires a subscription name to start with the role name, or by granting permission on a subscription to specific roles. Topic-level subscription operations, such as consuming, skipping messages, resetting a cursor or deleting a subscription, apply these subscription policies. Several admin REST API operations that act on subscriptions did not apply them: - the namespace-level and bundle-level clear backlog operations, both for a named subscription and for all subscriptions; - the namespace-level and bundle-level unsubscribe operations; - the topic-level operation that expires messages for all subscriptions of a topic. With the default authorization provider, the namespace-level operations require only consume permission granted on the namespace, and the expire-all operation requires only consume permission on the topic. As a result, when authentication and authorization are enabled and a namespace uses the Prefix subscription authentication mode or subscription permissions, an authenticated user with such consume permission can clear the backlog of, expire messages of, or delete a subscription that the policies would deny to that user. Tenant admin or superuser access is not required. The clear backlog operations can also be applied to the cursors used for geo-replication, so that messages not yet replicated to a remote cluster are skipped. Deployments that do not configure the Prefix subscription authentication mode or subscription permissions (the defaults) are not affected by this bypass, because consume permission on the namespace or topic already allows these operations there. This impact analysis assumes that Pulsar has been configured with the default authorization provider. For custom authorization providers, the impact could be different. The impact is on the integrity of subscription state. Clearing or expiring the backlog moves the subscription's cursor forward, so unacknowledged messages are no longer delivered to the subscription's consumers. The unsubscribe operation deletes a subscription that has no connected consumers. Skipped messages can be permanently deleted once no other subscription or the namespace's retention policy keeps them. The issue does not allow reading message contents, changing permissions or policies, or accessing other namespaces, and it does not affect the availability of the broker. In fixed versions, the namespace-level and bundle-level operations for a named subscription apply the subscription policies to that subscription, and the operations for all subscriptions, when requested by a user who is not a tenant admin or superuser, only act on the subscriptions that the user may access. Clearing the backlog of geo-replication cursors requires tenant admin or superuser access. Topic-level subscription operations apply the subscription policies for all types of authentication data, including that of the anonymous user role. This issue affects Apache Pulsar versions from 2.8.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. All brokers in a cluster must be upgraded before relying on the namespace-level restrictions, because brokers that have not been upgraded do not apply them when they receive the internal per-bundle requests. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/. Users who cannot upgrade immediately can reduce exposure by restricting network access to the broker's admin REST API to trusted administrative clients.	N/A	More Details
CVE-2026-113795	pypdf is a free and open-source pure-python PDF library. Prior to 6.20.0, accessing a crafted PDF stream with a large number of filters or multiple filters that pass around large amounts of data can cause large memory consumption. This issue is fixed in version 6.20.0.	N/A	More Details
	When a configured proxy role sends a request to the broker on behalf of another user (the original principal), the broker is designed to authorize both the proxy role and the original principal and to allow the request only if both are authorized, so that the original principal cannot exceed its own permissions. Two flaws weakened this check. First, the original principal was authorized with the proxy's		

CVE-2026-105942	authentication data instead of its own. On the admin REST API of the broker and the functions worker, this applied to every proxied request. On the Pulsar binary protocol, it applied when the proxy did not forward the client's credentials to the broker, which was the default of the Pulsar proxy in the affected versions. Authorization providers that resolve permissions from the role name only, such as the default authorization provider, are not affected by this flaw. The multi-role token authorization provider (MultiRolesTokenAuthorizationProvider), added in Apache Pulsar 2.8.1, resolves roles from the role claim of the token in the authentication data it is given. When a proxy authenticates to the broker with its own token, the original principal was therefore authorized with all roles in the proxy's token. If any of those roles is a superuser or tenant admin role, any user whose requests are forwarded by that proxy can perform operations that need those permissions without holding them. Custom authorization providers that use the authentication data could be affected in a similar way. Second, independently of the authorization provider, the tenant admin check of a proxied admin REST API request succeeded without checking the original principal when the proxy role was itself an admin role of the tenant. In deployments where a proxy role is a tenant admin, including with the default authorization provider, any user whose requests are forwarded by that proxy could perform tenant administration operations on that tenant. This behavior is present from 2.8.4 in the 2.8.x releases, from 2.9.3 in the 2.9.x releases, and from 2.10.1 onwards. Deployments that use neither the multi-role token authorization provider nor a custom provider that uses the authentication data are not affected by the first flaw, and with the multi-role token authorization provider, neither are deployments whose proxies do not authenticate to the broker with a token. Deployments in which no proxy role is listed as an admin role of a tenant are not affected by the second flaw. A proxy role that is a superuser but is not listed as a tenant admin does not trigger it, because the tenant admin check of the built-in authorization providers only considers the tenant's admin roles; this can differ for custom authorization providers that treat superusers as tenant admins. In fixed versions, both the proxy role and the original principal are authorized for each proxied request, and the original principal is authorized only with its own authentication data. With the multi-role token authorization provider, the roles of a forwarded token are used for the original principal only when that token authenticates as the original principal; otherwise only the original principal's role name is used. For tenant admin checks, the proxy role and the original principal must each be a superuser or a tenant admin of the tenant. This issue affects Apache Pulsar versions from 2.8.1 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/.	N/A	More Details
CVE-2026-105944	Apache Pulsar Manager is a web console for administering Apache Pulsar clusters. It has its own user accounts and login, and it can be configured with an HTTP authentication credential, such as a token, that it uses to call the Pulsar admin API on behalf of its users. Pulsar Manager has the following security issues: - The operation for creating the initial administrator (superuser) account is available without authentication, and it remains available after the instance has been initialized. Anyone who can reach Pulsar Manager can create an additional administrator account at any time and log in with it. This issue affects versions 0.3.0 and later. - Requests to the Pulsar admin API that Pulsar Manager proxies to the configured Pulsar cluster are not subject to Pulsar Manager's authentication. Version 0.2.0 checks the caller's permissions for some of these requests, but not for all of them. An unauthenticated user can therefore use Pulsar Manager to send requests to the Pulsar admin API, which the cluster processes with the credential configured in Pulsar Manager. - The destination of a proxied request can be chosen by the requester, and Pulsar Manager attaches its configured Pulsar admin credential to the request. An unauthenticated user can make Pulsar Manager send this credential to a host of the user's choosing, and the credential can then be used directly against the Pulsar cluster. - In versions 0.3.0 and later, Pulsar Manager does not distinguish between privilege levels of its accounts. Every account, including one created through the first issue, effectively has full administrative privileges in Pulsar Manager. The impact depends on the version and the deployment. In versions 0.3.0 and later, any user who can reach Pulsar Manager over the network can take full control of Pulsar Manager. In all versions, when Pulsar Manager is configured with an HTTP authentication credential for a Pulsar cluster, any user who can reach Pulsar Manager can perform the Pulsar admin operations that this credential permits and can obtain the credential itself. When Pulsar Manager is configured with the token signing key of a Pulsar cluster, an administrator of Pulsar Manager can create tokens for the cluster. Apache Pulsar Manager is no longer maintained and is completely unsupported. The Apache Pulsar PMC has retired it. There are no supported versions, and no further releases, fixes or security updates will be made for any version. These issues will not be fixed. Using Pulsar Manager is not recommended. Users should stop using Pulsar Manager, remove existing deployments, and move to one of the alternatives described in the Administration UI documentation (https://pulsar.apache.org/docs/administration-ui). Users who have exposed Pulsar Manager configured with a Pulsar credential or token signing key should also consider replacing that credential or key. Until a deployment is removed, users should not expose Pulsar Manager to untrusted networks, and should restrict network access to it to trusted administrators.	N/A	More Details
	HTTP::Server::PSGI versions before 1.0056 for Perl allow a single request to crash the server by		

CVE-2026-115825	declaring a large Content-Length. The request body is read with sysread() calls sized by the remaining Content-Length. Perl grows the read buffer to the requested length before any data arrives, so the server tries to allocate as much memory as the client declares. A remote attacker can send one request with a large Content-Length and a few body bytes, and perl exits with an out-of-memory error before the application is called. HTTP::Server::PSGI runs as a single process, so the server stays down until restarted.	N/A	More Details
CVE-2026-115824	Starman versions before 0.4019 for Perl allow memory exhaustion via unbounded buffering of chunked request bodies. The chunked body decoder keeps each chunk in memory until the whole chunk and its CRLF have arrived, and sets no limit on the length of a chunk-size line. A client that declares a huge chunk size and streams data, or sends a chunk-size line that never ends, grows the buffer for as long as it keeps sending. A remote attacker can exhaust worker memory with a single connection.	N/A	More Details
CVE-2026-115823	Starman versions before 0.4019 for Perl send response bodies for HEAD requests, desynchronizing responses on keep-alive connections. For a HEAD request Starman omits chunked framing but still writes the body returned by the application, on both the array and streaming response paths. Responses with a 1xx, 204 or 304 status also carry the body. On a keep-alive or pipelined connection the client reads those bytes as the start of the next response. A front-end proxy that reuses upstream connections may then pair responses with the wrong requests.	N/A	More Details
CVE-2026-115822	Starman versions before 0.4019 for Perl allow CPU exhaustion via quadratic regex backtracking when parsing chunked request bodies. Starman decodes chunked request bodies before calling the application, matching each chunk-size line with <code>`s/^([0-9a-fA-F]+).*(\015\012)//`</code> . A line of hex digits followed by a bare LF and then CRLF never matches, because <code>`.</code> does not match LF, and each failed attempt backtracks in time quadratic in the length of the hex run. Nothing limits the length of the line. A remote attacker can exhaust worker CPU by sending such requests to any route.	N/A	More Details
CVE-2026-115817	Starman versions before 0.4019 for Perl allow a stalled request body to hold a worker indefinitely, as <code>read_timeout</code> covers only the headers. Content-Length and chunked request bodies are read with blocking reads and no deadline. A client that sends complete headers and then stalls or trickles the body keeps the worker waiting until it disconnects. A remote attacker can take every worker with a few such connections, five with the default worker count, so the server stops answering requests. Version 0.4019 applies <code>read_timeout</code> to each wait for body data and adds a <code>body_timeout</code> limit, 300 seconds by default, on the total time to read a body.	N/A	More Details
CVE-2026-105945	The Pulsar Packages Management Service can be configured to store packages on the local filesystem of the broker by using the filesystem storage provider. Package names consist of a type, tenant, namespace, name and version. Authorization of package operations is based only on the tenant and namespace, while the storage location of a package is derived from all of these components. The version component of a package name was not validated, so it could contain URL-encoded path separators and parent directory references. The filesystem storage provider resolved such a version to a location belonging to a package in a different tenant or namespace. Because that location is inside the configured storage directory, the containment check that the filesystem storage provider performs in recent releases did not prevent the access. An authenticated user who has the packages permission on any namespace could use this to download, overwrite or delete packages stored in other tenants or namespaces for which the user has no permission. Packages overwritten or deleted in this way may be used by Pulsar Functions, Sources or Sinks, so the issue can affect the integrity and availability of these components. If authentication is disabled, the operations are available to any client that can reach the broker's admin REST API. This issue only affects deployments in which package management is enabled with the broker setting <code>enablePackagesManagement</code> and the filesystem storage provider (<code>FileSystemPackagesStorageProvider</code>) is configured with the setting <code>packagesManagementStorageProvider</code> . Package management is disabled by default, and the default storage provider, which stores packages in Apache BookKeeper, is not affected. This issue affects Apache Pulsar versions from 2.10.0 before 4.0.14 and from 4.1.0 before 4.2.5. - 4.0.x version users should upgrade to at least 4.0.14. - 4.1.x and 4.2.x version users should upgrade to at least 4.2.5. Users operating versions prior to those listed above should upgrade to the aforementioned patched versions or newer versions. The currently supported Apache Pulsar release lines are 4.0.x and 5.0.x. The supported versions schedule is available at https://pulsar.apache.org/contribute/release-policy/ .	N/A	More Details
	Vaadin Charts renders charts with the Highcharts library and pins an exact version of it, 9.2.2 in Vaadin 22 to 24 and 12.2.0 in Vaadin 25, so an application cannot remove these vulnerabilities by upgrading Highcharts on its own. Highsoft disclosed several vulnerabilities in that library that are reachable through the chart configuration: writes to <code>Object.prototype</code> through <code>dotted.series.keys</code> paths, and cross-site scripting through <code>credits.href</code> , through <code>breadcrumbs.style</code> , through HTML text parsing while printing, and through <code>xlink:href</code> attributes in exported charts. An application is affected where data an attacker controls reaches the chart configuration, for example through <code>additionalOptions</code> , through <code>updateConfiguration()</code> , or through text options that accept markup. The fixed Vaadin versions apply Highsoft's recommended workarounds to the bundled Highcharts when it loads, and keep the same		

CVE-2026-107588	Highcharts version. Users of affected versions should apply the following mitigation or upgrade. Releases that have fixed this issue include: Product version Vaadin 22.0.0 - 23.6.15 Vaadin 24.0.0 - 24.4.23 Vaadin 24.5.0 - 24.7.15 Vaadin 24.8.0 - 24.9.22 Vaadin 24.10.0 - 24.10.10 Vaadin 25.0.0 - 25.2.9 Vaadin 25.3.0 Mitigation Upgrade to 23.6.16 Upgrade to 24.4.24 Upgrade to 24.7.16 Upgrade to 24.9.23 Upgrade to 24.10.11 Upgrade to 25.2.10 Upgrade to 25.3.1 or newer Please note that Vaadin versions 10-13 and 15-22 are no longer supported and you should update either to the latest 23, 24, 25 version. Artifacts Maven coordinates / Vulnerable versions / Fixed version com.vaadin:vaadin / 22.0.0 - 23.6.15 / >=23.6.16 com.vaadin:vaadin / 24.0.0 - 24.4.23 / >=24.4.24 com.vaadin:vaadin / 24.5.0 - 24.7.15 / >=24.7.16 com.vaadin:vaadin / 24.8.0 - 24.9.22 / >=24.9.23 com.vaadin:vaadin / 24.10.0 - 24.10.10 / >=24.10.11 com.vaadin:vaadin / 25.0.0 - 25.0.13 / >=25.2.10 com.vaadin:vaadin / 25.1.0 - 25.1.13 / >=25.2.10 com.vaadin:vaadin / 25.2.0 - 25.2.9 / >=25.2.10 com.vaadin:vaadin / 25.3.0 / >=25.3.1 com.vaadin:vaadin-charts-flow / 22.0.0 - 23.6.15 / >=23.6.16 com.vaadin:vaadin-charts-flow / 24.0.0 - 24.4.22 / >=24.4.24 com.vaadin:vaadin-charts-flow / 24.5.0 - 24.7.14 / >=24.7.15 com.vaadin:vaadin-charts-flow / 24.8.0 - 24.9.22 / >=24.9.23 com.vaadin:vaadin-charts-flow / 24.10.0 - 24.10.10 / >=24.10.11 com.vaadin:vaadin-charts-flow / 25.0.0 - 25.0.13 / >=25.2.10 com.vaadin:vaadin-charts-flow / 25.1.0 - 25.1.13 / >=25.2.10 com.vaadin:vaadin-charts-flow / 25.2.0 - 25.2.9 / >=25.2.10 com.vaadin:vaadin-charts-flow / 25.3.0 / >=25.3.1 com.vaadin:vaadin-charts-flow-svg-generator / 21.0.0 - 23.6.15 / >=23.6.16 com.vaadin:vaadin-charts-flow-svg-generator / 24.0.0 - 24.4.22 / >=24.4.24 com.vaadin:vaadin-charts-flow-svg-generator / 24.5.0 - 24.7.14 / >=24.7.15 com.vaadin:vaadin-charts-flow-svg-generator / 24.8.0 - 24.9.22 / >=24.9.23 com.vaadin:vaadin-charts-flow-svg-generator / 24.10.0 - 24.10.10 / >=24.10.11 com.vaadin:vaadin-charts-flow-svg-generator / 25.0.0 - 25.0.13 / >=25.2.10 com.vaadin:vaadin-charts-flow-svg-generator / 25.1.0 - 25.1.13 / >=25.2.10 com.vaadin:vaadin-charts-flow-svg-generator / 25.2.0 - 25.2.9 / >=25.2.10 com.vaadin:vaadin-charts-flow-svg-generator / 25.3.0 / >=25.3.1 npm packages npm package / Vulnerable versions / Fixed version @vaadin/charts / 22.0.0 - 23.6.6 / >=23.6.7 @vaadin/charts / 24.0.0 - 24.4.18 / >=24.4.19 @vaadin/charts / 24.5.0 - 24.7.12 / >=24.7.13 @vaadin/charts / 24.8.0 - 24.9.18 / >=24.9.19 @vaadin/charts / 24.10.0 - 24.10.6 / >=24.10.7 @vaadin/charts / 25.0.0 - 25.0.13 / >=25.2.14 @vaadin/charts / 25.1.0 - 25.1.14 / >=25.2.14 @vaadin/charts / 25.2.0 - 25.2.13 / >=25.2.14 @vaadin/charts / 25.3.0 / >=25.3.1	N/A	More Details
CVE-2026-110803	VaadinService.requestEnd() in Vaadin Framework 8 locks the session and then reads the v-uiled request parameter outside the block that releases the lock. A request carrying a non-numeric v-uiled makes Integer.parseInt throw while the session lock is held, so the lock is never released and the container thread returns to its pool still owning it. Every later request to that session then blocks forever in VaadinService.lockSession(), and repeating the request exhausts the thread pool and takes the application down. No authentication is required: a session cookie, which any client obtains with a single request, is enough. Vaadin 10 and later are not affected, because Flow releases the lock in a finally block and does not read that parameter there, and Vaadin Framework 7 is not affected either. Users of affected versions should apply the following mitigation or upgrade. Releases that have fixed this issue include: Product version Vaadin Framework 8.16.0 - 8.32.1 Mitigation Upgrade to 8.32.2 Vaadin Framework 7 and Vaadin 10 and later do not carry this behaviour and need no upgrade for this issue. Artifacts Maven coordinates / Vulnerable versions / Fixed version com.vaadin:vaadin-server / 8.16.0 - 8.32.1 / >=8.32.2 com.vaadin:vaadin-server-mpr-jakarta / 8.18.0 - 8.32.1 / >=8.32.2	N/A	More Details
CVE-2026-113461	Local privilege escalation in Checkmk 2.5.0p14 allows a local user with write access to the Oracle client library directory of an Oracle home listed in the 'oratab' file or to the 'oratab' file itself to execute code as root if an agent has the 'mk-oracle' plugin enabled.	N/A	More Details
CVE-2026-113794	pypdf is a free and open-source pure-python PDF library. Prior to 6.20.0, an attacker can craft an incomplete PDF that ends inside a standard cross-reference table, causing pypdf to mishandle the unexpected end of file and enter an infinite loop. This issue is fixed in version 6.20.0.	N/A	More Details
CVE-2026-113793	pypdf is a free and open-source pure-python PDF library. Prior to 6.20.1, removing an object from a crafted cyclic PDF tree, such as an outline item, can enter an infinite loop and block the current process. This issue is fixed in version 6.20.1.	N/A	More Details
CVE-2026-114111	sharp before 0.35.5 bundles a libsvg dependency affected by a memory-related vulnerability (CVE-2026-96889) that, under certain runtime-specific conditions, can lead to code execution on glibc-based Linux. An attacker can supply a crafted SVG image that, when decoded by sharp through libsvg, may result in arbitrary code execution.	N/A	More Details
CVE-2026-113792	pypdf is a free and open-source pure-python PDF library. Prior to 6.20.1, calling PdfWriter.find_outline_item on a crafted PDF with a cyclic outline hierarchy can enter an infinite loop and block the current process. This issue is fixed in version 6.20.1.	N/A	More Details
CVE-2026-113811	Algorithm::UrataniTakeda versions from 0.1.5 before 0.1.7 for Perl allow memory exhaustion via a never-cleared failure shift cache in search. The failure shift is cached for each pair of trie state and text character in a field of the matcher object, and nothing clears it. A character not seen before at a given state adds an entry that lasts for the life of the matcher. An attacker who can submit text with many	N/A	More Details

	distinct characters to a long-lived matcher can grow its memory with each search.		
CVE-2026-99088	In the Linux kernel, the following vulnerability has been resolved: net: openvswitch: validate transport header presence in set_ipv6_addr When executing IPv6 address rewrite actions on IPv6 fragments, set_ipv6_addr() calls update_ipv6_checksum(). If parse_ipv6hdr() processes a non-first IPv6 fragment, it sets key->ip.proto to NEXTHDR_FRAGMENT and returns early without calling skb_set_transport_header(). update_ipv6_checksum() unconditionally evaluates skb_transport_offset() on entry before checking l4_proto. Because skb->transport_header is uninitialized, this triggers a warning under CONFIG_DEBUG_NET=y although it is completely harmless. Fix this by returning early in update_ipv6_checksum() if l4_proto is NEXTHDR_FRAGMENT. This avoids reading the uninitialized transport offset for fragments while preserving the debug warning for any other protocol where the transport header is unexpectedly missing. See the syzbot trace: !skb_transport_header_was_set(skb) WARNING: ./include/linux/skbuff.h:3075 at skb_transport_header include/linux/skbuff.h:3075 [inline], CPU#1: syz-executor463/5635 WARNING: ./include/linux/skbuff.h:3075 at skb_transport_offset include/linux/skbuff.h:3250 [inline], CPU#1: syz-executor463/5635 WARNING: ./include/linux/skbuff.h:3075 at update_ipv6_checksum net/openvswitch/actions.c:361 [inline], CPU#1: syz-executor463/5635 WARNING: ./include/linux/skbuff.h:3075 at set_ipv6_addr+0x462/0x660 net/openvswitch/actions.c:399, CPU#1: syz-executor463/5635 [...] RIP: 0010:skb_transport_header include/linux/skbuff.h:3075 [inline] RIP: 0010:skb_transport_offset include/linux/skbuff.h:3250 [inline] RIP: 0010:update_ipv6_checksum net/openvswitch/actions.c:361 [inline] RIP: 0010:set_ipv6_addr+0x462/0x660 net/openvswitch/actions.c:399 [...] Call Trace: <TASK> set_ipv6 net/openvswitch/actions.c:531 [inline] do_execute_actions+0x557e/0x8600 net/openvswitch/actions.c:1366 ovs_execute_actions+0xde/0x520 net/openvswitch/actions.c:1592 ovs_packet_cmd_execute+0xb4f/0xf10 net/openvswitch/datapath.c:705 genl_family_rcv_msg_doit+0x233/0x340 net/netlink/genetlink.c:1114 genl_rcv_msg+0x614/0x7a0 net/netlink/genetlink.c:1209 netlink_rcv_skb+0x226/0x4a0 net/netlink/af_netlink.c:2572 genl_rcv+0x28/0x40 net/netlink/genetlink.c:1218 netlink_unicast+0x7bd/0x940 net/netlink/af_netlink.c:1361 netlink_sendmsg+0x813/0xb40 net/netlink/af_netlink.c:1916 sock_sendmsg_nosec+0x14e/0x190 net/socket.c:800	N/A	More Details
CVE-2026-113791	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.0, an attacker can craft a PDF containing a content stream with many small operations, causing parsing to consume large amounts of memory. This issue is fixed in version 6.21.0.	N/A	More Details
CVE-2026-114139	Sentry-Javascript is an official Sentry SDKs for JavaScript. From 9.36.0 until 9.47.2, 10.75.2, and 11.0.0-rc.1, the webpack integration for @sentry/nextjs skips application middleware for every request path under a configured tunnelRoute, while the tunnel rewrite serves only the exact tunnel path with valid o and p query parameters. When an application route directly collides with the tunnel path or a root dynamic segment resolves under it, an unauthenticated request can reach that route without middleware authorization. Exploitation requires tunnelRoute, a webpack build, a colliding route, and middleware being the only authorization control. This issue is fixed in versions 9.47.2, 10.75.2, and 11.0.0-rc.1.	N/A	More Details
CVE-2026-99095	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: drain the hangcheck timer and works on V3D unbind Nothing stops the hangcheck timer that submitted jobs arm, or the job_done_work that the render-done interrupt queues, at teardown: vc4_irq_disable() cancels only overflow_mem_work, and vc4_gem_destroy() runs from the drm-managed release, after vc4_v3d_unbind() has already uninstalled the V3D interrupt and cleared vc4->v3d. A timer still armed by then reads V3D registers through the NULL vc4->v3d pointer, and the late callbacks run on the vc4_dev embedding them after it has been freed. Drain them in vc4_v3d_unbind(): shut the timer down and cancel reset_work before the interrupt is taken down, because its vc4_irq_reset() re-enables it, then flush job_done_work so that completions queued at teardown still release their jobs. This issue was found by an in-house static analysis tool.	N/A	More Details
CVE-2026-98890	In the Linux kernel, the following vulnerability has been resolved: wifi: libertas: reject short monitor TX frames In monitor mode, lbs_hard_start_xmit() casts skb->data to a radiotap TX header, skips that header, and then copies the 802.11 destination address from offset 4 in the remaining frame. The generic length check only rejects zero-length and oversized skbs, so a short monitor frame can be read past the end of the skb data. Require enough bytes for the radiotap TX header and the destination address field before using the monitor-mode header layout.	N/A	More Details
CVE-2026-116162	Parse Server before 8.6.104 and 9.0.0 before 9.10.5-alpha.5 saves the entire login request with master privileges when issuing an SMS multi-factor authentication verification code. An SMS MFA-enrolled user can include extra fields in that request to modify master-key-restricted fields on their own account, such as the email-verification flag or beforeSave-protected values.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: namespace: refuse an automount below a mount that is in no namespace It's possible to add autmounts even when the parent mount isn't in the mount namespace of the caller. The only requirement we have is that the parent mount		

CVE-2026-98979	namespace must not be NULL, i.e., unmounted. Problem is that clone_private_mount() has MNT_NS_INTERNAL which makes that trivially true. So that passes the test and attach_recursive_mnt() accepts that as a mount point and funny enough, count_mounts() dereferences MNT_NS_INTERNAL. The problem is it is an error pointer... So we can reach this in userspace via fanotify. A filesystem mark on the lower filesystem of an overlay reports paths on the layer clone and reading the event hands out a descriptor on it. For example with debugfs as the lower layer it goes kaboom: openat(evfd, "tracing", O_DIRECTORY) Oops: general protection fault KASAN: null-ptr-deref in range [0x1d0-0x1d7] RIP: 0010:count_mounts+0x35/0x200 attach_recursive_mnt finish_automount And since that sleeping beauty happens under namespace_sem held for writing every mount operation on the system blocks from then on. Congrats. Use is_mounted() instead which rejects unmounted and internal mounts alike. The open fails with EINVAL just as it did before clone_private_mount() used MNT_NS_INTERNAL.	N/A	More Details
CVE-2026-98972	In the Linux kernel, the following vulnerability has been resolved: cgroup/cpuset: Don't access cpuset_cgrp_subsys.root in is_in_v2_mode() After seeing the patch [1] to guard is_in_v2_mode() with RCU, it makes me realize that is_in_v2_mode() may be called in a context where a new cgroup filesystem is being rebound with stale cpuset_cgrp_subsys.root pointer. Avoid this potential UaF situation by adding a new cpuset_v2_mode flag which is set when the cpuset_v2_mode mount option is used. This flag is written into only when cpuset_bind() is being called with a stable cpuset_cgrp_subsys.root value. The is_in_v2_mode() helper is modified to read the new cpuset_v2_mode flag instead of accessing cpuset_cgrp_subsys.root directly. [1] https://lore.kernel.org/lkml/20260929084124.626693-2-arighi@nvidia.com	N/A	More Details
CVE-2026-98973	In the Linux kernel, the following vulnerability has been resolved: drm/bridge: Fix unlocked list_del in drm_bridge_add() When re-adding a bridge that was previously removed, drm_bridge_add() drops it from bridge_lingering_list without holding bridge_lock. Both bridge_list and bridge_lingering_list are protected by bridge_lock, as they are concurrently modified by drm_bridge_remove() and __drm_bridge_free(), and walked by the debugfs 'bridges' file. Running the list_empty() test and the list_del_init() outside of the lock may therefore corrupt either list. Perform both under bridge_lock.	N/A	More Details
CVE-2026-89317	gose parsed the JWS `crit` header into `JwsHeader.Crit` and then never read it. RFC 7515 §4.1.11 requires a recipient to reject a JWS whose listed critical extension parameters it does not understand. gose implements no header extensions, so every token carrying a non-empty `crit` should have been rejected — and none were.	N/A	More Details
CVE-2026-89316	`LoadPublicKey` resolved an EC key's curve with an unchecked type assertion, `algToOptsMap[alg].(*ECDSAOptions)`. A JWK declaring `"kty":"EC"` together with a non-EC `"alg"` — which RFC 7517 §4.4 permits — reaches that assertion and panics .	N/A	More Details
CVE-2026-89315	`JwtClaims.MarshalJSON` built a Go struct at runtime with `reflect.StructOf`, synthesising a field named `"A" + claimName` for each untyped claim. `reflect.StructOf` panics when a field name is not a valid Go identifier, or when it duplicates another field. Both conditions are reachable from ordinary JWT claim names.	N/A	More Details
CVE-2026-98977	In the Linux kernel, the following vulnerability has been resolved: namespace: handle mount locking for automounts correctly If mounts are propagated across user namespaces, attach_recursive_mnt() locks every copy of the source mount to protect overmounts from vanishing and revealing the underlying files or directories. The user namespace is taken from the caller's mount namespaces since this is where the mounts end up. Except, that's not always true. Automounts may legitimately get popped in by tasks located in a different mount and user namespace during path lookup. Then check doesn't make sense at that point. The copy in the namespace of the parent - which may be the host's - is now locked and the host cannot change the flags of its own mounts anymore. Here's the reproducer: - host has debugfs mounted nosuid,nodev,noexec and shared - tracefs automount below it is not yet active - hand a child process a directory descriptor on that mount - child process enters auser and mount namespace - child process' namespace now holds a locked copy of the host's debugfs mount that receives propagation from it - child process stats "tracing/" through the descriptor - lookup runs on the host's mount so the automount lands below the host's mount - propagation puts a copy below the child's copy - both try to clear the flags on the mount they got, with a bind remount: host, on its own automount: MS_REMOUNT MS_BIND = EPERM child, on the copy in its namespace: MS_REMOUNT MS_BIND = 0 So the lock landed on the host's mount instead of the child's copy. Congrats. So we need to compare with the owner of the namespace the mount actually gets mounted on. For all regular cases that is the caller's mount namespace and so nothing changes. Detached trees in anonymous mount namespaces by be handed over via SCM_RIGHTS or inherited in other ways on purpose so the attaching task's mount namespace is authoritative, not the creator of the detached tree.	N/A	More Details
CVE-2026-89314	`jwkFields.CheckConsistency` detected duplicate `key_ops` entries with a pairwise nested loop, making validation cost grow with the square of the entry count. Because the check runs on every JWK and JWKS unmarshal, a modestly sized but hostile JWKS document could occupy a CPU core for many seconds.	N/A	More Details

CVE-2026-89313	`JwksTrustStore.Get` assigned each parsed key to the function's **named return value** while refreshing the JWKS cache. When the requested `kid` was not present, the trailing naked `return` handed the caller the last key in the JWKS document together with a `nil` error — a confidently wrong answer instead of "not found".	N/A	More Details
CVE-2026-98989	In the Linux kernel, the following vulnerability has been resolved: namespace: check a recursive bind mount for mount namespace loops do_loopback() refuses to bind mount the file of a mount namespace that is as old as the caller's or older because a mount namespace that holds a mount of its own file, or of an ancestor's, would cause a cycle. But it only checks the dentry the bind mount starts from. With MS_REC everything below it is copied and do_loopback() passes CL_COPY_MNT_NS_FILE so mount namespace files below the source are copied. That's fine for a source in the caller's own mount namespace. Every mount namespace file in there passed the same check when it was bind-mounted. It isn't fine for a source in another mount namespace. may_copy_tree() accepts a bind mount of any nsfs or pidfs file no matter what mount namespace it lives in so that /proc/<pid>/ns/<ns> and pidfds can be bind mounted. If we create a bind-mount stack of mount namespaces file descriptors the kernel will copy them irrespective of their ancestral relationship to the mount namespace in question: 151 146 0:7 net:[4026531833] /tmp/nrc/y rw - nsfs nsfs rw 152 151 0:7 mnt:[4026532293] /tmp/nrc/y rw - nsfs nsfs rw Mount 152 is a mount of the mount namespace file of mount namespace 4026532293 inside mount namespace 4026532293. That namespace and every mount in it are leaked... All it takes is a process in an older mount namespace that stacks the file and repeating it leaks without limit: after control: Shmem: 380 kB child: mount(/proc/self/fd/6, MS_BIND MS_REC): ok after cycle: Shmem: 65916 kB do_move_mount() runs check_for_nsfs_mounts() over a detached tree for exactly that reason. Do the same for the copy before it is grafted.	N/A	More Details
CVE-2026-89312	`AesGcmCryptor.Open` passed a caller-supplied nonce straight to `cipher.AEAD.Open` without checking its length. The Go standard library **panics** rather than returning an error when a GCM nonce is not exactly `NonceSize()` bytes, so a JWE carrying a wrong-sized IV crashes the decrypting goroutine.	N/A	More Details
CVE-2026-16762	An attacker able to supply a crafted XML document to an application using QXmlStreamReader can cause disproportionate CPU consumption during parsing, making the application unresponsive (denial of service). Exploitation requires only that the application parse attacker-influenced XML; no authentication, privileges, or user interaction beyond ingesting the input are needed. There is no loss of confidentiality or integrity — no data disclosure, data modification, or code execution — and the impact is contained to the affected process.	N/A	More Details
CVE-2026-98983	In the Linux kernel, the following vulnerability has been resolved: unshare: don't drop active namespace references that were never taken Active references on the namespaces of an nsproxy are taken when the nsproxy is installed into a task in switch_task_namespaces() and copy_namespaces() and dropped again by put_nsproxy() through deactivate_nsproxy(). But ksys_unshare() calls put_nsproxy() on an nsproxy that was never installed when set_cred_ucounts() fails. The new namespaces of that nsproxy go from zero to minus one and the namespaces shared with the caller lose a reference that belongs to the nsproxy the caller keeps using: WARNING: kernel/nscommon.c:171 at __ns_ref_active_put+0x1cd/0x230 nsproxy_ns_active_put deactivate_nsproxy ksys_unshare Afterwards the namespaces the caller lives in aren't listed by listns() anymore. set_cred_ucounts() only fails when alloc_ucounts() can't allocate and it only allocates when the real uid of the caller differs from its effective uid. Commit cefd55bd2159 ("nsproxy: fix free_nsproxy() and simplify create_new_namespaces()") separated the two cases on purpose. nsproxy_free() frees an nsproxy that was prepared but never installed and that's what a failed setns() uses in put_nsset(). Export it and use it for a failed unshare() as well.	N/A	More Details
CVE-2026-98984	In the Linux kernel, the following vulnerability has been resolved: dcache: don't put a mountpoint on a dentry that's being removed rmdir(), unlink() and rename() call dont_mount() on the victim and then detach_mounts() with the victim's inode locked. do_lock_mount() takes the inode lock of the mountpoint and checks cant_mount() so a mount can't show up after detach_mounts(). But attach_recursive_mnt() makes a second mountpoint for the root of the source mount so that the mounts already located at the destination can be put on top of it. No inode is locked for that one and d_set_mounted() only refuses a dentry that is unlinked. Between detach_mounts() and d_delete() the victim is still hashed: rmrace: b passed dont_mount() and detach_mounts(), sleeping T2: move_mount(S, /tmp/plcant/x, BENEATH) = 0 errno 0 () T1: rmdir(/tmp/plcant/d/b) = 0 errno 0 109 107 0:61 /d/b//deleted /tmp/plcant/x rw,relatime - tmpfs tmpfs rw 108 109 0:63 / /tmp/plcant/x rw,relatime - tmpfs T rw A bind mount of the directory that's being removed is moved beneath an existing mount while the rmdir() is located between the two calls. The mount ends up on the removed directory and nothing will ever detach it. Check cant_mount() in d_set_mounted() as well. dont_mount() raises the flag under d_lock before detach_mounts() runs so either the mountpoint is set first and detach_mounts() finds the mount or the flag is seen and the mount is refused.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: namespace: check the mounts before		

CVE-2026-98985	reading their parents in pivot_root() path_pivot_root() reads the parents of new_root and of the caller's root and checks whether they are shared before it checks that either mount is in the caller's mount namespace. Only namespace_sem is held. That's fine for a mount that is in the caller's mount namespace. But new_root can be a file descriptor to a mount that has been unmounted and that only the file descriptor keeps alive. If that mount stayed attached to its parent when it was unmounted nothing pins the parent for it. The final mntput() of the parent unhooks the children under mount_lock alone and frees the parent afterwards: pivot_root() close(fd), last ref on the parent ----- ----- ex_parent = new_mnt->mnt_parent mntput_no_expire_slowpath() __umount_mnt(new_mnt) cleanup_mnt() call_rcu() IS_MNT_SHARED(ex_parent) BUG: KASAN: slab-use-after-free in path_pivot_root+0xf1a/0x1840 Read of size 4 at addr ffff8881047382f8 by task pivot_widen/157 path_pivot_root+0xf1a/0x1840 __x64_sys_pivot_root+0x165/0x190 The same goes for the caller's root via chroot(). Both outcomes of the check end in EINVAL so nothing but the read itself goes wrong. It's the same thing commit bb4609405752 ("statmount: read the parent of an unmounted mount under mount_lock") fixed for statmount(). Check that both mounts are in the caller's namespace before their parents are read. Every path returns EINVAL either way.	N/A	More Details
CVE-2026-98986	In the Linux kernel, the following vulnerability has been resolved: namespace: remove the fsnotify marks of a mount namespace in process context mnt_ns_release() drops the last passive reference of a mount namespace and removes its fanotify marks via fsnotify_mntns_delete(). That takes the mutex of every group with a mark on the namespace and the spinlock of the connector. Fine from process context. But mnt_ns_tree_remove() hands the reference the namespace was allocated with to call_rcu() and so the marks are removed from the RCU softirq: BUG: sleeping function called from invalid context at kernel/locking/mutex.c:623 in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 0, name: swapper/3 __mutex_lock+0x113/0x24b0 fsnotify_destroy_marks+0x11b/0x3d0 mnt_ns_release_rcu+0x57/0xa0 rcu_core+0x6b6/0x1e40 and lockdep complains about the connector lock being taken from softirq context. A fanotify group with a FAN_MARK_MNTNS mark on the mount namespace of another task is all that's needed. Root in a user namespace can do that for a mount namespace it owns. The task exits and the namespace is freed with the mark still on it. Remove the marks in free_mnt_ns() before the namespace is handed to RCU. That runs in process context once the last active reference is gone. A mark is added through a file descriptor to the namespace which holds an active reference so no mark can show up after that.	N/A	More Details
CVE-2026-98987	In the Linux kernel, the following vulnerability has been resolved: namespace: look at the topmost mount for a mount namespace file attach_recursive_mnt() preallocates a mountpoint for the root of the topmost mount of the source so that a mount that already sits at the destination can be put on top of the source or on top of one of its propagated copies. The copies are made without CL_COPY_MNT_NS_FILE so their chain of overmounts is shorter than the source's. It ends below the first bind mount of a mount namespace file. So while the loop walks up the chain of the source it remembers the mountpoint of that mount in "shorter" and the existing mount is put there for the copies. But the loop ends at the topmost mount without ever looking at it. If the topmost mount is the only mount namespace file in the chain "shorter" stays NULL and mnt_change_mountpoint() attaches the existing mount of the copy below the root of the mount namespace file. That's a dentry of nsfs. No path walk in the copy's mount namespace ever gets there. The mount is gone until its parent goes and the copy that took its place can't be unmounted synchronously because it has a child: S bind mount of a file N bind mount of a mount namespace file on top of S A shared mount, B a slave of A, Q mounted on B/file move_mount(S, A/file) cat B/file -> the content of S instead of Q umount B/file -> EBUSY Look at every mount of the chain including the topmost one.	N/A	More Details
CVE-2026-98971	In the Linux kernel, the following vulnerability has been resolved: gpio: exar: initialize the ID before registering its cleanup devm_add_action_or_reset() invokes the cleanup action immediately when registering the action fails. gpio_exar_probe() currently stores the allocated ID in exar_gpio only after registering the action, so that failure makes exar_devm_ida_free() release ID 0. If the allocated ID is nonzero, this leaks that ID and can free ID 0 belonging to another device. Store the ID before registering the cleanup action so both the immediate failure path and normal device teardown release the correct allocation.	N/A	More Details
CVE-2026-98970	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: Free the BO cache size list on teardown vc4_bo_cache_destroy() purges the cached BOs but never frees the size list, leaking it on every unbind.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: net/sched: em_text: reject unterminated algo before autoload This is a follow-up to commit 9c572a83037a ("net/sched: fix potential stack infoleak in em_text_dump()"), which zero-initialised the dump-side struct. Sashiko pointed at a pre-existing bug that exposed the change-side input-validation gap open. em_text_change() casts the raw netlink attribute payload to struct tcf_em_text and passes conf->algo, a 16-byte array with no enforced NUL terminator, to textsearch_prepare(). On the TS_AUTOLOAD retry textsearch_prepare() calls request_module("ts_%s", algo); vsnprintf() walks %s until it finds a NUL, so		More

2026-98969	an algo[] with no NUL reads past the array into the adjacent struct fields and payload and feeds those bytes to the usermode helper command line. This is an out-of-bounds read and a kernel memory disclosure. Reject the attribute when no NUL exists within TC_EM_TEXT_ALGOSIZ bytes, matching the check xt_string has had since 3ab720881b6e. Conditions to recreate the bug: - CAP_NET_ADMIN on the target netns (namespace-local via unshare -Urn is enough) - install clsact on an interface, then add a basic filter with one text ematch whose algo[] is 16 bytes with no NUL (e.g. "A"*16) and pattern_len at least 1 - the add fails with -ENOENT and the kernel invokes modprobe with a module name made of "ts_" plus the 16 bytes and the adjacent payload	N/A	Details
CVE-2026-89318	`JweRsaKeyEncryptionDecryptorImpl.Decrypt` passed the JWE's initialization vector to `cipher.AEAD.Open` without checking its length. The Go standard library **panics** rather than returning an error when a GCM nonce is not exactly `NonceSize()` bytes, so a JWE carrying a wrong-sized IV crashes the decrypting goroutine.	N/A	More Details
CVE-2026-105963	Externally-Generated Error Message Containing Sensitive Information vulnerability in plug_crypto in elixir-plug plug_crypto allows an unauthenticated remote attacker to expose the secret key base through application logs via a wrong-typed argument. Every public function in Plug.Crypto validates argument types in a function-head guard. A failed guard makes the Erlang VM raise FunctionClauseError with the actual argument values in the top stack frame, so the secret key base, a derived key, or the value an API key is compared against lands in that frame. The library's rescue blocks and Plug.Crypto.prune_args_from_stacktrace/1 only reach exceptions raised inside a function body, so they never see a guard failure. Plug.Cowboy and Bandit write the frame to the error log and attach it as crash_reason Logger metadata, which error trackers forward. Two entry points take request data. Plug.Crypto.verify/4 exposes the secret key base when the token is neither a binary nor nil, which Phoenix.Token.verify/4 reaches because it forwards the token unguarded. Plug.Crypto.secure_compare/2 exposes the other operand when one side is not a binary, as happens when a missing request header yields nil. Reading the leaked value needs separate access to the logs or the error tracker, so the request alone does not disclose the secret to the attacker. This issue affects plug: from 1.6.3 before 1.7.0; plug_crypto: from 1.0.0 before 2.2.1.	N/A	More Details
CVE-2026-98953	In the Linux kernel, the following vulnerability has been resolved: devlink: fix devlink_rel reference leak when notify work is pending devlink_rel_nested_in_notify_work_schedule() takes a reference on the devlink_rel for the notify work and then queues the work, ignoring the return value of schedule_delayed_work(). If the work is already pending, nothing new is queued, the work runs only once and drops only one reference, so the extra one is leaked together with the devlink_rel and its index in devlink_rels. This is easy to hit. devl_register() of a nested instance queues the work, and if the instance is unregistered before the work has run, devlink_rel_put() queues it again while it is still pending. The work also keeps rescheduling itself for as long as the parent devlink lock cannot be taken, which widens the window. Registering and unregistering a nested devlink instance 100 times while holding the parent lock leaks all 100 devlink_rel objects. The reschedule path in devlink_rel_nested_in_notify_work() has the same problem: if the work was queued again while it was running, the reference it holds is never dropped. Drop the reference in both places when the work was already pending. The pending work holds its own reference, so this can not be the last one.	N/A	More Details
CVE-2026-98954	In the Linux kernel, the following vulnerability has been resolved: net: always dissect GSO packets in __virtio_net_hdr_to_skb() Commit 9e8db5913264 ("net: avoid false positives in untrusted gso validation") added a '&& skb->network_header' check before flow-dissecting GSO packets without VIRTIO_NET_HDR_F_NEEDS_CSUM in __virtio_net_hdr_to_skb(), because some callers (such as tun_get_user(), tun_xdp_one(), virtnet_receive_done(), and raw_verify_header()) called virtio_net_hdr_*_to_skb() before initializing skb->network_header and skb->dev. Because __alloc_skb() and __build_skb_around() zero-initialize skb->network_header to 0 (unlike mac_header and transport_header which are initialized to ~0U), those four callers always had skb->network_header == 0 and bypassed flow dissection in __virtio_net_hdr_to_skb(). More generally, skb->network_header is an offset from skb->head (where 0 is also a valid offset whenever skb_headroom(skb) is 0), not a boolean flag. Whenever the 'if (gso_type && skb->network_header)' branch was skipped, the fallback 'else if (gso_type)' only pulled nh_min_len + thlen (40 bytes for TCPv4) without dissecting the packet, without validating ip_proto or n_proto, and without setting skb->transport_header. If the packet has a malformed network header, it is not rejected and a subsequent skb_probe_transport_header() also fails, leaving skb->transport_header at ~0U (0xfffff). Similarly, if an IPv4 packet carries IP options (ihl > 5) or an IPv6 packet carries extension headers, pulling only nh_min_len + thlen can leave the TCP header outside skb->head. In both cases, tcp_hdrlen(skb) in skb_gso_transport_seglen() reads out-of-bounds: BUG: KASAN: slab-out-of-bounds in skb_gso_transport_seglen Read of size 2 by task poc/133 skb_gso_transport_seglen (net/core/gso.c:155) skb_gso_validate_mac_len (net/core/gso.c:270) tbf_enqueue (net/sched/sch_tbf.c:260) dev_qdisc_enqueue (net/core/dev.c:4227) __dev_queue_xmit (net/core/dev.c:4884) In addition, checking virtio_net_hdr_match_proto() only inside 'if (!skb->protocol)' before flow dissection both skipped validation when skb->protocol was pre-set by the caller and rejected VLAN-tagged frames whose outer L2 protocol is ETH_P_8021Q or ETH_P_8021AD. Fix this by: 1.	N/A	More Details

	Initializing <code>skb->dev</code> and <code>skb->network_header</code> (plus <code>skb->protocol</code> for IFF_TUN) before <code>virtio_net_hdr_*_to_skb()</code> in <code>tun_get_user()</code>, <code>tun_xdp_one()</code>, <code>virtnet_receive_done()</code>, and <code>raw_verify_header()</code>. In <code>tun_get_user()</code>, drop the redundant <code>skb_reset_mac_header(skb)</code> in the IFF_TUN case since <code>__virtio_net_hdr_to_skb()</code> unconditionally resets <code>mac_header</code>. 2. Removing '&& <code>skb->network_header</code>' and the unvalidated 'else if (<code>gso_type</code>)' fallback in <code>__virtio_net_hdr_to_skb()</code> so all GSO packets without <code>VIRTIO_NET_HDR_F_NEEDS_CSUM</code> are flow-dissected, have their transport header pulled into linear data, and have <code>skb->transport_header</code> set. 3. Moving the <code>virtio_net_hdr_match_proto()</code> check to after <code>skb_flow_dissect_flow_keys_basic()</code>, validating <code>keys.basic.n_proto</code> against <code>hdr_gso_type</code>.		
CVE-2026-98955	In the Linux kernel, the following vulnerability has been resolved: <code>flow_dissector</code>: avoid u16 truncation of <code>skb->len</code> when computing <code>thoff</code> <code>__skb_flow_dissect()</code> computes <code>key_control->thoff</code> using: <code>key_control->thoff = min_t(u16, nhoff, skb ? skb->len : hlen); min_t(u16, ...)</code> casts both arguments to u16 before comparing them. Whenever <code>skb->len</code> (or <code>hlen</code>) modulo 65536 is smaller than <code>nhoff</code>, the truncated length wins and <code>thoff</code> is set to a bogus small value, even when the dissection succeeded. For example, an IPv4/TCP frame with <code>skb->len == 65540</code> and <code>nhoff == 34</code> gets <code>thoff == 4</code>, so callers such as <code>skb_probe_transport_header()</code> point the transport header inside the Ethernet header. Such skbs are not exotic: - At the time of commit <code>d0c081b49137</code> ("flow_dissector: properly cap thoff field"), <code>AF_PACKET</code> with <code>PACKET_VNET_HDR</code> could already build GSO skbs larger than 64KB (MTU checks are skipped for GSO, and <code>alloc_skb_with_frags()</code> accepted up to <code>MAX_SKB_FRAGS</code> (17) order-0 pages on top of the linear part). - BIG TCP now makes skbs larger than 64KB common. - The following patch makes <code>tun_get_user()</code> dissect IFF_TAP frames before <code>eth_type_trans()</code> pulls the Ethernet header, so a GSO frame carrying a 65522..65535 byte L3 packet will be dissected with <code>skb->len</code> in [65536, 65549]. Compare as u32 instead. If the resulting offset cannot be represented in the u16 <code>key_control->thoff</code>, cap it to <code>U16_MAX</code> and report the dissection as failed rather than silently returning a wrong transport offset. <code>thoff</code> is still set on failure, as some callers (such as <code>eth_get_headlen()</code>) use it regardless of the return value.	N/A	More Details
CVE-2026-98956	In the Linux kernel, the following vulnerability has been resolved: ASoC: <code>adau1372</code>: allow sleepy powerdown GPIOs <code>gpiod_set_value()</code> emits a WARN when the GPIO provider indicates that it can sleep, as is typical for I2C controlled GPIO expanders. Since the powerdown GPIO is only manipulated in process context, use the <code>_cansleep</code> variant instead. This lets the driver use sleepy GPIOs to power off the chip.	N/A	More Details
CVE-2026-98957	In the Linux kernel, the following vulnerability has been resolved: wifi: <code>mt76</code>: <code>mt7915</code>: disable rx napi when removing device Unloading the <code>mt7915e</code> module with a <code>MT7916</code> triggers multiples WARN in <code>__netif_napi_del_locked()</code> and in <code>page_pool_disable_direct_recycling()</code> because the driver does not disable the napi before destroying it. This is troublesome since on <code>MT7916</code> it is required to unload the module and reinsert it with a different <code>enable_6ghz</code> parameter to change the frequency. The system generally becomes unstable after reinserting the module. Fix it by disabling rx napi inside <code>mt7915e</code> itself, even through napi rx is enabled in the <code>mt76</code> generic code. Disabling it in the <code>mt76</code> generic code was previously attempted and reverted because napi was disabled too late and rx handlers could access freed structures. Note that the tx side is enabled and disabled by the <code>mt76</code> generic code. It is disabled in <code>mt7915_stop_hardware()</code> since, unlike <code>mt7915_unregister_device()</code>, is it also called in the probe path if an error happens after rx napi is enabled.	N/A	More Details
CVE-2026-98958	In the Linux kernel, the following vulnerability has been resolved: wifi: <code>mt76</code>: account non-AQL frames per peer rather than per link The driver increments <code>wcid->non_aql_packets</code> when a frame goes to the hardware queue. The completion path decrements the counter again. For an MLO peer, the hardware reports the completion on the link that it used for the transmission. This link can differ from the link that the driver queued the frame on. In that case, the increment and the decrement go to different link <code>wcids</code>, and the count drifts upwards until it blocks the station. Keep the count on the <code>wcid</code> of the primary link, like the other peer-wide state.	N/A	More Details
CVE-2026-98960	In the Linux kernel, the following vulnerability has been resolved: btrfs: don't stash <code>io_uring</code> encoded data across -EAGAIN Returning -EAGAIN while leaving <code>btrfs_uring_encoded_data</code> in the cmd PDU leaks if the request is cancelled or the ring exits before reissue. <code>io_uring</code> does not free driver PDU allocations on cleanup. Write: <code>io_queue_sqe()</code> always issues with <code>IO_URING_F_NONBLOCK</code> first, so return -EAGAIN before allocating and free data on every exit. Read: free on <code>nowait</code> -EAGAIN too; only -EIOCBQUEUED keeps the allocation for <code>btrfs_uring_read_finished()</code>.	N/A	More Details
CVE-2026-98961	In the Linux kernel, the following vulnerability has been resolved: btrfs: unlock inode and extent in caller when <code>io_uring</code> read extent fails <code>btrfs_uring_read_extent()</code> runs only after <code>btrfs_encoded_read()</code> has taken the inode shared lock and the extent lock. On failure it used to unlock in <code>out_fail</code>, and a pages-array allocation failure returned -ENOMEM without unlocking at all. Unlock in the caller instead on all failure returns, matching the <code>copy_to_user()</code> error path. The deferred -EIOCBQUEUED path still unlocks in <code>btrfs_uring_read_finished()</code>.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: btrfs: free iov when <code>btrfs_uring_read_extent()</code> fails After <code>btrfs_uring_read_extent()</code>, the caller always jumped to <code>out_acct</code>.	N/A	More

98962	That skips kfree(data->iov), which is only correct for -EIOCBQUEUED where the deferred path owns the iov. On failure, fall through to out_free instead.		Details
CVE-2026-98963	In the Linux kernel, the following vulnerability has been resolved: btrfs: always return -EIOCBQUEUED after btrfs_uring_read_extent_endio If all bios finish before btrfs_encoded_read_regular_fill_pages() returns, it calls btrfs_uring_read_extent_endio() and previously returned the I/O status. A negative errno then made btrfs_uring_read_extent() unlock and free while btrfs_uring_read_finished() did the same again. Return -EIOCBQUEUED so only the deferred path cleans up.	N/A	More Details
CVE-2026-89320	`HmacShaCryptor.Hash` never hashed its input. For any input it returned `HMAC_K("")` — the MAC of the empty message under the key. Every AES-CBC-HMAC JWE produced or verified by gose therefore carried an authentication tag that is constant per key and independent of the message, removing ciphertext integrity entirely.	N/A	More Details
CVE-2026-89319	`JweDirectDecryptorBlock.Decrypt` read a 64-bit plaintext length from the `_thales_aad` protected-header field and used it without any validation. A truncated field panics `binary.BigEndian.Uint64`, an oversized value drives an unbounded allocation, and a value larger than the decrypted block panics with a slice-bounds error.	N/A	More Details
CVE-2026-98966	In the Linux kernel, the following vulnerability has been resolved: bonding: fix the program leak when XDP is replaced We need to release the reference on the old BPF prog, whether we are going to no-prog, or replacing the prog with another. Spotted by AI while running tests for the propagation series (leaked programs were accumulating). Selftest will come in the net-next series to avoid conflicts.	N/A	More Details
CVE-2026-98967	In the Linux kernel, the following vulnerability has been resolved: cxgb4/ch_ktls: disable softirqs around tid_list erases chcr_ktls_cpl_act_open_rpl() inserts into tid_list from the receive path using xa_insert_bh(). The array is shared by the adapter's connections and initialized with XA_FLAGS_LOCK_BH. However, chcr_ktls_dev_del() and the chcr_ktls_dev_add() error path use xa_erase(), which takes the XArray lock without disabling softirqs. When cleanup runs with softirqs enabled, a receive softirq for another connection on the same CPU can interrupt the erase and spin on the lock held by the interrupted task. Use xa_erase_bh() at both sites to match the insertion path.	N/A	More Details
CVE-2026-115657	When used as an executor in either KNIME Server or KNIME Business Hub, KNIME Analytics Platform LTS versions before 5.8.4 log configured credentials of proxies during startup. An attacker with access to log files can steal the credentials and use the proxy as the user who configured the proxy settings. The issue has been fixed in LTS versions 5.8.4 and 5.12.0 and later.	N/A	More Details
CVE-2026-98990	In the Linux kernel, the following vulnerability has been resolved: namespace: check a submount for references right before unmounting it shrink_submounts() and mark_mounts_for_expiry() first collect all the mounts they are allowed to unmount and then unmount them. Whether a mount is busy is decided by propagate_mount_busy() on the tree. But unmounting the first mount changes the tree that the second one propagates into. propagate_umount() moves a surviving overmount off a copy it unmounts and mounts it where the copy was mounted. If that's where the propagated copy of the second victim is looked up the overmount becomes a candidate of the second umount_tree(). It's childless and so trim_one() commits it without looking at its reference count. So a synchronous umount pulls out a busy mount that is in use somewhere else and marks it MNT_SYNC_UMOUNT while it has users: umount2(/tmp/plshrink/p1, 0) = 0 errno 0 cwd is (unreachable) The same two umounts requested one after the other fail with EBUSY. This needs a mount with MNT_SHRINKABLE, i.e., automounted submounts of NFS, AFS and CIFS or the tracefs mount below debugfs. Bind mounts inherit the flag. Check each mount right before it is unmounted under the same hold of namespace_sem and mount_lock as the umount itself. Make sure that the algorithm stays linear.	N/A	More Details
CVE-2026-116161	Parse Server before 8.6.103 and 9.0.0 before 9.10.5-alpha.4 contains an information disclosure vulnerability in LiveQuery. Its subscription hash omits constraints next to a \$or operator, including the _Session owner filter and beforeSubscribe constraints, so distinct subscriptions are grouped together. Authenticated clients subscribing with \$or queries can receive other users' events, including session tokens usable for impersonation.	N/A	More Details
CVE-2026-75779	Improper Neutralization of Special Elements used in an OS Command (“OS Command Injection”) and Cross-Site Request Forgery (CSRF) in an ASUS router model allow an attacker to trick a logged-in administrator into visiting a phishing website, causing the router to execute arbitrary commands with root privileges. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2025-30049	The CGM CLININET application stores user passwords as unsalted MD5 hashes. This issue was fixed in versions 2026.MS1 and 2026.MS3	N/A	More Details
	CGM CLININET is vulnerable to Remote Code Execution caused by deserialization of untrusted data. The		

CVE-2025-30047	server exposes an HTTP-based API, used by the administrative module, that exchanges data using Java object serialization. An attacker within the same local network is able to send requests to this API and transmit a crafted serialized object. When it is deserialized, it allows arbitrary code execution with root privileges on the server. This issue was fixed in version 2026.MS3	N/A	More Details
CVE-2025-30046	In the CGM CLININET application server, multiple processes responsible for HL7 communication with external systems run with root privileges while processing untrusted data received from those external systems. An attacker in adjacent network able to send data to these HL7 communication processes can achieve remote code execution with root privileges. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
CVE-2025-30045	CGM CLININET is vulnerable to Code Injection in the ConvExpr function, which processes Perl expressions into RTF document content. Due to improper control of the code evaluated from these expressions, an authenticated, high-privileged attacker can cause arbitrary code to be executed with root privileges. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
CVE-2026-99015	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_sock: Serialize dead-device detachment Rebinding an HCI socket after its controller is unregistered can race monitor control replay and crash the kernel. hci_sock_bind() clears hci_pi(sk)->hdev and drops the device reference under only the socket lock, whereas send_monitor_control_replay() holds hci_sk_list.lock for reading. Replay can observe a non-NULL hdev, then bind can clear it and drop the last reference before create_monitor_ctrl_open() reads hdev->id. This can cause a NULL pointer dereference if the pointer is reloaded, or a use-after-free if the old pointer is retained. The kernel reported: KASAN: null-ptr-deref in range [0x0000000000000068-0x000000000000006f] RIP: 0010:create_monitor_ctrl_open+0x579/0x800 Call Trace: hci_sock_bind+0xd96/0x1190 __sys_bind+0x166/0x200 __x64_sys_bind+0x6d/0xb0 do_syscall_64+0xdd/0x4a0 entry_SYSCALL_64_after_hwframe+0x77/0x7f The same missing serialization lets hci_send_to_sock() select a socket for the old device and queue its frame after bind has detached the socket or installed a new binding. Take hci_sk_list.lock for writing when clearing the device pointer and resetting the socket state. This waits for existing replay and delivery readers to finish before detachment and prevents later readers from using the old binding. Keep hci_dev_put() outside the critical section because the final device release can sleep.	N/A	More Details
CVE-2025-30043	CGM CLININET is vulnerable to Missing Authentication for Critical Function in the /DocManagment path. The endpoint is accessible without authentication, allowing an unauthenticated attacker in the adjacent network to retrieve the list of patient documents and download those documents. This issue was fixed in version 2026.MS3	N/A	More Details
CVE-2026-93489	Improper Authentication in ASUS router modules allows an unauthenticated nearby user to send an HTTP request and forge an access ticket, subsequently enabling the user to perform administrator operations. Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-99019	In the Linux kernel, the following vulnerability has been resolved: scsi: ch: Do not keep references to data transfer element devices ch_readconfig() looks up the scsi_device of every data transfer element whose SCSI id the changer reports in READ ELEMENT STATUS, and stores it in ch->dt[]. scsi_device_lookup() takes a reference, and nothing ever drops it: ch_destroy() frees the array with kfree(). ch->dt[] is read nowhere else - it only supplies the vendor, model and revision printed in the same loop. Once such a drive is removed, its scsi_device can never be released. It stays on the host's device list at its address, so a new device there is refused by anything that walks the list - target_core_pscsi reports "scsi_device_get() failed for H:C:T:L" - and the low-level driver's module can no longer be unloaded. Only a reboot recovers. It shows with any changer that reports its drives' ids; with the mhvtl virtual library (IBM 3573-TL personality), each create and remove of a library with four drives leaves four references behind, counted by the module's use count in lsmmod. With ch not bound the count is unchanged, and with this patch applied it is unchanged too. Drop the reference as soon as the name has been printed, and remove the now unused dt[] array. That also removes the array leaked when ch_probe() fails after ch_readconfig(). The same leak was reported with an RFC patch in 2022, which was not merged.	N/A	More Details
CVE-2025-58403	CGM CLININET is vulnerable to Reflected Cross-Site Scripting (XSS) due to improper neutralization of user-supplied data. A remote unauthenticated attacker can craft a malicious URL containing JavaScript or HTML code, which is executed in the victim's browser when the URL is visited, allowing the attacker to steal session cookies and perform actions on behalf of the victim within the application. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: regulator: tps6594-regulator: Fix n_linear_ranges for TPS65224 BUCK2-4 The TPS65224 BUCK2, BUCK3 and BUCK4 descriptors pass 4 as		

CVE-2026-99020	n_linear_ranges, but tps65224_bucks_2_3_4_ranges[] only has two entries, covering selectors 0x00-0x45. The linear range helpers loop over n_linear_ranges entries, so any lookup of a selector above 0x45 reads past the end of the array and treats whatever follows it as two more ranges. This happens on the normal registration path when the core checks the voltage constraints against every selector. TPS652G1 uses the same descriptors and is affected too. Use ARRAY_SIZE() so the count always matches the table.	N/A	More Details
CVE-2026-99021	In the Linux kernel, the following vulnerability has been resolved: regulator: of: fill in supply names in of_regulator_bulk_get_all() of_regulator_bulk_get_all() returns an array it allocated itself, and fills in only the consumer of each entry. Every other way of getting a bulk array has the supply name set, because the caller provides it, and the core expects it to be there: regulator_bulk_enable() prints it when a supply fails to enable, so a caller that hands such an array to it dereferences uninitialised memory on that path. Copy each name into the array's own allocation, right behind the entries, so that it shares the array's lifetime and callers still have nothing extra to free. That also retires the fixed 64 byte stack buffer the names were assembled in, which is_supply_name() never bounded the copy against.	N/A	More Details
CVE-2026-99022	In the Linux kernel, the following vulnerability has been resolved: regulator: ltc3676: don't read the write-only HRST and CLIRQ registers HRST (0x1e) and CLIRQ (0x1f) are write-only command strobes. Writing HRST triggers a hard reset of the chip and writing CLIRQ clears the interrupt; reading them returns nothing useful, and the read itself fires the command. The result is that a plain register dump resets the board: cat /sys/kernel/debug/regmap/2-003c/registers walks every readable register, hits HRST, and the board reboots. Split the callback back in two so HRST and CLIRQ stay writable but are no longer readable. The readable range is left unchanged (LTC3676_BUCK1 ... LTC3676_IRQSTAT).	N/A	More Details
CVE-2026-99023	In the Linux kernel, the following vulnerability has been resolved: power: sequencing: pcie-m2: Fix leaking array from of_regulator_bulk_get_all() of_regulator_bulk_get_all() allocates the consumers array and hands it to the caller, who is expected to free it at an appropriate time. pwrseq-pcie-m2.c never did, leading to a memory leak. Add an explicit kfree() for the consumers array.	N/A	More Details
CVE-2026-99024	In the Linux kernel, the following vulnerability has been resolved: optee: register TEE devices only once fully initialized optee_probe() called tee_device_register() on both the client and the supplicant device before the rest of struct optee had been initialized. tee_device_register() calls cdev_device_add(), which does two things at once: it creates /dev/tee0 and /dev/teepriv0, and it links the device into the tee class so that class_find_device() can find it. From that moment on the device is reachable both from user space via tee_open() and from kernel space via tee_client_open_context(). The only gate in teedev_open() is tee_device_get(), which merely checks that teedev->desc is non-NULL, which was already set by tee_device_alloc(). Therefore, there is effectively no gate at all. A context opened in that window can run against a struct optee where - optee->call_queue.mutex is not initialized by optee_cq_init() - optee->supp mutex and completions is not initialized by optee_supp_init() - optee->rpm_dev_mutex is not initialized yet, - the message argument cache is not initialized by optee_shm_arg_cache_init() - optee->ctx is still NULL Any open session or invoke during this window takes uninitialized mutexes and dereferences a NULL pointer. Fix it by moving both tee_device_register() calls down to the point where all of struct optee is set up. [JW: adding fixes and cc]	N/A	More Details
CVE-2026-99025	In the Linux kernel, the following vulnerability has been resolved: openvswitch: fix soft lockup in the netlink flow dump A production compute node carrying a few thousand datapath flows hit a soft lockup inside a single netlink flow dump and panicked. ovs_flow_cmd_dump() calls ovs_flow_stats_get() for every flow it emits, and that releases stats->lock with spin_unlock_bh() once per CPU that has touched the flow. Every release is a local_bh_enable(), and each one runs the pending softirq backlog in the dumping thread's own context. The skb bounds the flows one callback emits, but not the softirq work it absorbs. On a CPU that carries the box's packet load the backlog refills as fast as it drains, so the dumping thread becomes that CPU's softirq engine. It never sleeps and it has no reschedule point, so under voluntary preemption nothing can take the CPU away from it: neither the ksoftirqd the kernel woke to take the work over, nor the stopper thread the softlockup detector dispatches to refresh its timestamp. Hold BH off across the whole callback instead, the way ctnetlink_dump_table() does, so the nested spin_unlock_bh() stop draining softirqs. The loop already runs under rcu_read_lock() and cannot sleep. What it gives up is preemption under CONFIG_PREEMPT, since a BH-off region is not preemptible outside PREEMPT_RT. The region stays short: the skb caps the flows one callback emits, and empty buckets cost no skb space but are each visited once per dump, as the cursor only moves forward. The table grows on insert and shrinks only on flush, so the walk is bounded by the largest flow count the datapath has held. The softirq backlog the callback used to absorb has no bound at all.	N/A	More Details
CVE-2026-99026	In the Linux kernel, the following vulnerability has been resolved: net: team: free skb when broadcast has no txable port bc_transmit() sends the original skb through the last txable port. When no port is txable, it returns false without consuming the skb, but team_xmit() still returns NETDEV_TX_OK. AF_PACKET sends then retain the skb and its socket write-memory charge. A TX-enabled, link-down port keeps the broadcast transmit op installed. Carrier can stay up through a second link-up, TX-disabled	N/A	More Details

	port or be forced on by userspace. Free the original skb on the no-port path while leaving the existing drop accounting intact. Discovered by Sashiko during review of previous commit.		
CVE-2026-99010	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: RFCOMM: free the skb when the DLC has no owner rfcmm_rcv_data() hands the skb to d->data_ready() and returns 0 without freeing it, so the callback owns the skb from that point on. The tty implementation frees it when the DLC has been detached from its device: static void rfcmm_dev_data_ready(struct rfcmm_dlc *dlc, struct sk_buff *skb) { struct rfcmm_dev *dev = dlc->owner; if (!dev) { kfree_skb(skb); return; } The socket implementation returns without freeing it, and the skb is leaked. A connected DLC can have no owner: rfcmm_sock_destruct() clears d->owner, while the DLC itself stays on the session until it is closed. Data frames that arrive in that window still find d->state == BT_CONNECTED in rfcmm_rcv_data() and are handed to the callback, one leaked skb each. The leak does not stop on its own either. The remote runs out of credits only once RFCOMM_RX_THROTTLED is set, and for a socket the only place that sets it is rfcmm_sk_data_ready() itself, below the return. rfcmm_process_dlc() keeps granting credits, so the remote keeps sending and the leak keeps growing. Free the skb, as the tty side does.	N/A	More Details
CVE-2026-99009	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: RFCOMM: connect the session socket without rfcmm_mutex An RFCOMM connect() issued while a BR/EDR link is being authenticated makes lockdep report a circular dependency, and the reported cycle is a real AB/BA between rfcmm_mutex and hdev->lock. rfcmm_security_cfm() is called from the HCI event path, which already holds hdev->lock: hci_rx_work() hci_event_packet() hci_cc_read_enc_key_size() [hdev->lock] hci_encrypt_cfm() [hci_cb_list_lock] rfcmm_security_cfm() [rfcmm_mutex] while an RFCOMM connect() from userspace takes the same two locks the other way round: rfcmm_sock_connect() rfcmm_dlc_open() [rfcmm_mutex] __rfcmm_dlc_open() rfcmm_session_create() kernel_connect() l2cap_sock_connect() l2cap_chan_connect() [hdev->lock] WARNING: possible circular locking dependency detected kworker/u131:1/1128 is trying to acquire lock: rfcmm_mutex, at: rfcmm_security_cfm+0x31/0x3e0 [rfcmm] but task is already holding lock: hci_cb_list_lock, at: hci_cc_read_enc_key_size+0x1d2/0xcc0 Chain exists of: rfcmm_mutex --> &hdev->lock --> hci_cb_list_lock hci_auth_complete_evt() and hci_encrypt_change_evt() reach the callback the same way. Both orders have to be seen in the same boot, which is why a BR/EDR connection alone is not enough to show it: a session set up by the remote side is created by rfcmm_accept_connection() in krfcommd, which calls kernel_accept() and never takes hdev->lock under rfcmm_mutex. Connecting a device that authenticates and encrypts the link and then calling connect() on an RFCOMM socket towards any address - the connect does not have to succeed, the order is recorded before the page timeout - reports it every time. Only the session socket has to be connected with the lock held, and it does not: nothing else can see the socket before it is put on the session list. So connect it first and take rfcmm_mutex afterwards, which removes the rfcmm_mutex -> hdev->lock order for good, rather than keeping the HCI event path out of rfcmm_mutex. rfcmm_session_create() becomes rfcmm_session_connect(), which returns the connected socket without touching the session list, and rfcmm_dlc_open() adds the session once it holds the lock again. If another opener added a session for the same pair while this socket was connecting, that session is used and this socket is dropped. __rfcmm_dlc_open() now takes the session it should use, and its state check runs after the lock is re-acquired, so a DLC that was opened or closed in the meantime is still handled. Over an existing ACL link the connection can complete before the session reaches the list, and the wakeup from the socket callback is then lost, so krfcommd is woken once the session is visible. Connecting without the lock opens a window in which the socket can be closed. The DLC is not on a session yet, so rfcmm_dlc_close() finds nothing to do and returns without touching it, and rfcmm_dlc_open() would then attach a DLC whose owner is gone. rfcmm_dlc_close() now marks such a DLC with RFCOMM_CLOSED, and rfcmm_dlc_open() checks the mark once it holds rfcmm_mutex again and drops the socket instead of attaching the DLC. The mark is cleared when an open starts, so a DLC that is reused is not affected. This covers both callers of rfcmm_dlc_open(), the socket and the tty layer.	N/A	More Details
CVE-2026-99008	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: RFCOMM: Fix NULL tty_dev dereference in rfcmm_dev_shutdown syzbot hit a NULL pointer dereference in rfcmm_dev_shutdown() when an RFCOMM tty is closed. rfcmm_dev_add() puts the new device on rfcmm_dev_list first and only then registers the tty device and saves it in dev->tty_dev. If that registration fails, dev->tty_dev is never set at all. The device can be looked up in the meantime, so an open() of /dev/rfcommN gets through rfcmm_tty_install() and attaches to it. When that tty is closed again, rfcmm_dev_shutdown() uses dev->tty_dev without checking it: Oops: general protection fault, probably for non-canonical address KASAN: null-ptr-deref in range [0x0000000000000040-0x0000000000000047] RIP: 0010:rfcmm_dev_shutdown+0x48/0xb0 net/bluetooth/rfcomm/tty.c:131 Call Trace: tty_port_shutdown+0x1bf/0x220 drivers/tty/tty_port.c:372 tty_port_close+0x4e/0x150 drivers/tty/tty_port.c:705 tty_release+0x359/0x1670 drivers/tty/tty_io.c:1718 __fput+0x418/0xa50 fs/file_table.c:512 rfcmm_dev_destruct() already copes with a NULL dev->tty_dev. Add the same check to rfcmm_dev_shutdown(), so the tty device is only moved when it was actually registered.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: drm/fb-helper: defer setup when		

CVE-2026-99007	fbdev_probe() fails, not just on -EAGAIN __drm_fb_helper_initial_config_and_unlock() calls drm_client_modeset_probe() first, which already fills in each mode_set->mode/num_connectors for every connected output, and only afterwards calls drm_fb_helper_single_fb_probe() to create fb_helper->fb and wire it into those mode_sets via drm_setup_crtcs_fb(). If drm_fb_helper_single_fb_probe() fails with anything other than -EAGAIN (e.g. dev->driver->fbdev_probe() itself fails), the function bails out before drm_setup_crtcs_fb() runs and before setting fb_helper->deferred_setup. The client is left with mode_sets that have a mode and connectors but mode_set->fb == NULL, and nothing marks this fb_helper as needing a retry. __drm_fb_helper_restore_fbdev_mode_unlocked() only skips committing when fb_helper->deferred_setup is set, so a later .restore() call (drm_lastclose() on process exit, in this report) commits the stale mode_set through drm_client_modeset_commit() -> drm_client_modeset_commit_atomic() -> __drm_atomic_helper_set_config(), which hits WARN_ON(!set->fb). Mark fb_helper->deferred_setup on any drm_fb_helper_single_fb_probe() failure, not just -EAGAIN, so restore keeps skipping the commit until a later hotplug event successfully re-probes and re-populates mode_set->fb. The real error code is still returned to the caller unchanged.	N/A	More Details
CVE-2026-98992	In the Linux kernel, the following vulnerability has been resolved: namespace: don't copy an unmounted mount tree may_copy_tree() accepts a mount whose root is an nsfs or pidfs file before it checks whether that mount is mounted. This exception exists for nsfs and pidfs mounts. They are part of no mount namespace so neither check_mnt() nor is_mounted() would let them through and bind-mounting a nsfsd or a pidfd wouldn't work. Problem is that exception also admits a bind mount of such a file when it was lazily unmounted and is only kept alive by a file descriptor: mount --bind /proc/self/ns/net /tmp/x exec 3 < /tmp/x umount -l /tmp/x mount --rbind /proc/self/fd/3 /tmp/y With MS_REC or AT_RECURSIVE copy_tree() doesn't stop at the unmounted mount. It copies every mount that stayed attached to it as well. So if you have a bunch of overmounts over the nsfsd or the pidfd the whole unmounted stack is copied. That's a bug. So refuse to walk an unmounted tree. Copying its root alone stays possible. So a bind mount of a namespace that is gone can still be recreated from a file descriptor on the old one. The file descriptor pins that mount and its root and __has_locked_children() only reads the direct children which don't leave the parent's list without namespace_sem. The internal mounts carry MNT_NS_INTERNAL instead of NULL and are still walked. Copying an unmounted tree recursively fails with EINVAL.	N/A	More Details
CVE-2025-30054	CGM CLININET is vulnerable to local privilege escalation. The service that processes incoming HL7 messages runs with root privileges, but due to incorrect permissions, its scripts and configuration files can be modified by the uhcapache user. An attacker with uhcapache privileges can modify these files to execute arbitrary code as root. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
CVE-2026-98994	In the Linux kernel, the following vulnerability has been resolved: statmount: read the parent of an unmounted mount under mount_lock statmount_mnt_basic() reads the ids of the parent through m->mnt_parent with namespace_sem held for reading. That's fine for a mount that is in a namespace. Everything above it is pinned by that namespace. STATMOUNT_BY_FD also accepts a file descriptor on a mount that has been unmounted on purpose and that only the file descriptor keeps alive. If that mount stayed attached to its parent when it was unmounted the parent is unmounted as well and nothing pins it for the child. The final mntput() unmounts the children under mount_lock alone and frees the parent afterwards. Concurrently, statmount_mnt_basic() reads m->mnt_parent twice while that can happen and can easily dereference a mount that's already gone. Read the parent of an unmounted mount under mount_lock. The parent can't be put while it is held and both ids come from the same mount. A mounted mount's parent is pinned by the namespace and a mount can't leave its namespace while namespace_sem is held, so the common case takes no global lock.	N/A	More Details
CVE-2025-30053	The CGM CLININET application runs the cgm-updater process listening on an open network port without any authentication. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fs: don't silently unmount busy mounts The propagate_mount_busy() helper exists to decide whether a synchronous can succeed. It does that by looking at the copies of the victims in the mounts its parent propagates to. One of those functions from propagation hell. It skips a copy that has child mounts except for the case of a single child mount covering the copy's root. That algorithm used to be exact. It isn't anymore. Originally only childless copies were unmounted during propagation. And 1064f874abc0 ("mnt: Tuck mounts under others instead of creating shadow/side mounts.") ended up adding covered mounts to both sides in one go. Along came 99b19d16471e ("mnt: In propagate_umount handle visiting mounts in any order") and that started to unmount more than before: [1]: A copy of a mount is unmounted when each of its children is either its overmount or a copy of the victim itself. In the former case the overmount gets reparented. In the latter case the copy of the victim will end up being unmounted. So propagate_umount() ended up being rewritten without adjusting propagate_mount_busy(). And currently propagate_mount_busy() doesn't look at copies of the victim that have two or more children		

CVE-2026-98996	nor a single child that is not the overmount. So any such copy in [1] is not checked for references. The result is that a synchronous unmount succeeds, pulling out a mount that is still in use. A container can run into this with its own mounts. The host shares a tree with the container and the host has a mount in that tree: mount -t tmpfs none /mnt mount --make-shared /mnt mount -t tmpfs none /mnt/a The container's /mnt is a slave copy of the host mount. The container now takes a detached copy of the tree and mounts that copy beneath its propagated copy of /mnt/a and keeps the file descriptor to the tree open: unshare -m --propagation unchanged mount --make-rslave / fd = open_tree(AT_FDCWD, "/mnt", OPEN_TREE_CLONE AT_RECURSIVE) move_mount(fd, "", AT_FDCWD, "/mnt/a", MOVE_MOUNT_BENEATH) grep /mnt /proc/self/mountinfo 775 355 ... /mnt ... master:646 777 775 ... /mnt/a ... master:646 776 777 ... /mnt/a ... master:647 778 777 ... /mnt/a/a ... master:647 The now attached tree (777) is now mounted beneath the container's copy of /mnt/a. That's also where the host mount sits (as seen from the host mount namespace). So the propagated copy (776) of the host's mount is on top of its root and the copy of that (778) is inside it. Now the host runs: umount /mnt/a and it succeeds. The moved tree tree and the copy inside of it are gone from the container. Now only the propagated copy is left and it got reparented to where it was before: grep /mnt /proc/self/mountinfo 775 355 ... /mnt ... master:646 776 775 ... /mnt/a ... But the file descriptor still refers to the moved tree. So a synchronous umount detached a mount that is still in use. Teach propagate_mount_busy() the same checks that propagate_umount() does. It now checks for a single victim mount without children for which it can assume that MNT_LOCKED is cleared on every copy like propagate_umount() does. The copies form chains. IOW, when a copy itself receives propagation from the victim's parent then the mount at the same mountpoint inside that copy is a copy as well. Whether a mount receives propagation from the victim's can simply be checked by walking its masters. Walk each chain once from its top so we are sure that nested copies stay linear. Check every reference once of every copy that gets unmounted.	N/A	More Details
CVE-2026-98997	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/userq: only accept doorbell BOs as queue doorbell amdgpu_userq_get_doorbell_index() pins any BO passed as doorbell_handle into the doorbell domain. For a regular BO, TTM moves it with ttm_bo_move_null() and its contents are lost. Since an importer on the same device gets the exporter's GEM object, a client can do this to a buffer another client shared with it. Reject BOs that were not created in the doorbell domain. v2: drop the error message, use a local abo variable (Christian) (cherry picked from commit f7497bbc19db4821c62b9d97b31c9086feb2d917)	N/A	More Details
CVE-2026-98998	In the Linux kernel, the following vulnerability has been resolved: drm/bridge: th1520-dw-hdmi: Fix remove() callback This driver stores struct th1520_hdmi * in platform drvdata, but th1520_dw_hdmi_remove() was reading it back as struct dw_hdmi * and passing it to dw_hdmi_remove(), so teardown runs on the wrong pointer. Retrieve struct th1520_hdmi * from platform drvdata and pass hdmi->dw_hdmi to dw_hdmi_remove().	N/A	More Details
CVE-2025-30052	The rsyncd service of the CGM CLININET is vulnerable to Authentication Bypass via IP spoofing. Access is granted based solely on the client's source IP address. An attacker able to spoof an allowed IP address, can bypass this check and gain the ability to edit files on the affected system. The vulnerability was confirmed to affect versions up to 2026.MS3 but later versions may also be vulnerable.	N/A	More Details
CVE-2026-99000	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Fix shrinker accounting double-subtraction on nested external pins xe_bo_pin_external() and xe_bo_unpin_external() support nested pins, since the underlying ttm_bo_pin()/ttm_bo_unpin() maintain a pin_count refcount rather than a boolean. However, the shrinker accounting calls xe_ttm_tt_account_subtract() and xe_ttm_tt_account_add() are invoked unconditionally on every pin_external()/unpin_external() call, instead of only on the transition into and out of the pinned state. An external BO (dma-buf) can be pinned more than once while already pinned, for example when it has multiple attachments/importers, each independently calling xe_bo_pin_external(). Each such nested pin subtracts the BO's pages from the shrinker's accounting again, even though the BO's pages were already removed from consideration by the first pin. Symmetrically, each nested unpin adds them back again before the BO is actually unpinned. This desynchronizes the shrinker's page and object counts from reality, and since they are declared as signed long but interpreted as unsigned long in xe_shrinker_count(), the underflow can turn into an enormous shrinkable/purgeable page count, causing the shrinker to be invoked excessively under memory pressure. Guard the accounting calls with the same pin-count transition checks already used to guard the pinned_link list maintenance, so accounting is only updated on the outermost pin and the final unpin. While centralizing this logic into xe_bo_account_pin()/xe_bo_account_unpin(), also close a second, related accounting gap affecting imported dma-bufs. For an imported bo's ttm_tt, xe's ttm_tt_populate hook (xe_ttm_tt_populate()) early-returns without calling xe_ttm_tt_account_add(), since TTM TT_FLAG_EXTERNAL is set without TTM TT_FLAG_EXTERNAL_MAPPABLE. However, TTM core's ttm_tt_populate() wrapper unconditionally marks the tt as populated on driver-hook success, so ttm_tt_is_populated() still reports true for it. xe_bo_account_pin()/xe_bo_account_unpin() trusted ttm_tt_is_populated() as a proxy for "these pages are accounted for by the shrinker", and would subtract/add pages for such a tt that were never added in the first place, again underflowing the	N/A	More Details

	shrinker's counts. This is reachable today via <code>xe_bo_pin_external()</code>, and would have become newly reachable via the fb-pin path once this patch routes <code>__xe_pin_fb_vma()</code> through <code>xe_bo_pin_account()</code> below, so fix it here rather than in a later patch, to avoid a commit in this series that introduces the regression before fixing it. Skip the accounting for imported bos, using the same <code>EXTERNAL && !EXTERNAL_MAPPABLE</code> test <code>xe_ttm_tt_populate()/xe_ttm_tt_unpopulate()</code> use to decide whether to act, via a new <code>xe_ttm_bo_is_imported()</code> helper (hoisted up from its only prior use site). v2: - Fix the critical issue caught by Sashiko AI review: the pin-count transition check introduced here was still bypassed by the fb-pin path (<code>__xe_pin_fb_vma()/__xe_unpin_fb_vma()</code>), which could leave a stale pinned_link entry and a permanent shrinker accounting leak when the fb-unpin performed the true last unpin. Addressed in a separate follow-up patch in this series rather than folded in here, since it is a distinct bug with its own Fixes: tag. - Drop the unneeded local "last_unpin" variable in <code>xe_bo_unpin_external()</code>, restoring the original inline <code>bo->ttm.pin_count == 1</code> check; it was a no-op rename not used by any of the new accounting helpers, and was going to be deleted again by the follow-up patch anyway. v3: - Exclude imported dma-bufs from the pin/unpin shrinker accounting, fixing another underflow caught by Sashiko AI review. - Add a short comment on the backup-object pin site noting it's already populated/accounted at pin time, addressing (as a false-positive clarification, no code fix needed) a separate Medium-severity accounting-leak concern also caught by Sashiko AI ---truncated---		
CVE-2026-99001	In the Linux kernel, the following vulnerability has been resolved: <code>drm/xe/i2c</code>: cancel the client work on remove <code>xe_i2c_notifier()</code> stores the DesignWare adapter in <code>i2c->adapter</code> and schedules <code>i2c->work</code> when the adapter is registered under the xe I2C platform device, and <code>xe_i2c_client_work()</code> then instantiates the AMC client device on that adapter. <code>xe_i2c_remove()</code> tears down the AMC, unregisters the client devices, the bus notifier and the adapter platform device, but it never drains <code>i2c->work</code>. A work item that is still queued or running when the adapter is unregistered dereferences <code>i2c->adapter</code> in <code>i2c_new_client_device()</code> after <code>platform_device_unregister()</code> has released the adapter. The work item is also embedded in the devm-allocated struct <code>xe_i2c</code>, so a work item still queued after the drm device devm unwind frees that allocation runs its callback on freed memory. The bus notifier is the only thing that schedules this work, and it is unregistered after the client devices. An instance that is still queued when the teardown runs can therefore write <code>i2c->client[XE_I2C_CLIENT_AMC]</code> while the loop is unregistering and clearing the same array, and an AMC client it instantiates late is only cleaned up by the adapter's own child sweep in <code>i2c_del_adapter()</code>. Move <code>bus_unregister_notifier()</code> in front of the client teardown loop and cancel the work right after it, so no new instance can be scheduled and a queued instance is drained before the client array is touched. A running instance still finds a live adapter, since the adapter is unregistered later. (cherry picked from commit <code>a23f27d5ec11c3f418c5e5531eea34ccda2b5b20</code>)	N/A	More Details
CVE-2025-30051	The CGM CLININET application stores the production database login and password, as well as RabbitMQ broker credentials, in the <code>SysIO.pl</code> configuration file. Insecure file permissions make this file readable by any local account. This issue was fixed in version 2026.MS3	N/A	More Details
CVE-2026-99003	In the Linux kernel, the following vulnerability has been resolved: <code>drm/panthor</code>: validate userspace queue count against initialized firmware slot count In <code>panthor_drv.c:panthor_ioctl_group_create()</code>, the queue count is validated against the hardcoded maximum: <code>if (!args->queues.count args->queues.count > MAX_CS_PER_CSG) return -EINVAL</code>; However, if the firmware reports fewer slots than <code>MAX_CS_PER_CSG</code>, <code>panthor_fw_init_ifaces()</code> will only initialize that smaller number of stream interfaces. If unprivileged userspace requests more queues than the firmware reported (but no more than <code>MAX_CS_PER_CSG</code>), <code>cs_slot_prog_locked()</code> fetches an uninitialized stream interface and may crash when dereferencing <code>cs_iface->input->ringbuf_base</code>.	N/A	More Details
CVE-2026-99004	In the Linux kernel, the following vulnerability has been resolved: <code>drm/imagination</code>: Size page table preallocation by device address Commit <code>0a8224058a58</code> ("drm/imagination: Fix page count for page table for <code>map()</code> interface") passed the device address to <code>pvr_mmu_op_context_create()</code>, but the preallocation is still sized from <code>device_addr + sgt_offset</code> with an exclusive end. The offset into the object's pages has no place in a device-virtual range, and the exclusive end allocates one table too many when the range ends on a table boundary. Count the tables from <code>device_addr</code> and the inclusive end of the range. The new calculation does not take page table indices into account. This fixes another issue with the previous implementation, where the difference between start and end indices could underflow when a mapping crossed a page table boundary. This would lead to the driver preallocating a huge number of MMU pages, exhausting system memory and triggering an out of memory panic. [Alessio: add paragraph to point out OOM bug being fixed, cc stable]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: <code>drm/i915/gt</code>: Unmask interrupts only when ACTIVE is true There is a small race window between <code>intel_rps_boost()</code> reading a stale ACTIVE state and queuing <code>rps->work</code>, where ACTIVE may be changed via asynchronous calls triggered by wakeref drop. This means that <code>rps_work()</code> could be triggered after the GT has been parked, so unmasking RPS PM interrupts would happen without holding PM wakelock. This sequence of events may be observed during <code>live_rc6_ctx_wa()</code> selftest: <code>1) live_rc6_ctx_wa() -> intel_engine_pm_put() -></code>		

CVE-2026-99005	wakeref_put() ASYNC1) wakeref_put() callback -> __engine_park() -> intel_gt_pm_put_async() -> intel_gt_pm_put_async_untracked() -> intel_wakeref_put_async() ASYNC2) wakeref_put_async() callback -> __gt_park() -> intel_rps_park() -> intel_rps_clear_active() [clear ACTIVE] WITH INTERRUPTS: -> rps_disable_interrupts() -> cancel_work_sync(rps->work) -> rps_reset_interrupts() [rps->pm_iir = 0] WITHOUT INTERRUPTS (TIMER mode): -> rps_stop_timer() -> cancel_work_sync(rps->work) 2) live_rc6_ctx_wa() -> intel_gt_wait_for_idle() -> intel_gt_retire_requests_timeout() -> dma_fence_wait_timeout() -> ops->wait() [i915_fence_wait()] -> i915_request_wait_timeout() -> intel_rps_boost() -> intel_rps_is_active() [check ACTIVE] [OPEN RACE WINDOW] -> queue_work(rps->work) -> retire_requests() -> i915_request_retire() -> intel_rps_dec_waiters() [from 1 to 0] If ASYNC1 and ASYNC2 would execute during [OPEN RACE WINDOW], then: 3) rps_work() executes from the workqueue and may satisfy "(!pm_iir && !client_boost)" test (in TIMER mode rps->pm_iir is zeroed by an earlier rps_work() call or never gets a new value after previous intel_rps_unpark() event), so it jumps to "out" label, bypassing intel_rps_is_active() guard and executes gen6_gt_pm_unmask_irq(); 4) GT is parked and runtime PM wakelock is dropped, so only a raw reference in __gt_park() -> intel_display_power_put_async() remains for 100ms, then the MMIO write in gen6_gt_pm_unmask_irq() detects missing wakelock via assert_rpm_wakelock_held() and emits "RPM wakelock ref not held during HW access" tainting kernel. Commit 3e7abf814193 ("drm/i915: Extract GT render power state management") dropped the "if (rps->interrupts_enabled)" guard that gen6_pm_rps_work() had around this unmask in i915_irq.c. INTEL_RPS_ACTIVE brackets the same window, since it is set in intel_rps_unpark() before rps_enable_interrupts() and rps_start_timer() and cleared in intel_rps_park() before rps_stop_timer() and rps_disable_interrupts(). Skip unmasking when RPS is no longer active (reaching unmask operation is independent of INTERRUPT mode). This does not prevent RPS operation from resuming on the next unpark. In interrupt mode, intel_rps_unpark() calls rps_enable_interrupts(), which enables and unmask the PM interrupts. In timer mode, it calls rps_start_timer(), restarting periodic engine-busyness evaluation. The timer generates software RPS events and queues rps_work() directly, independently of PM interrupt delivery. Neither mode requires the late worker to unmask interrupts while the GT is parked (cherry picked from commit ca4a6deadad15b91ee083afe445fa6bdbfb68206)	N/A	More Details
CVE-2025-30050	CGM CLININET is vulnerable to Code Injection in the RunCGI.pm component, which dynamically imports the module specified as a query parameter and executes code from untrusted data provided by the user. An attacker from adjacent network can supply a crafted query parameter to have an attacker-controlled module imported which results in arbitrary code execution with root privileges. This issue was fixed in version 2026.MS3	N/A	More Details
CVE-2026-98951	In the Linux kernel, the following vulnerability has been resolved: riscv: ptrace: Zero-initialize regset buffers before copyin user_regset_copyin() only copies `count` bytes, so a PTRACE_SETREGSET request with a short iov_len leaves the rest of the stack buffer uninitialized. In riscv_vr_set() that garbage is validated and copied into the target's vstate, and in riscv_cfi_set() it can be stored as the target's shadow stack pointer. Zero both buffers, as the get path already does.	N/A	More Details
CVE-2026-98950	In the Linux kernel, the following vulnerability has been resolved: flow_dissector: avoid u16 truncation of hlen in bpf_flow_dissect() bpf_flow_dissect() sanitizes the offsets returned by the BPF program with clamp_t(u16, ..., hlen). If hlen is bigger than 65535, the upper bound is truncated before the comparison. For instance, with hlen == 65536, both nhoff and thoff are forced to zero, even if the program returned sane values. This is the same class of bug as the one fixed in commit 99bda1ecbd56 ("flow_dissector: avoid u16 truncation of skb->len when computing thoff"). That fix does not cover the BPF path, because __skb_flow_dissect() returns early when a BPF flow dissector handles the packet. struct bpf_flow_keys stores nhoff and thoff as u16, so cap hlen to U16_MAX before clamping. Found by sashiko while reviewing the commit above.	N/A	More Details
CVE-2026-98949	In the Linux kernel, the following vulnerability has been resolved: drm/imagination: Fix reference and vm_bo handling in remap() When a map overlaps part of an existing mapping, pvr_vm_gpuva_remap() splits the mapping into prev/next parts covering what the request did not take, instead of creating something new. It gets two things wrong. - A GEM reference is taken for each part but it's never needed and never dropped, so it leaks one reference per split (remap-next-in-2m in tests/imagination/pvr_vm_map.c): CRITICAL: 33558528 bytes of shmem still held after close - The parts still belong to the original object being split, but pvr_vm_gpuva_remap() links them to ctx->gpuvvm_bo, the new object's vm_bo: prev_va --obj--> BO_A BO_B <--obj-- vm_bo (ctx->gpuvvm_bo) _____link_____ / (mismatch: BO_A != BO_B) drm_gpuva_link() catches the mismatch: WARNING: drivers/gpu/drm/drm_gpuvvm.c:2108 at drm_gpuva_link+0x2ec/0x310 drm_WARN_ON(obj != vm_bo->obj) Call trace: drm_gpuva_link pvr_vm_gpuva_remap __drm_gpuvvm_sm_map pvr_vm_map pvr_ioctl_vm_map Link them to op->remap.unmap->va->vm_bo instead. The locking of the GPUVA lists of the other objects touched by a split is not addressed here; it is handled by switching the GPUVm to immediate mode.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mmc: core: Add validation for host-provided max_segs The max_segs field is of type unsigned short, and if a host driver sets an		

CVE-2026-98884	excessively large value, it may be truncated to zero. This can cause mmc_alloc_sg() to call kmalloc_objs() with a zero size allocation request, which leads to undefined behavior. Under the SLUB allocator, kmalloc(0) returns a special pointer (ZERO_SIZE_PTR). The subsequent 'if (sg)' check will evaluate to true, and sg_init_table() will then attempt to access invalid memory, resulting in a crash: dwmmc_rockchip 2a310000.mmc: Successfully tuned phase to 133 mmc1: new UHS-I speed SDR104 SDHC card at address aaaa Unable to handle kernel paging request at virtual address 0000001fffffff0 Mem abort info: ESR = 0x0000000096000004 EC = 0x25: DABT (current EL), IL = 32 bits SET = 0, FnV = 0 EA = 0, S1PTW = 0 FSC = 0x04: level 0 translation fault Data abort info: ISV = 0, ISS = 0x000000004, ISS2 = 0x00000000 CM = 0, WnR = 0, TnD = 0, TagAccess = 0 GCS = 0, Overlay = 0, DirtyBit = 0, Xs = 0 user pgtbl: 4k pages, 48-bit VAs, pgdp=0000000102c88000 [0000001fffffff0] pgd=0000000000000000, p4d=0000000000000000 Internal error: Oops: 0000000096000004 [#1] SMP Modules linked in: CPU: 2 UID: 0 PID: 102 Comm: kworker/2:1 Not tainted 7.0.0-rc6-next-20260331-00013-g4d93c25963c5-dirty #80 PREEMPT Hardware name: Rockchip RK3576 EVB V10 Board (DT) Workqueue: events_freezable mmc_rescan pstate: 80000005 (Nzcv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : sg_init_table+0x2c/0x50 lr : sg_init_table+0x24/0x50 sp : ffff8000837db710 x29: ffff8000837db710 x28: 000000000000c000 x27: 0000000000000300 x26: 0000000000000000 x25: 0000000000000040 x24: ffff0000c46a0000 x23: 0000000000000000 x22: ffff0000c0c73c00 x21: 0000000000000010 x20: 0000000000000010 x19: 0000000000000000 x18: 000000000000002c x17: 0000000000000000 x16: 0000000000000001 x15: 0000000000000000 x14: 0000000000000400 x13: ffff8000837dc000 x12: 0000000000000000 x11: ffff0000c0c73ca0 x10: 0000000000000040 x9 : 459ec1f0abbdbb00 x8 : 0000001ffffffe0 x7 : 0000000000000000 x6 : 000000000000003f x5 : 0000000000035579 x4 : 0000000000000901 x3 : 0000000000000000 x2 : 0000000000000000 x1 : 0000000000000000 x0 : 0000000000000010 Call trace: sg_init_table+0x2c/0x50 (P) mmc_mq_init_request+0x64/0x90 blk_mq_alloc_map_and_rqs+0x3ac/0x480 blk_mq_alloc_set_map_and_rqs+0x98/0x1e0 blk_mq_alloc_tag_set+0x1c0/0x290 mmc_init_queue+0x120/0x370 mmc_blk_alloc_req+0x150/0x420 To prevent this, add a validation check in mmc_mq_init_request() to detect when sg_len (derived from max_segs) is zero. If sg_len is zero, we return an error and print an error message, allowing host driver developers to identify and fix incorrect max_segs configuration. This is a defensive measure that ensures the MMC core fails gracefully when host drivers provide invalid max_segs values, rather than crashing with a page fault.	N/A	More Details
CVE-2026-98886	In the Linux kernel, the following vulnerability has been resolved: drm/gma500: return errors from Oaktrail HDMI I2C reads xfer_read() waits for the HDMI I2C transaction to reach I2C_TRANSACTION_DONE, but it ignores both timeout and signal returns from wait_for_completion_interruptible_timeout(). If the interrupt never advances the transaction state, the loop can wait forever. Return -ETIMEDOUT when the completion wait expires, propagate interrupted waits, and make the I2C master_xfer callback return the first transfer error instead of reporting a successful message count.	N/A	More Details
CVE-2026-98912	In the Linux kernel, the following vulnerability has been resolved: ALSA: seq: oss: Reject reads that cannot fit the next event snd_seq_oss_read() checks whether the next queued OSS sequencer event fits in the remaining userspace buffer before removing it from the read queue. The check is inverted. It currently stops when the event is smaller than the remaining buffer, so a normal 4-byte event is not copied for an 8-byte read buffer. Conversely, an 8-byte event can be copied for a smaller read count. Break only when the remaining userspace buffer is smaller than the next event, and report -EINVAL if no complete event has been copied. This prevents an undersized read from looking like end-of-file while leaving the event queued for a later read with a large enough buffer.	N/A	More Details
CVE-2026-98913	In the Linux kernel, the following vulnerability has been resolved: net: thunderx: fix PTP device ref leak in nicvf_probe() cavium_ptp_get() acquires a reference to the PTP PCI device through pci_get_device(). If any initialization step fails after cavium_ptp_get(), the PTP PCI device reference is leaked. Add a common error path to release the PTP reference before returning from probe failures.	N/A	More Details
CVE-2026-98914	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: cap ATOM command table nesting depth Cap nesting at 32 levels with execute_depth and return -ELOOP when exceeded.	N/A	More Details
CVE-2026-98915	In the Linux kernel, the following vulnerability has been resolved: media: video-i2c: use vb2_video_unregister_device on driver removal The driver uses vb2_fop_release() as its file release operation, so vb2_video_unregister_device() should be used instead of video_unregister_device() during driver removal. This ensures that the vb2 queue is properly disconnected before the video device is unregistered.	N/A	More Details
CVE-2026-98916	In the Linux kernel, the following vulnerability has been resolved: wifi: nl80211: check link is beaconing for color change When trying to do a color change, validate that the link is beaconing first, to avoid calling the driver with an invalid link.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ima: return error early if file xattr cannot be changed During early boot, the filesystem is read-only and any changes to xattrs are not		

CVE-2026-98885	allowed. This fails in case of ext4 because changing xattr starts an ext4 transaction which fails with the following warning. WARNING: fs/ext4/ext4_jbd2.c:75 at ext4_journal_check_start+0x63/0xa0 [ext4], CPU#1: systemd-sysroot/561 CPU: 1 UID: 0 PID: 561 Comm: systemd-sysroot Not tainted 6.19.12-1-default #1 PREEMPT(voluntary) openSUSE Tumbleweed c2dfc3c9d9f6f1233251c5d4410574fe82a348ee Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS unknown 02/02/2022 RIP: 0010:ext4_journal_check_start+0x63/0xa0 [ext4] Call Trace: __ext4_journal_start_sb+0x3e/0x180 [ext4 6d025f3bc52c89a957b89a89d211fadf5e9434e1] ext4_xattr_set+0x104/0x150 [ext4 6d025f3bc52c89a957b89a89d211fadf5e9434e1] __vfs_setxattr+0x9a/0xd0 __vfs_setxattr_noperm+0x76/0x1f0 ima_appraise_measurement+0x23e/0xe40 ima_d_path+0x5a/0xd0 process_measurement+0xb29/0xc40 ? copy_from_kernel_nofault+0x21/0xe0 ? fscrypt_file_open+0xc0/0xe0 ? ext4_file_open+0x60/0x490 [ext4 6d025f3bc52c89a957b89a89d211fadf5e9434e1] ? bpf_prog_31efb7c56239148b_restrict_filesystems+0xab/0x126 ? __bpf_prog_exit+0x23/0xd0 ? __bpf_trampoline_exit+0xd/0x50 ? bpf_trampoline_6442530367+0x9f/0xea ima_file_check+0x57/0x80 security_file_post_open+0x50/0xf0 path_openat+0x493/0x1650 do_filp_open+0xc7/0x170 Detect the state of the file early and return the error.	N/A	More Details
CVE-2026-98883	In the Linux kernel, the following vulnerability has been resolved: serial: 8250_port: recognize UPIO_AU My MIPS Alchemy systems generate the following warning during probe of the 8250 driver: WARNING: drivers/tty/serial/8250/8250_port.c:462 at set_io_from_upio+0xfc/0x124, CPU#0: swapper/0/1 Unsupported UART type 4 [...] [<80521d40>] set_io_from_upio+0xfc/0x124 [<80521dfc>] serial8250_set_defaults+0x94/0xe0 [<80520fb4>] serial8250_register_8250_port+0x288/0x51c [<805214ec>] serial8250_probe+0x160/0x1e8 [<8053b5f0>] platform_probe+0x64/0x90 The least invasive fix is to recognize UPIO_AU (type 4) in set_io_from_upio() and do nothing, since all parameters have already been set up in 8250_rt288x.c::au_platform_setup(). Run-tested on Alchemy Au1300 platform.	N/A	More Details
CVE-2026-98948	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_flower: Fix port range register leak in tmplt_create() mlxsw_sp_flower_tmplt_create() parses a flow_cls_offload template into a stack-local struct mlxsw_sp_acl_rule_info purely to compute rulei.values.elusage. Parsing can acquire port range registers via mlxsw_sp_flower_parse_ports_range(), but since this rulei never goes through mlxsw_sp_acl_rulei_destroy(), those registers were never released, including through chain template deletion. Factor out of mlxsw_sp_acl_rulei_destroy() the code to actually release the necessary resources and call from mlxsw_sp_flower_tmplt_create() to plug the leak. The issue was found during a review of Wentao Liang's patch referenced below.	N/A	More Details
CVE-2026-98882	In the Linux kernel, the following vulnerability has been resolved: media: chips-media: wave5: Add range checks for dec_output_info If the driver's dec_output_info contains invalid data the driver can write in invalid memory. Add a range check for that. This fixes this smatch error: drivers/media/platform/chips-media/wave5/wave5-vpuapi.c:588 wave5_vpu_dec_get_output_info() error: buffer overflow 'inst->frame_buf' 64 <= 127	N/A	More Details
CVE-2026-98881	In the Linux kernel, the following vulnerability has been resolved: net: bridge: remove stale rcu_barrier() in br_multicast_dev_del() This rcu_barrier() came from a time call_rcu() calls were used in net/bridge/br_multicast.c. Now kfree_rcu() is there, we can remove this problematic rcu_barrier() which causes extreme RTNL pressure in many syzbot reports. INFO: task syz-executor:77945 is blocked on a mutex likely owned by task kworker/u1024:5:36537. task:kworker/u1024:5 state:D stack:24616 pid:36537 tgid:36537 ppid:2 task_flags:0x4208060 flags:0x00080000 last_sleep:612797637337 Workqueue: netns cleanup_net Call Trace: <TASK> [<ffffffff81914eaa>] context_switch+0xf2a/0x1730 kernel/sched/core.c:6483 [<ffffffff81201143>] __schedule+0x1133/0x43a0 kernel/sched/core.c:8411 [<ffffffff8120446b>] __schedule_loop kernel/sched/core.c:8514 [inline] [<ffffffff8120446b>] schedule+0xab/0x260 kernel/sched/core.c:8529 [<ffffffff8121a093>] schedule_timeout+0xc3/0x2b0 kernel/time/sleep_timeout.c:75 [<ffffffff81205347>] do_wait_for_common kernel/sched/completion.c:100 [inline] [<ffffffff81205347>] __wait_for_common kernel/sched/completion.c:121 [inline] [<ffffffff81205347>] wait_for_common kernel/sched/completion.c:132 [inline] [<ffffffff81205347>] wait_for_completion+0x2c7/0x5d0 kernel/sched/completion.c:153 [<ffffffff81b8f27f>] rcu_barrier+0x49f/0x620 kernel/rcu/tree.c:3888 [<ffffffff860091b3>] br_multicast_dev_del+0x303/0x350 net/bridge/br_multicast.c:4459 [<ffffffff85fb5dbc>] br_dev_uninit+0x1c/0x40 net/bridge/br_device.c:157 [<ffffffff8568058c>] unregister_netdevice_many_notify+0x1c1c/0x2300 net/core/dev.c:12599 [<ffffffff8562be43>] ops_exit_rtnl_list net/core/net_namespace.c:187 [inline] [<ffffffff8562be43>] ops_undo_list+0x3d3/0x940 net/core/net_namespace.c:248	N/A	More Details
CVE-2026-98880	In the Linux kernel, the following vulnerability has been resolved: ntfs3: handle set_blocksize failures ntfs3 uses buffer_heads, which don't handle block size > PAGE_SIZE well. Without this, mounting we will hit the BUG_ON(offset >= folio_size(folio)); in folio_set_bh on the first __bread_gfp call.	N/A	More Details
CVE-	pghoard is a PostgreSQL backup daemon and restore tooling that stores backup data in cloud object		

2026-54711	stores. Prior to version 2.7.1, when using .pgpass, database connection information including the username and password will be logged at the debug level. This issue has been patched in version 2.7.1.	N/A	More Details
CVE-2026-98879	In the Linux kernel, the following vulnerability has been resolved: isofs: handle set_blocksize failures isofs uses buffer_heads, which don't handle block size > PAGE_SIZE well. Without this, mounting we will hit the BUG_ON(offset >= folio_size(folio)); in folio_set_bh on the first __bread_gfp call.	N/A	More Details
CVE-2026-55888	XDG Desktop Portal is a desktop integration portal. Prior to version 1.22.1, malicious sandboxed applications can use the FileChooser.SaveFiles method to gain arbitrary file access to non-existing files on the host (limited to files the user has access to). The FileChooser.SaveFiles method takes optional filenames which get appended to the selected directory by backends. The backends are supposed to check that the file does not exist yet. The filenames are not validated by the frontend and also not by common backends. This issue has been patched in version 1.22.1.	N/A	More Details
CVE-2026-55889	XDG Desktop Portal is a desktop integration portal. Prior to version 1.22.1, malicious sandboxed applications can use the FileTransfer.RetrieveFiles method to redirect drag-and-drop and copy-paste data targeted at other applications to itself. Some portals generate tokens which provide access to resources. The sharing of the token transfers access to the resource and is supposed to happen out of band. The tokens were generated with g_random_int or g_uuid_string_random, both of which use GLib's GRandom based on the Mersenne Twister PRNG which created a stream of reproducible random numbers. By generating enough tokens, the state of the GRandom can be recover and with it, the next tokens can be predictedinsufin. This issue has been patched in version 1.22.1.	N/A	More Details
CVE-2026-98910	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Unwind debug trap enable on copy_to_user failure If kfd_dbg_trap_enable() fails while copying runtime_info to userspace, it had already activated the trap, set debug_trap_enabled, taken an extra process reference, and opened the debug event file. Return -EFAULT without unwinding that state, leaving inconsistent trap state and a refcount imbalance that could break later DISABLE/ENABLE. On copy_to_user failure, deactivate the trap and undo the rest of the enable setup before returning.	N/A	More Details
CVE-2026-98909	In the Linux kernel, the following vulnerability has been resolved: net: cpsw_new: unregister devlink on port registration failure cpsw_probe() registers devlink before registering the CPSW ports. If cpsw_register_ports() fails, the error path only unregisters the notifiers and then releases the lower level resources. It does not undo the successful cpsw_register_devlink() call, leaving the devlink instance and its parameters registered after probe has failed. Add a devlink cleanup label for the path where devlink registration has already succeeded, and use it when port registration fails.	N/A	More Details
CVE-2026-98908	In the Linux kernel, the following vulnerability has been resolved: net: microchip: sparx5: clean up PSFP resources on flower setup failure sparx5_tc_flower_psfp_setup() allocates PSFP stream gate, flow meter and stream filter resources before adding VCAP actions. If a later step fails, the resources allocated earlier in the function are not unwound. Add error paths to release the stream filter, flow meter and stream gate when setup fails after they have been acquired. Also make sparx5_psfp_fm_add() return the acquired flow-meter id before the existing-flow-meter early return. When an existing flow meter is reused, sparx5_psfp_fm_get() increments its pool reference count, but the caller previously kept psfp_fm_id as 0. If a later setup step failed, the error path could try to delete flow-meter id 0 instead of the reused flow meter, leaving the incremented reference behind.	N/A	More Details
CVE-2026-98907	In the Linux kernel, the following vulnerability has been resolved: btrfs: protect sb_write_pointer() with invalidate lock sb_write_pointer() reads the super block from the block device page cache using read_cache_page_gfp(). This has the same race with BLKBSZSET as the one fixed by commit 3f29d661e568 ("btrfs: sync read disk super and set block size"). Take the mapping invalidate lock around read_cache_page_gfp() to serialize the read against block size changes.	N/A	More Details
CVE-2026-98892	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scm: Validate BASE_ERROR_EVENT payload size BASE_ERROR_EVENT carries a variable number of message reports, with the count encoded in error_status. The notification parser used that count without checking whether the received payload contained all reported entries. Reject truncated payloads before copying the report array.	N/A	More Details
CVE-2026-98893	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scm: Validate SENSOR_UPDATE payload size SENSOR_UPDATE carries one or more sensor readings after the fixed notification header. The parser derives the expected reading count from the sensor description, but it did not verify that the received payload contains those entries before parsing them. Reject truncated update notifications before reading the variable array.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: drm/amd/ras: reset CPER ring on corrupt entry size When CPER ring overflow handling advances the read pointer, it trusts the parsed entry size from the current ring contents. Corrupt CPER data can produce an entry size that does not advance rp_ptr after dword conversion and pointer masking. In that case the recovery loop keeps testing	N/A	More

98894	the same location while holding the CPER ring mutex. This can hang the worker that is writing the next CPER record. Detect a no-progress rprr update and reset the CPER ring to an empty state instead. This drops the corrupt contents and lets the writer leave the recovery path without spinning.		Details
CVE-2026-98888	In the Linux kernel, the following vulnerability has been resolved: arm64: kprobes: Allow reentering kprobes while single-stepping A kprobe can be hit while another kprobe is in KPROBE_HIT_SS state. This can happen when tracing or perf code runs from the debug exception path while the first kprobe is preparing or executing its out-of-line single-step instruction. Currently arm64 treats a kprobe hit in KPROBE_HIT_SS as unrecoverable, the same as a hit in KPROBE_REENTER. This is too strict. A hit in KPROBE_HIT_SS is still a one-level reentry and can be handled by saving the current kprobe state and setting up single-step for the new probe, just like reentry from KPROBE_HIT_ACTIVE or KPROBE_HIT_SSDONE. The truly unrecoverable case is hitting another kprobe while already in KPROBE_REENTER, because the reentry save area has already been consumed. Move KPROBE_HIT_SS to the recoverable reentry cases and leave KPROBE_REENTER as the unrecoverable nested reentry case. This change also requires saving saved_irqflag in struct prev_kprobe. When a nested kprobe calls kprobes_save_local_irqflag(), it overwrites kcb->saved_irqflag with the currently masked DAIF value, losing the outer kprobe's original DAIF state. Without this fix, when the outer kprobe's single-step finishes, kprobes_restore_local_irqflag() applies the wrong DAIF mask and leaves interrupts permanently disabled. Extend struct prev_kprobe with a saved_irqflag field and save/restore it alongside kp and status. This ensures the outer kprobe's original interrupt state is preserved across reentry. This mirrors the x86 fix in commit 6a5022a56ac3 ("kprobes/x86: Allow to handle reentered kprobe on single-stepping").	N/A	More Details
CVE-2026-98896	In the Linux kernel, the following vulnerability has been resolved: thunderbolt: Keep the domain reference while processing hotplug We process hotplug events in a workqueue that may run after the domain has been removed by tb_domain_remove(). For example if user unloads the driver while at the same time plugging a device router we may have scheduled tb_handle_hotplug() to run. Avoid possible UAF in this case by taking the domain reference before scheduling the hotplug handler in tb_queue_hotplug().	N/A	More Details
CVE-2026-98897	In the Linux kernel, the following vulnerability has been resolved: ksmbd: don't hold ci->m_lock while waiting for a lease break ack When a cifs.ko client caches a read-handle (RH) lease via deferred close and a conflicting open arrives, ksmbd breaks the lease and waits for the acknowledgment in wait_for_break_ack() for up to OPLOCK_WAIT_TIME (35s). __smb_break_all_level_oplock() runs that wait while holding ci->m_lock for read. cifs.ko reacts to a handle-lease break by closing the deferred handle rather than sending a lease break acknowledgment. That close path (close_id_del_oplock() -> opinfo_del()) takes ci->m_lock for write and is exactly what would wake the waiter, but it blocks on the read lock held by the waiting thread. The break is then resolved only by the 35s timeout, so xfstests generic/001 takes ~78s with leases enabled versus ~4s with oplocks only. Collect the target opinfos (each pinned with a reference) while holding ci->m_lock, then break them after releasing it, matching how smb_grant_oplock() already breaks a conflicting lease using only a reference. The reference keeps the opinfo (and its conn and lease) alive across the unlocked window, and a close racing the break is handled by the existing OPLOCK_CLOSING state check. Apply the same fix to the parent lease break paths.	N/A	More Details
CVE-2026-98898	In the Linux kernel, the following vulnerability has been resolved: drm/amd/pm/si: Don't schedule thermal work when queue isn't initialized When DPM is turned off with the amdgpu.dpm=0 module parameter, the thermal work queue isn't initialized so we shouldn't schedule any work on it. (cherry picked from commit bd018d36171a695952c6d391471c279c9e05c8b2)	N/A	More Details
CVE-2026-98887	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211_hwsim: reject undersized HWSIM_ATTR_TX_INFO hwsim_tx_info_frame_received_nl() casts the HWSIM_ATTR_TX_INFO payload to a struct hwsim_tx_rate * and unconditionally reads IEEE80211_TX_MAX_RATES entries (8 bytes) from it. The policy only bounds the attribute from above (NLA_BINARY .len is a maximum) and the op sets GENL_DONT_VALIDATE_STRICT, so a short or zero-length attribute is accepted and the loop reads past the payload. Require the exact length in the policy, so a malformed attribute is rejected before the handler runs.	N/A	More Details
CVE-2026-98900	In the Linux kernel, the following vulnerability has been resolved: xen/gntalloc: validate grant count before allocation gntalloc_ioctl_alloc() allocates the grant-id array before checking whether the requested count fits within the global grant limit. Counts above that limit cannot succeed, so reject them before the user-controlled allocation reaches kcalloc(). Use a subtraction-based check while holding gref_mutex so adding the requested count cannot wrap. Also cast the count before advancing the per-file index so the page-size multiplication is performed in 64-bit arithmetic.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix lease break and ack state handling Do not skip valid lease states containing WRITE_CACHING when breaking level-II/read leases		More

2026-98901	for writes and truncates. Handle lease break acknowledgments according to the SMB2 rule that the acknowledged state must be a subset of the server's break target. Apply the acknowledged state directly and keep the break pending on failed ACKs.	N/A	Details
CVE-2026-98902	In the Linux kernel, the following vulnerability has been resolved: md/raid5: let stripe batch bm_seq comparison wrap-safe Once the 32-bit seq wraps, a newer bm_seq can look smaller than old, so .. covert to wrap-safe calculate way.	N/A	More Details
CVE-2026-98903	In the Linux kernel, the following vulnerability has been resolved: mailbox: imx: Add a channel shutdown field sashiko complained about possible teardown problem. The scenario CPU 0 CPU 1 imx_mu_isr() imx_mu_shutdown() imx_mu_xcr_rmw(priv, IMX_MU_RCR, 0, IMX_MU_xCR_RIE(priv->dcfg->type, cp->idx)); imx_mu_specific_rx() imx_mu_xcr_rmw(priv, IMX_MU_RCR, IMX_MU_xCR_RIE(priv->dcfg->type, 0), 0); free_irq() The RX event remains enabled because in this short window the RX event was disabled in ->shutdown() while the interrupt was active and then enabled again by the ISR while ->shutdown waited in free_irq(). This race requires timing and if happens can be problematic on shared handlers if the "removed" channel triggers an interrupt. In this case the irq-core will shutdown the interrupt with the "nobody cared" message. Introduce imx_mu_con_priv::shutdown to signal that the channel is shutting down. This flag is set with the lock held (by imx_mu_xcr_clr_shut()). The unmask side uses imx_mu_xcr_set_act() which only enables the event if the channel has not been shutdown and serialises on the same lock.	N/A	More Details
CVE-2026-98904	In the Linux kernel, the following vulnerability has been resolved: mailbox: imx: Use devm_pm_runtime_enable() sashiko complained about early usage of the device while probe isn't completed. This can be mitigated by delaying the pm_runtime_enable() into the removal path instead doing it early. This ensures that in an error case the device is removed (and imx_mu_shutdown()) before pm_runtime_disable() so we don't have to do this manually. For the order to work, lets move devm_mbox_controller_register() until after the pm-runtime part. So the reverse order will be mbox-controller removal followed by disabling pm runtime. Use devm_pm_runtime_enable(), remove manual pm_runtime_disable() invocations and move the pm_runtime handling in probe before devm_mbox_controller_register().	N/A	More Details
CVE-2026-98905	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Properly acquire queue buffers in CRIU restore When kfd_queue_acquire_buffers() was split off from set_queue_properties_from_user(), set_queue_properties_from_criu() was missed. Thus, set_queue_properties_from_criu() is not filling out the buffer fields of queue_properties, which can come up when subsequent code expects them to be non-null. Add the proper call to kfd_queue_acquire_buffers(), and also use the right cast types in set_queue_properties_from_criu() (which were missed at the same time) (cherry picked from commit 88ed96abbb27b70193544fbc1ee06448c274714)	N/A	More Details
CVE-2026-98906	In the Linux kernel, the following vulnerability has been resolved: net: lan966x: restore RX state on reload failure lan966x_fdma_reload() backs up rx->page_pool and rx->fdma before reallocating the RX resources for the new MTU. If the allocation fails, the restore path puts these fields back before restarting RX. However, the reload path also updates rx->page_order and rx->max_mtu before calling lan966x_fdma_rx_alloc(). These fields are not restored on failure, so RX can be restarted with the old pages, old FDMA state and old page pool, but with the page geometry from the failed new MTU. This can make the XDP path advertise a frame size derived from the new page_order while the actual RX pages still come from the old allocation. For example, after a failed reload to a jumbo MTU, xdp_init_buff() may be called with a frame size larger than the restored RX pages. lan966x_fdma_rx_alloc_page_pool() also registers the newly allocated page pool with each port's XDP RXQ before fdma_alloc_coherent() is called. If fdma_alloc_coherent() fails, the new page pool is destroyed, but the rollback path does not restore the per-port XDP RXQ mem model registration either. Save and restore rx->page_order and rx->max_mtu, and restore the old page pool registration for each port's XDP RXQ before RX is started again. This keeps the restored RX state consistent after a failed reload.	N/A	More Details
CVE-2026-98878	In the Linux kernel, the following vulnerability has been resolved: irqchip/gic-v5: Immediately exec priority drop following activate With GICv5 an interrupt of equal or lower priority cannot be signalled until there has been a priority drop. This is done via the GIC CDEOI system instruction. Once this has been executed, the hardware is able to signal the next interrupt if there is one. As all interrupts are programmed to have the same priority, no new interrupts can be signalled until the priority drop has happened. This can cause issues when, for example, an interrupt remains active while a long running process takes place, such as when injecting a physical interrupt into a guest VM in software. The GICv5 driver has so far done the priority drop as part of irq_eoi(), i.e., at the same time as deactivating the interrupt. This means that any long running process (or VM) could block incoming interrupts, effectively causing a denial of service for all other interrupts. Rather than doing the EOI as part of irq_eoi() (which the name would suggest would be a good place for it), move it to happen immediately after acknowledging an interrupt in the main GICv5 interrupt handler. The deactivation of interrupts (GIC CDDI) remains implemented as part of irq_eoi(), which means that the same interrupt cannot be	N/A	More Details

	signalled a second time until deactivated by software.		
CVE-2026-98877	In the Linux kernel, the following vulnerability has been resolved: irqchip/gic-v4: Don't advertise VLPIs if no ITS is probed When accidentally setting "kvm-arm.vgic_v4_enable=1" on a system that has no MSI controller device tree node and GICv4, it results a panic as "gic_domain" is NULL and the kernel attempts to access it. Unable to handle kernel NULL pointer dereference at virtual address 0000000000000028 Mem abort info: ESR = 0x0000000096000006 CPU: 1 UID: 0 PID: 295 Comm: lkvm-static Not tainted 7.1.0-rc4-ge3f15ad3970e #5 PREEMPT Hardware name: linux,dummy-virt (DT) pstate: 81402005 (Nzcv daif +PAN -UAO -TCO +DIT -SSBS BTYPE=--) pc : __irq_domain_instantiate+0x1d4/0x578 lr : __irq_domain_instantiate+0x1cc/0x578 Set vLPI support to false at init time if the host has no ITS, so it propagates properly to kvm_vgic_global_state.has_gicv4.	N/A	More Details
CVE-2026-98876	In the Linux kernel, the following vulnerability has been resolved: host1x: bus: Fix missing ops null check in error teardown In host1x_device_init(), the error teardown paths do not check client->ops before dereferencing it, unlike the forward init paths which correctly guard with 'client->ops &&'. This can result in a NULL pointer dereference if client->ops is NULL. Fix by adding the missing client->ops check in both the teardown and teardown_late labels.	N/A	More Details
CVE-2026-98928	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Derive GUEST_HAS_SVE from the SVE feature bit at EL2 pkvm_init_features_from_host() takes KVM_ARCH_FLAG_GUEST_HAS_SVE and KVM_ARM_VCPU_SVE from the host separately, but pkvm_vcpu_init_sve() tests the bit while vcpu_has_sve() reads the flag. A host that sets the flag without the bit gets a vCPU with a NULL sve_state that the world switch loads the guest's SVE state from. Derive the flag from the bit, and drop the protected path's copy of the host's flag, which is dead code since protected VMs are not allowed SVE.	N/A	More Details
CVE-2026-98930	In the Linux kernel, the following vulnerability has been resolved: KVM: s390: Fix IRQ injection with SIGP Stop and Store Status When __inject_sigp_stop() is called for a Stop and Store Status operation, if the vCPU is running, the interrupt is marked as pending and the status is stored by the thread performing the KVM_RUN IOCTL. If the vCPU is already stopped, the status is stored immediately. Storing the status means writing into userspace, which might fault, and __inject_sigp_stop() is called from do_inject_vcpu() which in turn is always called holding a spinlock, which is obviously an issue. Fix this by returning -EWOULDBLOCK from __inject_sigp_stop(), and adding a bool flag to indicate whether a store status is needed. The callers of do_inject_vcpu() are modified to pass the pointer to the bool flag; whenever a Store Status operation is needed, the callers can now perform it outside the spinlock. Opportunistically refactor kvm_s390_set_irq_state() to use scoped_guard() and __free(). [Added Fixes tag while picking -- Claudio]	N/A	More Details
CVE-2026-98931	In the Linux kernel, the following vulnerability has been resolved: mtd: spinand: fix NULL pointer dereference with no ECC engine When "nand-no-ecc-engine" is set in DT, nanddev_get_ecc_engine() takes the NAND_ECC_ENGINE_TYPE_NONE path and returns success while leaving nand->ecc.engine NULL. The SPI-NAND code nevertheless dereferences it unconditionally to test for a pipelined engine, so probing such a device oopses immediately. Rather than open-coding the test three times, add a nand_ecc_is_pipelined() helper to the NAND core that folds the NULL check into the integration comparison, and use it everywhere. Future callers then cannot reintroduce the problem.	N/A	More Details
CVE-2026-98933	In the Linux kernel, the following vulnerability has been resolved: usb: hub: use shorter 120ms post resume hold for SS root hubs Holding a USB3 hub PM runtime reference for 200ms at hub resume triggers an AMD platform issue. Users running Android adb report crashes after adb has been polling and waking up the USB subsystem once a second for some time. Vendors are working on a solution. Disabling USB runtime PM is one way to prevent this issue, but it is also proven that reducing the hold time to 120ms in resume also mitigates it. See Link for more details. Reducing the hold time to 120ms for the USB3 roothub is in itself a valid change and optimization, as the current 200ms is excessive: a root hub has no upstream suspended hub whose wake propagation needs to be accounted for, but still need some time for USB3 link training to complete. Keep the 200ms hold for external hubs that commit 8f5b7e2bec1c ("usb: hub: fix detection of high tier USB3 devices behind suspended hubs") is intended for.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: ipv6: sr: use skb_get_hash_net() in seg6_make_flowlabel() Since commit d58e468b1112 ("flow_dissector: implements flow dissector BPF hook") __skb_flow_dissect() needs a net pointer, either from skb->dev, skb->sk, or since commit 3cbf4ffba5ee ("net: plumb network namespace into __skb_flow_dissect") a caller provided pointer. syzbot was able to reach seg6_make_flowlabel() with an skb having neither skb->dev nor skb->sk set: a TIPC UDP bearer sends a discovery message through an IPv4 route using seg6 encap, while net.ipv6.seg6_flowlabel is set to 1. seg6_make_flowlabel() already has a net pointer, use skb_get_hash_net(). WARNING: net/core/flow_dissector.c:1131 at __skb_flow_dissect+0x910/0x5368 net/core/flow_dissector.c:1126, CPU#0: syz.0.17/4930 Call trace: __skb_flow_dissect+0x910/0x5368 net/core/flow_dissector.c:1126 (P) __skb_get_hash_net+0xe0/0x29c net/core/flow_dissector.c:1903		

2026-98934	skb_get_hash include/linux/skbuff.h:1663 [inline] seg6_make_flowlabel+0xcc/0x1ec net/ipv6/seg6_ip tunnel.c:132 __seg6_do_srh_encap+0x320/0xbf4 net/ipv6/seg6_ip tunnel.c:161 seg6_do_srh+0x4b4/0xa44 net/ipv6/seg6_ip tunnel.c:431 seg6_output_core+0x164/0x688 net/ipv6/seg6_ip tunnel.c:681 seg6_output+0x44/0x1ac net/ipv6/seg6_ip tunnel.c:744 lwtunnel_output+0x3d0/0x664 net/core/lwtunnel.c:356 dst_output include/net/dst.h:470 [inline] ip_local_out+0x110/0x148 net/ipv4/ip_output.c:131 iptunnel_xmit+0x50c/0xd38 net/ipv4/ip_tunnel_core.c:97 udp_tunnel_xmit_skb+0x220/0x348 net/ipv4/udp_tunnel_core.c:187 tipc_udp_xmit+0x75c/0x9c0 net/tipc/udp_media.c:202 tipc_udp_send_msg+0x214/0x374 net/tipc/udp_media.c:274 tipc_bearer_xmit_skb+0x260/0x3b0 net/tipc/bearer.c:576 tipc_enable_bearer net/tipc/bearer.c:366 [inline] __tipc_nl_bearer_enable+0xc90/0xfb0 net/tipc/bearer.c:1048 tipc_nl_bearer_enable+0x2c/0x48 net/tipc/bearer.c:1057 genl_family_rcv_msg_doit+0x1e4/0x2d4 net/netlink/genetlink.c:1114	N/A	More Details
CVE-2026-98935	In the Linux kernel, the following vulnerability has been resolved: rtc: dev: zero-initialize struct rtc_wkalrm to prevent information leak The struct rtc_wkalrm has padding holes after the 'pending' member. When returned to userspace via the RTC_WKALRM_RD ioctl, uninitialized kernel stack data residing in these holes may be leaked. Zero-initialize the struct at declaration to eliminate this information leak risk. Fixes smatch warnings: - drivers/rtc/dev.c:392 rtc_dev_ioctl() warn: check that 'alarm' doesn't leak information (struct has a hole after 'pending')	N/A	More Details
CVE-2026-98936	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: reset ar->num_stations on hardware start When (multiple) hw restart occurs, we end up in a situation where we cannot accept new stations / mesh peers. This seems to be the 'num_stations' in 'struct ath11k' that did not get reset properly in this case. In ath11k_mac_op_start(), it was indeed the only accounting var that did not get reset so let's clear it too. This avoids problems like those: [1248828.088023] WARNING: CPU: 0 PID: 12281 at net/mac80211/util.c:1682 ieee80211_reconfig_stations+0x98/0xc0 [1248828.088115] Call trace: [1248828.088116] ieee80211_reconfig_stations+0x98/0xc0 (P) [1248828.088119] ieee80211_reconfig+0x774/0x1438 [1248828.088123] ieee80211_restart_work+0x10c/0x178 [1248828.088127] process_one_work+0x178/0x3c0 [1248828.088130] worker_thread+0x280/0x440 [1248828.088133] kthread+0xe0/0xf8 [1248828.088137] ret_from_fork+0x10/0x20 [1248828.088141] ---[end trace 0000000000000000]--- [1248836.686604] ath11k_pci 0000:04:00.0: refusing to associate station: too many connected already (128) [1248836.686609] ath11k_pci 0000:04:00.0: Failed to add station: xx:xx:xx:xx:xx:xx for VDEV: 1 Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.9.0.1-01977-QCAHKSUPL_SILICONZ-1	N/A	More Details
CVE-2026-98937	In the Linux kernel, the following vulnerability has been resolved: perf/core: Check kernel access when kernel callchains are requested perf_event_open() currently gates perf_allow_kernel() only on !attr.exclude_kernel. However, users can still request kernel callchain collection with attr.exclude_callchain_kernel == 0 even when attr.exclude_kernel == 1. That still requires kernel profiling privilege, but the existing check does not enforce it. Update the permission check to call perf_allow_kernel() when either kernel sampling is requested or kernel callchains are requested. This keeps permission checks aligned with requested data and prevents unprivileged use of kernel callchain capture.	N/A	More Details
CVE-2026-98938	In the Linux kernel, the following vulnerability has been resolved: rust: irq: pass RegistrationInner as cookie to request_irq With commit ba268514ea14 ("rust: devres: fix race condition due to nesting"), Devres::new() returns Result<Self> by value instead of initializing in-place via PinInit. Because request_irq() is called inside Devres::new(), registration.inner is still uninitialized when the IRQ is enabled, so accessing registration.inner.device() in the IRQ callback can read uninitialized memory. Fix this by storing the Device and handler pointer in RegistrationInner and passing RegistrationInner as the IRQ cookie after its fields are initialized. This is not needed in mainline because commit 98c63ce4d760 ("rust: irq: make Registration compatible with lifetime-bound drivers") removed Devres from irq::Registration and moved request_irq() after all fields are initialized.	N/A	More Details
CVE-2026-98939	In the Linux kernel, the following vulnerability has been resolved: x86/mm: Drop unnecessary PMD page copy when freeing On a box with a discrete GPU, lockdep reports a possible deadlock as soon as kswapd shrinks the TTM page pool. The immediate cause is an x86 commit that added an mmap_read_lock() to kernel page protection munging code. The huge vmmap code holds the same lock over a GFP_KERNEL allocation, which is a no-no now that reclaim can take it. That allocation is in a page table *free* path and ends up being for dubious purposes[1]. Basically, it tries to avoid hardware setting Accessed=1 in page table entries that are unreachable by the hardware, a non-issue. Remove the PMD copy. Detach the original PMD page at the PUD, flush the mid-level caches, and free the PTE tables straight from the detached PMD page. With no allocation left, the locking issue is gone. Lockdep splat/analysis: WARNING: possible circular locking dependency detected 7.3.0-rc3-f6e7b42bf05b+ #183 Tainted: G U ----- kswapd0/269 is trying to acquire lock: ((init_mm).mmap_lock){++++}-{4:4}, at: change_page_attr_set_clr+0x29a/0x4a0 but task is already holding lock: (pool_shrink_rwsem){.+.-}{4:4}, at: ttm_pool_shrink+0xb2/0x330 [ttm] Chain exists of: (init_mm).mmap_lock --> fs_reclaim --> pool_shrink_rwsem The cycle is built from three edges: 1)	N/A	More Details

	pool_shrink_rwsem -> (init_mm).mmap_lock The TTM shrinker restores the caching attribute of every page it frees, while holding pool_shrink_rwsem: ttm_pool_shrink() -> ttm_pool_dispose_list() -> ttm_pool_free_page() -> set_pages_wb() -> change_page_attr_set_clr() [init_mm mmap read lock] 2) fs_reclaim -> pool_shrink_rwsem The same shrinker, called from reclaim. 3) (init_mm).mmap_lock -> fs_reclaim ioremap() installing a huge PUD mapping over an existing PMD table: ioremap_page_range() -> vmmap_range_noflush() -> vmmap_try_huge_pud() [init_mm mmap read lock] -> pud_free_pmd_page() -> __get_free_page(GFP_KERNEL) [enters reclaim] [dhansen: Lots of changelog munging/trimming and merged comments from my version of the fix.]		
CVE-2026-98940	In the Linux kernel, the following vulnerability has been resolved: ASoC: qcom: lpass-cpu: use DAI table index for playback constraints The DAI ID is a hardware port ID and is not necessarily the index of the corresponding entry in variant->dai_driver. When configuring a DAI with LPAIF_I2SCTL_MODE_QUAD01, the probe code currently uses dai_id to index variant->dai_driver. This is incorrect for platforms where DAI IDs are sparse. For example, the IPQ806x MI2S DAI has ID 4 while it is the only entry in the DAI driver table. If that DAI is configured for QUAD01 and probe reaches this branch, the code writes beyond that single entry via dai_driver[4] instead of updating the DAI being processed. Use the loop index i when updating the current DAI's playback channel constraints, while retaining dai_id for indexing the hardware-port specific playback SD-line mode array. This fixes the incorrect DAI table access for platforms where the DAI ID does not match its position in the driver table.	N/A	More Details
CVE-2026-98942	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Build the cid tables privately and publish them with RCU The cid tables are visible to the cid kfuncs while being modified: the first enable publishes the global pointers before filling them, ops.init_cids() overrides rewrite them in place, and re-enables rebuild them in place. A racing TRACING or SYSCALL program can read unfilled entries, including uninitialized memory in the kmalloc'd tables, or torn topo updates. Tie the tables' lifetimes to the root sched instead: each root enable builds a fresh set privately and publishes the per-table __rcu globals once the layout is final, and root disable unpublishes and RCU-frees the set. A non-NULL global is now always a fully built table which stays valid for the reader's RCU read section, and lookups stay two loads. Kfuncs treat NULL as no-mapping, also after the scheduler exits instead of reporting the stale last mapping. The cid kfuncs are available whether the root scheduler is cid-form or cpu-form, the latter to allow gradual migration to cids. Every root therefore builds and publishes a default mapping. Every reader must either be gated on scheduler liveness or NULL-check inside an RCU read section. Fix the two kfuncs that were neither: scx_bpf_this_cid() read the table with no RCU or preemption protection and scx_bpf_task_cid() relied on KF_RCU, which doesn't put a sleepable program in an RCU read section. The hotplug callbacks are instead serialized by retiring the tables inside the cpus_read_lock() section that clears scx_root. v2: Document why every root builds the tables (desc + cid.c comment).	N/A	More Details
CVE-2026-98943	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-107685	Improper Certificate Validation vulnerability in the WebSocket streaming client in agentjido req_llm allows an on-path attacker to read and modify OpenAI WebSocket sessions, including the authorization bearer token, via a TLS certificate the client does not trust. ReqLLM.Streaming.WebSocketSession.Client.start/3 in lib/req_llm/streaming/web_socket_session/client.ex starts WebSockex without TLS options, and WebSockex defaults to insecure: true, which connects with verify: :verify_none. The option allowlist also drops any caller-supplied cacerts or ssl_options, so an application cannot turn verification on. The OpenAI Responses, Realtime, and Codex WebSocket paths reach this client and send the API key or OAuth access token in the authorization header. HTTP streaming through Finch and the HTTP/2 duplex session through Mint verify certificates and are not affected. This issue affects req_llm: from 1.9.0 before 1.29.0.	N/A	More Details
CVE-2026-98946	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: Disable the V3D interrupt across runtime suspend vc4_irq_disable() masks the V3D interrupt sources and then calls synchronize_irq() before the V3D is powered down. However, by itself, this is not enough to quiesce the interrupt handler. synchronize_irq() waits for handlers that have already set IRQD_IRQ_INPROGRESS, and for irqchips reporting IRQCHIP_STATE_ACTIVE. A GIC interrupt chip is able to mark an interrupt active as soon as a CPU acknowledges it, so there these checks cover the whole dispatch path. On RPi 0-3, however, the interrupt controller is ARMCTRL, which has no active state. A CPU that has read the hwirq out of the pending register but has not yet reached handle_level_irq() stays invisible to synchronize_irq(). During a power transition, vc4_irq() may therefore run after the power domain is off, where every V3D register read will return 0xdeadbeef. 0xdeadbeef has FLDONE, FRDONE and OUTOMEM set. This problem doesn't trigger NULL pointer dereference issues only because the functions vc4_irq_finish_bin_job() and vc4_irq_finish_render_job() return early when there is no job pointer, but OUTOMEM is still able to schedule vc4_overflow_mem_work() after vc4_irq_disable() has already cancelled it. Therefore, fix the spurious interrupts by disabling the interrupt line across the PM	N/A	More Details

	transition. To preserve the enable/disable balance, move vc4_irq_install() and request the IRQ line disabled before the first runtime resume.		
CVE-2026-98947	In the Linux kernel, the following vulnerability has been resolved: vsock: Fix memory leak in vmci_transport_recv_dgram_cb() During the closure of a datagram socket, the vmci_transport_recv_dgram_cb() function may be called, which will add the packets to the socket's backlog; then, after the receive queue is cleared, the __release_sock() function will move the packet back from the socket's backlog to the receive queue, which will lead to a memory leak. sock_close __sock_release __vsock_release // take ownership by user-space lock_sock_nested sock_set_flag(sk, SOCK_DEAD) vmci_transport_release vmci_dispatch_dgs vmci_datagram_invoke_guest_handler vmci_transport_recv_dgram_cb sk_receive_skb if (!sock_owned_by_user(sk)) [...] (!) else if sk_add_backlog() [...] (!) skb_queue_purge(&sk->sk_receive_queue) release_sock __release_sock // move all packets from the backlog // to the receive queue (!) sk_backlog_rcv vsock_queue_rcv_skb [...] __sock_queue_rcv_skb if (!sock_flag(sk, SOCK_DEAD)) // Since the socket is "dead", the function will not // be called sk->sk_data_ready(sk) sock_release_ownership(sk); sock_put Add a receive queue cleanup in the socket destructor to fix this. syzkaller report: unreferenced object 0xffff8880117bfb80 (size 240): comm "irq/56-vmw_vmci", pid 205, jiffies 4295064448 (age 73.992s) hex dump (first 32 bytes): 78 dc 26 1e 80 88 ff ff 78 dc 26 1e 80 88 ff ff x.&.....x.&..... 00 00 00 00 00 00 00 00 c0 da 26 1e 80 88 ff ff&..... backtrace: [<ffffff82be8d97>] __alloc_skb+0x287/0x330 net/core/skbuff.c:505 [<fffffffa114bcd>] vmci_transport_recv_dgram_cb+0xbd/0x210 [vmw_vsock_vmci_transport] [<fffffffa05072c6>] vmci_datagram_invoke_guest_handler+0x366/0x450 [vmw_vmci] [<fffffffa050ad75>] vmci_dispatch_dgs+0x235/0x490 [vmw_vmci] [<fffffffa050b025>] vmci_interrupt+0x55/0x290 [vmw_vmci] [<ffffff8137ad2b>] irq_thread_fn+0x8b/0x1a0 kernel/irq/manage.c:1205 [<ffffff8137c7ae>] irq_thread+0x28e/0x530 kernel/irq/manage.c:1314 [<ffffff81253dd6>] kthread+0x2e6/0x3a0 kernel/kthread.c:376 [<ffffff81004392>] ret_from_fork+0x22/0x30 arch/x86/entry/entry_64.S:295 Found by InfoTeCS on behalf of Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2026-98929	In the Linux kernel, the following vulnerability has been resolved: macsec: initialize SecY before registering the netdevice Creating a MACsec device with MAC offload over an LRO-capable lower device triggers a warning in rtmsg_ifinfo_build_skb() when IPv4 forwarding is enabled by default. register_netdevice() invokes inetdev_init(), which disables LRO and emits a NETDEV_FEAT_CHANGE notification. This reaches macsec_fill_info() before macsec_add_dev() initializes the SecY. key_len is still zero, so macsec_fill_info() returns -EMSGSIZE and trips the WARN_ON in rtmsg_ifinfo_build_skb(), even though the skb has enough space. Even without the warning, notifications during registration can report uninitialized SecY attributes, including the SCI. This ordering has existed since the driver was introduced. Initialize the SecY and apply the new-link attributes before registration. Move MAC address inheritance into macsec_newlink() so the SCI can also be initialized before registration-time notifications report it. Move the per-CPU statistics and metadata destination allocation into ndo_init(), and release partial allocations on failure.	N/A	More Details
CVE-2026-95620	A use-after-free vulnerability in the Qt QNetworkAccessManager HTTP/2 client support of the Qt Network module allows a malicious or compromised HTTP/2 server to cause the client application to crash, affecting availability.	N/A	More Details
CVE-2026-98875	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/gsp: add SEC2 to GA100 chip table The booter-load and booter-unload firmware run on the SEC2 falcon. During tu102_gsp_oneinit(), the booter constructor needs device->sec2 to access the SEC2 falcon. Without the .sec2 entry, device->sec2 is NULL and this dereference crashes during GSP-RM boot.	N/A	More Details
CVE-2026-98926	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Remove sink usage from DPMS [Why] stream->sink is optional and can be null, so should always be checked before dereference. Additionally, most of its usage in DPMS sequences is for stream->sink->link, which can be replaced with stream->link, as the two should always be the same. [How] * Replace stream->sink->link in DPMS on/off * Add assert to USB4 BW allocation where sink is required * Avoid inconsistencies in resource access, e.g. don't repeat stream->link after it was already saved to a local variable * Pull out effective VPG calculation to helper getter * Formatting fixes	N/A	More Details
CVE-2026-98874	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mld: don't WARN on WoWLAN suspend w/o netdetect Clearly, from a user perspective, it must be valid to configure WoWLAN and then suspend while not connected to a network. Since mac80211 doesn't distinguish these cases and simply calls the driver to suspend whenever WoWLAN is configured, the driver has to cleanly handle the case where it's called for WoWLAN, it's not connected but there's also no netdetect configured. Remove the WARN_ON() and keep returning 1 to disconnect and then suspend.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_conntrack_expect: zero at allocation time There are occasional LLM hints wrt. leaking uninitialized data to userspace via ctnetlink. Just zero at allocation time, expectations are not frequently used these days. Intentionally keeps _init	N/A	More Details

98873	as-is because we could theoretically support re-init, so add the missing exp->dir there.		
CVE-2026-98872	In the Linux kernel, the following vulnerability has been resolved: freevxfs: don't BUG() on unknown typed-extent type vxfs_bmap_typed() handles four typed-extent types and calls BUG() in its default case, so an on-disk typed extent with any other type value crashes the kernel. It is reachable from ioctl(FIBMAP) on a regular file: kernel BUG at fs/freevxfs/vxfs_bmap.c:230! RIP: vxfs_bmap_typed fs/freevxfs/vxfs_bmap.c:230 [inline] vxfs_bmap1+0x128a/0x12d0 fs/freevxfs/vxfs_bmap.c:257 Replace the BUG() with WARN_ON_ONCE() and return 0 -- the value vxfs_bmap_typed() already returns on failure (and from the DEV4 case above); vxfs_getblk() maps 0 to -EIO, so the ioctl fails cleanly.	N/A	More Details
CVE-2026-57209	Heimdall is a cloud native Identity Aware Proxy and Access Control Decision service. Prior to version 0.17.17, when Heimdall operates in proxy mode, it constructs the Forwarded HTTP header after executing the matched rule pipeline by inserting the incoming request's Host header value directly into the header string without sanitizing commas or semicolons. This allows an attacker to inject additional parameters into the Forwarded header, potentially spoofing IP addresses for upstream services. This issue has been patched in version 0.17.17.	N/A	More Details
CVE-2026-57210	Heimdall is a cloud native Identity Aware Proxy and Access Control Decision service. Prior to version 0.17.17, when the trusted_proxies option is configured, heimdall extracts client IP addresses from the Forwarded (for= parameter) and X-Forwarded-For headers and exposes them as Request.ClientIPAddresses to the rule pipeline. However, extracted values are not validated to be syntactically valid IP addresses. Arbitrary strings, malformed IP literals, and RFC 7239 unknown values and obfuscated identifiers are accepted without further checks. In addition, the Forwarded header parser splits on , and ; without accounting for RFC 7239 quoted strings, which can cause a single quoted value to be parsed as multiple entries, with fragments — including trailing quote characters — treated as independent addresses. This issue has been patched in version 0.17.17.	N/A	More Details
CVE-2026-98871	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Fix possible deadlock in SEV init failure path __sev_platform_init_handle_init_ex_path() calls rmp_mark_pages_firmware() with locked=false while the parent function of init_ex_path already acquired the sev_cmd_mutex. In the case of an RMPUPDATE failure for any page after the first, the cleanup path would invoke reclaim pages which would result in a deadlock in sev_do_cmd. Pass locked=true to honor the lock status of the parent function.	N/A	More Details
CVE-2026-98917	In the Linux kernel, the following vulnerability has been resolved: parisc: Increase kernel stack size to 32kb For 64-bit Linux kernels, increase the default kernel stack size (THREAD_SIZE_ORDER) to 32 kB, in order to avoid kernel crashes which have been triggered recently when building the debian vtk9 package with gcc 17: stackcheck: kworker/u128:0 will most likely overflow kernel stack (sp:179a83af0, stk bottom-top:179a80000-179a84000) Kernel panic - not syncing: low stack detected by irq handler - check messages CPU: 2 UID: 0 PID: 30760 Comm: kworker/u128:0 Tainted: G W 6.18.46-dirty #1 NONE Tainted: [W]=WARN Hardware name: 9000/800/rp3440 Workqueue: writeback wb_workfn (flush-259:0) Backtrace: [<000000004022f050>] show_stack+0x70/0x90 [<000000004022378c>] dump_stack_lvl+0x124/0x190 [<000000004022382c>] dump_stack+0x34/0x48 [<000000004020212c>] vpanic+0x204/0x648 [<00000000402025c4>] panic+0x54/0x58 [<0000000040232230>] do_cpu_irq_mask+0x3f8/0x440 [<0000000040227070>] intr_return+0x0/0xc	N/A	More Details
CVE-2026-98918	In the Linux kernel, the following vulnerability has been resolved: mm/damon/ops-common: use a page-aligned address in damon_ptep_mkold() __damon_va_prepare_access_check() picks a random byte address within the region and stores it in r->sampling_addr. damon_va_mkold() passes it into a page table walk, which hands it to damon_ptep_mkold() as the address of the page to sample: damon_va_mkold(mm, r->sampling_addr) damon_va_walk_page_range(mm, addr, addr + 1) damon_mkold_pmd_entry() damon_ptep_mkold(pte, vma, addr) ptep_test_and_clear_young(vma, addr, pte) mmu_notifier_clear_young(mm, addr, addr + PAGE_SIZE) For arm64, before commit 6f0e1142173a ("arm64: mm: support batch clearing of the young flag for large folios"), the contpte helper walked exactly CONT_PTES entries from the aligned-down page table pointer and used @addr only to pass down to each entry, so an unaligned value was harmless: ptep = contpte_align_down(ptep); addr = ALIGN_DOWN(addr, CONT_PTE_SIZE); for (i = 0; i < CONT_PTES; i++, ptep++, addr += PAGE_SIZE) Now the range to walk is derived from @addr instead: end = addr + nr * PAGE_SIZE, rounded up to CONT_PTE_SIZE. For a sample in the last page of a contpte block, the sub-page offset puts end just past the block boundary, so the round-up lands a whole block further and the walk clears PTE_AF in CONT_PTES entries beyond the sampled block. For the last block in a page table page, those entries are past the end of that page, so the walk writes into the page that follows. Triggered by the full 7.1/7.2 kernel selftest suite on arm64 (EC2 c/m6g.4xlarge). The kernel sometimes crashes at or shortly after the DAMON test. What the overrun does depends on the page that happens to follow the page table, so there is no single signature. If that page is read-only, the write faults in the sampling path itself: Unable to handle kernel write to read-only memory at virtual address ffff0003c5d2d000 FSC = 0x0f: level 3 permission fault CM = 0, WnR = 1, TnD = 0, TagAccess = 0 CPU: 10 UID: 0 PID: 3487 Comm:	N/A	More Details

	kdamond.2 pc : contpte_test_and_clear_young_ptes+0x70/0xc0 lr : damon_ptep_mkold+0x1e8/0x1f8 Call trace: contpte_test_and_clear_young_ptes+0x70/0xc0 (P) damon_mkold_pmd_entry+0x150/0x170 walk_pmd_range+0x110/0x2b0 walk_pud_range+0x10c/0x208 walk_pgd_range+0x134/0x258 __walk_page_range+0x98/0x1b0 walk_page_range_vma_unsafe+0x90/0x148 walk_page_range_vma+0x28/0x40 damon_va_walk_page_range+0x114/0x2b8 damon_va_prepare_access_checks+0xec/0x1a8 kdamond_fn+0x534/0x770 kthread+0x128/0x138 ret_from_fork+0x10/0x20 Otherwise the page is writable, the PTE_AF clearing succeeds silently and the damage only surfaces later, in whatever happened to own the page, so the backtrace is unrelated to DAMON and differs between runs. Pass a page-aligned address to the ptep_test_and_clear_young() call in damon_ptep_mkold(), which is the only place DAMON can reach contpte_test_and_clear_young_ptes() from. Nothing else sees the aligned address, and r->sampling_addr itself is left as is, so the sampling and region bookkeeping semantics are unchanged.		
CVE-2026-98919	In the Linux kernel, the following vulnerability has been resolved: drm/xe: harden adjust_idledly() against divide-by-zero and overflow adjust_idledly() has several corner-case issues flagged during review: 1. If xe_gt_clock_init() failed to recognise the crystal clock, gt->info.timestamp_base is 0, which makes idledly_units_ps also 0. The subsequent DIV_ROUND_CLOSEST(..., idledly_units_ps) is then a divide-by-zero and panics the kernel. 2. The tick-to-ns conversions are done in u32: idledly * idledly_units_ps, (maxcnt - 1) * 1000 Both overflow u32 before DIV_ROUND_CLOSEST() sees them. 3. If IDLE_WAIT_TIME reads back as 0, maxcnt evaluates to 0 and the maxcnt - 1 clamp wraps to 0xFFFFFFFF in u32. 4. The register only stores whole ticks, so the clamped ns value has to be converted to ticks and back. DIV_ROUND_CLOSEST() can round that conversion up past maxcnt: maxcnt = 640 ns, one tick = 666664 ps clamp: maxcnt - 1 = 639 ns ns -> ticks: 639000 / 666664 = 0.958 -> rounds to 1 tick tick -> ns: 1 * 666664 / 1000 = 667 ns 667 ns is programmed into RING_IDLEDLY, but 667 >= maxcnt (640), so xe_gt_WARN_ON() fires again on every subsequent init. Return early if timestamp_base is 0 (the unknown-crystal path). Do the conversions in u64 via the *_ULL() helpers so they cannot wrap. Clamp with a floor (DIV_ROUND_DOWN_ULL) so the programmed delay stays strictly below maxcnt, and guard the maxcnt == 0 case with a zero delay while still writing RING_IDLEDLY so INHIBIT_SWITCH_UNTIL_PREEMPTED is cleared. v2: Drop the redundant warn on the timestamp_base == 0 path; xe_gt_clock_init() already warns on an unrecognised crystal clock. Keep the early return to avoid the divide-by-zero. - Vinay v3: Field-mask the RING_IDLEDLY write with REG_FIELD_PREP(IDLE_DELAY, ...) instead of writing the raw tick count, which could clobber INHIBIT_SWITCH_UNTIL_PREEMPTED and reserved bits. Split the inhibit-switch clear from the maxcnt clamp so a set inhibit bit no longer forces a needless delay overwrite when the delay itself is already valid. Use gt_to_xe(gt) instead of gt_to_xe(hwe->gt). (cherry picked from commit d864065ea25e9d12897c175de9176ce46677e176)	N/A	More Details
CVE-2026-98920	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix missing alloc tagging of direct mempool allocations Commit 1d78d56c43ef ("netfs: Fix folio_queue ENOMEM in writeback by adding a mempool") added a mempool for the folio_queues and made the request, subrequest and folio_queue allocations distinguish between writeback and everything else. Writeback is part of memory reclaim and must not fail due to ENOMEM, so it allocates under GFP_NOFS through mempool_alloc(), which may dip into the pool's reserve and, if that runs empty, wait for elements to be returned. The GFP_KERNEL paths, which can return -ENOMEM to their callers, invoke the pool's ->alloc() callback directly instead. The direct call, however, skips the alloc_hooks() wrapper that the mempool_alloc() macro provides. The pool callbacks, mempool_alloc_slab() and mempool_kmalloc(), call kmem_cache_alloc_noprof() and kmalloc_noprof() and rely on current->alloc_tag having been set by the caller. With CONFIG_MEM_ALLOC_PROFILING_DEBUG=y this leads to current->alloc_tag not set WARNING: ./include/linux/alloc_tag.h:161 at __alloc_tagging_slab_alloc_hook alloc_tag was not set WARNING: ./include/linux/alloc_tag.h:166 at __alloc_tagging_slab_free_hook at allocation and free time respectively, as reported when reading files on a CIFS mount. The allocations are also missing from /proc/allocinfo. Wrap the direct ->alloc() invocations in alloc_hooks() with a new mempool_alloc_noreserve() helper in include/linux/mempool.h, next to the other alloc_hooks()-wrapped macros such as mempool_alloc(). The GFP_KERNEL paths keep their failable allocation semantics, they just get tagged now.	N/A	More Details
CVE-2026-98921	In the Linux kernel, the following vulnerability has been resolved: ovl: fix UAF in ovl_do_mkdir() debug print ovl_do_mkdir() prints the input dentry with %pd after vfs_mkdir(). Since commit fe497f0759e0 ("VFS: change vfs_mkdir() to unlock on failure."), vfs_mkdir() calls end_creating() on the input dentry on failure and may replace it on success, so the post-call %pd can use-after-free the dentry when CONFIG_OVERLAY_FS_DEBUG is enabled. Print the dentry before the call and only the result afterward.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net: emac: move setting of netops to fix crash fixes the following crash on driver initialization: BUG: Kernel NULL pointer dereference on read at 0x00000158 Faulting instruction address: 0xc0566b40 Oops: Kernel access of bad area, sig: 11 [#1] BE PAGE_SIZE=4K PowerPC 44x Platform Modules linked in: CPU: 0 UID: 0 PID: 1 Comm: swapper/0 Tainted: GW 7.3.0-rc3+ #1 Tainted: [W]=WARN Hardware name: MyBook Live APM821XX		

CVE-2026-98922	0x12c41c83 PowerPC 44x Platform NIP: c0566b40 LR: c05648f8 CTR: c04c1e9c REGS: c1053a20 TRAP: 0300 Tainted: GW (7.3.0-rc3+) MSR: 0002b000 <CE,EE,FP,ME> CR: 24008808 XER: 00000000 DEAR: 00000158 ESR: 00000000 GPR00: c05648f8 c1053b10 c1063600 c1030000 c5ab3000 00000000 [...] GPR08: 00000002 00000000 00000000 c1053b40 84002808 00000000 [...] GPR16: cfffd210 00000002 c0beafcc cfffc960 00000000 c1030644 [...] GPR24: c0beafbc c1037000 00000000 0000000a 00000000 c1030000 [...] NIP [c0566b40] phy_link_topo_add_phy+0x2c/0x1d0 LR [c05648f8] phy_attach_direct+0x1a4/0x368 Call Trace: [c1053b10] [c0811e04] klist_put+0x54/0xb4 (unreliable) [c1053b40] [c05648f8] phy_attach_direct+0x1a4/0x368 [c1053b70] [c0564ae8] phy_connect_direct+0x2c/0x60 [c1053b90] [c056c7a8] of_phy_connect+0x50/0x74 [c1053bc0] [c0572a88] emac_probe+0xd50/0x119c [c1053c90] [c04cb770] platform_probe+0x74/0xa4 [c1053cb0] [c04c8f08] really_probe+0x120/0x2b0 [c1053cd0] [c04c9254] __driver_probe_device+0x1bc/0x1fc [c1053d00] [c04c9334] driver_probe_device+0x38/0xa8 [c1053d30] [c04c9560] __driver_attach+0xf4/0x10c [c1053d50] [c04c6af8] bus_for_each_dev+0x68/0xd0 [c1053d90] [c04c7c34] bus_add_driver+0xcc/0x1ec [c1053dc0] [c04c9f6c] driver_register+0xcc/0x110 [c1053de0] [c0aa3f40] emac_init+0x1c4/0x200 This bug showed up starting with v7.3-rc1. At the NIP in phy_link_topo_add_phy() is a netdev_need_ops_lock() check. This was added by the following commit ded86da4bbb7 ("net: ethtool: relax ethnl_req_get_phydev() locking assertion") The bug shows up because at the time of phy_connect() was called, the netdev_ops were *not yet* determined. My fix is to move the code that sets netdev_ops+commac.ops+ethtool_ops further up as emac_init_config() derives that by looking at the device-tree and sets the required dev->phy_mode accordingly. During review, the Sashiko bot's AI stated that the commac assignment became a dead store. Great catch! To keep the original behavior as-is, one mentioned option "set dev->commac.ops = &emac_commac_ops only in the non-gige case?" sounded like a great plan. So the dev->commac.ops assignment for the non-gige-case moves into the else block. This patch was tested on a WD MyBook Live (RGMII). the device now works again.	N/A	More Details
CVE-2026-98923	In the Linux kernel, the following vulnerability has been resolved: arm64: io: Reject non-user protection in ioremap_prot() Mapping a stack-top page via /dev/mem with PROT_NONE and then reading that process's /proc/<pid>/cmdline triggers a spurious WARN in ioremap_prot() through generic_access_phys(): WARNING: ./arch/arm64/include/asm/io.h:275 at generic_access_phys Call trace: generic_access_phys+0x1c8/0x228 (P) __access_remote_vm+0x2b4/0x398 access_remote_vm+0x14/0x30 get_mm_cmdline+0xf8/0x2a0 proc_pid_cmdline_read+0x68/0x120 generic_access_phys() passes the protection derived from the user PTE to ioremap_prot(). On arm64, a PROT_NONE mapping is represented by a present-invalid PTE, so pte_present() still returns true and the protection reaches ioremap_prot(). A PROT_NONE mapping does not have PTE_USER, causing the existing WARN_ON_ONCE() in ioremap_prot() to fire even though this is a valid user mapping. Execute-only mappings have the same issue and must not be readable through this path either. ioremap_prot() should therefore reject protection values without PTE_USER without warning. This makes the access fail cleanly for PROT_NONE and execute-only mappings while retaining the existing user-protection contract.	N/A	More Details
CVE-2026-98924	In the Linux kernel, the following vulnerability has been resolved: fs: avoid repeated scans in evict_inodes() We observed hung tasks when users attempted to unmount a filesystem after its disk had been removed while still in use. During device removal, fs_bdev_mark_dead() calls evict_inodes() while holding s_umount. Each time evict_inodes() drops s_inode_list_lock to reschedule, it restarts the walk from the head of s_inodes. With many referenced inodes at the head of the list, these restarts repeatedly scan the same inodes without reclaiming them. This can keep s_umount held for a long time, blocking concurrent umount attempts and triggering hung-task reports. Keep the current inode, already marked I_FREEING, out of the disposal batch until s_inode_list_lock is reacquired. Resume the walk from this inode and dispose of it in a later batch or at the end of the walk. The zero-refcount and state checks under i_lock allow this walker to claim the inode by setting I_FREEING and removing it from the LRU. Other reclaimers skip the inode, leaving this walker responsible for eviction. Only evict() removes it from s_inodes, so keeping it out of the disposal batch ensures that it remains on the list while the lock is dropped. After reacquiring the lock, reading its current next pointer accounts for concurrent removal of following inodes. The existing inode lifetime rules prohibit acquiring a reference to an inode marked I_FREEING or I_WILL_FREE. __iget() requires its caller to hold i_lock and establish that taking a reference is valid. Inode lookup and igrob() check these flags under i_lock when acquiring a reference from zero. ihold() requires an existing reference, which would keep i_count nonzero and prevent this walker from claiming the inode. These rules already allow iput_final() and the inode shrinker to release i_lock after setting I_FREEING and before eviction completes. A temporary __iget() reference would also keep the inode on the list, but its release must preserve last-reference handling. Another user can acquire a reference, update lazy timestamps and drop its reference while the pin is held. If the pin becomes the last reference, dropping it with atomic_dec_and_test() and evicting directly bypasses iput()'s lazytime handling and can lose those timestamp updates. Releasing the pin with iput() preserves that handling, but does not guarantee eviction. fs_bdev_mark_dead() runs with SB_ACTIVE set, so iput() may retain the inode in cache, whereas evict_inodes() must evict eligible zero-reference inodes. The inode may also have been freed when iput() returns, so the walker cannot then use it to force eviction. Using I_FREEING preserves the existing eviction behavior without introducing an	N/A	More Details

	additional last-reference transition. The xfstests auto group passed on ext4 and XFS with known unrelated failures excluded. No new issues were observed, and the previously reproducible hung task no longer occurs with this patch.		
CVE-2026-98925	In the Linux kernel, the following vulnerability has been resolved: bpf, sockmap: Fix self-redirect copied_seq double-counting When a BPF stream_verdict program redirects an skb back to the same socket (self-redirect with BPF_F_INGRESS), sk_psock_verdict_apply() calls tcp_eat_skb() which advances tcp_sk->copied_seq. However, the skb is then delivered to the socket's psock ingress queue and later read by tcp_bpf_recvmmsg_parser(), which also advances copied_seq via the copied_from_self accounting path. This double-counting causes copied_seq to advance by 2x the actual data length, triggering: TCP recvmmsg seq # bug 2: copied BF2E806, seq BF2E7FD, \ rcvnxt BF2E806, fl 0 WARNING: net/ipv4/tcp.c:2745 at tcp_recvmmsg_locked+0x72b/0x2640 Call Trace: tcp_recvmmsg+0x10a/0x500 sock_recvmmsg+0x168/0x1d0 __sys_recvfrom+0x19a/0x2a0 __x64_sys_recvfrom+0xe4/0x1f0 do_syscall_64+0xf7/0x530 entry_SYSCALL_64_after_hwframe+0x77/0x7f cleanup rbuf bug: copied BF2E806 seq BF2E806 rcvnxt BF2E806 WARNING: net/ipv4/tcp.c:1609 at tcp_cleanup_rbuf+0xf2/0x1c0 Call Trace: tcp_recvmmsg_locked+0x8d1/0x2640 tcp_recvmmsg+0x10a/0x500 sock_recvmmsg+0x168/0x1d0 __sys_recvfrom+0x19a/0x2a0 __x64_sys_recvfrom+0xe4/0x1f0 do_syscall_64+0xf7/0x530 entry_SYSCALL_64_after_hwframe+0x77/0x7f Fix this by converting self-redirect verdict to __SK_PASS at the beginning of sk_psock_verdict_apply(). This bypasses the __SK_REDIRECT case entirely (which calls sk_psock_eat_skb), letting the __SK_PASS path queue the skb to the psock ingress queue. The data is then read via tcp_bpf_recvmmsg_parser(), which advances copied_seq exactly once through copied_from_self. Cross-socket redirects continue through __SK_REDIRECT with sk_psock_eat_skb() unchanged.	N/A	More Details
CVE-2026-99027	In the Linux kernel, the following vulnerability has been resolved: net: team: stop reusing skb after queue override A nonzero dev_queue_xmit() result does not return ownership of the skb. When an override port's qdisc drops it, the queue override helper used to try another port or the mode transmit op with the freed skb. Return the handoff decision separately from transmit success. Select only the first override port under RCU and stop after handing it the skb, regardless of the lower transmit result. Keep the existing success and drop accounting. BUG: KASAN: slab-use-after-free in sk_skb_reason_drop Write of size 4 by task poc/5040 sk_skb_reason_drop (net/core/skbuff.c:1240) dev_kfree_skb_any_reason (net/core/dev.c:3470) ab_transmit (drivers/net/team/team_mode_activebackup.c:49) team_xmit (drivers/net/team/team_core.c:1849) __dev_direct_xmit (net/core/dev.c:4934) Freed by task 5040: __tcf_kfree_skb_list (net/sched/sch_generic.c:43) __dev_queue_xmit (net/core/dev.c:4831) team_xmit (drivers/net/team/team_core.c:1846) Kernel panic - not syncing: KASAN: panic_on_warn set ...	N/A	More Details
CVE-2026-99028	In the Linux kernel, the following vulnerability has been resolved: net: xilinx: axienet: Free outstanding DMA buffers on dmaengine stop In the dmaengine path the driver pre-submits RX buffers and holds in-flight TX buffers whose SKBs are DMA-mapped by the driver and freed only in the completion callbacks. On ndo_stop() dmaengine_terminate_sync() aborts these descriptors without running their callbacks, and the driver then frees only the ring shells, leaking every SKB still owned by the engine and its DMA mapping on each ifdown. With 128 RX buffers pre-posted per channel, the mapping leak can eventually exhaust a limited IOMMU aperture. Clear the slot's skb in the TX and RX callbacks so a non-NULL skb marks a slot that still owns a live, DMA-mapped buffer, and on stop unmap and free every such buffer. axienet_dma_rx_cb() runs from the DMA tasklet and re-arms the RX ring on each completion, so it can race axienet_stop(): a completion may submit a fresh buffer after dmaengine_terminate_sync() has returned, leaving the channel armed with a buffer the teardown then frees while the engine may still write into it (dma_release_channel() does not stop it either). Add a lock that axienet_dma_rx_cb() holds across the @stopping check and the resubmit, and axienet_stop() holds to set @stopping before terminating. Once @stopping is set no callback can arm a new buffer, and any armed just before is aborted by the terminate, so teardown only frees buffers the engine no longer owns.	N/A	More Details
CVE-2026-99029	In the Linux kernel, the following vulnerability has been resolved: net: sparx5: free the matchall entry on destroy sparx5_tc_matchall_replace() allocates a struct sparx5_mall_entry for every offloaded matchall filter and adds it to sparx5->mall_entries. sparx5_tc_matchall_destroy() removes the entry from the list, but never frees it, so the entry of every deleted mirror and goto matchall filter is leaked. Free the entry after unlinking it. The leak was discovered by an AI code review agent, and reproduced with kmemleak on a lan969x EV board (EV23X71A) by repeatedly adding and deleting matchall mirror and goto filters. With the fix, kmemleak no longer reports the leak.	N/A	More Details
CVE-2026-115775	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.10.0rc1, the twisted.words.xish.domish ExpatElementStream parser allows an unauthenticated remote peer to keep a container element open while continuously adding child elements. The parser retains the unfinished stanza as an in-memory tree without a byte or element-count limit, and XMPP pre-authentication stream parsing exposes this behavior before credentials are accepted. A single connection can therefore exhaust server memory and terminate the process. This issue is fixed in version 26.10.0rc1.	N/A	More Details

CVE-2026-115502	A missing authorization check allows anyone having access to a "link without login" for a DataApp to list and potentially access jobs of other users created by the same link in KNIME Business Hub prior to 1.19.0. As workaround remove all "links without login" to DataApps.	N/A	More Details
CVE-2026-114608	Pylint is a static code analyser for Python. From 0.15.2 until 4.0.9, Pylint stores module statistics in predictable files under PYLINT_HOME and loads those files with pickle.load on a later run. An attacker who can write to a shared, group-writable, or otherwise untrusted cache directory can replace a stats cache with crafted serialized data, causing arbitrary code to execute with the privileges of the user who next runs Pylint. The default per-user cache is not normally writable by another user. This issue is fixed in version 4.0.9.	N/A	More Details
CVE-2026-114606	draw.io Desktop is the official Electron build of draw.io. Prior to 31.4.4, draw.io Desktop is vulnerable to arbitrary local-file disclosure through the Electron main process's read-side IPC handlers. The readFile, getFileDrafts, getBkpFile, fileStat, isFileWritable, checkFileExists, and watchFile actions operate on a renderer-chosen filesystem path with no blessed-path (assertWritablePath) authorization. The only gates are validateSender — a URL-prefix origin check — plus, for readFile alone, a content "sniff" and a lexical appBaseDir prefix block. The write-side handlers (saveFile, writeFile, saveDraft, deleteFile) are all protected by assertWritablePath; the read side is not. The result is a read/write authorization asymmetry: any frame executing in the app's file:// origin can read the contents of local files that the same origin is explicitly forbidden to write. This issue is fixed in version 31.4.4.	N/A	More Details
CVE-2026-94482	Next.js is a React framework for building full-stack web applications. From 13.0.0 until 15.5.28 and 16.4.1, crafted requests targeting App Router applications with at least one Server Action can cause excessive CPU consumption and block processing of further requests in the same process. Applications using the Pages Router and App Router applications without Server Actions are not affected. This issue is fixed in versions 15.5.28 and 16.4.1.	N/A	More Details
CVE-2026-113877	HTTP::Headers::ActionPack versions before 0.10 for Perl allow CPU exhaustion via Accept-Language ranges matched as regular expressions. _language_match interpolates each Accept-Language range into the pattern `/\^\$range\-/i` without escaping it, so regular expression syntax in the header reaches the regex engine. A 12-byte range such as `(a? b?){26}z` backtracks for seconds, and each increment of the quantifier doubles the time. An invalid range such as `(` fails to compile. A remote attacker can exhaust CPU, or make choose_language die, in any application that passes an untrusted Accept-Language header to it.	N/A	More Details
CVE-2026-113951	If an attacker can write to the filesystem of a container k8s cluster container, they can plant an exec credential plugin in /etc/kubernetes/admin.conf. kubectl will run that plugin as the macOS host user the next time someone pulls the cluster's kubeconfig onto the host with container k8s write-config or create and then runs kubectl against it — turning filesystem access inside an untrusted Linux container into arbitrary command execution on the host. This vulnerability is addressed in Container version 1.5.0.	N/A	More Details
CVE-2026-51954	An issue in Webasyst Framework 4.0.3 and before allows a remote attacker to execute arbitrary code via a crafted OAuth authorization URL with an arbitrary redirect_uri.	N/A	More Details
CVE-2026-115496	Overly informative error messages in the catalog functionality of KNIME Business Hub prior to version 1.18.5 may allow an authenticated remote attacker to see workflow names if their ID is known. A workaround is not possible.	N/A	More Details
CVE-2026-105483	A flaw in Node.js TLS hostname verification allows a certificate without a matching IP address SAN to authenticate a connection to an IPv4 address when the host is written with non-ASCII characters, such as fullwidth digits (` 1 2 7 .0.0.1`). `tls.checkServerIdentity()` did not recognize such hosts as IP addresses, even though they normalize and resolve to a regular IPv4 address, and matched them against the certificate DNS names and Common Name instead. A certificate containing `DNS:127.0.0.1` or `DNS:*0.0.0.1` could therefore be accepted for an IP endpoint it was not issued for. This is a regression introduced after the fix for CVE-2026-48618. This vulnerability affects Node.js **26.x** and **24.x** .	N/A	More Details
CVE-2026-115494	A flaw in the cross-site request forgery (CSRF) prevention of KNIME Business Hub prior to version 1.19.0 allows an authenticated attacker to perform a CSRF attack with passive user interaction. This means the state of a resource could be changed without the user noticing. As a workaround, an API gateway for firewall in front of KNIME Business Hub could remove the X-Forwarded-For header from all incoming requests.	N/A	More Details
CVE-2026-115454	An excessive resource allocation issue in the remote workflow editing capabilities of KNIME Business Hub prior to version 1.19.0 allows a authenticated attack to exhaust resources on the server side. This leads to intermittent unavailability of workflow execution functionality. A workaround is not possible.	N/A	More Details
	Authorization Bypass Through User-Controlled Key (CWE-639) in the cart and wishlist repositories of		

CVE-2026-115453	MarcosCamara01 Ecommerce Template allows any authenticated user to delete and modify the cart and wishlist rows of any other user, because <code>cartRepository.delete()</code> , <code>cartRepository.updateQuantityInternal()</code> and <code>wishlistRepository.delete()</code> in <code>src/lib/db/drizzle/repositories/</code> match rows by primary key alone, with no predicate on the owning user, and the per user authorization for those three operations rests on PostgreSQL row level security. The migrations run <code>ALTER TABLE ... ENABLE ROW LEVEL SECURITY</code> and create policies comparing <code>app.current_user_id()</code> with <code>user_id</code> , but never <code>ALTER TABLE ... FORCE ROW LEVEL SECURITY</code> , and PostgreSQL does not apply policies to the role that owns a table. The template uses one <code>DATABASE_URL</code> both to run the migrations, which create the tables, and for the application at runtime, so the application connects as the owner of the tables and the policies do not apply to it. <code>cart_items.id</code> and <code>wishlist.id</code> are sequential bigserial values, so the identifiers of rows belonging to other users can be enumerated, and the update path returns the modified row, which carries the identifier of the user it belongs to. This results in the destruction and modification of cart and wishlist contents belonging to other accounts of the installation.	N/A	More Details
CVE-2026-57849	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-115450	External Control of Assumed-Immutable Web Parameter (CWE-472) in the cart and checkout flow of MarcosCamara01 Ecommerce Template allows an authenticated user with no special privileges to be charged the price of one product while placing an order for another, because <code>POST /api/user/cart</code> accepts <code>variantId</code> , which determines what the order records, and <code>stripeId</code> , which determines what Stripe charges, as two independent fields of the request body, stores both on the cart row, and no check requires the <code>stripeId</code> to be the one that belongs to the <code>variantId</code> . <code>POST /api/stripe/payment</code> builds the Stripe Checkout line items from the stored value, as <code>{ price: item.stripeId, quantity }</code> , so the amount charged is the amount of that Stripe price. The <code>checkout.session.completed</code> webhook then creates the order products from the <code>variantId</code> of the same cart row, so the order records a product that was not paid for. Every <code>stripeId</code> in the catalogue is readable by any visitor, so no value has to be guessed. This results in orders for goods at the price of the cheapest item in the catalogue, repeatable without limit, and in order records whose amount does not correspond to the product they list.	N/A	More Details
CVE-2026-114487	Missing authorization in the PAM AD computer import endpoint in Devolutions Server 2026.3.5.0 through 2026.3.10.0 and 2026.2.20.0 and earlier allows an authenticated non-administrative user with a PAM license to create PAM sub-providers under existing AD providers via a crafted import request.	N/A	More Details
CVE-2026-114486	Incorrect authorization in the Devolutions Gateway token generation in Devolutions Server 2026.3.10.0 and earlier allows an authenticated non-administrative attacker to bypass the Network Access Rules of a virtual gateway and reach denied hosts via wildcard additional destinations signed into the gateway token.	N/A	More Details
CVE-2026-67597	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-97631	EDK2 contains a vulnerability in BIOS where a user may cause “Out-of-bounds Write”, Classic Buffer Overflow, and “Incorrect Calculation of Buffer Size” by remote access. Successful exploitation of this vulnerability can lead to arbitrary code execution. It could also impact Confidentiality, Integrity, and Availability.	N/A	More Details
CVE-2026-115586	<code>virtualenv</code> is a tool for creating isolated virtual python environments. From 21.14.9 until 21.14.13, the <code>bash bin/activate</code> script can copy an attacker-controlled <code>--prompt</code> value into <code>PS1</code> without escaping dollar signs or backticks when both <code>promptvars</code> and <code>POSIX</code> mode are disabled during activation. Re-enabling <code>promptvars</code> , running <code>set -o posix</code> , or starting an interactive child <code>bash</code> that inherits the exported <code>PS1</code> can then expand command substitution in the stored <code>VIRTUAL_ENV_PROMPT</code> value and execute commands as the user who activates the environment. This issue is fixed in version 21.14.13.	N/A	More Details
CVE-2026-115585	<code>pypdf</code> is a free and open-source pure-python PDF library. Prior to 6.21.2, an incomplete remediation for repeated malformed cross-reference entries in GHSA-55h5-xmcq-c37v could still trigger a complete PDF stream reload during each recovery attempt even when a recovery cache was present, causing crafted cross-reference streams to produce long runtimes. This issue is fixed in version 6.21.2.	N/A	More Details
CVE-2026-106070	A flaw in the Node.js <code>zlib.ZipFile`</code> API can cause it to return bytes that do not belong to the requested entry after the local and central header validation fails. When <code>zipEntry.modified`</code> is read on an archive whose local header does not match its central directory, either directly or while compacting the archive, the validation error is caught but the parsed local header is left cached without a content offset. A later content read then returns attacker-controlled bytes from the start of the file instead of rejecting the archive, and CRC-32 verification still passes because the attacker controls the central directory CRC. A Node.js-based scanner can therefore approve different content from what another ZIP implementation extracts from the same archive. This vulnerability affects Node.js <code>**26.x**</code> .	N/A	More Details

CVE-2026-106071	A flaw in Node.js Permission Model filesystem enforcement can grant access to paths that do not match a wildcard <code>--allow-fs-read`</code> or <code>--allow-fs-write`</code> rule. When the requested path is shorter than the prefix of a wildcard grant, the radix tree lookup stops comparing early and treats the path as a match. For example, <code>--allow-fs-read=/etc/passwd.*`</code> also allows reading <code>/etc/passwd`</code> . This requires the wildcard grant to sit directly under the root of the permission tree, for example when it is the only grant for that scope. This is a different root cause from CVE-2026-58043. This vulnerability affects Node.js 26.x .	N/A	More Details
CVE-2026-106072	A flaw in Node.js TLS hostname verification allows a certificate with only a DNS SAN to authenticate an IP endpoint when the client connects using a non-canonical IPv4 spelling, such as <code>0177.0.0.1`</code> , <code>0x7f.0.0.1`</code> , <code>2130706433`</code> or <code>127.1`</code> . <code>tls.checkServerIdentity()</code> did not classify these hosts as IP addresses, but normalized them to their canonical form and matched the result against the certificate DNS names. A certificate containing <code>DNS:127.0.0.1`</code> and no <code>IP Address:127.0.0.1`</code> SAN was therefore accepted, while the same certificate is correctly rejected for <code>127.0.0.1`</code> . Exploitation requires a certificate with a numeric DNS SAN issued by a CA trusted by the client, and a platform resolver that accepts these spellings, such as glibc. This vulnerability affects Node.js 26.x , 24.x , and 22.x .	N/A	More Details
CVE-2026-106073	A flaw in Node.js Permission Model filesystem enforcement leaves <code>fs.realpathSync()</code> without the required read permission checks. This is an incomplete fix of CVE-2026-21715, which only covered <code>fs.realpathSync.native()</code> . As a result, code running under <code>--permission`</code> with restricted <code>--allow-fs-read`</code> can still use <code>fs.realpathSync()</code> to check file existence, resolve symlink targets, and enumerate filesystem paths outside of permitted directories. This vulnerability affects Node.js 26.x , 24.x , and 22.x .	N/A	More Details
CVE-2026-106074	A flaw in Node.js <code>crypto`</code> causes the internal entropy caches used by <code>crypto.randomUUID()</code> , <code>crypto.randomUUIDv7()</code> , and <code>crypto.randomInt()</code> to be saved into user-land startup snapshots ( <code>--build-snapshot`</code>) and single executable applications built with <code>useSnapshot`</code> . If any of these functions, or an API that relies on them such as <code>URL.createObjectURL()</code> , is called while the snapshot is built, every process started from that snapshot returns the same values until the cache is refilled: up to 127 UUIDs and about 1,000 <code>randomInt()</code> values. Identifiers expected to be unique can collide across processes, and values issued by one process can be predicted by observing another. This vulnerability affects Node.js 26.x , 24.x , and 22.x .	N/A	More Details
CVE-2026-87901	Improper sanitization in the registration module allows reflected Cross-Site Scripting.	N/A	More Details
CVE-2026-93336	Missing authorization in the unit resources download route in Canonical Juju 4.0 through 4.0.15 allows a controller-authenticated principal (a unit agent, or a user with only login permission) to read any application's charm resource in any model, including container-registry credentials, via a crafted request naming an arbitrary unit.	N/A	More Details
CVE-2026-115567	Tokio is a runtime for writing applications with Rust. In 1.54.0, a race condition between <code>tokio::time::Interval::poll_tick</code> (including <code>.tick().await</code>) and the multi-threaded runtime's time driver could drop an internal timer entry without synchronizing with the driver thread that was firing it. The race caused a use-after-free on the timer's waker and task state in safe Rust, which could produce assertion panics, task reference-count corruption, or heap memory corruption. This issue is fixed in 1.54.1.	N/A	More Details
CVE-2026-115516	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-108359	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. Prior to 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0, <code>Ods::processTableRow</code> in <code>src/PhpSpreadsheet/Reader/Ods.php</code> accepts an unbounded <code>table:number-rows-repeated</code> attribute from an ODS table row. When the row contains a value, the reader materializes a cell for every repetition without checking the spreadsheet row limit, so a small attacker-controlled ODS file can consume substantial CPU and memory and terminate the PHP worker. Column repetition is separately bounded and is not part of this vulnerability. This issue is fixed in versions 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0.	N/A	More Details
CVE-2026-108358	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. Prior to 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0, <code>XmlScanner::scan</code> in <code>src/PhpSpreadsheet/Reader/Security/XmlScanner.php</code> uses a DOCTYPE detection pattern that accounts for NUL-byte interleaving but not ISO-2022 escape sequences in XLSX XML parts. The escape sequences prevent the scanner from matching the declaration, and XML charset normalization then removes the escape bytes before parsing, exposing a valid DOCTYPE and entity definitions. Processing an attacker-controlled XLSX can therefore cause entity-expansion denial of service on supported PHP versions and can permit local file access or SSRF on PHP before 8.0 or when <code>LIBXML_NOENT</code> is enabled. This issue is fixed in versions 1.30.7, 2.1.19, 2.4.8, 3.10.8, and 5.10.0.	N/A	More Details

CVE-2026-108357	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. From 3.3.0 until 3.10.8 and 5.10.0, Reader\Xml passes stored SpreadsheetML XML through XmlScanner::scan in Xml::canRead and Xml::trySimpleXMLLoadStringPrivate, while Xml::unentity decodes numeric entities only after each DOCTYPE pattern check. A double-encoded less-than character is reduced by the first callback to another encoded character and by the second callback to a literal less-than sign, so SimpleXML receives an unchecked DOCTYPE and entity declaration. Processing attacker-controlled SpreadsheetML XML can cause entity-expansion denial of service and can permit local file access or SSRF on PHP before 8.0 or when LIBXML_NOENT is enabled. This issue is fixed in versions 3.10.8 and 5.10.0.	N/A	More Details
CVE-2026-91093	A potential security vulnerability has been identified in HP One Agent for certain HP PC products, which might allow for escalation of privilege and/or arbitrary code execution. HP is releasing software updates to mitigate the potential vulnerability.	N/A	More Details
CVE-2026-115572	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.2, an attacker can craft a PDF that repeatedly references the same nonexistent indirect object, causing relevant file parts such as pages to trigger repeated full-file recovery loads and long runtimes when accessed. This issue is fixed in version 6.21.2.	N/A	More Details
CVE-2026-115580	pypdf is a free and open-source pure-python PDF library. Prior to 6.21.2, an attacker can craft a PDF with a standard cross-reference table whose declared entry count exceeds what the file can contain, causing excessive processing and long runtimes when the file is opened. This issue is fixed in version 6.21.2.	N/A	More Details
CVE-2026-114485	Insufficient protection of sensitive information in the user list endpoint in Devolutions Server 2026.3.10.0 and earlier allows an authenticated attacker with delegated user view permissions to obtain the MFA secrets of all users via the user list API response.	N/A	More Details
CVE-2026-16477	CleanShot X for macOS insecurely deserializes untrusted data embedded in a .cleanshot project file (editable screenshot annotations). Exploitation requires an attacker to craft a malicious .cleanshot file, deliver it to the target, and have the target open it; the payload then executes automatically during parsing, resulting in arbitrary code execution in the context of the current user. Users who have never opened a .cleanshot file from an untrusted source are not affected by this vulnerability. This vulnerability was fixed in versions: 4.5.2, 4.6.3, 4.7.8 and 4.8.10	N/A	More Details
CVE-2026-115883	LangGraph checkpoint savers provide Postgres- and SQLite-backed state persistence for LangGraph workflows and agents. Prior to langgraph-checkpoint-postgres 3.1.3 and langgraph-checkpoint-sqlite 3.1.2, PostgresStore, AsyncPostgresStore, SqliteStore, and AsyncSqliteStore flatten namespace labels with ".".join(namespace), while BaseStore validates labels only for put and aput. Unvalidated dotted labels passed through get, search, delete, list_namespaces, their asynchronous equivalents, or direct batch and abatch operations can resolve to another tenant's stored namespace, allowing cross-namespace reads, overwrites, or deletion. InMemoryStore is not affected because it compares namespace tuples. This issue is fixed in langgraph-checkpoint-postgres 3.1.3 and langgraph-checkpoint-sqlite 3.1.2.	N/A	More Details
CVE-2026-96293	A cleartext storage of sensitive information vulnerability in WatchGuard Endpoint Security's Full Disk Encryption feature for macOS allows a local, unprivileged user or process to obtain a managed computer's FileVault Personal Recovery Key by reading a persistently written diagnostic log file left readable by all local users.	N/A	More Details
CVE-2026-86696	Allocation of Resources Without Limits or Throttling vulnerability in esaml in dropbox esaml allows an unauthenticated attacker to exhaust the Erlang VM atom table and crash the node via unsigned LogoutResponse messages. esaml:status_code_map/1 in src/esaml.erl maps any StatusCode value it does not recognize to an atom with list_to_atom/1 on the last colon-separated segment of the URN. esaml_sp:validate_logout_response/2 verifies a signature only when the message carries one, so an unsigned LogoutResponse posted to the single-logout endpoint reaches esaml:decode_logout_response/1 and mints one permanent atom per distinct status value. Atoms are never reclaimed and the VM caps the atom table at about one million entries, so an attacker who posts that many distinct values aborts the node. This issue affects esaml: from 0.1.0 onward.	N/A	More Details
CVE-2026-115253	An incorrect permission assignment vulnerability in the WatchGuard Agent's log directory allows a local, authenticated non-administrative user to create, modify, append to, or delete the agent's log files and write other files into the directory. This allows an attacker to tamper with or destroy the agent's security logs and stage further local attacks by placing executable content in a directory owned by a privileged service.	N/A	More Details
	The org.json parser in com.vaadin.external.google:android-json recurses once per nesting level with no depth limit: JSONTokener.nextValue() calls readArray() and readObject(), which call nextValue() again. A deeply nested JSON document, around 20 KB of nested arrays with the default thread stack, makes the JSONArray and JSONObject constructors throw a raw java.lang.StackOverflowError instead of a		

CVE-2026-110974	JSONException. Any application that parses untrusted JSON with this library can be made to fail the thread doing the parsing. The artifact is a 2013 repackage of third-party org.json classes, published once and never maintained, and no fixed version will be released. Users of the affected version should apply the following mitigation. Product version com.vaadin.external.google:android-json 0.0.20131108.vaadin1 Mitigation No fixed version will be released. Replace the artifact with an org.json implementation that limits the nesting depth, such as org.json:json 20231013 or later. Artifacts Maven coordinates / Vulnerable versions / Fixed version com.vaadin.external.google:android-json / 0.0.20131108.vaadin1 / none	N/A	More Details
CVE-2026-105270	Allocation of Resources Without Limits or Throttling vulnerability in esaml_binding in dropbox esaml allows an unauthenticated attacker to exhaust the Erlang VM atom table and crash the node via a crafted SAMLResponse. esaml_binding:decode_response/2 parses the decoded SAMLResponse with xmerl_scan:string/2, which interns every element name, attribute name, and namespace prefix it reads as a permanent atom. The SAML consume (ACS) and single-logout endpoints call it on unauthenticated requests before any signature check, so one well-formed document with thousands of distinct names mints thousands of atoms while staying inside the 128 KB request body limit. Atoms are never reclaimed and the VM caps the atom table at about one million entries, so a few dozen requests abort the node. This issue affects esaml: from 0.1.0 onward.	N/A	More Details
CVE-2026-96391	Missing Authorization vulnerability in Drupal Webform REST allows Forceful Browsing. This issue affects Webform REST versions: from 0.0.0 to 4.2.1.	N/A	More Details
CVE-2026-98891	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: validate deauth frame length before reason access ieee80211_rx_mgmt_deauth() reads the deauth reason code before checking that the fixed field is actually present in the received frame. Validate the deauth frame length first and only then read the reason code.	N/A	More Details
CVE-2026-96379	Improper Neutralization of Input During Web Page Generation ("Cross-site Scripting") vulnerability in Drupal CookieCuttr allows Stored XSS. This issue affects CookieCuttr versions: from 2.0.0 to 2.0.3.	N/A	More Details
CVE-2026-114629	Web::Simple versions before 0.035 for Perl allocate a buffer sized by the Content-Length header for urlencoded body parameters. Web::Dispatch::ParamParser::get_unpacked_body_from reads an application/x-www-form-urlencoded body using a single read call, with no cap and the return value ignored. Perl's read grows the buffer to the requested length before reading, so the declared Content-Length determines the allocation. An attacker can trigger this through any route whose dispatch specification has a % body-parameter match, such as sub (POST + %name=), sending a Content-Length far larger than the body actually sent. No authentication is involved.	N/A	More Details
CVE-2026-115557	Python has added a denylist for known compromised certificates to the "ssl" module to ensure user safety even without certificate revocation configured. Known compromised certificates will automatically be rejected regardless of the SSLContext policy configuration.	N/A	More Details
CVE-2026-96363	Privilege Defined With Unsafe Actions vulnerability in Drupal Webform allows Stored XSS. This issue affects Webform versions: from 0.0.0 to 6.2.12, from 6.3.0 to 6.3.1.	N/A	More Details
CVE-2026-109111	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-111326	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-89284	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-87941	Vulnerability in Drupal Patreon. This issue affects Patreon versions: 7.x.	N/A	More Details
CVE-2026-87938	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Drupal CSP log allows SQL Injection. This issue affects CSP log versions: from 0.0.0 to 1.0.2.	N/A	More Details
	Improper Handling of Highly Compressed Data (Data Amplification) in esaml_binding in dropbox esaml		

CVE-2026-93471	allows an unauthenticated attacker to exhaust server memory and crash the Erlang VM via a crafted compressed SAMLResponse. esaml_binding:decode_response/2 base64-decodes the attacker-supplied SAMLResponse, inflates it with zlib:unzip/1 with no bound on the output, then converts the result to a charlist with binary_to_list/1 before parsing. This runs on the SAML consume (ACS) and single-logout endpoints, which accept unauthenticated requests by design, and it runs before any signature check. A small deflate stream reaches a compression ratio near 1000 to 1, so a request body of tens of kilobytes inflates to hundreds of megabytes and then to several gigabytes of list heap, so a few requests exhaust memory and take the node down. This issue affects esaml: from 0.1.0 onward.	N/A	More Details
CVE-2026-106017	Improper certificate validation in the OpenVPN 3 Core library version 3.4 through 3.11.7 allows attackers to circumvent extended key usage restrictions.	N/A	More Details
CVE-2026-115885	The LangGraph CLI is the official command-line interface for LangGraph. From 0.4.15 until 0.4.32, langgraph deploy --verbose passes a temporary registry push token on standard input to docker login --password-stdin, but the CLI subprocess helper prints command input in verbose mode. Anyone who can read resulting CI logs or terminal recordings before the token's maximum 12-hour expiration can pull deployment images or push images to that deployment's image repository. Deployments without --verbose and remote image builds are not affected, and the token alone cannot update a deployment to execute a pushed image. This issue is fixed in version 0.4.32.	N/A	More Details
CVE-2026-102785	Joomla Extension - yootheme.com - Unauthenticated remote code execution in YOOtheme Pro 5.0.0-5.0.52 - YOOtheme Pro's publicly accessible image transformation handler incompletely binds normalized image operations and the output destination to its request signature. An unauthenticated remote attacker can reuse a signature disclosed in a legitimate public image URL, inject additional image operations, traverse out of the intended cache directory, and write an attacker-controlled PNG/PHP polyglot to an executable public media directory.	N/A	More Details
CVE-2026-94502	In EmberZNet 9.1.1 and earlier, a malformed Publish Price Matrix command can result in an out-of-bounds write potentially leading to remote code execution. These messages must come from a device that has already joined the network. Only devices supporting the Price Cluster are impacted.	N/A	More Details
CVE-2026-89410	In Z-Wave SDK 8.1.1 and earlier, a malformed NCP SPI command can trigger out-of-bounds reads that leak sensitive information.	N/A	More Details
CVE-2026-115892	Wasmtime is a runtime for WebAssembly. From 46.0.0 until 48.0.7 and 49.0.4, lowering a host-created component-model stream or future into a guest can omit a required payload-type check after the value passes through guest code and returns through StreamAny, FutureAny, Val::Stream, Val::Future, or an equivalent typed reader. A malicious guest that controls exported function types or returned values can cause the host to pass the type-erased value to a guest expecting a different payload type, which panics the Wasmtime host. This issue is fixed in versions 48.0.7 and 49.0.4.	N/A	More Details
CVE-2026-89409	In Z-Wave SDK 8.1.1 and earlier (for 700/800 series devices) and 6.85.02 and earlier (for 500 series devices), missing checks in the ZWAVE_CMD_NODES_EXIST handler can result in nodes being removed from the network. These messages must come from a device that has already joined the network.	N/A	More Details
CVE-2026-115897	Serialize JavaScript serializes JavaScript values to a superset of JSON that includes regular expressions and functions. Prior to 7.1.3, the function-source escaping logic does not neutralize a bare closing-script prefix when attacker-influenced function source text ends exactly at that prefix. Because each function body is escaped independently, a delimiter supplied by surrounding serialized output can terminate an HTML script element, and a later serialized value supplies injected markup that becomes a live element in the page. Exploitation requires control of a serialized function's source text, such as through a spoofed toString(). This issue is fixed in version 7.1.3.	N/A	More Details
CVE-2026-85080	In EmberZNet 9.1.1 and earlier, a malformed GetGroupMembership message can result in an out-of-bounds write leading to a crash. A limited number of values can be written using this vulnerability, so no route to remote code execution was observed. These messages must come from a device that has already joined the network. Only devices supporting the Group Cluster are impacted.	N/A	More Details
CVE-2026-85079	In EmberZNet 9.1.1 and earlier, a malformed Simple_Desc_req response can result in an out-of-bounds write potentially leading to remote code execution. This response must come from a device that has already joined the network, and the attacker must wait for a Simple_Desc_req before the attack can be performed.	N/A	More Details
CVE-2026-66282	In EmberZNet 9.1.1 and earlier, a malformed Publish Billing Period command can result in a recursive function call that will never reach its end condition. These messages must come from a device that has already joined the network. Only devices supporting the Price Cluster are impacted.	N/A	More Details
CVE-			

2026-108537	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-108793	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-110814	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-66281	In EmberZNet 9.1.1 and earlier, a malformed Discover Attributes command can result in a loop with unreachable exit condition (infinite loop). These messages must come from a device that has already joined the network. Only devices supporting manufacturer-specific clusters are impacted.	N/A	More Details
CVE-2026-66279	In Connect SDK 4.3.1 and earlier and Flex SDK 3.7.7 and earlier, a message with invalid security settings may bypass authentication resulting in network disruption. This vulnerability does not allow an attacker to gain the Network Key or otherwise break confidentiality.	N/A	More Details
CVE-2026-111497	A buffer overflow vulnerability in the Brotli decoder (c/dec/decode.c) in Google Brotli versions 1.0.3 through 1.2.0 on all platforms allows a remote unauthenticated attacker to corrupt adjacent decoder state memory, disclose sensitive memory contents, cause a denial of service, or potentially execute arbitrary code using a crafted Brotli-compressed stream. Users need to update any Brotli libraries or binaries between 1.0.3 through 1.2.0 to the latest available version: 1.0.10 or 1.1.1 or 1.2.1.	N/A	More Details
CVE-2026-107079	An issue in Systrome Validation & Certification 1.0 allows an attacker to cause a denial of service via the t URL parameter in the qrcode_auth.html component	N/A	More Details
CVE-2026-106069	A flaw in Node.js TLS hostname verification can cause a server certificate to be validated against a different IP address than the one the client connected to, when the host is a non-canonical IPv4 literal with a leading zero, such as `0127.0.0.1`. `tls.checkServerIdentity()` normalizes the host following the WHATWG URL rules, which read a leading-zero octet as octal (`0127.0.0.1` becomes `87.0.0.1`), while the connection is made through the operating system resolver, which on some platforms such as macOS reads it as decimal (`127.0.0.1`). A server at the connected address can then be authenticated with a certificate issued for another IP address. This vulnerability affects Node.js 26.x , 24.x , and 22.x .	N/A	More Details
CVE-2026-105482	A flaw in Node.js DNS resolution allows hostnames containing a NUL byte to be silently truncated by the dns.resolve*() functions and dns.reverse(). This is an incomplete fix of CVE-2026-48930, which only covered dns.lookup() and the net/tls host options. Because the name is passed to c-ares as a C string, resolving attacker.example\0.trusted.example queries attacker.example and returns its records without any error. Applications that validate a hostname, for example against a domain allowlist, and then resolve it with dns.resolve*() can act on records for a different name than the one they validated. This vulnerability affects Node.js 26.x, 24.x, and 22.x.	N/A	More Details
CVE-2026-75775	Improper Neutralization of Special Elements used in an OS Command in certain ASUS routers allows remote user to execute arbitrary commands with root privileges via crafted parameters sent to the router while the device is in its factory-default state.Refer to the ' Security Update for ASUS Router Firmware ' section on the ASUS Security Advisory for more information.	N/A	More Details
CVE-2026-99063	In the Linux kernel, the following vulnerability has been resolved: dm ioctl: don't let data_start run past the buffer struct dm_ioctl has seven bytes of trailing padding, so the smallest buffer userspace is allowed to ask for, offsetof(struct dm_ioctl, data), is seven bytes shorter than sizeof(struct dm_ioctl). get_result_buffer() sets data_start to sizeof(struct dm_ioctl) without looking at how big the buffer actually is. Ask for the minimum and data_start ends up seven bytes past the end. Callers add it to data_size, and ctl_ioctl() copies that much back to userspace: BUG: KASAN: slab-out-of-bounds in _copy_to_user+0xad/0xd0 Read of size 312 at addr ffff8880169e4c00 by task repro-dm/140 ctl_ioctl+0x5e3/0xcb0 dm_ctl_ioctl+0x25/0x40 allocated 305-byte region [fff8880169e4c00, ffff8880169e4d31) A DM_LIST_DEVICES with data_size set to 305 is enough to hit it, and the seven bytes that leak are whatever sat next to the allocation. Clamp data_start to the buffer size. Note that this is not a security bug because the out-of-bounds read can only be triggered by root.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: selinux: preserve NATIVE_LABELS on already-initialized sb in set_mnt_opts When an NFSv4 filesystem exported with security_label is mounted a second time, sget_fc() returns the existing superblock and nfs_get_root() calls security_sb_set_mnt_opts() again on it. Since the superblock is already SB_INITIALIZED, selinux_set_mnt_opts() takes an early exit without setting SECURITY_LSM_NATIVE_LABELS in		

2026-99052	*set_kern_flags. nfs_get_root() then sees kflags without SECURITY_LSM_NATIVE_LABELS set and clears NFS_CAP_SECURITY_LABEL on the shared nfs_server, silently disabling per-file labeling for both the new and the original mount. Commit 3815a245b501 ("security/selinux: fix SECURITY_LSM_NATIVE_LABELS on reused superblock") fixed the same problem in selinux_sb_clone_mnt_opts(). The reused-superblock path for a plain second mount now goes through sb_set_mnt_opts rather than sb_clone_mnt_opts, so it needs the same treatment: on both SE_SBINITIALIZED early exits, set SECURITY_LSM_NATIVE_LABELS in *set_kern_flags when the superblock was already using SECURITY_FS_USE_NATVE.	N/A	More Details
CVE-2026-99055	In the Linux kernel, the following vulnerability has been resolved: tcp: reject net_iov in zerocopy receive mapping hints After copying a readable prefix, receive_fallback_to_copy() asks tcp_zerocopy_set_hint_for_skb() where page mapping can resume. If the next skb is unreadable, find_next_mappable_frag() passes its net_iov fragment to can_map_frag(). skb_frag_page() returns NULL for a net_iov, but can_map_frag() dereferences it in PageCompound(). A v7.2 KASAN run on a connected TCP socket with 64 readable bytes followed by a 4096-byte NET_IOV_DMABUF fragment reported: BUG: KASAN: null-ptr-deref in can_map_frag tcp_zerocopy_receive -> can_map_frag Kernel panic - not syncing: KASAN: panic_on_warn set The diagnostic inserted the net_iov directly because the test host has no devmem-capable NIC. Hardware end-to-end reachability remains untested and requires CONFIG_NET_DEVMEM plus a supported DMA-buf-bound RX queue. Reject all net_iov fragments before skb_frag_page(). This covers both DMABUF and IOURING net_iov types while leaving page-backed checks unchanged. With the guard, the same queue copied the readable prefix, returned a 4096-byte skip hint, and completed without a fault.	N/A	More Details
CVE-2026-99057	In the Linux kernel, the following vulnerability has been resolved: xen/netfront: don't leak the skb when xennet_fill_frags() fails When a response chain has more slots than fit in the skb's frags, xennet_fill_frags() returns an error and xennet_poll() jumps to its error path. That path moves what's left on tmpq to errq to be freed, but the skb being filled was already dequeued from tmpq, so it's never freed. Each chain that overflows leaks the skb and the pages attached to it as frags, and the backend decides how many slots it sends. Put the skb back on tmpq before taking the error path, like the xennet_set_skb_gso() failure just above it does.	N/A	More Details
CVE-2026-99058	In the Linux kernel, the following vulnerability has been resolved: workqueue: Fix NULL current_pwq deref in chained work check current_wq_worker() only tells us that %current is a kworker. It does not guarantee that the worker is currently executing a work item, as worker->current_pwq is populated only while process_one_work() is running the work function. is_chained_work() can be reached while queueing on a draining or destroying workqueue. If a kworker outside work-item execution reaches that path, current_wq_worker() returns a worker but worker->current_pwq is NULL, and the chained-work test can fault before it emits the intended warning. Guard the chained-work test with worker->current_pwq. A kworker without current_pwq is not executing work on the target workqueue, so the helper should return false and let the draining/destroying warning fire.	N/A	More Details
CVE-2026-99060	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7603: initialize global station WCID mt7603 uses a driver-private global station WCID for frames that have neither a station nor a vif, including frames injected through a monitor interface. Unlike the core global WCID initialized by mt76_register_device(), this WCID is published without ever being passed to mt76_wcid_init(). Since commit 0335c034e726 ("wifi: mt76: fix race condition related to checking tx queue fill status"), mt76_tx() queues every such frame on the WCID's tx_pending or tx_offchannel queue and links its tx_list. The first injected frame therefore reaches __skb_queue_tail() with an uninitialized sk_buff_head and dereferences a NULL prev pointer. Initialize the WCID before publishing it. The fix was tested on an MT7603 PCIe device. Three wildcard probe requests were injected through a monitor interface. The hardware queue head and tail advanced from 0 to 3 and a second radio captured the third request over the air. No oops or reset occurred. After restoring AP mode, a 150 MB transfer also completed without a reset.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: wifi: brcmfmac: fix P2P device removal race in brcmf_detach() When the driver is removed while the user space process that created the P2P device (e.g. wpa_supplicant) is still exiting, the P2P device interface is removed twice and the kernel crashes. The two paths are: - wpa_supplicant, on exit, sends NL80211_CMD_DEL_INTERFACE for the P2P device. nl80211_del_interface() holds RTNL and the wiphy mutex and calls brcmf_p2p_del_vif(), which disables discovery and waits up to 1.5 s for BRCMF_E_IF_DEL. Then it calls brcmf_remove_interface(), which unregisters the wdev and frees the vif and the ifp. - brcmf_detach() sets the bus down, so BRCMF_E_IF_DEL never arrives, and calls brcmf_remove_interface(ifp, false) for the same ifp. brcmf_p2p_ifp_removed() reads ifp->vif and then blocks on rtnl_lock(). When brcmf_p2p_del_vif() releases RTNL, it calls cfg80211_unregister_wdev() on the wdev that was already unregistered and freed. The second list_del_rcu() of wdev->list faults on LIST_POISON2: brcmfmac: brcmf_p2p_del_vif delete P2P vif brcmfmac: brcmf_p2p_deinit_discovery enter brcmfmac: brcmf_p2p_del_vif P2P: GO_NEG_PHASE status cleared brcmfmac: brcmf_detach Enter brcmfmac: brcmf_bus_change_state 1 -> 0 brcmfmac: brcmf_remove_interface Enter, bsscfgidx=1, ifidx=0 brcmfmac: brcmf_del_if Enter, bsscfgidx=1, ifidx=0 brcmfmac: brcmf_p2p_ifp_removed P2P: device		

CVE-2026-99061	interface removed [1.5 s later, brcmf_p2p_del_vif() timed out] brcmfmac: brcmf_remove_interface Enter, bsscfgidx=1, ifidx=0 brcmfmac: brcmf_del_if Enter, bsscfgidx=1, ifidx=0 brcmfmac: brcmf_p2p_ifp_removed P2P: device interface removed Unable to handle kernel paging request at virtual address dead00000000122 Internal error: Oops: 96000044 [#1] PREEMPT_RT SMP pc : _cfg80211_unregister_wdev+0x6c/0x264 [cfg80211] Call trace: _cfg80211_unregister_wdev+0x6c/0x264 [cfg80211] cfg80211_unregister_wdev+0x10/0x1c [cfg80211] brcmf_p2p_ifp_removed+0x84/0xb0 [brcmfmac] brcmf_remove_interface+0x1f4/0x254 [brcmfmac] brcmf_detach+0x88/0x180 [brcmfmac] brcmf_sdio_remove+0x90/0x7b0 [brcmfmac] brcmf_sdiod_remove+0x20/0xa0 [brcmfmac] brcmf_ops_sdio_remove+0xbc/0x12c [brcmfmac] sdio_bus_remove+0x38/0x144 device_release_driver_internal+0x1e8/0x2c0 device_release_driver+0x14/0x20 brcmf_sdio_bus_remove+0x2c/0x40 [brcmfmac] brcmf_fwvid_unregister_vendor+0xd8/0x170 [brcmfmac] __exit_compat+0x18/0x418 [brcmfmac_cyw] __arm64_sys_delete_module+0x1ac/0x244 The log was taken with debug=0x406 on an i.MX 8M Plus board with a CYW55513 (Sona IF513) on SDIO. The driver is the Ezurio backport of brcmfmac from v6.18.22 on a 5.4-rt kernel. The code involved is the same in mainline. To reproduce: 1. Start wpa_supplicant on wlan0 and let it associate. It creates the P2P device (iw dev shows "type P2P-device"). 2. Stop wpa_supplicant and remove the driver without waiting for wpa_supplicant to exit: kill \$(pidof wpa_supplicant) rmmod brcmfmac_cyw brcmfmac The crash is reproducible on that board. For the interface without a netdev (the P2P device), take RTNL and the wiphy mutex in brcmf_detach() before reading iflist, and remove it with locked=true. If brcmf_p2p_del_vif() runs first, brcmf_detach() finds the slot empty. If brcmf_detach() runs first, nl80211 does not find the wdev any more. The interfaces with a netdev are removed as before, because brcmf_del_if() takes RTNL on its own for the primary interface.	N/A	More Details
CVE-2026-99062	In the Linux kernel, the following vulnerability has been resolved: vsock/virtio: account only unread bytes in read_skb() After a partial stream receive, rx_bytes tracks the unread suffix while buf_used and peer credit still cover the whole packet. virtio_transport_read_skb() dequeues that packet but passes pkt_len for both counters, underflowing rx_bytes. On a connected CID_LOCAL stream with an SK_SKB verdict, a 40-byte read from a 100-byte packet made SIOCINQ report -40, an empty rcv return ELOOP, and poll report the empty socket readable. After verdict detach, receiving 40 bytes wrapped the counter to zero and hid those queued bytes until one more byte arrived. The fixed run kept the counter balanced and exposed all queued data immediately. This was found during an LLM-assisted manual source audit of VSOCK receive accounting while re-evaluating virtio_transport_read_skb() after CVE-2024-50169. Subtract only skb->len minus the VSOCK offset from rx_bytes. Retain the full packet length for buf_used and peer credit.	N/A	More Details
CVE-2026-107858	Missing Origin Validation in WebSockets (CWE-1385) and Permissive Cross-domain Policy with Untrusted Domains (CWE-942) in the development web server (ADK Web) in Google Agent Development Kit (ADK) for Java (com.google.adk:google-adk-dev versions 0.1.0 before 1.13.0 and com.google.adk:google-adk-maven-plugin versions 0.3.0 before 1.13.0) allow a remote attacker to control a locally running agent and read its data via a malicious web page visited by a developer while the server is running. The /run_live WebSocket endpoint accepts connections from any origin by default (hardcoded via setAllowedOrigins("*") in versions 0.1.0 through 1.5.0, and defaulting to * via adk.web.cors.origins in versions 1.6.0 through 1.12.0), and since version 0.4.0 the REST API returns Access-Control-Allow-Origin: * by default. Because the server has no authentication and requests originate from the developer's browser (so binding to localhost does not prevent exploitation), a cross-origin page can send messages to the agent (triggering registered tools), stream agent responses, and list, read, create, or delete sessions and artifacts. In versions 1.6.0 through 1.12.0, users can mitigate the issue by setting the adk.web.cors.origins property to an explicit allowlist of trusted origins (such as http://localhost:8000). The vulnerability is fixed in version 1.13.0, which enforces a same-origin default for CORS and the /run_live WebSocket endpoint, and rejects DNS-rebinding requests over loopback.	N/A	More Details
CVE-2026-105481	A flaw in Node.js HTTP Agent connection reuse can cause a client to accept as valid a response that was sent before the client has sent the request. This is an incomplete fix of CVE-2026-48931. The guard that destroys idle keep-alive sockets when unexpected data arrives only runs when the event loop polls for I/O. If an unsolicited response has already reached the socket but has not been read yet when the socket is handed to a new request, for example because the event loop was busy with synchronous work, the guard is removed without seeing it and the data is parsed as the response to the new request. This vulnerability affects Node.js **26.x**, **24.x**, and **22.x**.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: cgroup/cpuset: Handle cpu hotplug race in guarantee_active_cpus() With commit 2125c0034c5d ("cgroup/cpuset: Make cpuset hotplug processing synchronous"), the cpuset hotplug operation becomes synchronous. That commit also removes the code that handles race between cpuset_hotplug_work and cpu hotplug notifier with the assumption that race is now gone. Later commit 7a0aabd9ce69 ("cgroup/cpuset: Always use cpu_active_mask") updates the cpuset code to always use cpu_active_mask instead of cpu_online_mask in some places including guarantee_online_cpus() which is renamed to guarantee_active_cpus() in that		More

2026-99065	commit. In the case of CPU offline operation, <code>cpu_active_mask</code> is updated first in <code>sched_cpu_deactivate()</code> to remove the offline CPU before <code>cpuset_handle_hotplug()</code> is called to update the effective cpus of the affected cpusets. The <code>cpu_online_mask</code> is updated after that near the end of the offline operation to remove the offline CPU. As a result, the race comes back and the top cpuset may not have any active CPU leading to NULL pointer dereference during the race window when <code>guarantee_active_cpus()</code> is called after <code>cpu_active_mask</code> is updated to remove the CPU to be torn down but before <code>cpuset_handle_hotplug()</code> is able to properly update the effective cpus of the top cpuset. Fix this by adding back the NULL cs check to avoid this problem.	N/A	Details
CVE-2026-99066	In the Linux kernel, the following vulnerability has been resolved: can: dev: <code>init_can_skb()</code> : restore skb header initialization Commit 9f10374bb024 ("can: remove private CAN skb headroom infrastructure") removed the <code>skb_reset_mac_header()/skb_reset_network_header()/skb_reset_transport_header()</code> calls from <code>init_can_skb()</code> . As a result, RX skbs from <code>alloc_can_skb()</code> and friends again carry <code>mac_header = 0xFFFF</code> . When such an skb reaches <code>packet_rcv_spkt()</code> (SOCK_PACKET), the push length calculation overflows and triggers <code>skb_under_panic -> kernel BUG -> full machine panic</code> . The same issue was originally reported in 2014 on linux-can and fixed by commit 969439016d2c ("can: add missing initialisations in CAN related skbuffs"). <code>packet_rcv_spkt()</code> itself has never been hardened: only <code>packet_rcv()</code> and <code>tpacket_rcv()</code> gained <code>dev_has_header()</code> checks in commit d549699048b4 ("net/packet: fix packet receive on L3 devices without visible hard header").	N/A	More Details
CVE-2026-99067	In the Linux kernel, the following vulnerability has been resolved: btrfs: scrub: fix local_root reference leak in <code>scrub_print_warning_inode()</code> When <code>paths_from_inode()</code> fails, <code>scrub_print_warning_inode()</code> jumps to err without dropping the reference taken by <code>btrfs_get_fs_root()</code> , leaking a reference to the root every time path resolution fails while printing scrub warnings. Every other error and success path of the function drops the reference. Drop the reference on the <code>paths_from_inode()</code> failure path too.	N/A	More Details
CVE-2026-99068	In the Linux kernel, the following vulnerability has been resolved: ALSA: ctxfi: Fix memory leakage when clearing DAO input mappers Commit e576e7843c0d ("ALSA: ctxfi: Simplify <code>dao_clear_{left,right}_input()</code> functions") centralized the DAO input mapper clearing logic in <code>dao_clear_input()</code> , but omitted the <code>kfree()</code> calls. It removes the mapper entries from the list and sets their pointers to NULL without freeing the allocated array. This leaks memory each time <code>dao_clear_input()</code> is called. Store the base pointer of the allocation before clearing the mapper entries and free the array after removing entries from the list.	N/A	More Details
CVE-2026-99069	In the Linux kernel, the following vulnerability has been resolved: Input: <code>ati_remote2</code> - fix type confusion in device attribute handlers Previously, the <code>channel_mask</code> and <code>mode_mask</code> attributes were registered on the parent <code>usb_device</code> via <code>sysfs_create_group()</code> . When commit b20d6bf8014b ("Input: <code>ati-remote2</code> - use driver core to instantiate device attributes") converted the driver to let the driver core manage them via <code>dev_groups</code> in struct <code>usb_driver</code> , the attributes moved to the bound <code>usb_interface</code> devices, but the show and store callbacks were not updated accordingly. As a result, the callbacks still invoked <code>to_usb_device()</code> on a device pointer belonging to a struct <code>usb_interface</code> , producing an invalid <code>usb_device</code> pointer. Use <code>to_usb_interface()</code> directly in the attribute handlers to obtain the interface and its driver data.	N/A	More Details
CVE-2026-99070	In the Linux kernel, the following vulnerability has been resolved: ipv6: do not warn on route notification size race <code>fib6_info_hw_flags_set()</code> sizes the skb with <code>rt6_nlmsg_size()</code> and later fills it with <code>rt6_fill_node()</code> . With nexthop compatibility mode enabled, a concurrent replacement can grow the group between these independent snapshots. <code>rt6_fill_node()</code> can then legitimately return <code>-EMSGSIZE</code> , so the warning does not prove a sizing bug. Remove the warning. The existing error path still frees the skb and reports the error to listeners.	N/A	More Details
CVE-2026-99071	In the Linux kernel, the following vulnerability has been resolved: ipv4: do not warn on route notification size race Commit 680aea08e78c ("net: ipv4: Emit notification when fib hardware flags are changed") added asynchronous route notifications for hardware flag changes. <code>fib_alias_hw_flags_set()</code> sizes the skb with <code>fib_nlmsg_size()</code> and later fills it with <code>fib_dump_info()</code> while holding only RCU. With nexthop compatibility mode enabled, a concurrent replacement can grow the group between these independent snapshots. <code>fib_dump_info()</code> can then legitimately return <code>-EMSGSIZE</code> , so the warning does not prove a sizing bug. Remove the warning. The existing error path still frees the skb and reports the error to listeners.	N/A	More Details
CVE-2026-99049	In the Linux kernel, the following vulnerability has been resolved: gpio: xilinx: fix runtime PM leak on request error path <code>pm_runtime_get_sync()</code> leaves the usage counter incremented even when it fails, and <code>xgpio_request()</code> returns the error without dropping it. <code>gpiolib</code> does not call <code>->free()</code> when <code>->request()</code> fails, so <code>xgpio_free()</code> , which holds the only matching <code>pm_runtime_put()</code> , never runs. The reference is leaked and the controller can no longer runtime-suspend, so its clock stays enabled. Switch to <code>pm_runtime_resume_and_get()</code> , which only increments the usage counter on success.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: gve: fix napi_disable deadlock when		

CVE-2026-99048	attempting to disable XSK pools When disabling XSK pools, GVE calls the unlocked versions of napi_disable and napi_enable. However, the netdev lock has already been acquired before ndo_bpf is called because GVE supports queue management ops. Calling the unlocked versions of napi_disable/enable results in a deadlock when attempting to disable XSK pools, as the thread attempts to re-acquire a lock it already holds. Update the NAPI calls to use the locked versions. This issue was found in production.	N/A	More Details
CVE-2026-99047	In the Linux kernel, the following vulnerability has been resolved: gve: fix XSK buffer leak when rings are stopped GVE does not free XSK buffers when resetting ring state as a part of stopping queues. This causes all XSK buffers which are posted to the NIC to be leaked. Free XSK buffers attached to an allocated buf_state when stopping rings. This issue was hit in production.	N/A	More Details
CVE-2026-4950	The CustomerAdministration portal (KundenManager) of VertiGIS WebOrderSystem allows privileged users to manage the employees of defined companies (user groups). User properties, for example the "secondName" parameter, can be filled with JavaScript code while editing a user. The code is stored and executed in the browser of any other privileged user who opens the edited user entry or the user overview of the corresponding user group. An authenticated attacker with access to the portal can use this to run script code in the session of other portal users. Version 7.0.8.6, the latest available at the time of testing, is affected. The vendor did not provide information about further affected versions.	N/A	More Details
CVE-2026-99031	In the Linux kernel, the following vulnerability has been resolved: net: bridge: avoid recursive multicast port cleanup br_multicast_sg_del_exclude_ports() removes automatically added STAR_EXCL port groups by calling br_multicast_del_pg(). For S,G entries, that deletion calls back into br_multicast_sg_del_exclude_ports(), so a large number of STAR_EXCL ports consumes one kernel stack frame per port and can hit the stack guard page. Keep the complete port-group deletion path for cleanup, but suppress the S,G exclude-port cleanup while that helper is already walking the same entry. Normal deletion callers retain the existing behavior.	N/A	More Details
CVE-2026-4952	The "ASMobile" web interface, which is part of the VertiGIS software package "GIS Portal", allows authenticated users to define a business reference. JavaScript code can be inserted into this field. On its own this only affects the user who entered it (self-XSS). However, the business reference is included in the order details when the user places an order through the web interface. The order is later listed in the OrderAdministration portal (AuftragsManager) of VertiGIS WebOrderSystem, where the stored code is executed as soon as a highly privileged user opens the order. Version 7.0.8.6 of WebOrderSystem, the latest available at the time of testing, is affected. The vendor did not provide information about further affected versions.	N/A	More Details
CVE-2026-99033	In the Linux kernel, the following vulnerability has been resolved: net/packet: guard the ll header push in packet_rcv_spkt() packet_rcv_spkt() restores the link layer header with skb_push(skb, skb->data - skb_mac_header(skb)); That subtraction is only meaningful when the device actually has a link layer header. packet_rcv() and tpacket_rcv() both wrap it in dev_has_header(), which is also the predicate the block comment at the top of the file states the restore in terms of; packet_rcv_spkt() does not. commit d549699048b4 ("net/packet: fix packet receive on L3 devices without visible hard header") introduced the helper and changed the two call sites, and this one stayed behind. A device without a visible ll header can leave skb->mac_header at the 0xFFFF sentinel that __alloc_skb() initialises it to. A CAN skb does: init_can_skb() sets pkt_type and ip_summed but does not reset the headers, and commit 9f10374bb024 ("can: remove private CAN skb headroom infrastructure") dropped the skb_reset_*_header() calls that used to be there. skb_mac_header() is then 0xFFFF, the length becomes a large negative number and skb_push() reports it through skb_under_panic() -- from softirq context, so it is a full system panic even with panic_on_oops=0: skbuff: skb_under_panic: text:ffffffff8bd21bc1 len:-65455 put:-65471 head:... data:... tail:0x50 end:0x180 dev:can0 kernel BUG at net/core/skbuff.c:214! RIP: 0010:skb_panic+0x50/0x60 Call Trace: <IRQ> skb_push+0x38/0x40 packet_rcv_spkt+0xe1/0x170 __netif_receive_skb_core.constprop.0+0x7e8/0xd30 ... Kernel panic - not syncing: Fatal exception in interrupt The socket type is reachable: packet_create() accepts SOCK_PACKET alongside SOCK_RAW and SOCK_DGRAM behind the same CAP_NET_RAW check, and neither the socket length nor a capability check keeps it away from a CAN interface. The missing skb_reset_*_header() calls in init_can_skb() are a regression in their own right and are being fixed separately, but a packet socket should not turn a link layer that did not initialise its mac header into a kernel panic. Guard the push the way the other two receive paths do. Tested on v7.3-rc5 under QEMU with a slcan device on a pty, which is the driver RX path: vcan does not reproduce it, because can_send() resets the headers on the way out. One SOCK_PACKET socket bound to can0 and one frame written into the line discipline panics an unpatched kernel with the trace above; the same image with this patch prints no panic and powers off normally. Both kernels are this tree, defconfig plus CONFIG_CAN_SLCAN=y, differing only in this hunk.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: mmc: sdhci-sprd: disable runtime PM on remove sdhci_sprd_probe() enables runtime PM, while sdhci_sprd_remove() does not perform the corresponding runtime PM cleanup. The probe failure path disables runtime PM before disabling the		More

2026-99034	clocks, but the normal remove path directly disables clocks that are also managed by the runtime PM callbacks. If the device is runtime suspended, those clocks may already be disabled. Resume the device before removal, disable runtime PM and drop the temporary runtime PM reference before disabling the clocks. This issue was found by manual code inspection.	N/A	Details
CVE-2026-99035	In the Linux kernel, the following vulnerability has been resolved: mmc: mtk-sd: Cancel request timeout work on remove The driver queues req_timeout while a request is active. msdc_request_done() uses cancel_delayed_work(), which does not wait for a timeout callback that has already started. The timeout callback calls mmc_request_done(), which wakes the request waiter, and then continues to use host->dev_comp and check the SDIO IRQ. During unbind, msdc_drv_remove() can return and the managed mmc_host can be freed before the callback finishes. KASAN reported use-after-free accesses in msdc_request_done() and msdc_recheck_sdio_irq() in each of three unbind tests. The same tests completed without a kernel diagnostic after this change. Call cancel_delayed_work_sync() after mmc_remove_host() has stopped new requests and before the driver releases the host resources.	N/A	More Details
CVE-2026-99036	In the Linux kernel, the following vulnerability has been resolved: dm: fix reading free memory in do_resume When do_resume calls dm_table_get_mode(new_map), the call is done without holding any locks (it only holds a reference to the md). It may be possible that another concurrent ioctl on the same device will swap table after dm_swap_table and before dm_table_get_mode. In this case, dm_table_get_mode(new_map) reads freed memory. This race condition was triggered by syzbot. This commit fixes the race by reading mode of the new table before dm_swap_table. Note that new_map is also passed to dm_ima_need_measure, but this function doesn't attempt to dereference it, so this call should be safe. Note that this is not a security bug because only root can trigger it and the commonly used tools such as lvm or cryptsetup do not call the ioctls concurrently, so that they can't trigger it accidentally.	N/A	More Details
CVE-2026-99037	In the Linux kernel, the following vulnerability has been resolved: dm-crypt: reject the lmk IV mode with AEAD ciphers A "capi:" table line that pairs an AEAD cipher with the lmk IV mode, for example "capi:gcm(aes)-lmk" with the integrity:28:aead option, causes an oops when it is loaded. lmk sets tfms_count to 64, but crypt_alloc_tfms_aead only allocates one transform, so crypt_setkey reads past the end of the tfms_aead array and calls crypto_aead_setkey on the values it finds there. lmk exists for compatibility with Loop-AES, which uses CBC, so refuse it with an AEAD cipher rather than allocate 64 AEAD transforms. The check cannot go in crypt_iv_lmk_ctr with the other lmk checks, because the IV constructor runs after crypt_set_key. Note that this is not a security bug because the memory corruption can only be triggered by root.	N/A	More Details
CVE-2026-99038	In the Linux kernel, the following vulnerability has been resolved: memstick: rtsx_usb_ms: complete requests after eject instead of dropping them memstick_check() can pass its host->removing check just before rtsx_usb_ms_drv_remove() sets eject/removing. Its next request is then silently dropped: rtsx_usb_ms_request() skips schedule_work() once eject is set, and drv_remove's cancel_work_sync() can also cancel a queued handle_req before it picks the request up. Nobody completes card->mrq_complete. Once memstick core waits for requests without a timeout ("memstick: core: wait for request completion before freeing card"), this hangs removal: memstick_check() never returns and memstick_remove_host() blocks in flush_workqueue(): INFO: task kworker/u10:3:65 blocked for more than 20 seconds. Workqueue: kmemstick memstick_check __wait_for_common memstick_check INFO: task kworker/1:1:33 blocked for more than 20 seconds. Workqueue: usb_hub_wq hub_event __flush_workqueue memstick_remove_host rtsx_usb_ms_drv_remove Never drop a request. rtsx_usb_ms_request() always schedules handle_req, and handle_req fails requests with -ENOMEDIUM once eject is set, without touching the device. drv_remove flushes handle_req instead of cancelling it, and cancels it only after memstick_remove_host(), when no new request can arrive, so it cannot run on a freed host. The host_mutex drain in drv_remove is removed: handle_req always leaves host->req NULL when it finishes, so the drain never had anything to do, and it would now race with handle_req.	N/A	More Details
CVE-2026-99039	In the Linux kernel, the following vulnerability has been resolved: memstick: core: wait for request completion before freeing card In memstick_alloc_card() and memstick_set_rw_addr(), wait_for_completion_timeout() is used to wait for device identification and address configuration requests to complete. The memstick subsystem does not provide an abort or cancellation callback in struct memstick_host for in-flight requests. If wait_for_completion_timeout() expires (for instance, when a USB transfer or command retry is delayed), memstick_alloc_card() returns an error and frees the memstick_dev structure via kfree(card). However, the host controller driver (such as rtsx_usb_ms) is still running the request asynchronously in worker context. When the transfer eventually finishes, the driver writes into host->req->data, which points to memory inside the already freed card structure. This leads to a KASAN slab-use-after-free write: BUG: KASAN: slab-use-after-free in ms_read_bytes+0x8fe/0x990 Write of size 1 at addr ffff88813c188079 by task kworker/1:5/5239 Workqueue: events rtsx_usb_ms_handle_req Call Trace: ms_read_bytes+0x8fe/0x990 drivers/memstick/host/rttsx_usb_ms.c:450 rtsx_usb_ms_issue_cmd drivers/memstick/host/rttsx_usb_ms.c:477 [inline] rtsx_usb_ms_handle_req+0x557/0x1a00 drivers/memstick/host/rttsx_usb_ms.c:531 process_one_work+0xac7/0x1b10 kernel/workqueue.c:3396	N/A	More Details

	worker_thread+0x5ef/0xe50 kernel/workqueue.c:3560 kthread+0x373/0x450 kernel/kthread.c:436 Allocated by task 1662: memstick_alloc_card drivers/memstick/core/memstick.c:383 [inline] memstick_check+0x2c1/0x10d0 drivers/memstick/core/memstick.c:452 process_one_work+0xac7/0x1b10 kernel/workqueue.c:3396 Freed by task 1662: kfree+0x20c/0x650 mm/slab.c:6792 memstick_alloc_card drivers/memstick/core/memstick.c:420 [inline] memstick_check+0x72b/0x10d0 drivers/memstick/core/memstick.c:452 process_one_work+0xac7/0x1b10 kernel/workqueue.c:3396 Host controller drivers (rtsx_usb_ms, jmb38x_ms, tfm_ms, r592) already implement their own hardware and bus timeouts, and callers in mspro_block and ms_block already rely on wait_for_completion() unconditionally. Replace wait_for_completion_timeout() with wait_for_completion() in memstick_set_rw_addr() and memstick_alloc_card() so that the card structure is never freed while an asynchronous request is in flight.		
CVE-2026-99040	In the Linux kernel, the following vulnerability has been resolved: KEYS: Fix add_key() race with keyring restriction __key_create_or_update() snapshots keyring->restrict_link before taking the destination keyring's semaphore. keyring_restrict() installs a restriction while holding that semaphore. This allows a writer to observe no restriction, wait for the keyring owner to install a reject-all restriction and return successfully, and then link a key using the stale NULL snapshot. The writer only needs write permission on the destination keyring. Move the restrict_link read after __key_link_lock() and __key_link_begin(). The read and the subsequent restriction check are then serialized with restriction installation by keyring->sem. The race was reproduced on v7.2.8 in 19 executions where restriction installation returned before the link completed. All 19 linked the key despite the reject-all restriction. With this change, 312 executions reached the same ordering and every add_key() call failed with -EPERM. The issue was found by manual concurrency analysis assisted by AI-based analysis and independently verified with a QEMU reproducer and kernel instrumentation.	N/A	More Details
CVE-2026-99041	In the Linux kernel, the following vulnerability has been resolved: ip6_gre: remove obsolete ERSPAN PMTU update A tc tunnel_key action can attach a METADATA_IP_TUNNEL dst to an skb, and mirrored can redirect it to a native ip6erspan device. The legacy PMTU update in ip6erspan_tunnel_xmit() treats that metadata dst as a route. Since it has no output device, dst_dev(dst)->mtu dereferences NULL. On v7.2 with KASAN, an initial UID/GID 65534 process with CapEff 0 used self-created user and network namespaces to trigger: KASAN: null-ptr-deref RIP: ip6erspan_tunnel_xmit+0x10fc/0x2cc0 Kernel panic - not syncing: Fatal exception in interrupt The update is obsolete. Commit fe1a4ca0a2b7 ("ip6_gre: process too big in a better way") changed IPv6 GRE Too Big handling to update the underlay route and removed the equivalent block from __gre6_xmit(). ip6_tnl_xmit() then propagates the underlay PMTU to the overlay route. Remove the stale ERSPAN copy and its now-unused dst variable. This eliminates the metadata-dst dereference and keeps ERSPAN aligned with IPv6 GRE PMTU handling.	N/A	More Details
CVE-2026-4951	The CustomerAdministration portal (KundenManager) of VertiGIS WebOrderSystem allows privileged users to search for employees across all defined companies (user groups). The search term can be passed in the "searchValue" URL parameter, so the search is performed when a logged-in user opens such a URL. JavaScript code placed in this parameter is executed in the victim's browser. An unauthenticated attacker who knows the URL structure can craft a link that contains script code. The code runs if a privileged user who is currently logged into the portal opens the link. Because administrative users can change the passwords of other users, including administrators, without entering the old password, an attacker may be able to take over accounts or lock out other users. Version 7.0.8.6, the latest available at the time of testing, is affected. The vendor did not provide information about further affected versions.	N/A	More Details
CVE-2026-99043	In the Linux kernel, the following vulnerability has been resolved: ice: fix metadata_dst refcount handling on representor teardown ice_eswitch_release_repr() uses metadata_dst_free() to release the representor's metadata_dst. metadata_dst_free() directly frees the underlying memory without checking the dst_entry refcount. When ice_eswitch_port_start_xmit() processes a packet, it takes a reference via dst_hold() and attaches the metadata_dst to the skb. If the representor is torn down while packets are still queued on the lower device (e.g. in a qdisc), the metadata_dst is freed while references are still held. Use dst_release() instead, which correctly decrements the refcount and only frees the object when all references are dropped. The dst subsystem already handles metadata_dst cleanup in dst_destroy() when DST_METADATA is set. Other drivers sharing this pattern (nfp, airoha, bnxt) already use dst_release() for their metadata_dst lifecycle.	N/A	More Details
CVE-2026-99044	In the Linux kernel, the following vulnerability has been resolved: ice: fix use-after-free in dynamic port cleanup ice_dealloc_dynamic_port() uses dyn_port->vsi->idx to erase the dynamic port from pf->dyn_ports. However, it frees the VSI before reading the index for the erase, resulting in a use-after-free. Follow the reverse of the allocation order in ice_alloc_dynamic_port() by erasing the xarray entry before freeing the VSI.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: ibmveth: h_free logical LAN on open-fail after register ibmveth_open() registers the logical LAN, then allocates the RX buffer pools. A pool		

2026-99045	failure jumps to out_free_buffer_pools without h_free_logical_lan(), so PHPY still owns the buffer-list page when it is unmapped. request_irq() failure already issued the hypercall before taking the same label. Move that h_free to the shared post-register unwind so both paths deregister before the buffer list is unmapped.	N/A	More Details
CVE-2026-99072	In the Linux kernel, the following vulnerability has been resolved: ipv4: stop route notification sizing when nexthop group shrinks Commit 680aea08e78c ("net: ipv4: Emit notification when fib hardware flags are changed") added an RCU-only fib_nlmsg_size() call for asynchronous hardware flag notifications. fib_nlmsg_size() checks fib_info_num_path() before each iteration, then fib_info_nhc() independently reloads nh->nh_grp. If RTM_NEWNEXTHOP replaces a group with fewer paths between those loads, fib_info_nhc() returns NULL and fib_nexthop_nlmsg_size() dereferences it. Stop sizing when the indexed path is absent. Nexthop groups are dense, so the current snapshot has no later path.	N/A	More Details
CVE-2026-99073	In the Linux kernel, the following vulnerability has been resolved: ipv4: stop exception dump when nexthop group shrinks Commit 4ce5dc9316de ("inet: switch inet_dump_fib() to RCU protection") allowed IPv4 FIB dumps to run concurrently with nexthop replacement. fib_dump_info_fnhe() uses fib_info_num_path() as the loop bound, then fib_info_nhc() independently reloads nh->nh_grp. If RTM_NEWNEXTHOP replaces a group with fewer paths between those loads, fib_info_nhc() returns NULL and the dump dereferences it through nhc_flags. Stop the walk when the indexed path is absent. Nexthop groups are dense, so the current snapshot has no later path.	N/A	More Details
CVE-2026-99074	In the Linux kernel, the following vulnerability has been resolved: ipv4: stop PMTU walk when nexthop group shrinks Commit 7d3f3b4367f3 ("net: ipv4: Cache pmtu for all packet paths if multipath enabled") made __ip_rt_update_pmtu() update every path. For nexthop objects, fib_info_num_path() and fib_info_nhc() independently load nh->nh_grp. RCU protects each group's lifetime but does not make the loads observe the same group. If replacement shrinks the group after the loop accepts an index, fib_info_nhc() returns NULL and update_or_create_fnhe() dereferences it. A deterministic probe only widened the existing window between the real operations. On v7.2, an RTM_NEWNEXTHOP replacement published a one-member group after the PMTU reader accepted index 1 from a two-member group, producing: KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] RIP: update_or_create_fnhe+0x45/0x15b0 Stop when the indexed path is absent. Groups are dense, so no later index exists in that snapshot. The same real-writer window completed 114,423 PMTU iterations without an oops with this check. A reproducer is available on request. Replacement requires CAP_NET_ADMIN in the network namespace. Where unprivileged user namespaces are permitted, a local user can obtain it in a new user and network namespace.	N/A	More Details
CVE-2026-19678	A server-side request forgery (SSRF) vulnerability in the unauthenticated form service in version 11.0.14 in onPhase Forms allows an unauthenticated remote attacker to access sensitive internal Apache Solr data via an attacker-controlled URL supplied to a publicly accessible form.	N/A	More Details
CVE-2026-115481	baserCMS is a website development framework. Starting in version 5.0.0-beta1 and prior to versions 5.3.2 and 5.4.2, `PermissionGroupsService::build()` reads a plugin's `config/permission.php` to decide what each user group may reach. When that file is absent it writes a wildcard allow rule instead: `url` `/baser/<prefix>/<plugin>/*`, `method` `*`, `auth` true, for every non-admin group. BcCustomContent has no such file in any release. After an administrator installs it, the default `operators` group can store and execute arbitrary PHP under the document root. Versions 5.3.2 and 5.4.2 patch the issue.	N/A	More Details
CVE-2026-53918	odhcpd is the DHCPv6, DHCPv4, and router-advertisement server used by OpenWrt. Prior to version 25.12, odhcpd's DHCPv6 IA handler keeps a per-message `first` lease pointer so multiple IA options in one message reuse the same Reconfigure authentication key. A DHCPv6 REQUEST from an adjacent-network client can make the first IA set `first`, then make a later same-IAID IA fail the on-link check and free that same lease, leaving `first` dangling. A following IA in the same packet allocates a new lease and copies `first->key`, producing a 16-byte heap-use-after-free read. ASan confirms the UAF; native tracing shows the packet shape can deterministically reclaim the freed lease as the destination allocation, producing a zeroed Reconfigure key copy. The validated impact is unauthenticated adjacent-network denial of service and DHCPv6 Reconfigure/ForceRenew key confusion. Arbitrary PC control is not proven for this UAF alone. This issue has been patched in version 25.12.	N/A	More Details
CVE-2026-53920	odhcpd is the DHCPv6, DHCPv4, and router-advertisement server used by OpenWrt. From version 17.01.0 to before version 25.12.3, odhcpd's DHCPv6 IA parser accepts an IA_NA or IA_PD option whose body is shorter than struct dhcpv6_ia_hdr, casts the option header to that full structure, and reads ia->iaid before proving that the IA body contains the IAID field. If an unauthenticated network-adjacent client places a zero-length or otherwise short IA option at the end of odhcpd's 8192-byte receive stack buffer, the IAID read crosses from data_buf into adjacent stack memory. The disclosed bytes are then reflected in the DHCPv6 reply IAID; prior ARM validation showed data_buf[8192:8196] disclosing the saved stack canary, and wider stack scans can leak ASLR-useful stack/process contents. This issue has been patched in version 25.12.3.	N/A	More Details

CVE-2026-55826	The GTM Server Side plugin by stape.io implements server-side tracking on WordPress websites. From version 2.0.0 to before version 2.1.48, the WooCommerce purchase event tracking mechanism allows an unauthenticated attacker to trigger the rendering of another customer's full order data — including personally identifiable information (PII) — into the page's JavaScript dataLayer. This issue has been patched in version 2.1.48.	N/A	More Details
CVE-2026-114583	tarfile.tar_filter could be bypassed using crafted link entries, including symlinks with empty or directory-like names, to redirect later archive members outside the intended extraction directory. This allowed a malicious tar archive to cause tarfile.extractall() with filter="tar" to write files outside the destination directory, subject to the permissions of the extracting process. This is the tarfile.tar_filter counterpart of CVE-2026-7774, which covered tarfile.data_filter.	N/A	More Details
CVE-2026-57208	Entire CLI hooks into Git workflows to capture AI agent sessions. Prior to version 0.7.7, a path traversal vulnerability in Entire CLI allows an attacker with push access to the checkpoints repository to craft malicious checkpoint metadata that causes entire session resume or entire checkpoint rewind to write attacker-controlled transcript data outside of the expected session directory. The issue occurs because checkpoint metadata is fetched from the remote entire/checkpoints/v1 branch and the SessionID field was used to construct filesystem paths without validation in the restore path. A malicious SessionID containing absolute paths or path traversal sequences could cause arbitrary files on the victim's machine to be overwritten. This issue has been patched in version 0.7.7.	N/A	More Details
CVE-2026-115519	Missing security attributes on sensitive cookies in KNIME Business Hub before version 1.19.0 allow an unauthenticated remote attacker to compromise user sessions if other security issues such as a cross-site scripting vulnerability or a man-in-the-middle attack are present. A workaround is not possible.	N/A	More Details
CVE-2026-115517	Unrestricted network access between execution contexts in KNIME Business Hub before version 1.19.0 allows an authenticated remote user to access open network ports on executors in different execution contexts. This only poses a security issue if the IP address of the attacked executor is known and there are open network ports that provided access without further authentication. By default no network ports are open. However, depending on the running workflows, ports may be opened and provide an attack surface. As a workaround, network policies can be created manually.	N/A	More Details
CVE-2026-102786	Joomla Extension - phoca.cz - Authenticated stored cross-site scripting through HTML attribute injection in Phoca Gallery 4.5.0 - 6.1.5	N/A	More Details
CVE-2026-115477	obsidian-tc is a governed, agent-ready Model Context Protocol server for Obsidian vaults. From 1.13.0 until 1.31.4, on a Streamable HTTP deployment with `auth.jwtSecret` configured and a 2026-07-28-era client declaring `elicitation.form` support, a dispatch-gated tool call offered an `inputRequired` confirmation could still complete after the human declined or cancelled, or the elicitation response was omitted or malformed, provided the client echoed the validly bound `requestState` on resubmission. The server trusted an authentic, call-bound state as proof of approval even though that state was minted before the `elicitation/create` leg was asked and proved only that the resubmission belonged to the same call. A client behaving incorrectly or under an attacker's control could therefore execute dispatch-gated destructive vault operations without the human's approval, while still requiring the caller's ordinary tool permissions. Only tools marked `destructive: true` or requiring a scope with a human-in-the-loop floor were reachable through this bypass. The sixteen handler-side conditionally gated tools, including `write_note` overwrite and `move_note`, did not read elicitation state on the modern wire and let neither approval nor decline through that path. Affected deployments using stdio or leaving `auth.jwtSecret` unset instead used the CLI `elicit_token` fallback and did not offer this `requestState` round trip. The fix does not prevent a malicious client from reporting a fabricated approval, because the server cannot verify that the client displayed the form or that a human saw it. This issue is fixed in version 1.31.4.	N/A	More Details
CVE-2026-115478	obsidian-tc is a governed, agent-ready Model Context Protocol server for Obsidian vaults. From 1.0.1 until 1.31.4, `create_entity`, `link_entities`, `rename_entity`, `unlink_entities` and `delete_entity` regenerate memory notes from SQLite at the computed `memory/<type>/<name>.md` path without reliably preserving note ownership or the consistency of a failed lifecycle operation. During shared materialization, a hand-authored note or an orphan left by a different entity whose name sanitized to the same path could have its entire body overwritten on the next write naming that type and name, with non-owned frontmatter preserved. `rename_entity` also seeded the destination path with a direct atomic write without validating ownership and could overwrite a foreign note there. Separately, several tools mutated SQLite before fallible filesystem work: `create_entity` and `link_entities` inserted the row or relation before materialization; `rename_entity` updated the entity, including a retired status, before updating its note; `unlink_entities` removed a relation before rematerializing the source note; and `delete_entity`'s cascade removed the target and its relations before rematerializing neighbouring notes. A later filesystem failure could therefore leave a row or relation changed and the memory graph inconsistent with the notes on disk even when the operation failed. An authorized caller with the	N/A	More Details

	required memory-tool permissions could, without intending to, destroy the body of a vault note outside the entity being written; `delete_entity` additionally requires its deletion permission and destructive confirmation, which this memory flaw does not bypass. Snapshots-before-write through `restore_note` mitigate the overwrite only where snapshots are enabled and the note had previously been written through obsidian-tc. This issue is fixed in version 1.31.4.		
CVE-2026-54668	Kitodo.Production is a workflow management tool for mass digitization and is part of the Kitodo Digital Library Suite. Prior to version 3.9.2, org.kitodo.production.helper.XMLUtils exposes several XML parsing helpers that construct a JAXP parser with the default, non-secure configuration. The methods parseXMLString(String), getNumberOfEADElements(String, String) and checkIfXmllsWellFormed(String) build a DocumentBuilderFactory / XMLInputFactory / SAXParserFactory via newInstance() and parse the attacker-supplied string without disabling DOCTYPE declarations or external entity resolution. These helpers sit directly on the EAD-import path: an authenticated user who uploads an EAD/XML file through the "Create process" file-upload dialog has the uploaded bytes parsed by these methods. A crafted document with an external entity therefore yields XML External Entity (XXE) injection: local file disclosure and outbound SSRF requests from the server. This issue has been patched in version 3.9.2.	N/A	More Details
CVE-2026-58497	Uni-CLI gives people and agents one command interface for websites, browsers, desktop applications, local tools, and MCP. Prior to version 0.225.2, Uni-CLI exposed the legacy JSON-RPC-over-HTTP MCP transport on loopback without validating browser Origin headers before routing requests. A malicious web page could send a CORS simple POST request, such as text/plain, to the local /mcp endpoint and deliver a JSON-RPC body to the dispatcher. If the user had started the local MCP HTTP transport, that page could drive tools/call requests against the user's local Uni-CLI server. The Streamable HTTP transport already enforced this browser-to-localhost boundary. The legacy stateless HTTP path did not, so the two HTTP transports had drifted. This issue is about the browser-to-localhost boundary; it does not change Uni-CLI's local-code-execution trust model. This issue has been patched in version 0.225.2.	N/A	More Details
CVE-2025-45528	An issue in Nitro Software, Inc Nitro PDF Pro for macOS v.14.7 allows a local attacker to execute arbitrary code via a crafted PDF	N/A	More Details
CVE-2026-114519	Missing Authorization (CWE-862) in the user API of Roskus Prospero Flow CRM from 2.0.0 before 5.26.0 allows a user with a basic role to read the accounts of other users of its own company, including accounts with higher privileges, and, in versions from 3.5.0, to delete them, because the routes GET /api/user/{id} and DELETE /api/user/{id} in routes/api.php carry only the auth:api middleware, which checks that a JSON Web Token is valid and nothing more, and neither UserReadController::read() nor UserDeleteController::delete() checks a permission, a policy or a role before acting. Object access is scoped by company: the read path limits the query to the company of the authenticated user, and the delete path compares the company of the target account with the company of the authenticated user. The delete path skips that comparison when the company of the authenticated user is Company::DEFAULT_COMPANY, so a user with a basic role in that company can delete any account of any company of the installation. Deleting an account removes the data held on it, and the last administrative account of a company can be deleted. This gives the attacker the personal data held on the other accounts, including email address, telephone number, postal address and IBAN, and the ability to delete accounts it has no permission to administer.	N/A	More Details
CVE-2026-19677	An authentication bypass vulnerability in the public onPhase/ frevvo registration form version 11.0.14 allows an unauthenticated remote attacker to authenticate as an arbitrary tenant user, obtain TenantAdmin and Designer privileges, and execute arbitrary operating system commands via manipulation of registration parameters containing a valid tenant name and existing user identifier.	N/A	More Details
CVE-2026-54615	Part-DB is an open source inventory management system for electronic components. Prior to version 2.12.1, LogEntryExtraFormatter builds the HTML shown in the "Extra" column of the activity-log table (/log) and in element history tables. Every branch of LogEntryExtraFormatter::getInternalFormat() runs user-controlled strings through htmlspecialchars() before placing them into the HTML it returns — except the CollectionElementDeleted branch, which interpolates getOldName() (and getCollectionName()) into a sprintf() template with no escaping. The returned string is emitted by LogEntryExtraColumn::render() and inserted as raw HTML into the DataTables cell (the bundle does not escape column output; sibling columns call htmlspecialchars() themselves). CollectionElementDeleted.getOldName() is the name of a deleted collection sub-element — a part Parameter, Attachment, PartLot, Orderdetail, etc. — all freely set by any user with edit rights on the owning element. A low-privileged content editor can therefore store an HTML/script payload as the name of such a sub-element, delete it (which writes the log entry), and have that markup parsed as live DOM in the browser of a higher-privileged victim (an administrator/auditor with system.show_logs, or any user viewing that element's history). This is a stored cross-site-scripting / HTML-injection defect caused by a single missed sanitizer in one sibling branch. CSP policies block inline javascript execution under normal conditions. This issue has been patched in version 2.12.1.	N/A	More Details

CVE-2026-99075	In the Linux kernel, the following vulnerability has been resolved: rust_binder: reschedule node refcount update on thread exit When a thread exits via BINDER_THREAD_EXIT, its pending work items are cancelled. If a thread exits while holding a pending node refcount increment (e.g. pushed as deferred work to that thread), the refcount increment was previously dropped because Node::cancel() and NodeWrapper::cancel() were no-ops. Dropping the refcount update leaves the node's delivery state and count state desynchronized, and userspace will not receive the notification, which can cause the node to never be freed from the process's nodes tree when all external references are dropped. Fix this by implementing DeliverToRead::cancel() for Node and NodeWrapper to move the pending refcount update to the process's work queue on thread exit so another thread can deliver it to userspace.	N/A	More Details
CVE-2026-55600	Part-DB is an open source inventory management system for electronic components. Prior to version 2.12.2, Part-DB renders BOM validation information messages using Twig's raw filter. These messages include translated placeholders populated with user-controlled Part names. A user who can edit a Part name can store HTML markup in the Part name. When another user imports a BOM CSV referencing that Part ID, the malicious Part name is rendered as HTML in the victim's BOM validation result page. This results in a cross-user stored HTML injection issue. JavaScript execution appears to be blocked by the current Content Security Policy, but arbitrary HTML links, forms, images, and CSS-based UI elements can still be injected into a trusted Part-DB page. This issue has been patched in version 2.12.2.	N/A	More Details
CVE-2026-99076	In the Linux kernel, the following vulnerability has been resolved: ipv6: udp: Create exceptions before socket matching Currently, when ICMPv6 Packet Too Big and Redirect Message packets are locally delivered and quote a UDP packet, an exception is only created in the IPv6 exception cache if the kernel can match the UDP packet to an existing socket. This behavior allows off-path attackers to conduct a side-channel attack on the exception cache in order to discover the ephemeral port used by a connected UDP socket. Commit 4785305c05b2 ("ipv6: use siphash in rt6_exception_hash()") and commit a00df2caffed ("ipv6: make exception cache less predictable") tried to mitigate such attacks by making it harder for attackers to discover hash collisions in the exception cache and by randomizing the number of exceptions a hash bucket can hold, respectively. Unfortunately, both of the mitigations can be bypassed. Instead, mitigate such attacks by always creating an exception, even before trying to find a matching socket. Do that by calling ip6_update_pmtu() and ip6_redirect(), the helpers used when the quoted packet did not originate from a socket. This means that guesses (right or wrong) from an off-path attacker will always result in an exception being created or updated in the cache that the attacker can observe. Pass the ifindex of the ingress device and the default uid, in a similar fashion to icmpv6_err(). Unlike IPv4, an oif of 0 would not match any nexthop in ip6_redirect_nh_match() and no exception would be created in response to a Redirect Message. Note that this does not allow attackers to create exceptions that they could not create before, as both helpers can already be reached with little to no validation. For example, by sending an ICMPv6 error that quotes an ICMPv6 Echo Reply or one that quotes a UDP source port that matches a wildcard socket. Also note that in the good case (matched socket) the above scheme comes at the cost of an extra route lookup, as the no socket helpers perform their own lookup before the one performed by ip6_sk_update_pmtu() / ip6_sk_redirect(). When the two resolve to different nexthops, it also results in two exceptions being created for the same destination IP. One in the exception cache of the nexthop resolved by the no socket helpers and another in the exception cache of the nexthop used by the socket.	N/A	More Details
CVE-2026-86055	Part-DB is an open source inventory management system for electronic components. Prior to version 2.13.3, Part-DB builds the absolute URL of the password-reset link it emails to a user from the client-supplied HTTP Host header, with no trusted-host validation. An unauthenticated attacker who submits a forgot-password request for a known account while sending a spoofed Host header causes Part-DB to email that account a reset link whose hostname is attacker-controlled, while still embedding the genuine reset token. When the victim opens the link (or an automated inbound mail-security scanner dereferences it), the valid token is disclosed to the attacker, who replays it against the real instance to set a new password and take over the account. Because the seeded admin account can be targeted, this yields full pre-authentication administrative account takeover. This issue has been patched in version 2.13.3.	N/A	More Details
CVE-2026-99080	In the Linux kernel, the following vulnerability has been resolved: fs: Free any excess xarray nodes in clear_inode() For many years we've had a hard to hit leak of xarray nodes. Hugh documented it well in commit 786b31121a2c. Recently people and syzbot have found ways to force it to happen with madvise. Rather than fix the leaks where they happen, just call xa_destroy() which has the side-effect of cycling the i_pages lock.	N/A	More Details
CVE-2026-86019	Part-DB is an open source inventory management system for electronic components. Prior to version 2.13.2, Part-DB allows authenticated API users to bypass the configured maximum attachment file size by uploading attachments through the API using base64-encoded file data. The application enforces the configured attachment size limit for normal browser-based file uploads, but the API base64 upload path does not apply the same size validation. As a result, an authenticated user with API access and	N/A	More Details

	permission to create attachments can upload files larger than the configured maximum attachment size, potentially exhausting disk space and causing a denial of service. This issue has been patched in version 2.13.2.		
CVE-2026-86018	Part-DB is an open source inventory management system for electronic components. Prior to version 2.13.2, Part-DB allows authenticated users to bypass the configured maximum attachment file size when using the server-side "Download external file" attachment feature. When server-side attachment downloads are enabled, Part-DB downloads the remote file and streams it directly to disk without enforcing the configured MAX_ATTACHMENT_FILE_SIZE limit. As a result, an authenticated user who can create or edit attachments can cause Part-DB to store files larger than the configured maximum attachment size, potentially exhausting disk space and causing a denial of service. This issue has been patched in version 2.13.2.	N/A	More Details
CVE-2026-99084	In the Linux kernel, the following vulnerability has been resolved: readdir: take no inode lock on an immutable directory iterate_dir() takes the directory's i_rwsem shared and holds it across ->iterate_shared(). For a directory that never has an entry and is never removed the lock keeps nothing still, it only orders every reader and every writer of that inode behind each other. For the directory of a nullfs instance that matters. The instance of the initial mount namespace is the root of every empty mount namespace and the private instance is the root of every kernel thread, so one inode is shared across users who have nothing else in common. And a reader can hold the lock for as long as it likes: back the getdents() buffer with a mapping of a file on a FUSE mount of your own, let the copy of "." and ".." fault and let the server wait. Queue an exclusive taker behind it, a mkdir() in that directory goes through start_dir() before the read-only mount is reported, and from then on every lookup that misses the dcache in that directory, every create and every mount on it waits until the server answers. One user of an empty mount namespace stalls all the others. Add FOP_IMMUTABLE for the file operations of a directory that never changes and is never removed and let iterate_dir() skip the lock for it. The flag never changes for a file, ->f_pos is protected by f_pos_lock since directories are FMODE_ATOMIC_POS, IS_DEADDIR can't be set on such a directory and neither touch_atime() nor fsnotify take i_rwsem. Set it on the nullfs directory. The placeholder directories of libfs never have an entry either but their owners remove them, so they keep the lock.	N/A	More Details
CVE-2026-99085	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Fix CPU hotplug hang when a dying CPU's tasks sit in the BPF scheduler A CPU going down has to empty its own rq. Its hotplug thread waits in sched_cpu_wait_empty() until nothing else is left, and the only thing that wakes it is balance_push(), which runs from __schedule() on the dying CPU and pushes off the migratable tasks the CPU picks. sched_ext breaks this. A task sitting on a user DSQ or held by the BPF scheduler is still counted on its rq, but an inactive CPU no longer calls ops.dispatch() and can't pull the task back. The dying CPU goes idle with the task still counted, and one of two things happens: - Another CPU consumes the task. The rq empties without a __schedule() on the dying CPU, so the hotplug thread is never woken and cpu_down() hangs holding cpu_hotplug_lock. - The task is affine only to the dying CPU. Nothing can move it, and the offline stalls until the watchdog ejects the BPF scheduler. When the rq goes offline, re-enqueue every task on it that isn't already on the local DSQ. An enqueue on an offline rq lands on the local DSQ, so the dying CPU itself picks the tasks and balance_push() pushes them off, the same as for the other sched classes. From then on no sched_ext path on another CPU can pull a task off the rq. ops.dispatch() currently stops as soon as the CPU goes inactive, and CPU hotplug then waits for an RCU grace period before taking the rq offline. A BPF-held task affine only to the dying CPU and preempted inside an RCU read-side critical section would block that grace period, and the rq would never go offline. Test SCX_RQ_ONLINE directly so that ops.dispatch() keeps running until the rq goes offline. Only the dying CPU's own hotplug thread clears the flag during teardown, and task_can_run_on_remote_rq() still keeps other rqs' tasks off the CPU.	N/A	More Details
CVE-2026-99086	In the Linux kernel, the following vulnerability has been resolved: dm-integrity: validate the superblock on resume dm_integrity_resume() re-reads the superblock from the device so that it picks up the flags and the recalculate position. It performs no validation on the result, while the constructor validates the superblock it reads and sizes all the in-memory structures according to it. The user may modify the superblock on the underlying device while the dm-integrity device is suspended, so that the two no longer agree. In particular, access_journal_data() shifts the journal entry index by ic->sb->log2_sectors_per_block, while ic->journal_section_sectors and the journal page list were computed with the value that was present at constructor time. Increasing log2_sectors_per_block makes the index run past the end of the journal, and dm-integrity then writes 512 bytes through lowmem_page_address(NULL). Snapshot the validated superblock in the constructor and refuse to resume if any of the fields that describe the on-disk geometry changed.	N/A	More Details
CVE-2026-99087	In the Linux kernel, the following vulnerability has been resolved: dm-integrity: fix NULL pointer dereference when the 'R' flag is used If the dm-integrity device has the SB_FLAG_DIRTY_BITMAP flag set and the user activates the device in the 'R' mode, a crash in dm_integrity_resume happens because the function attempts to read the journal containing the bitmap. This patch makes dm-integrity skip any writes to the device in dm_integrity_resume if the device is activated in the 'R' mode.	N/A	More Details

CVE-2026-55835	Lookyloo is a web interface that allows users to capture a website page and then display a tree of domains that call each other. Prior to version 1.39.4, an unauthenticated user can trigger Lookyloo's "report a capture" email on demand (this is expected). The /json/<uuid>/report API endpoint (CaptureReport) is incorrectly passing the email and comment parameters, overwriting as_admin. When True, as_admin will trigger a search query on a MISP instance (if configured) and add a list of matching URLs in the email sent to the analyst (not the submitter, so there is no risk of leaking information) . It could potentially cause a DoS on the MISP instance. This issue has been patched in version 1.39.4.	N/A	More Details
CVE-2026-114353	Hidden functionality (undocumented backdoor) in the root-privileged 'assistant' daemon in DJI CrystalSky ZS600C firmware v0.3.9.7 through v0.6.6.9 on Android allows a physically-present attacker with USB access to execute arbitrary OS commands as root by sending the hard-coded magic string 'dji26656677' to the USB-gadget serial interface (/dev/ttyGS0), after which each subsequent line is passed to system().	N/A	More Details
CVE-2026-114354	Missing authorization in the native binder service 'android.dji.server' (transaction code 1018) in DJI CrystalSky ZS600C firmware v0.3.9.7 through v0.6.6.9 on Android allows a local application with no permissions to execute arbitrary OS commands as root, because the handler passes a caller-supplied string to system() in the root-privileged 'assistant' daemon without any caller-identity or permission check.	N/A	More Details
CVE-2026-114355	OS command injection in the root-privileged 'hardware_monitor' helper in DJI CrystalSky ZS600C firmware v0.3.9.7 through v0.6.6.9 on Android allows a local unprivileged process to execute arbitrary commands as root by writing a crafted line to the Android log (logcat), which the helper reads and passes to a shell eval.	N/A	More Details
CVE-2026-114356	Missing authorization in the system-privileged 'DJIUpdateService' binder in DJI CrystalSky ZS600C firmware v0.3.9.7 through v0.6.6.9 on Android allows a local application with no permissions to invoke privileged operations, including factory reset with user-data wipe (transaction 153), shutdown (105), clock/timezone change (101/123), and the firmware-apply hand-off (151), because onTransact() performs no caller-identity or permission check.	N/A	More Details
CVE-2026-114357	Missing authentication for a critical function in the Android Debug Bridge daemon (adbd) in DJI CrystalSky ZS600C firmware v0.3.9.7 through v0.6.6.9 allows an attacker to obtain an unauthenticated shell (uid 2000), a foothold for privilege escalation to root. The device ships with ro.adb.secure=0 (no host key authorization) and USB ADB enabled by default; ADB over TCP is not enabled by default but, once enabled (e.g. via the shipped /system/bin/adbd_tcp.sh or 'setprop persist.adb.tcp.port 5555'), the same unauthenticated shell is reachable over the network (CVSS 4.0 base 9.3).	N/A	More Details
CVE-2026-106676	qiujie mall-server 0.0.1-SNAPSHOT contains an integer overflow in CartServiceImpl that allows an authenticated user to wrap a cart quantity to a negative value. The order creation flow does not validate the resulting quantity and passes it to the stock update query, increasing inventory and creating negative-value orders.	N/A	More Details