SingCERT
Singapore Cyber Emergency Response Team

Security Bulletin 29 July 2026

Generated on 29 July 2026

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

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

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

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

CRITICAL VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2026-66012	SiYuan before v3.7.2 contains a missing authorization vulnerability in the POST /mcp kernel endpoint, which is gated only by a general auth check (model.CheckAuth) with no admin-role or read-only enforcement. This exposes 31 MCP tools, including a file tool with list/read/write/delete/rename/copy actions across the entire workspace. When the Publish server is enabled in anonymous mode (Conf.Publish.Enable=true and Conf.Publish.Auth.Enable=false), the Publish reverse proxy attaches an anonymous RoleReader JWT to proxied requests, allowing a remote unauthenticated attacker to reach /mcp. The attacker can read conf/conf.json to extract accessAuthCode, api.token, and cookieKey in plaintext, write arbitrary files in the workspace, and plant a plugin into data/plugins/ that executes with nodeIntegration:true and no contextIsolation on the next desktop launch, leading to administrator takeover.	10.0	More Details
CVE-2026-60366	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Platform Security for Java. While the vulnerability is in Oracle Platform Security for Java, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 10.0 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H).	10.0	More Details
CVE-2026-58630	Improper access control in Azure App Service allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-64812	In JetBrains IntelliJ IDEA before 2026.2 unauthorized input injection was possible in a Remote Development session	10.0	More Details
CVE-2026-64813	In JetBrains IntelliJ IDEA before 2026.2 unauthorized settings modification was possible in a Remote Development session	10.0	More Details
CVE-2026-57106	Server-side request forgery (ssrf) in Data Quality allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-56163	Missing authentication for critical function in Microsoft Azure Kubernetes Service allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-62825	Improper authentication in Azure Key Vault allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
CVE-2026-58275	Missing authorization in Azure DNS allows an unauthorized attacker to elevate privileges over a network.	10.0	More Details
	A Deserialization of Untrusted Data vulnerability affecting Station Launcher App in 3DEXPERIENCE platform		

CVE-2026-11756	from Release 3DEXPERIENCE R2023x through Release 3DEXPERIENCE R2026x could lead to an unauthenticated remote code execution.	10.0	More Details
CVE-2026-56191	Improper authentication in Microsoft Exchange Online allows an unauthorized attacker to perform tampering over a network.	10.0	More Details
CVE-2026-59555	Unauthenticated Arbitrary File Deletion in Participants Database <= 2.7.8.3 versions.	10.0	More Details
CVE-2026-42933	Pronetiqs IntraVUE versions 3.2.1a14 and prior have an unintended proxy or intermediary vulnerability which could allow an attacker to use an active proxy, which would bypass OT segmentation.	10.0	More Details
CVE-2026-47668	DbGate is cross-platform database manager. In versions 7.1.8 and prior, DbGate's JSON script runner (`POST /runners/start`) allows remote code execution via code injection in the `functionName` parameter of JSON script `assign` commands. The `functionName` value is interpolated directly into dynamically generated JavaScript source code via string concatenation. The generated code is then executed in a forked Node.js child process. Version 7.1.9 contains a patch.	10.0	More Details
CVE-2026-16812	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. This functionality was intended to be for internal use only and is not intended to be remotely accessible. Hosted and Dedicated versions of VCO have already been patched in advance of this notice going out. This issue was discovered externally and is known to be actively exploited.	10.0	More Details
CVE-2025-71389	Cal.com (calcom/cal.diy) before 5.9.9 is vulnerable to unauthenticated remote code execution because it bundles a version of Next.js whose React Server Components (RSC) request handling deserializes attacker-controlled input. A remote attacker can send a crafted RSC request to the server and cause arbitrary code to be executed during server-side processing, without authentication or user interaction. The flaw derives from the upstream Next.js vulnerability CVE-2025-55182 and is resolved in 5.9.9 by updating the affected dependency.	10.0	More Details
CVE-2026-16498	The terraform-mcp-server before version 1.1.0 is vulnerable to a cross-tenant credential reuse issue in the streamable-HTTP stateless transport mode that may allow one user's Terraform token to be used to execute tool calls on behalf of subsequent users. This vulnerability, CVE-2026-16498, is fixed in terraform-mcp-server 1.1.0.	10.0	More Details
CVE-2026-6516	Zohocorp ManageEngine ADAudit Plus versions before 8606 are affected by Unauthenticated Remote code execution due to the vulnerable agent API.	10.0	More Details
CVE-2026-47724	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.4, the `/api/v1/*` route surface trusts the bearer token alone for authorisation on most endpoints. The codebase itself admits this at `internal/api/hosts.go:384`: "API trusts the bearer token for authorisation; per-CA ownership is enforced only in the Web layer." The Web UI gates state-changing routes through `loadAccessibleCA` (`internal/web/cas.go`); CA-management endpoints in `internal/api/cas.go` ALSO have proper `canAccessCA` gates. The gap is on the host, network, firewall, mobile-bundle, and most operator endpoints. Combined with the per-operator CA model from ADR 0002, this gives any non-admin operator API key broad cross-tenant access — instant privilege escalation in the worst case. Version 0.3.4 fixes the issue.	9.9	More Details
CVE-2026-54120	Improper input validation in Microsoft Surface allows an authorized attacker to execute code over a network.	9.9	More Details
CVE-2026-47752	Tugtainer is a self-hosted app for automating updates of Docker containers. Versions prior to 1.30.2 are vulnerable to Server-Side Template Injection (SSTI) in the notification template feature. The `title_template` and `body_template` fields are rendered using an unsandboxed `jinja2.Environment`, allowing any authenticated user to execute arbitrary OS commands as root inside the container. Version 1.30.2 fixes the issue.	9.9	More Details
CVE-2026-63732	9router 0.4.59 (fixed in 0.4.60) contains a chain of vulnerabilities: a hardcoded default password (123456) that authenticates any fresh installation, a bypass of the LOCAL_ONLY network gate via a spoofed Host header, and unvalidated arguments passed to child_process.spawn() when registering MCP plugins. A remote, unauthenticated attacker can log in with the default credential, spoof the Host header to reach local-only routes, and register a malicious MCP plugin (e.g. node -e <payload>) to achieve arbitrary code execution on the host operating system when the plugin's SSE endpoint is triggered.	9.9	More Details
CVE-2024-58354	cal.com (calcom repository, later renamed cal.diy) is affected by a repository takeover vulnerability in its GitHub Actions workflows. The workflow pr.yml uses the pull_request_target trigger with the repository's default write permissions and passes them down to check-types.yml. check-types.yml then performs a 'dangerous' checkout of the attacker-submitted pull request code (via the dangerous-git-checkout action) and subsequently executes it (through yarn install and package.json scripts). An attacker can open a pull request whose code runs arbitrary commands with the repository's write-scoped GITHUB_TOKEN, allowing them to push commits, merge or mutate pull requests, add or delete comments, and delete or force-push branches, thereby compromising the repository. The main branch is affected; no patched version is available.	9.9	More Details
CVE-2026-15630	A non-global organization admin in one tenant can bypass tenant boundaries to delete, create, or modify resources in any other tenant by exploiting a mismatch between authorization (based on ?id=) and action (based on request body).	9.9	More Details

CVE-2026-50517	Deserialization of untrusted data in M365 Copilot allows an authorized attacker to execute code over a network.	9.9	More Details
CVE-2026-59543	Subscriber Remote Code Execution (RCE) in Advanced Views <= 3.8.11 versions.	9.9	More Details
CVE-2026-60369	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. While the vulnerability is in Oracle Platform Security for Java, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 9.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H).	9.9	More Details
CVE-2026-48030	Pheditor is a single-file editor and file manager written in PHP. From version 2.0.1 to before version 2.0.4, an OS Command Injection vulnerability in the terminal action handler allows any authenticated user to execute arbitrary OS commands by injecting shell metacharacters into the 'dir' POST parameter, completely bypassing the TERMINAL_COMMANDS whitelist and achieving full Remote Code Execution with web server privileges. This issue has been patched in version 2.0.4.	9.9	More Details
CVE-2026-43750	A buffer overflow was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to execute arbitrary code out of its sandbox or with certain elevated privileges.	9.8	More Details
CVE-2026-63077	In JetBrains TeamCity before 2026.1.3, 2025.11.7 unauthenticated remote code execution was possible via the agent polling protocol	9.8	More Details
CVE-2026-28911	The issue was addressed with improved memory handling. This issue is fixed in macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious app may be able to corrupt memory of a system process.	9.8	More Details
CVE-2026-43805	A race condition was addressed with improved state handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, watchOS 26.6. An app may be able to cause unexpected system termination or write kernel memory.	9.8	More Details
CVE-2026-43807	A buffer overflow was addressed with improved bounds checking. This issue is fixed in iOS 26.5.2 and iPadOS 26.5.2, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.5.2, tvOS 26.6, visionOS 26.6, watchOS 26.6. A malicious accessory may be able to cause unexpected app termination.	9.8	More Details
CVE-2026-55579	Pheditor is a single-file editor and file manager written in PHP. From version 2.0.1 to before version 2.0.6, Pheditor ships with a hardcoded default password admin (SHA-512 hash stored at pheditor.php:11). There is no mechanism to force a password change on first login. Any deployment using the default credentials grants an attacker full access to the file editor, file upload, and terminal features, enabling arbitrary file read/write and remote code execution. This issue has been patched in version 2.0.6.	9.8	More Details
CVE-2026-51303	A use-after-free (UAF) vulnerability was discovered in the core parsing component of SQLite 3.41. The flaw occurs because the program frees an ExprList object via sqlite3ExprListDelete and then subsequently accesses the dangling pointer of the released object. A remote adversary can supply specially crafted SQL queries to trigger this vulnerability during SQL statement parsing. Successful exploitation may result in application crash (denial of service), sensitive memory information leakage, and in some scenarios, arbitrary code execution on the affected host.	9.8	More Details
CVE-2026-43757	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-51302	SQLite 3.41 has a use-after-free vulnerability exists in the expression evaluation logic. The sqlite3ReleaseTempReg function improperly releases temporary register resources, and the subsequent exprComputeOperands function continues to access the already freed register memory. By supplying a malicious SQL statement, a remote attacker can exploit this flaw to cause denial of service, leak sensitive information, or potentially execute arbitrary code on the affected system.	9.8	More Details
CVE-2026-65879	Joomla Extension - joomshaper.com - Unauthenticated mail relay via a hardcoded, product-wide secret in SP Page Builder < 6.7.1 - A hardcoded secret allowed attackers to forge the mail from address of forms.	9.8	More Details
CVE-2026-61511	vBulletin 5.x through 5.7.5 and 6.x through 6.2.1 contains an eval injection vulnerability in the vB5_Template_Runtime::runMaths() method within the template runtime that allows unauthenticated remote attackers to execute arbitrary PHP code by supplying crafted input through the pagenav[pagenumber] parameter. Attackers can exploit the insufficiently restrictive regex filter by using phpfuck-style encoding with permitted characters to inject and execute arbitrary PHP code via the unauthenticated ajax/render template route without any authentication.	9.8	More Details
CVE-2026-55971	Heap-based Buffer Overflow vulnerability in Apache Thrift C++ bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	9.8	More Details
CVE-2026-	The Realtyna Organic IDX plugin + WPL Real Estate WordPress plugin before 5.3.0 does not validate the type of uploaded files, and its file upload functionality is gated only by an API that is enabled by default and	9.8	More Details

13714	authenticated with hardcoded credentials shipped identically across all installations. This makes it possible for unauthenticated attackers to upload arbitrary PHP files and achieve remote code execution.		
CVE-2026-12394	The MemberGlut WordPress plugin before 1.1.5 does not validate the role chosen during front-end registration, allowing unauthenticated users to register an account with an arbitrary role, including administrator, leading to full site compromise.	9.8	More Details
CVE-2026-28928	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43803	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. A remote attacker may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43802	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43799	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-28982	A race condition was addressed with improved locking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote user may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-39873	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. Connecting to a malicious SMB server may lead to unexpected system termination.	9.8	More Details
CVE-2026-43682	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote user may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-43694	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination or write kernel memory.	9.8	More Details
CVE-2026-43793	An issue existed in the handling of environment variables. This issue was addressed with improved validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43710	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An attacker may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-43779	A logic issue was addressed with improved restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to intercept network connections intended for another process.	9.8	More Details
CVE-2026-43778	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-43730	A permissions issue was addressed with additional restrictions. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to fingerprint the user.	9.8	More Details
CVE-2026-43773	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. Mounting a maliciously crafted disk image may cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-43748	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43769	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43764	An integer overflow was addressed with improved input validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43809	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43822	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6,	9.8	More Details

	watchOS 26.6. An app may be able to cause unexpected system termination.		
CVE-2026-43810	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. A remote user may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-15014	The SMS Alert – SMS & OTP for WooCommerce, Order Notifications & Abandoned Cart Recovery plugin for WordPress is vulnerable to Authentication Bypass leading to Account Takeover in all versions up to, and including, 3.9.7 via the `billing_phone` parameter. This is due to the `processRegistration()` function using a phone-unbound `$_SESSION['sa_mobile_verified']` boolean flag as the sole gate before issuing an authentication cookie — the flag is set to `true` after any successful OTP validation without being bound to the specific phone number that was verified. This makes it possible for unauthenticated attackers to complete OTP verification for a phone number they control, then resubmit the registration request with a victim's `billing_phone` value to have `wp_set_auth_cookie()` called for the resolved victim account, enabling full authentication as any existing WordPress user whose registered phone number is known or guessable, including administrators.	9.8	More Details
CVE-2026-64770	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may be able to cause unexpected application termination or heap corruption.	9.8	More Details
CVE-2026-64771	A buffer overflow was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may be able to cause unexpected application termination or heap corruption.	9.8	More Details
CVE-2026-64772	An out-of-bounds write issue was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may be able to cause unexpected application termination or heap corruption.	9.8	More Details
CVE-2026-64774	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may be able to cause unexpected application termination or heap corruption.	9.8	More Details
CVE-2026-64775	A memory initialization issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2021-32084	An issue was discovered in Quest KACE Systems Deployment Appliance (SMA) 11.0.273. If a customer restricts access to the web console by IP address or subnets, the API endpoints are not restricted. If credentials/API keys are known to an attacker, the appliance can still be accessed via the API, leading to a potential compromise of the entire environment that is configured for KACE.	9.8	More Details
CVE-2021-32086	An issue was discovered in Quest KACE Systems Deployment Appliance (SMA) 11.0.273. It uses a hardcoded symmetric encryption key to encrypt secrets in the MySQL databases. (This key is not unique for each installation.) An attacker that gains access to the MySQL server or a backup files can decrypt the secrets. Often, the decrypted secrets can be used to escalate privileges within KACE, or gain privileged access to unrelated systems or services.	9.8	More Details
CVE-2021-32088	An issue was discovered in Quest KACE Systems Deployment Appliance (SMA) 11.0.273. Certain API endpoints contain a rate-limiting feature to minimize a brute-force attack. This protection can be bypassed by removing the kboxid cookie.	9.8	More Details
CVE-2026-14545	The TrueBooker WordPress plugin before 1.2.4 does not validate account ownership when resetting a user's password through one of its front-end account handlers, allowing unauthenticated attackers to set an arbitrary password on any account, including an administrator, and take over the site.	9.8	More Details
CVE-2026-16462	In PROCON-WEB SCADA the endpoint 'GetGridData' is not properly sanitized. This allows a remote unauthenticated attacker to execute arbitrary SQL commands.	9.8	More Details
CVE-2026-64767	A buffer overflow was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote attacker may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-51252	schreibfaul1 ESP32-audioI2S 3.4.5 has a buffer overflow vulnerability in the MP3Decoder::UnpackSFMPEG1 function due to missing input validation on attacker-controlled MP3 metadata.	9.8	More Details
CVE-2026-51259	Unchecked unsigned integer overflow in buffer size calculation in schreibfaul1 ESP32-audioI2S 3.4.5 leads to undersized PSRAM buffer allocation. Subsequent normal audio buffer read and write operations cause heap out-of-bounds access, memory corruption, denial of service, and potential code execution.	9.8	More Details
CVE-2026-66713	Deserialization of Untrusted Data (CWE-502) in the Tribes-based clustering component in Apache Software Foundation Apache Axis2/Java through 2.0.0 on Apache Tomcat (only when Tribes clustering is enabled, which is off by default) allows an unauthenticated remote attacker with network access to the clustering port to execute arbitrary code via a crafted serialized Java object delivered to the cluster channel and deserialized in org.apache.axis2.clustering.tribes.Axis2ChannelListener#messageReceived. Users are recommended to upgrade to version 2.0.1, which fixes this issue by removing the clustering feature entirely.	9.8	More Details

CVE-2026-51263	schreibfaul1 ESP32-audioI2S 3.4.5 is vulnerable to Buffer Overflow. The Audio::openai_speech function in the Audio library manually constructs JSON request bodies and HTTP request headers by directly concatenating externally controllable input and instructions strings without effective length restriction and boundary validation. An unauthenticated remote attacker can send oversized malicious string data to trigger a heap buffer overflow during string splicing, resulting in memory corruption.	9.8	More Details
CVE-2026-51266	schreibfaul1 ESP32-audioI2S 3.4.5 has a heap-based buffer overflow vulnerability in the HTTP request header construction logic. The application dynamically splices attacker-controlled host name, path, query string, and multiple HTTP header fields into a fixed ps_ptr heap buffer without proper size limitation and boundary validation. Remote attackers can use an oversized crafted network request parameter to trigger out-of-bounds heap write, leading to arbitrary code execution.	9.8	More Details
CVE-2026-51267	schreibfaul1 ESP32-audioI2S 3.4.5 has a heap-based buffer overflow vulnerability in the URL path concatenation and encoding module. The application splices untrusted extension path and attacker-controlled query string into a path buffer, then invokes urlencode without validating the final string length. Remote attackers can construct an oversized malicious URL path and query string to trigger out-of-bounds heap write, resulting in arbitrary code execution, information disclosure, service crash, or privilege escalation.	9.8	More Details
CVE-2026-14446	IBM WebSphere Application Server 9.0, and 8.5 is vulnerable to broken access control/privilege escalation in the administrative console.	9.8	More Details
CVE-2026-14512	IBM WebSphere Application Server 9.0, and 8.5 traditional is vulnerable to pre-authentication unsafe deserialization which could allow a remote attacker to bypass authentication or execute arbitrary code.	9.8	More Details
CVE-2026-54658	Hypequery is a TypeScript semantic layer for ClickHouse. Prior to 2.0.2, escapeValue() in packages/clickhouse/src/core/utils.ts did not escape backslashes before single quotes during parameter substitution, allowing attacker controlled query parameters with a trailing backslash to escape the closing quote and inject arbitrary SQL. This issue is fixed in version 2.0.2.	9.8	More Details
CVE-2026-64769	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may be able to cause unexpected application termination or heap corruption.	9.8	More Details
CVE-2026-64762	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-43812	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64702	An access issue was addressed with additional sandbox restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to break out of its sandbox.	9.8	More Details
CVE-2026-43814	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64523	In the Linux kernel, the following vulnerability has been resolved: net/handshake: Take a long-lived file reference at submit handshake_nl_accept_doit() needs the file pointer backing req->hr_sk->sk_socket to survive the window between handshake_req_next() and the subsequent FD_PREPARE() and get_file(). The submit-side sock_hold() does not provide that. sk_refcnt keeps struct sock alive, but struct socket is owned by sock->file: when the consumer fputs the last file reference, sock_release() tears the socket down regardless of any sock_hold. Add an hr_file pointer to struct handshake_req and acquire an explicit reference on sock->file during handshake_req_submit(). handshake_complete() and handshake_req_cancel() release the reference on the completion-bit-winning path. The submit error path must also release the file reference, but after rhashtable insertion a concurrent handshake_req_cancel() can discover the request and race the error path. Gate the error-path cleanup -- sk_destruct restoration, fput, and request destruction -- with test_and_set_bit(HANDSHAKE_F_REQ_COMPLETED), the same serialization handshake_complete() and handshake_req_cancel() already use. When cancel has already claimed ownership, the submit error path returns without touching the request; socket teardown handles final destruction. The accept-side dereferences are not yet retargeted; that change comes in the next patch.	9.8	More Details
CVE-2026-64691	A buffer overflow was addressed with improved size validation. This issue is fixed in macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64694	An integer overflow was addressed with improved input validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64695	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote user may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-64696	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote user may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details

CVE-2026-64697	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination or corrupt kernel memory.	9.8	More Details
CVE-2026-64698	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination or read kernel memory.	9.8	More Details
CVE-2026-64700	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64703	A use after free issue was addressed with improved memory management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause a denial-of-service.	9.8	More Details
CVE-2026-64751	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination or write kernel memory.	9.8	More Details
CVE-2026-64704	A type confusion issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64720	A race condition was addressed with improved state handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64726	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An attacker in physical proximity may be able to corrupt process memory.	9.8	More Details
CVE-2026-64727	A type confusion issue was addressed with improved memory handling. This issue is fixed in macOS Tahoe 26.6, tvOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64729	A use after free issue was addressed with improved memory management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	9.8	More Details
CVE-2026-64731	A path handling issue was addressed with improved validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Tahoe 26.6. A malicious app may be able to break out of its sandbox.	9.8	More Details
CVE-2026-64733	This issue was addressed with improved data protection. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to fingerprint the user.	9.8	More Details
CVE-2026-64738	A permissions issue was addressed with additional restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious app may be able to break out of its sandbox.	9.8	More Details
CVE-2026-64746	An authorization issue was addressed with improved validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, visionOS 26.6, watchOS 26.6. An app may be able to add contacts without user authorization.	9.8	More Details
CVE-2026-64530	In the Linux kernel, the following vulnerability has been resolved: net/sched: cls_api: Handle TC_ACT_CONSUMED in tcf_qevent_handle tcf_classify() can return TC_ACT_CONSUMED while the skb is held by the defragmentation engine (e.g. act_ct on out-of-order fragments). When that happens the skb is no longer owned by the caller and must not be touched again. tcf_qevent_handle() did not handle TC_ACT_CONSUMED: it fell through the switch and returned the skb to the caller as if classification had passed. The only qdisc that wires up qevents today is RED, via three call sites (qe_mark on RED_PROB_MARK/HARD_MARK, qe_early_drop on congestion_drop) red_enqueue() was continuing to operate on an skb it no longer owns in this case -- enqueueing it, dropping it, or updating statistics. Resulting in a UAF. tc qdisc add dev eth0 root handle 1: red ... qevent early_drop block 10 tc filter add block 10 ... action ct (with ct defrag enabled and traffic that produces out-of-order fragments, e.g. a fragmented UDP stream) Handle TC_ACT_CONSUMED in tcf_qevent_handle() the same way the ingress and egress fast paths do: treat it as stolen and return NULL without touching the skb. Unlike the TC_ACT_STOLEN case, the skb must not be dropped/freed here, as it is no longer owned by us.	9.8	More Details
CVE-2026-65590	n8n before 2.29.8 and 2.30.x before 2.30.1 does not enforce shell sandbox restrictions on Linux and Windows in the @n8n/computer-use package (sandboxing was applied only on macOS). Shell commands executed by the tool run without any filesystem or network restrictions, allowing unrestricted access to the host filesystem and network from within the computer-use agent process. This issue only affects deployments where the @n8n/computer-use package is explicitly installed and running; standard n8n installations are not affected.	9.8	More Details
CVE-2026-16280	An integer overflow when calculating physical offsets for sparse PMRs may result in 32-bit truncation of address computations for PMRs larger than 4 GB. This can lead to incorrect GPU MMU mappings and may allow a non-privileged user to trigger access to unintended physical memory, resulting in memory corruption or information disclosure.	9.8	More Details
CVE-2026-64384	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix change notify replay double-free A response-bearing attempt can return a replayable error and free its response buffer. If SMB2_notify_init() fails before the next send, cleanup retains the previous buffer type and frees that response again. Reset	9.8	More Details

	response bookkeeping before each attempt to prevent the stale free.		
CVE-2026-64355	In the Linux kernel, the following vulnerability has been resolved: bpf: Reject fragmented frames in devmap Devmap broadcast redirects clone the packet for all but the last destination. For native XDP, that clone path copies only the linear xdp_frame data, while fragmented frames keep skb_shared_info in tailroom outside the linear area. Cloning such a frame leaves XDP_FLAGS_HAS_FRAGS set but without valid frag metadata, and the later free path can interpret uninitialized tail data as skb_shared_info, leading to an out-of-bounds access during frame return. Reject fragmented native XDP frames in dev_map_enqueue_clone(). Add the same restriction to the generic XDP clone path in dev_map_redirect_clone(). Generic XDP represents fragmented packets as nonlinear skbs, and rejecting them here keeps clone-based broadcast support aligned between native and generic XDP.	9.8	More Details
CVE-2026-64874	Joomla Extension - regularlabs.com - CDN Credential leakage Cache Cleaner Pro extension - CDN credentials were exposed in administrator request URLs.	9.8	More Details
CVE-2026-64303	In the Linux kernel, the following vulnerability has been resolved: spi: fsl-lpspi: terminate the RX channel on TX prepare failure path When dmaengine_prep_slave_sg() fails for the TX channel, the error path terminates the TX DMA channel but leaves the RX channel running. Since the RX channel was already submitted and issued prior to preparing the TX descriptor, returning -EINVAL causes the SPI core to unmap the DMA buffers while the RX DMA engine continues writing to them, leading to potential memory corruption or use-after-free. Terminate the RX channel before returning on the TX prepare failure path.	9.8	More Details
CVE-2026-64796	Joomla Extension - regularlabs.com - various code injection vectors in Sourcerer extension - Free did not require both the article creator and last modifier to be Super Users before executing article PHP. Pro did not consistently enforce configured CSS, JavaScript and PHP permissions across tags, attributes, files and both article owners. PHP include attributes could also escape the configured include folder, and executable script/style variants could bypass detection.	9.8	More Details
CVE-2026-64268	In the Linux kernel, the following vulnerability has been resolved: RDMA/siw: bound Read Response placement to the RREAD length In drivers/infiniband/sw/siw/siw_qp_rx.c, siw_proc_rresp() places each inbound Read Response DDP segment at sge->laddr + wqe->processed and then accumulates wqe->processed, but it never checks the running total against the sink buffer length on continuation segments. siw_check_sge() resolves and validates the sink memory only on the first fragment (the if (!*mem) branch), and siw_rresp_check_ntoh() compares the cumulative length against wqe->bytes only on the final segment (the !frx->more_ddp_segs guard). A connected siw peer that answers an outstanding RREAD with Read Response segments that keep the DDP Last flag clear, carrying more total payload than the RREAD requested, drives wqe->processed past the validated sink buffer; the next siw_rx_data() call writes out of bounds at sge->laddr + wqe->processed. siw runs iWARP over ordinary routable TCP, so the peer is the remote end of an established RDMA connection and needs no local privilege. Bound every segment before placement, exactly as siw_proc_send() and siw_proc_write() already do for their tagged and untagged paths, and terminate the connection with a base-or-bounds DDP error when the Read Response would overrun the sink buffer. This is the second receive-path length fix for this file. A separate change rejects an MPA FPDU length that underflows the per-fragment remainder in the header decode; that guard does not cover this case, because here each individual segment length is self-consistent and only the accumulated placement offset overruns the buffer.	9.8	More Details
CVE-2026-65687	Bold Reports Standalone Report Designer before 14.1.12 contains a missing filepath validation vulnerability in its SVG processing feature that allows unauthenticated attackers to read arbitrary files from the server filesystem by supplying a crafted request. Attackers can exploit this path traversal weakness to disclose sensitive server files, including authentication credentials, enabling full unauthorized access to the application. The vulnerability is specific to the DataHub module, which was introduced in Bold Reports 6.3. Therefore, versions prior to 6.3 are not affected.	9.8	More Details
CVE-2026-16766	Catalyst::View::Wkhtmltopdf versions before 0.6.1 for Perl allow shell command injection (RCE) via PDF render options. Options are passed directly to the wkhtmltopdf command without sanitization. Any web application that passes user-controlled options such as the page_size, orientation or margins without validation allows shell command injection. Version 0.6.0 was released with an incomplete fix for this issue. Note that the wkhtmltopdf project is no longer being developed, and users of this package should migrate to alternative solutions.	9.8	More Details
CVE-2026-61884	The web management interface of Tycon Systems TPDIN-Monitor-WEB2 does not perform server-side validation of credentials during the login process. By submitting empty values for both credential fields, an unauthenticated remote attacker can bypass the authentication check and establish a valid administrative session. This grants full access to device controls including power relay management, device reboot, remote access service configuration, and network settings, which could allow an attacker to disrupt connected infrastructure or cause physical damage to equipment.	9.8	More Details
CVE-2026-65688	Bold Reports Standalone Report Designer before 14.1.12 contains a missing filepath validation vulnerability in its font processing feature that allows unauthenticated attackers to read arbitrary files from the server filesystem by supplying a crafted request. Attackers can exploit this path traversal weakness to disclose sensitive server files, including authentication credentials, enabling full unauthorized access to the application. The vulnerability is specific to the DataHub module, which was introduced in Bold Reports 6.3. Therefore, versions prior to 6.3 are not affected.	9.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: block: recompute nr_integrity_segments in blk_insert_cloned_request blk_insert_cloned_request() already recomputes nr_phys_segments against the bottom queue, because "the queue settings related to segment counting may differ from the original queue."		

CVE-2026-64232	The exact same reasoning applies to integrity segments: a stacked driver's underlying queue can have tighter virt_boundary_mask, seg_boundary_mask, or max_segment_size than the top queue, in which case blk_rq_count_integrity_sg() against the bottom queue produces a different count than the cached rq->nr_integrity_segments inherited from the source request by blk_rq_prep_clone(). When the cached count is lower than the bottom queue's actual count, blk_rq_map_integrity_sg() trips BUG_ON(segments > rq->nr_integrity_segments); on dispatch. The same families of stacked setups that motivated the existing nr_phys_segments recompute -- dm-multipath fanning out to nvme-rdma in particular -- can produce this. Mirror the nr_phys_segments handling: when the request carries integrity, recompute nr_integrity_segments against the bottom queue and reject the request if it exceeds the bottom queue's max_integrity_segments. blk_rq_count_integrity_sg() and queue_max_integrity_segments() are both already available via <linux/blk-integrity.h>, which blk-mq.c includes. This closes a latent gap in the stacking contract and brings the integrity-segment accounting in line with the existing phys-segment accounting.	9.8	More Details
CVE-2026-64216	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix potential UAF in netfs_unlock_abandoned_read_pages() netfs_unlock_abandoned_read_pages(rreq) accesses the index of the folios it is wanting to unlock and compares that to rreq->no_unlock_folio so that it doesn't unlock a folio being read for netfs_perform_write() or netfs_write_begin(). However, given that netfs_unlock_abandoned_read_pages() is called _after_ NETFS_RREQ_IN_PROGRESS is cleared, the one folio that it's not allowed to dereference is the one specified by ->no_unlock_folio as ownership immediately reverts to the caller. Fix this by storing the folio pointer instead and using that rather than the index. Also fix netfs_unlock_read_folio() where the same applies.	9.8	More Details
CVE-2026-58586	Image::WebP versions through 0.2 for Perl bundle a vulnerable version of libwebp. Image::WebP does not link to the system libwebp. Instead, it uses a bundled copy of libwebp 0.3.0 (released 2013-03-20). That version has multiple known vulnerabilities, including CVE-2023-4863. Any caller that decodes an untrusted WebP image reaches the bundled decoder. Because the library is compiled into the module, upgrading the system libwebp does not remediate this.	9.8	More Details
CVE-2026-16634	TOML::XS versions before 0.06 for Perl bundle an unsupported and vulnerable version of tomlc99. The tomlc99 library is no longer maintained, and has an uncontrolled recursion vulnerability publicly reported in the issue tracker. Any caller that passes untrusted TOML to from_toml risks a stack overflow from a deeply-nested document. TOML::XS version 0.06 or later uses the successor tomlc17 library.	9.8	More Details
CVE-2026-15704	In Eclipse BaSyx Go Components versions up to and including 1.0.0, ABAC-enabled deployments are vulnerable to an authorization bypass caused by inconsistent trailing-slash handling between the ABAC middleware and the HTTP router. The shared router configuration used Chi's `middleware.StripSlashes`, so a request such as `GET /shells/` was dispatched to the registered `GET /shells` route. However, the ABAC middleware evaluated the original request path including the trailing slash. If ABAC route lookup did not find a matching slash-suffixed route, the request was passed onward and the router then stripped the slash and executed the protected handler without the intended ABAC authorization decision and without the expected ABAC query filters. An unauthenticated or unauthorized network attacker could append a trailing slash to protected API routes to reach handlers that should have been denied by ABAC policy. Depending on the exposed component, HTTP method, and deployed policy, this could allow unauthorized read, create, update, delete, or upload operations. The issue affects ABAC-enabled deployments of services that use the shared router and ABAC middleware, including AAS Repository, Submodel Repository, AAS Registry, Submodel Registry, Concept Description Repository, Discovery, AAS Environment upload, and related services. The issue is fixed in Eclipse BaSyx Go Components v1.0.1.	9.8	More Details
CVE-2026-60367	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-56165	Heap-based buffer overflow in Microsoft Account allows an unauthorized attacker to execute code over a network.	9.8	More Details
CVE-2026-60372	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	9.8	More Details
CVE-2026-52439	An issue in xiandafu beetl 3.20.2 allows a remote attacker to execute arbitrary code via the type.new function and the property reflection mechanism	9.8	More Details
CVE-2026-15981	The SAML Single Sign On – SSO Login plugin for WordPress is vulnerable to Authentication Bypass in all versions up to, and including, 5.4.4. This is due to the mo_saml_validate_signature() function performing a loose boolean check on the raw tri-state integer returned by PHP's openssl_verify(), causing an error return value of -1 to be evaluated as truthy and therefore treated as a successful signature verification. This makes it possible for unauthenticated attackers to log in as any existing WordPress user, including administrators, by submitting a crafted SAMLResponse containing an attacker-controlled NameID and a deliberately malformed signature value that triggers an OpenSSL processing error — bypassing verification entirely and resulting in wp_set_auth_cookie() being called for the targeted account.	9.8	More Details

CVE-2026-63359	The Appriss Insights (Equifax) Victim Information Notification Exchange (VINE) applications allow an unauthenticated attacker to send a specially-crafted request to bypass the login page, access other users' credentials, take over other user accounts, access sensitive PII, and dump other information from the database.	9.8	More Details
CVE-2026-64459	In the Linux kernel, the following vulnerability has been resolved: tcp: restore RCU grace period in tcp_ao_destroy_sock Commit 51e547e8c89c ("tcp: Free TCP-AO/TCP-MD5 info/keys without RCU") removed the call_rcu() callback from tcp_ao_destroy_sock(), arguing that "the destruction of info/keys is delayed until the socket destructor" and therefore "no one can discover it anymore". That argument does not hold for the call site in tcp_connect() (net/ipv4/tcp_output.c:4327-4332). At that point the socket is in TCP_SYN_SENT, has already been inserted into the inet ehash by inet_hash_connect() in tcp_v4_connect(), and is therefore very much discoverable: any softirq running tcp_v4_rcv() on another CPU can take the socket out of the ehash, walk into tcp_inbound_hash(), and load tp->ao_info via implicit RCU before bh_lock_sock_nested() is taken on the destroying CPU. The reader path then enters __tcp_ao_do_lookup() (net/ipv4/tcp_ao.c:208) which re-loads tp->ao_info via rcu_dereference_check(); the re-load can still observe the (about-to-be-freed) pointer because there is no synchronize_rcu() between rcu_assign_pointer(tp->ao_info, NULL) and tcp_ao_info_free() in tcp_ao_destroy_sock(). The captured pointer is then walked at line 223: hlist_for_each_entry_rcu(key, &ao->head, node, ...) The writer's synchronous kfree() is free to complete between the line 218 re-fetch and the line 223 hlist iteration. The slab is reused (or simply LIST_POISON1-stamped if not yet reused) and the iteration walks attacker-controlled or poison memory in softirq context. Reproducer (no debug shim, stock x86_64 v7.1-rc2 SMP+KASAN, QEMU+KVM): an unprivileged uid=1000 process inside CLONE_NEWUSER|CLONE_NEWNET installs TCP_MD5SIG + TCP_AO_ADD_KEY on a TCP socket, sprays forged TCP-AO segments toward its eventual 4-tuple via raw sockets, then calls connect(). The md5-wins reconciliation in tcp_connect() fires tcp_ao_destroy_sock(); the softirq backlog reader on the loopback NAPI path crashes on the freed ao->head.first walk: Oops: general protection fault, probably for non-canonical address 0xfbd59c000000002f KASAN: maybe wild-memory-access in range [0xdead000000000178-0xdead00000000017f] CPU: 0 UID: 1000 PID: 100 Comm: repro_userns RIP: 0010:__tcp_ao_do_lookup+0x107/0x1c0 Call Trace: <IRQ> __tcp_ao_do_lookup+0x107/0x1c0 tcp_ao_inbound_lookup.constprop.0+0x12a/0x200 tcp_inbound_ao_hash+0x5ea/0x1520 tcp_inbound_hash+0x7ce/0x1240 tcp_v4_rcv+0x1e7a/0x3e10 ... Restore the RCU grace period: re-add struct rcu_head to tcp_ao_info and replace the synchronous tcp_ao_info_free() with a call_rcu() callback. Readers that captured tp->ao_info before rcu_assign_pointer NULLed it now see the object remain valid until rcu_read_unlock(). With the patch applied the reproducer runs cleanly for 2000 iterations on the same kernel build.	9.8	More Details
CVE-2026-65700	h2oGPT through 0.2.1 contains a path traversal vulnerability in the OpenAI-compatible files API that allows unauthenticated remote attackers to read, write, and delete arbitrary files accessible to the server process by supplying traversal sequences in the bearer token. The get_user_dir function in openai_server/backend_utils.py uses the bearer token string unsanitized as a path component via os.path.join, and because the default API key is EMPTY authentication is bypassed, enabling attackers to traverse outside the intended user directory through the file content, delete, and upload endpoints to achieve remote code execution by writing to startup hooks or application-loaded files.	9.8	More Details
CVE-2026-14282	The GoDAM – Organize WordPress Media Library & File Manager with Unlimited Folders for Images, Videos & more plugin for WordPress is vulnerable to arbitrary file uploads in versions up to, and including, 1.12.2. This is due to insufficient file type validation in the save_video_file() function hooked into WPForms' public wpforms_process_before_filter, which trusts the attacker-supplied multipart Content-Type header, preserves the original filename via wp_unique_filename(), and moves the raw upload with $wp_filesystem->move() into a web-served directory — bypassing wp_handle_upload()'s MIME/extension allowlist. This makes it possible for unauthenticated attackers to upload arbitrary files on the affected site's server which may make remote code execution possible.	9.8	More Details
CVE-2026-15011	The Customer Support Ticket System & Helpdesk plugin for WordPress is vulnerable to Code Injection via the 'path' parameter in all versions up to, and including, 6.0.5 due to the use of dynamic function invocation on an attacker-controlled value with insufficient validation. This makes it possible for unauthenticated attackers to invoke arbitrary parameterless PHP functions, which can be used to disrupt site functionality or expose sensitive information. The required nonce is publicly emitted via wp_localize_script whenever the plugin's [emd_form] shortcode is rendered on any public-facing page, making the endpoint reachable by unauthenticated visitors without any prior authentication or privilege.	9.8	More Details
CVE-2026-15015	The MountDev AI MCP Connector for WordPress plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 1.6.1. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for unauthenticated attackers to obtain an administrator-bound OAuth Bearer token via a self-registered client, granting full administrator-equivalent access to the plugin's MCP tool surface and all exposed WordPress content, users, and options. This is exploitable by combining the publicly accessible Dynamic Client Registration endpoint, which allows unauthenticated callers to register arbitrary OAuth clients with an attacker-controlled redirect_uri, with the unprotected authorization endpoint to complete the full OAuth flow without any administrator interaction.	9.8	More Details
CVE-2026-65689	Bold Reports Standalone Report Designer before 14.1.12 contains a missing filepath validation vulnerability in its database download feature that allows unauthenticated attackers to read arbitrary files from the server filesystem by supplying a crafted request. Attackers can exploit this path traversal weakness to disclose sensitive server files, including authentication credentials, enabling full unauthorized access to the application. The vulnerability is specific to the DataHub module, which was introduced in Bold Reports 6.3. Therefore, versions prior to 6.3 are not affected.	9.8	More Details

CVE-2026-64383	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix double-free in SMB2_flush() replay SMB2_flush() keeps its response buffer bookkeeping across replay attempts. If a replayable flush response is received and the retry then fails before cifs_send_recv() stores a replacement response, flush_exit will free the stale response pointer a second time. Reinitialize resp_buftype and rsp_iov at the top of the replay loop so cleanup only acts on response state produced by the current attempt. This fixes a double-free without changing replay handling for successful requests.	9.8	More Details
CVE-2026-64385	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix double-free in SMB2_ioctl() replay A response-bearing attempt can return a replayable error and free its response buffer. If SMB2_ioctl_init() fails before the next send, cleanup retains the previous buffer type and frees that response again. Reset response bookkeeping before each attempt to prevent the stale free.	9.8	More Details
CVE-2026-16606	A vulnerability in Fujitsu Software Linux openFT and Fujitsu Software Oracle Solaris openFT before version 12.1D00 allows for unauthenticated remote code execution (pre-auth RCE) on GNU/Linux or Oracle Solaris. The Fsas Technologies PSIRT obtained that intelligence internally and covers the CVE beyond its CNA scope under existing agreement with Fujitsu Germany.	9.8	More Details
CVE-2026-59540	Unauthenticated Privilege Escalation in SMS Alert Order Notifications <= 3.9.6 versions.	9.8	More Details
CVE-2026-64439	In the Linux kernel, the following vulnerability has been resolved: crypto: krb5 - filter out async aead implementations at alloc krb5_aead_encrypt(), krb5_aead_decrypt() in rfc3961_simplified.c and rfc8009_encrypt(), rfc8009_decrypt() in rfc8009_aes2.c set a NULL completion callback and treat any negative return from crypto_aead_{encrypt,decrypt}() as terminal, falling through to kfree_sensitive(buffer). When the encrypt_name resolves to an async AEAD instance the request returns -EINPROGRESS, the buffer is freed while the backend's worker still holds a pointer, and the worker dereferences the freed slab on completion. KASAN report under UML+SLUB with a synthetic async aead backend bound to krb5->encrypt_name: BUG: KASAN: slab-use-after-free in t5_stub_complete+0x7d/0xc7 The helpers were written synchronously, so filter the async instances out at allocation time instead of plumbing crypto_wait_req() through every call site. Reachable via net/rxrpc/rxgk.c, fs/afs/cm_security.c and net/ceph/crypto.c on systems with an async AEAD provider bound to the krb5 enctype name.	9.8	More Details
CVE-2026-59544	Unauthenticated PHP Object Injection in Thrive Quiz Builder <= 10.9.3.0 versions.	9.8	More Details
CVE-2026-64410	In the Linux kernel, the following vulnerability has been resolved: netfilter: flowtable: IPIP tunnel hardware offload is not yet support No driver supports for IPIP tunnels yet, give up early on setting up the hardware offload for this scenario. This patch adds a stub that can be enhanced to add more configuration that are currently not supported. As of now, the offload work is enqueued to the worker, then ignored if the hardware offload configuration is not supported. Check the NF_FLOW_HW flag to know if this entry was already tried once to be offloaded so this is not retried on refresh when unsupported. Move NF_FLOW_HW flag check to nf_flow_offload_add(). If this NF_FLOW_HW flag is unset the _del and _stats variants are never called. This can be updated later on to skip hardware offload work to be queued in case hardware offload does not support it.	9.8	More Details
CVE-2026-2395	Improper neutralization of special elements used in an SQL command ('SQL injection') vulnerability in Xpoda Türkiye Informatics Technology Inc. No Code Platform allows SQL Injection. This issue affects No Code Platform: from 4.3.1.0 through 20260722. NOTE: The vendor was contacted early about this disclosure but did not respond in any way.	9.8	More Details
CVE-2026-64399	In the Linux kernel, the following vulnerability has been resolved: ksmbd: add permission checks for FSCTL_DUPLICATE_EXTENTS_TO_FILE The FSCTL_DUPLICATE_EXTENTS_TO_FILE arm of smb2_ioctl() overwrites the destination file's data via vfs_clone_file_range() with neither the share-level KSMBD_TREE_CONN_FLAG_WRITABLE check nor a per-handle fp->daccess check that the other write-bearing arms carry. A client can overwrite destination data on a read-only share, or from a handle opened with only FILE_WRITE_ATTRIBUTES (which still yields an FMODE_WRITE filp). FILE_WRITE_ATTRIBUTES-only destination handle overwrote the file's data via the clone. Add both checks, matching the FSCTL_SET_SPARSE permission fix; require FILE_WRITE_DATA since this writes data.	9.8	More Details
CVE-2026-64397	In the Linux kernel, the following vulnerability has been resolved: ksmbd: serialize QUERY_DIRECTORY requests per file smb2_query_dir() stores a pointer to its stack-allocated private data in the ksmbd_file readdir_data. Concurrent QUERY_DIRECTORY requests using the same file handle can overwrite this pointer while an iterate_dir() callback is still using it, resulting in a stack use-after-free. Add a per-file mutex and hold it while accessing the shared directory enumeration state. The lock covers scan restart, dot entry state, readdir_data setup and iteration, and response construction. This prevents another request from replacing readdir_data.private before the current request has finished using it and also serializes the shared file position.	9.8	More Details
CVE-2026-61951	Unauthenticated Privilege Escalation in TrueBooker <= 1.2.3 versions.	9.8	More Details
CVE-2026-64386	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix query_info() replay double-free A response-bearing attempt can return a replayable error and free its response buffer. If SMB2_query_info_init() fails before the next send, cleanup retains the previous buffer type and frees that response again. Reset response bookkeeping before each attempt to prevent the stale free.	9.8	More Details
CVE-2026-65431	Joomla Extension - regularlabs.com - Zipslip in GeoIP extension - Geo IP database update archives have been broadly extracted without path validation, leading to unsafe file extractions.	9.8	More Details

CVE-2026-64873	Joomla Extension - regularlabs.com - SSRF in Cache Cleaner Pro extension - Custom query URLs could access internal or reserved network services.	9.8	More Details
CVE-2025-50329	An issue in ConeXware, Inc Power Archiver v.22.00.11 and before allows a remote attacker to escalate privileges and execute arbitrary code via the powerarc.exe.	9.8	More Details
CVE-2026-64387	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix query directory replay double-free A response-bearing attempt can return a replayable error and free its response buffer. If SMB2_query_directory_init() fails before the next send, cleanup retains the previous buffer type and frees that response again. Reset response bookkeeping before each attempt to prevent the stale free.	9.8	More Details
CVE-2026-64391	In the Linux kernel, the following vulnerability has been resolved: ksmbd: use opener credentials for ADS I/O Alternate data streams are stored as xattrs. Unlike regular file I/O, their read and write paths therefore call VFS xattr helpers which recheck inode permissions and LSM policy using the current task credentials. Run ADS I/O with the credentials captured when the SMB handle was opened.	9.8	More Details
CVE-2026-65606	SiYuan before v3.7.2 contains a cross-site scripting vulnerability in the siyuan:// protocol handler. When a siyuan://plugins/<name> link references a name that is not an installed plugin, the application opens a custom tab and inserts the link's icon parameter into the tab header via innerHTML without escaping it (app/src/layout/Tab.ts), allowing injection of an <img onerror=...> element. Because the SiYuan Desktop renderer runs with nodeIntegration:true, the injected JavaScript can access Node's require and call require('child_process').execSync(...), escalating the cross-site scripting into arbitrary operating-system command execution.	9.6	More Details
CVE-2026-16624	Cal.com OSS ships lacks authorization on webhook teamId creation, allowing any authenticated user to create a webhook on any team via unvalidated teamId injection, then steal booking data, including fields like organizer/attendee emails and custom responses, and conditionally video-call passwords, by triggering webhook delivery.	9.6	More Details
CVE-2026-51271	In schreibfaul1 ESP32-audioI2S 3.4.5, a heap-based buffer overflow vulnerability exists in the WAV header parsing function read_WAV_Header(). The function reads untrusted chunk size and bytes-to-skip value directly from malicious WAV files without reasonable range restriction. Abnormally large bts and headerSize values lead to out-of-bounds heap memory read/write during header parsing, which can be exploited to execute arbitrary code, disclose sensitive information, cause denial of service or escalate privileges.	9.6	More Details
CVE-2026-65605	SiYuan before v3.7.2 contains a stored cross-site scripting vulnerability in Attribute View (database) cell rendering. A Template column value is rendered as HTML via text/template without auto-escaping, and EscapeHTML is only applied when HasUnclosedHtmlTag returns true; because balanced self-closing tags such as <img> are skipped by that check, a payload like <img src=x onerror=...> is stored unescaped and later inserted into the page via innerHTML, executing when the database is viewed. Because the desktop renderer runs with nodeIntegration enabled, the injected script can reach require and escalate to arbitrary command execution.	9.6	More Details
CVE-2026-57784	Unauthenticated Cross Site Request Forgery (CSRF) in Ninja Forms File Uploads Extension <= 3.3.26 versions.	9.6	More Details
CVE-2026-65471	Unauthenticated Cross Site Request Forgery (CSRF) in Avada Core <= 5.15.6 versions.	9.6	More Details
CVE-2026-66395	SiYuan desktop before v3.7.2 contains a reflected cross-site scripting vulnerability in the bazaar plugin readme handler that allows attackers to execute arbitrary code by crafting a malicious siyuan:// deep link. Attackers can inject HTML payloads via the plugin name parameter that execute with full Node.js access through insertAdjacentHTML rendering in an insecurely configured Electron renderer.	9.6	More Details
CVE-2026-11841	An attacker may perform unauthenticated read and write operations on sensitive filesystem areas via the AppEngine Fileaccess over HTTP due to improper access restrictions. A critical filesystem directory was unintentionally exposed through the HTTP-based file access feature, allowing access without authentication. This includes device parameter files, enabling an attacker to read and modify application settings, including customer-defined passwords. Additionally, exposure of the custom application directory may allow execution of arbitrary Lua code within the sandboxed AppEngine environment.	9.4	More Details
CVE-2026-51260	Unsafe fixed-size memcpy operation in AudioBuffer::writeSpace() of schreibfaul1 ESP32-audioI2S 3.4.5 allows remote heap buffer overflow. The code copies a full UINT16_MAX bytes without validating destination available space, causing out-of-bounds memory write.	9.4	More Details
CVE-2026-61949	Unauthenticated SQL Injection in Bookly <= 27.7 versions.	9.3	More Details
CVE-2026-61950	Unauthenticated SQL Injection in TrueBooker <= 1.2.3 versions.	9.3	More Details
CVE-2026-61948	Unauthenticated SQL Injection in WPDM – Premium Packages <= 6.2.0 versions.	9.3	More Details
CVE-2026-64740	A parsing issue in the handling of directory paths was addressed with improved path validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6. A malicious app may be able to break out of its sandbox.	9.3	More Details

CVE-2026-59526	Unauthenticated SQL Injection in MapSVG <= 8.14.0 versions.	9.3	More Details
CVE-2026-59525	Unauthenticated SQL Injection in Participants Database <= 2.7.8.3 versions.	9.3	More Details
CVE-2026-59514	Unauthenticated SQL Injection in Buddyboss Platform <= 3.0.5 versions.	9.3	More Details
CVE-2026-62835	Improper authorization in Azure Portal allows an unauthorized attacker to disclose information over a network.	9.3	More Details
CVE-2026-59549	Unauthenticated SQL Injection in rtMedia for WordPress, BuddyPress and bbPress <= 4.7.10 versions.	9.3	More Details
CVE-2026-50252	In NLnet Labs Unbound 1.4.22 up to and including 1.25.1, UDP source port is randomized and intended to serve as a secret value that increases the entropy of DNS transactions. When resolver load balancing policies depend on the source port while their outcome is revealed this secrecy is undermined. The vulnerability arises when the load balancing policy is consistent with respect to the incoming source UDP port and IP address while heavily depending on the incoming source UDP port as a randomization source. When the SO_REUSEPORT configuration option is enabled ('so-reuseport: yes') in Unbound (by default), it meets these conditions, making it vulnerable for DNS cache poisoning attacks. Upon startup, Unbound randomly partitions the available UDP source port space into disjoint subsets of (almost) equal size, assigning each subset to a specific worker thread. When an incoming DNS query is received, the kernel's SO_REUSEPORT load balancing mechanism deterministically assigns the query to a socket associated with a particular thread. All outgoing DNS queries generated during the resolution of that request use source ports selected exclusively from the port subset assigned to the corresponding thread. Since these port subsets are disjoint across threads, the source port observed in a resolver's outgoing query to an authoritative name server serves as a reliable indicator of the worker thread that processed the original client query. A malicious actor can acquire the mapping between incoming UDP source ports (for a given fixed source IP address) and Unbound worker threads and leverage it to conduct DNS cache poisoning attacks by effectively lowering the random port population per thread.	9.3	More Details
CVE-2026-14973	IBM Aspera Desktop App 1.0.5 through 1.0.19 IBM Aspera for desktop can allow files to be written outside of the user's selected download destination.	9.3	More Details
CVE-2026-59527	Unauthenticated SQL Injection in MapSVG <= 8.14.0 versions.	9.3	More Details
CVE-2026-59538	Unauthenticated SQL Injection in GamiPress <= 7.9.7 versions.	9.3	More Details
CVE-2026-59533	Unauthenticated SQL Injection in Relevanssi Light <= 1.2.2 versions.	9.3	More Details
CVE-2026-59550	Unauthenticated SQL Injection in AWP Classifieds <= 4.4.7 versions.	9.3	More Details
CVE-2026-27064	Editor Arbitrary File Upload in Mailster <= 4.1.17 versions.	9.1	More Details
CVE-2026-64798	Joomla Extension - regularlabs.com - Insecure login URL keys in IP login extension - Persistent URL login keys were also generated using a non-cryptographic random generator with insufficient entropy.	9.1	More Details
CVE-2026-40712	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) an Improper Input Validation vulnerability in the REST API. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Elevation of privileges.	9.1	More Details
CVE-2026-62144	An authentication bypass vulnerability in Check Point Security Management and Multi-Domain Security Management allows an unauthenticated remote attacker to execute administrative commands on the Management Server. Successful exploitation may also allow command execution on managed Security Gateways. Exploitation requires network access to the Management Server without firewall protection or a configuration that does not restrict Trusted Clients.	9.1	More Details
CVE-2026-14958	IBM Aspera Faspex 5 5.0.0 through 5.0.15.4 could allow a remote authenticated attacker to execute arbitrary code due to unquoted shell interpolation.	9.1	More Details
CVE-2026-14959	IBM Aspera Faspex 5 5.0.0 through 5.0.15.4 could allow a remote authenticated attacker to execute arbitrary code due to shell command injection.	9.1	More Details
CVE-2026-46738	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) an Improper Input Validation vulnerability in the REST API. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Elevation of privileges.	9.1	More Details
CVE-2026-64793	Joomla Extension - regularlabs.com - Content access and publication bypass in Articles Anywhere and Modules Anywhere extensions - Content tags could use ignore flags or property overrides to render restricted or unpublished articles or modules. A content author could thereby expose content to visitors who lacked the required access.	9.1	More Details

CVE-2026-62325	goshs is a feature-rich single-binary file server for red teamers and developers. From 2.1.3 until 2.1.4, the sftpserver/sftpserver.go password handler used Username != "" && Password != "", so running goshs with -b 'admin:' -sftp and no -fkf left both SFTP authentication handlers unset and allowed unauthenticated file access. This issue is fixed in version 2.1.4.	9.1	More Details
CVE-2026-16232	An authentication bypass vulnerability in the Check Point SmartConsole login process allows an unauthenticated remote attacker to obtain an application login token and use it to authenticate with full administrative privileges. Successful exploitation allows the attacker to modify security policies and security configurations. Remote exploitation requires internet access to the Management Server IP address and a configuration that does not restrict Trusted Clients. Check Point is aware that this vulnerability is being exploited and has affected a very small number of customers.	9.1	More Details
CVE-2026-65701	SoftVC VITS Singing Voice Conversion through commit 730930d contains a path traversal vulnerability in the full-song inference server that allows unauthenticated remote attackers to read and exfiltrate arbitrary files by supplying attacker-controlled filesystem paths through the audio_path field of an unauthenticated POST request to the /wav2wav route. Attackers can pass arbitrary server-side paths verbatim to librosa.load, torchaudio.load, and soundfile.write sinks, causing the server to decode and return file contents via the HTTP response body while also writing attacker-specified .wav files to arbitrary locations on the filesystem.	9.1	More Details
CVE-2026-65455	Administrator Arbitrary File Upload in MapSVG <= 8.14.0 versions.	9.1	More Details
CVE-2026-65461	Administrator Arbitrary File Upload in Really Simple CSV Importer <= 1.3 versions.	9.1	More Details
CVE-2026-64450	In the Linux kernel, the following vulnerability has been resolved: tipc: fix out-of-bounds read in broadcast Gap ACK blocks A broadcast PROTOCOL/STATE_MSG can carry a Gap ACK blocks record in its data area. tipc_get_gap_ack_blks() only verifies that the record's len field is self-consistent with its ugack_cnt/bgack_cnt counts (sz == struct_size(p, gacks, ugack_cnt + bgack_cnt)); it does not check that the record actually fits in the message data area, msg_data_sz(). The unicast caller tipc_link_proto_rcv() bounds it ("if (glen > dlen) break;"), but the broadcast caller tipc_bcast_sync_rcv() discards the returned size, so tipc_link_advance_transmq() copies the record off the receive skb with an attacker-controlled count: this_ga = kmemdup(ga, struct_size(ga, gacks, ga->bgack_cnt), GFP_ATOMIC); A TIPC neighbour that negotiated TIPC_GAP_ACK_BLOCK triggers it with one ordinary broadcast STATE_MSG (msg_bc_ack_invalid() clear), sized so its data area is short, carrying a Gap ACK record with len = 0x400, bgack_cnt = 0xff and ugack_cnt = 0. len then equals struct_size(p, gacks, 255), so the consistency check passes and ga is non-NULL; kmemdup() reads struct_size(ga, gacks, 255) = 1024 bytes out of the much smaller skb: BUG: KASAN: slab-out-of-bounds in kmemdup_noprof+0x48/0x60 Read of size 1024 at addr ffff0000c7030d38 by task poc864/69 Call trace: kmemdup_noprof+0x48/0x60 tipc_link_advance_transmq+0x86c/0xb80 tipc_link_bc_ack_rcv+0x19c/0x1e0 tipc_bcast_sync_rcv+0x1c4/0x2c4 tipc_rcv+0x85c/0x1340 tipc_l2_rcv_msg+0xac/0x104 The buggy address belongs to the object at ffff0000c7030d00 which belongs to the cache skbuff_small_head of size 704 The buggy address is located 56 bytes inside of allocated 704-byte region [ffff0000c7030d00, ffff0000c7030fc0) The copied-out bytes are subsequently consumed as gap/ack values, but the read is already out of bounds at the kmemdup() regardless of how they are used. The unicast STATE path drops such a message: "if (glen > dlen) break;" skips the rest of STATE_MSG handling and the skb is freed. Make the broadcast path drop it too. tipc_bcast_sync_rcv() now bounds the record against msg_data_sz() and, when it does not fit, reports it back through tipc_node_bc_sync_rcv() to tipc_rcv() so the skb is discarded rather than processed. ga is not cleared on this path: ga == NULL already means "legacy peer without Selective ACK", a distinct legitimate state.	9.1	More Details
CVE-2026-13332	The Masteriyo LMS WordPress plugin before 2.3.1 does not correctly verify authorization on an unauthenticated AJAX action used to clear user sessions, allowing unauthenticated attackers to terminate the active sessions (force-logout) of any user on the site, including administrators.	9.1	More Details
CVE-2026-13597	The 微信二维码登陆 WordPress plugin through 1.3 does not properly validate WeChat webhook requests, as its signature check always passes, and it discloses the generated login code in the webhook response. This allows an unauthenticated attacker to forge a login event for any existing username, read the login code, and redeem it through an unauthenticated AJAX action to log in as that user, including an administrator, without a password.	9.1	More Details
CVE-2026-48144	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift c_glib bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	9.1	More Details
CVE-2026-58023	Out-of-bounds Read vulnerability in Apache Thrift c_glib bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	9.1	More Details
CVE-2026-58662	Improper Validation of Specified Quantity in Input, Out-of-bounds Read vulnerability in Apache Thrift C++ bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	9.1	More Details
CVE-2025-50455	SQL injection vulnerability exists in the order_by parameter of the /customers/search endpoint in Alex Tselegidis EasyAppointments <= 1.5.1. The vulnerability arises from unsanitized user input passed to the order_by method of the CodeIgniter Query Builder, enabling attackers to perform time-based queries and schema enumeration. Under certain MySQL configurations, the flaw may lead to remote code execution by writing a PHP shell using INTO OUTFILE.	9.1	More Details

CVE-2026-64393	In the Linux kernel, the following vulnerability has been resolved: ksmbd: run set info with opener credentials SMB2 SET_INFO handlers call path-based VFS helpers after checking the access mask granted to the SMB handle. Those helpers perform their owner, inode permission and LSM checks using the current ksmbd worker credentials. Run the complete SET_INFO dispatch with the credentials captured when the handle was opened. This also removes the separate security information credential setup and keeps all SET_INFO classes under one credential scope. Direct override_creds() is used because it can nest with the request credential overrides already used by rename and link helpers.	9.1	More Details
CVE-2026-51300	A use-after-free vulnerability exists in the expression parsing and memory management logic of SQLite 3.41. After invoking sqlite3ExprDelete to release an expression object, the program still retains the dangling pointer and subsequently accesses member fields of the already freed memory. By constructing malicious SQL queries, a remote attacker can trigger invalid memory access, leading to application crash and sensitive memory information leakage.	9.1	More Details
CVE-2026-64392	In the Linux kernel, the following vulnerability has been resolved: ksmbd: use opener credentials for delete-on-close Delete-on-close can be completed by deferred or durable handle teardown, where no request work is available. Both the base-file unlink and the ADS xattr removal consequently run with the ksmbd worker credentials and can bypass filesystem permission checks. Run both operations with the credentials captured in struct file when the handle was opened. This preserves the authenticated user's fsuid, fsgid, supplementary groups and capability restrictions at final close.	9.1	More Details
CVE-2026-17191	An input validation vulnerability exists in an API component of the orchestrator. An authenticated user can exploit this flaw to manipulate backend queries, which may result in unauthorized access to data beyond their intended privileges and cause the underlying system to initiate unintended outbound network connections. This issue was discovered internally by Arista and the company is not aware of any malicious uses of this issue in customer networks.	9.1	More Details
CVE-2026-17552	Plack::App::Prerender versions before 0.3.0 for Perl can proxy to an arbitrary host via unvalidated REQUEST_URI concatenation in call. When the rewrite base is a plain string, the REQUEST_URI is appended to it, with no check that the path starts with a forward slash ('/'). When the rewrite base does not contain a path (which is the standard given in the SYNOPSIS), an attacker can create a request that changes the hostname. A request target starting with an at-sign ('@') changes the base to a RFC 3986 userinfo component. For example, a rewrite base of "https://example.com" with the submitted request "GET @192.168.1.2/" will send a request to "https://example.com@192.168.1.2/", with the rendered content returned to the attacker. This allows an attacker to access internal or restricted hosts that only the webserver has access to.	9.1	More Details
CVE-2026-64320	In the Linux kernel, the following vulnerability has been resolved: nvmet: fix pre-auth out-of-bounds heap read in Discovery Get Log Page nvmet_execute_disc_get_log_page() validates only the dword alignment of the host-supplied Log Page Offset (lpo). The 64-bit offset is then added to a small kzalloc'd buffer that holds the discovery log page and the result is passed straight to nvmet_copy_to_sgl(), which memcpy()s data_len bytes out to the host with no source-side bound check: u64 offset = nvmet_get_log_page_offset(req->cmd); /* 64-bit host */ size_t data_len = nvmet_get_log_page_len(req->cmd); /* 32-bit host */ ... if (offset & 0x3) { ... } /* only check */ ... alloc_len = sizeof(*hdr) + entry_size * discovery_log_entries(req); buffer = kzalloc(alloc_len, GFP_KERNEL); ... status = nvmet_copy_to_sgl(req, 0, buffer + offset, data_len); The Discovery controller is unauthenticated -- nvmet_host_allowed() returns true unconditionally for the discovery subsystem -- so the call is reachable pre-authentication by any TCP/RDMA/FC peer that can reach the nvmet target. With a discovery log page of ~1 KiB, an attacker requesting up to 4 KiB starting at offset == alloc_len reads the next slab page out and gets its content returned over the fabric (an empirical run on a default nvmet-tcp loopback target leaked 81 canonical kernel pointers in one Get Log Page response). Pointing the offset at unmapped kernel memory faults the in-kernel memcpy and crashes (or panics, on panic_on_oops=1) the target host instead. The attacker-controlled source-side offset pattern "nvmet_copy_to_sgl(req, 0, buffer + ATTACKER_OFFSET, ...)" is unique to nvmet_execute_disc_get_log_page in the entire nvmet codebase: every other Get Log Page handler in admin-cmd.c either ignores lpo (and silently starts every response at offset 0) or tracks a local destination offset with a fixed source pointer. Validate the host-supplied offset against the log page size, cap the copy length to what is actually available, and zero-fill any remainder of the host transfer buffer. The zero-fill matches the existing short-response pattern in nvmet_execute_get_log_changed_ns() (admin-cmd.c) and prevents leaking transport SGL contents when the host asks for more bytes than the log page contains.	9.1	More Details
CVE-2026-64319	In the Linux kernel, the following vulnerability has been resolved: nvmet-auth: validate reply message payload bounds against transfer length nvmet_auth_reply() accesses the variable-length rval[] array using attacker-controlled hl (hash length) and dhvlen (DH value length) fields without verifying they fit within the allocated buffer of tl bytes. A malicious NVMe-oF initiator can craft a DHCHAP_REPLY message with a small transfer length but large hl/dhvlen values, causing out-of-bounds heap reads when the target processes the DH public key (rval + 2*hl) or performs the host response memcmp. With DH authentication configured, the OOB pointer is passed directly to sg_init_one() and read by crypto_kpp_compute_shared_secret(), reaching up to 526 bytes past the buffer. This is exploitable pre-authentication. Add bounds validation ensuring sizeof(*data) + 2*hl + dhvlen <= tl before any access to the variable-length fields. Discovered by Atuin - Automated Vulnerability Discovery Engine.	9.1	More Details
	In the Linux kernel, the following vulnerability has been resolved: RDMA/rtrs-srv: Bound RDMA-Write length to chunk size in rdma_write_sg When the server answers an RTRS READ, rdma_write_sg() builds the source scatter/gather entry for the IB_WR_RDMA_WRITE that returns data to the peer. Its length is taken directly from the wire descriptor: plist->length = le32_to_cpu(id->rd_msg->desc[0].len); rd_msg points into the chunk buffer that the remote peer filled via RDMA-WRITE-WITH-IMM (rtrs_srv_rdma_done() -> process_io_req() -> process_read()), so desc[0].len is attacker-controlled and, before this change, was only rejected when zero. The		

CVE-2026-64269	source address is the fixed chunk start (dma_addr[msg_id]) and the source lkey is the PD-wide local_dma_lkey, which is not tied to the chunk's MR mapping, so the verbs layer does not constrain the transfer length to max_chunk_size. msg_id and off are bounded against queue_depth and max_chunk_size in rtrs_srv_rdma_done(), but desc[0].len is a separate field that was not checked against the chunk size. A peer that advertises desc[0].len larger than max_chunk_size can make the posted RDMA write read past the chunk's mapped region. The resulting behaviour depends on the IOMMU configuration: with no IOMMU or in passthrough mode the read may extend into memory adjacent to the chunk and be returned to the peer, which can disclose host memory; with a translating IOMMU the out-of-range access is expected to fault and abort the connection. In either case the transfer exceeds what the protocol permits and is driven by a remote peer. Reject a descriptor length above max_chunk_size, mirroring the existing off >= max_chunk_size bound in rtrs_srv_rdma_done(). Legitimate clients do not exceed it: the client sets desc[0].len to its MR length, which is capped at the negotiated max_io_size (max_chunk_size - MAX_HDR_SIZE).	9.1	More Details
CVE-2026-64257	In the Linux kernel, the following vulnerability has been resolved: smb: client: reject overlapping data areas in SMB2 responses Commit 53b7c271f06b ("smb: client: restrict implied bcc[0] exemption to responses without data area") restricted the implied bcc[0] length exception to responses without a data area. However, the overlap handling in __smb2_calc_size() clears data_length, which can make an invalid response appear to have no data area and so qualify for the exception. Track data area overlap separately and reject such responses before applying the length compatibility exceptions.	9.1	More Details
CVE-2026-48021	In epa4all, prior to version 2026-05-20, an attacker who can intercept the TLS connection between epa4all and the ePA backend can complete the VAU handshake with attacker-controlled keys and obtain the session encryption keys. All inner HTTP traffic (patient consent decisions, medication data, document operations, authorization tokens, and entitlement queries) becomes readable and modifiable. The attacker can also inject arbitrary requests through the hijacked channel. This issue has been patched in version 2026-05-20.	9.1	More Details
CVE-2026-12877	The Project Management, Bug and Issue Tracking Plugin WordPress plugin before 5.1.0 does not sanitise and escape user supplied input before using it in a SQL query, allowing unauthenticated attackers to perform SQL injection attacks. This is exploitable in the Project Management, Bug and Issue Tracking Plugin WordPress plugin before 5.1.0's standard front-end issue-tracker configuration.	9.1	More Details
CVE-2026-56160	Improper authorization in Azure Red Hat OpenShift (ARO) allows an authorized attacker to elevate privileges over a network.	9.1	More Details
CVE-2026-15617	Logto performs principal lookup without normalizing email and identifier strings, enabling principal collision and unauthorized account access via case- or Unicode-different identities.	9.1	More Details
CVE-2026-15616	Logto does not enforce locally configured MFA during SSO authentication, allowing users to bypass second-factor requirements and grants unauthorized access.	9.1	More Details
CVE-2026-15612	Logto bypasses OIDC nonce validation when the nonce claim is absent from the id_token, enabling replay of authentication tokens and weakening session-binding.	9.1	More Details
CVE-2026-15611	Logto allows unverified email-based SSO account linking, enabling an attacker to register an identity at a permissive IdP using a victim’s email and gain unauthorized access to the victim’s account.	9.1	More Details
CVE-2026-65907	In JetBrains TeamCity before 2026.1.2, 2025.11.6 code execution in Git VCS roots was possible	9.1	More Details
CVE-2026-64863	goshs is a feature-rich single-binary file server for red teamers and developers. Prior to 2.1.4, the httpserver/server.go wdGuard handled WebDAV MOVE as a write-only method and did not enforce --no-delete, allowing WebDAV clients to delete or overwrite files via MOVE with Overwrite: T. This issue is fixed in version 2.1.4.	9.1	More Details
CVE-2026-16723	A remote code execution (RCE) vulnerability exists in fastjson 1.2.68 through 1.2.83. This vulnerability is exploitable under fastjson's stock default configuration — no AutoType enablement required, no classpath gadget required.	9.0	More Details
CVE-2026-14289	The FacturaONE para WooCommerce con VeriFactu WordPress plugin before 5.37 does not authenticate one of its request handlers, whose only protection is derived from a cryptographic key that is empty in the default, unconfigured state, allowing unauthenticated attackers to write an arbitrary file into a web-accessible directory and achieve remote code execution.	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2026-16496	The terraform-mcp-server before version 1.1.0 is vulnerable to an authorization bypass in the streamable-HTTP stateful transport mode that may allow a user who obtains another user's MCP session ID to have their tool calls executed using that user's Terraform credentials. This vulnerability, CVE-2026-16496, is fixed in terraform-mcp-server 1.1.0.	8.9	More Details
CVE-2024-	Cal.com (repository calcom/cal.diy) in versions <= 4.7.15 is vulnerable to cross-site scripting (XSS) on the publicly accessible single booking view (e.g., /booking/<id>). Booking question (form field) labels are rendered via React's dangerouslySetInnerHTML without proper input sanitization or CSP, so an attacker who can create an event type with	8.9	More

58353	a malicious booking question label can inject arbitrary HTML/JavaScript that executes when a victim visits the booking view URL. Self-hosted instances with open registration are particularly at risk. The issue is fixed in version 4.7.16.		Details
CVE-2024-58355	Cal.com (calcom/cal.diy) versions through 4.7.15 contain a stored cross-site scripting vulnerability. The single booking view (e.g., https://app.cal.com/booking/<id>) renders booking-question field labels via React's dangerouslySetInnerHTML without sanitizing or escaping user input. An attacker who can create an event type with a malicious booking-question label can inject arbitrary HTML/JavaScript that executes when a victim opens the crafted booking URL. The issue is fixed in v4.7.16.	8.9	More Details
CVE-2026-64739	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An attacker may be able to cause unexpected app termination.	8.8	More Details
CVE-2026-56747	Improper control of generation of code in the JSON Pointer-to-accessor compiler in Cribl Stream before 4.18.2 allows a remote authenticated attacker with edit privileges to execute arbitrary JavaScript on the server via a crafted database connection identifier or pack configuration value.	8.8	More Details
CVE-2026-59541	Subscriber Privilege Escalation in WP BASE Booking <= 6.3.1 versions.	8.8	More Details
CVE-2026-64475	In the Linux kernel, the following vulnerability has been resolved: vfio/pci: Release the VGA arbiter client on register_device() failure The re-order in the Fixes commit below displaced vfio_pci_vga_init() as the last failure point of what is now vfio_pci_core_register_device() without introducing an unwind for the VGA arbiter registration. In current kernels this is mostly benign because vfio_pci_set_decode() only uses pci_dev state, but the original failure path could leave a callback with a freed vdev cookie. The stale registration also becomes unsafe again once the callback follows drvdata to the vfio device. Add the required VGA unwind callout.	8.8	More Details
CVE-2025-44089	An issue in NCH Software ExpressZip v11.29 allows attackers to execute arbitrary code via downloading and executing a crafted archive file.	8.8	More Details
CVE-2026-56748	Improper validation of symbolic links in the Pack Git import feature in Cribl Stream before 4.18.2 allows a remote authenticated attacker with Pack import and pipeline preview permissions to execute arbitrary code as the Cribl server process via a crafted Git repository containing a symbolic link in the pack's functions directory.	8.8	More Details
CVE-2025-44090	An issue in OhSoft CoffeeZip v4.8.0.0 allows attackers to execute arbitrary code via downloading and executing a crafted archive file.	8.8	More Details
CVE-2026-65616	Incorrect authorization validation in refresh token signature allows non-admin users to obtain a signed JFrog administrator token.	8.8	More Details
CVE-2026-65617	A deserialization weakness in JFrog Artifactory package handling could allow a low-privileged user to impact confidentiality, integrity, and availability under specific repository conditions.	8.8	More Details
CVE-2026-57785	Unauthenticated Cross Site Request Forgery (CSRF) in ApusListing <= 1.2.63 versions.	8.8	More Details
CVE-2026-66041	FFmpeg 7.0 through 8.1.2, fixed in commit 4da9812, contains a heap out-of-bounds write vulnerability in the vf_quirc filter that allows an attacker to corrupt heap memory by supplying a crafted PGS/SUP subtitle file with mismatched frame dimensions. Attackers can provide a subtitle file whose second presentation has larger dimensions than its first, causing av_image_copy_plane() to copy data exceeding the initial allocation size into the undersized libquirc grayscale image buffer, resulting in heap corruption and process crash with potential for code execution.	8.8	More Details
CVE-2026-64757	A memory corruption issue was addressed with improved state management. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, visionOS 26.6, watchOS 26.6. Processing maliciously crafted web content may lead to an unexpected Safari crash.	8.8	More Details
CVE-2026-63727	Anchore Enterprise versions from 5.11.0 to 5.27.1 and 6.0.0 contain an improper privilege escalation vulnerability in the user management API. An authenticated attacker who is able to access the Anchore Enterprise API could issue an API call capable of modifying user permissions to gain access to additional resources and operations. It is not possible to grant the system-admin role, but a read only user could be granted write access. This issue is fixed in Anchore Enterprise 5.27.2 and 6.0.1.	8.8	More Details
CVE-2026-16870	Multiple security vulnerabilities in Snowflake libsnowflakeclient versions prior to 2.9.2 could allow remote code execution and credential exfiltration. A stack-based buffer overflow in the file download path could allow remote code execution on a victim host. An attacker could exploit this by uploading a file with a crafted encryption metadata field to a shared internal stage that a victim process later downloads, and impact would be limited to deployments where principals with different privilege levels share the same internal stage. A related out-of-bounds write in the same download path could allow memory corruption with attacker-controlled write primitives. An attacker may exploit this through a crafted initialization vector metadata field on a shared stage, and impact would be limited by the same stage-write precondition. Improper validation of connection parameters could allow an attacker-controlled input to	8.8	More Details

	redirect outbound authentication requests — including credentials and tokens — to an attacker-controlled endpoint. Impact is limited to embedding deployments where a lower-privileged principal can influence connection configuration while higher-privileged service credentials are in use. The fix is available in Snowflake libsnowflakeclient version 2.9.2. The Snowflake PHP PDO Driver and Snowflake ODBC Driver embed the affected library; fixes are available in versions 4.1.0 and 3.19.0 respectively. Users must manually upgrade.		
CVE-2026-51269	schreibfaul1 ESP32-audioI2S 3.4.5 has a heap-based buffer overflow vulnerability in the connecttospeech() function. The application accepts attacker-controlled long speech text input, performs URL encoding, and directly appends the encoded result into a fixed ps_ptr heap buffer when constructing HTTP TTS request headers. Lack of input length validation and boundary checking allows remote attackers to craft oversized input to trigger out-of-bounds heap write, resulting in arbitrary code execution.	8.8	More Details
CVE-2026-49499	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) a Generation of Incorrect Security Tokens vulnerability in the IAM. A low privileged attacker with remote access could potentially exploit this vulnerability, leading to Elevation of privileges.	8.8	More Details
CVE-2025-50324	An issue in Milos Paripovic OneCommander v.3.96.0.0 allows a remote attacker to execute arbitrary code via the OneCommander.exe component.	8.8	More Details
CVE-2026-64516	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vce1: Fix VCE 1 firmware size and offsets The VCPU BO contains the actual FW at an offset, but it was not calculated into the VCPU BO size. Subtract this from the FW size to make sure there is no out of bounds access. Make sure the stack and data offsets are aligned to the 32K TLB size. Check that the FW microcode actually fits in the space that is reserved for it. (cherry picked from commit c16fe59f622a080fc457a57b3e8f14c780699449)	8.8	More Details
CVE-2026-42017	An event-handling weakness in JFrog Artifactory could expose privileged authorization material to a lower-privileged user under specific conditions.	8.8	More Details
CVE-2026-61892	Weintek cMT3092X HMI allows a non-privileged user to modify tokens to escalate privileges.	8.8	More Details
CVE-2026-64467	In the Linux kernel, the following vulnerability has been resolved: rust_binder: use a u64 stride when cleaning up the offsets array Allocation's Drop walks the offsets array (binder_size_t = u64 entries), cleaning up the objects, but it used usize instead of u64 for both the stride and the per-entry read. On 64-bit kernels (usize == u64) this is harmless, but on 32-bit kernels it walks the 8-byte entries in 4-byte steps, iterating an N-entry array 2N times, and reads the always-zero high word as offset 0, cleaning up the object at offset 0 N extra times. As a result the referenced node or handle ends up with a lower reference count than it actually has (a refcount over-decrement), and binder's reference accounting is corrupted; for example, the owner can be notified of a strong reference release (BR_RELEASE) even though references still remain. Change the stride to u64, and read each entry as a u64, narrowing it to usize with try_into(). On 32-bit ARM, when this over-decrement would drive a count below zero, the driver's existing refcount guard refuses it and fires: rust_binder: Failure: refcount underflow!	8.8	More Details
CVE-2025-50327	An issue in Franco Corbelli ZPAQFRANZ v.61.3 and before allows a remote attacker to escalate privileges and execute arbitrary code via a bypass of the Mark-of-the-Web protection mechanism	8.8	More Details
CVE-2025-50330	An issue in ZipGenius Team ZipGenius v.6.3.2.3116 and before allows a remote attacker to escalate privileges and execute arbitrary code via the zipgenius.exe.	8.8	More Details
CVE-2026-16347	MikroTik RouterOS contains a weakness in its API authentication handling that lacks effective safeguards against excessive login attempts. The system does not enforce meaningful rate-limiting, account lockout, or source-based restrictions, allowing repeated authentication failures to proceed without defensive response. In some versions, a fixed per-connection delay is present, but it can be bypassed through concurrent sessions, resulting in continued high-volume attempts. This deficiency increases the risk that an attacker could eventually obtain valid credentials and gain unauthorized access to administrative services.	8.8	More Details
CVE-2026-64364	In the Linux kernel, the following vulnerability has been resolved: HID: multitouch: fix out-of-bounds bit access on mt_io_flags mt_io_flags is a single unsigned long, but mt_process_slot(), mt_release_pending_palms() and mt_release_contacts() use it as a per-slot bitmap indexed by the slot number. That slot number is only bounded by td->maxcontacts, which is taken from the device's ContactCountMaximum feature report and can be up to 255, not by BITS_PER_LONG. As a result, a multitouch device that advertises a large contact count makes set_bit()/clear_bit() operate past the mt_io_flags word and corrupt the adjacent members of struct mt_device. The sticky-fingers release timer is the easiest way to reach this. mt_release_contacts() runs for (i = 0; i < mt->num_slots; i++) clear_bit(i, &td->mt_io_flags); with num_slots == maxcontacts. For maxcontacts around 250 the loop clears the bits that overlap td->applications.next, zeroing that list head, and the list_for_each_entry() that immediately follows then dereferences NULL. The kernel panics from timer (softirq) context. On a KASAN build this shows up as a general protection fault in mt_release_contacts() with a null-ptr-deref at offset 0x58, which is offsetof(struct mt_application, num_received). The state is reachable from an untrusted USB or Bluetooth HID multitouch device; no local privileges are required. Store the per-slot active state in a separately allocated bitmap sized for maxcontacts, the same pattern already used for pending_palm_slots, and keep only MT_IO_FLAGS_RUNNING in mt_io_flags. The two "mt_io_flags & MT_IO_SLOTS_MASK" arming checks become bitmap_empty(td->active_slots, td->maxcontacts). Move MT_IO_FLAGS_RUNNING back to bit 0. It was bumped to bit 32 by the same commit to leave the low byte for the slot	8.8	More Details

	bits; with the slot bits gone it fits in bit 0 again, which also keeps it within the unsigned long on 32-bit.		
CVE-2026-51297	sqlite 3.41 has a use-after-free vulnerability in the JSON parsing logic. Remote adversaries can craft malicious JSON payload to trigger memory free followed by illegal memory access, which may lead to arbitrary code execution, sensitive information leakage and service denial.	8.8	More Details
CVE-2026-43818	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. Processing a maliciously crafted image may lead to arbitrary code execution.	8.8	More Details
CVE-2026-64366	In the Linux kernel, the following vulnerability has been resolved: HID: wacom: fix slab-out-of-bounds write in wacom_wac_queue_insert wacom_wac_queue_insert() calls kfifo_skip() in a loop when the kfifo doesn't have enough space for the incoming report. If the kfifo is empty, kfifo_skip() reads stale data left in the kmalloc'd buffer via __kfifo_peek_n() and interprets it as a record length, advancing fifo->out by that garbage value. This corrupts the internal kfifo state, causing kfifo_unused() to return a value much larger than the actual buffer size, which bypasses __kfifo_in_r()'s guard: if (len + recsize > kfifo_unused(fifo)) return 0; kfifo_copy_in() then performs an out-of-bounds memcpy, writing up to 3842 bytes past the 256-byte buffer. Add a !kfifo_is_empty() condition to the while loop so kfifo_skip() is never called on an empty fifo, and check the return value of kfifo_in() to reject reports that are too large for the fifo.	8.8	More Details
CVE-2026-65921	A path validation weakness in archive extraction/write handling allows entries with traversal sequences to be written outside the intended build artifacts location.	8.8	More Details
CVE-2026-60368	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via SOAP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	8.8	More Details
CVE-2026-61246	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	8.8	More Details
CVE-2026-7187	Missing authentication for critical function vulnerability in Universal Software Inc. UKBS allows Accessing Functionality Not Properly Constrained by ACLs. This issue affects UKBS: through 28072026. NOTE: The vendor was contacted and it was learned that the product is not supported.	8.8	More Details
CVE-2026-57510	SuperPlane before 0.27.0 contains a broken object-level authorization vulnerability in the CanvasService gRPC handlers that allows authenticated users with viewer-level access to one organization to access resources belonging to other organizations by supplying arbitrary canvas or queue UUIDs without organization scoping. Attackers can read cross-tenant execution history and event payloads containing sensitive secrets, write queue items and canvas events into victim organizations, delete arbitrary canvases, and disrupt automation workflows across tenant boundaries.	8.8	More Details
CVE-2026-66032	libssh2 through 1.11.1, fixed in commit 5e47761, contains a double-free vulnerability in the sftp_open() function in src/sftp.c that allows a malicious SSH server to corrupt the heap of any authenticated client opening an SFTP session. When a server responds to SSH_FXP_OPEN with SSH_FXP_STATUS containing FX_OK, the response data buffer is freed, and if a subsequent sftp_packet_require() call returns a specific error such as LIBSSH2_ERROR_CHANNEL_PACKET_EXCEEDED, the same pointer is freed a second time, enabling tcache dup conditions on glibc systems that allow overlapping allocations and function pointer overwrites.	8.8	More Details
CVE-2026-66036	FFmpeg through 8.1.2, fixed in commit 5d7112c, contains a heap out-of-bounds write vulnerability in the vf_hqdn3d filter that allows attackers to corrupt heap memory by supplying a crafted video whose frame resolution increases between frames when filtergraph reinitialization is disabled via the -reinit_filter 0 option. Attackers can provide a malicious video input where vf_hqdn3d.config_input() allocates undersized per-plane line-history buffers based on the initial frame width, and subsequent larger frames cause denoise_spatial() to write beyond the allocation boundary, resulting in heap memory corruption.	8.8	More Details
CVE-2026-64255	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mld: validate sta_mask before ffs() in BA session handlers Three BA session handlers use ffs(ba_data->sta_mask) - 1 to derive a station ID without checking that sta_mask is non-zero. When sta_mask is zero, ffs() returns 0 and the subtraction wraps to 0xFFFFFFFF, causing an out-of-bounds access on fw_id_to_link_sta[]. Add WARN_ON_ONCE(!ba_data->sta_mask) guards before each ffs() call, consistent with the existing check in iwl_mld_ampdu_rx_start().	8.8	More Details
CVE-2026-64280	In the Linux kernel, the following vulnerability has been resolved: fpga: dfl-afu: validate DMA mapping length in afu_dma_map_region() afu_ioctl_dma_map() accepts a 64-bit length from userspace via DFL_FPGA_PORT_DMA_MAP ioctl without an upper bound check. The value is passed to afu_dma_pin_pages() where npages is derived as length >> PAGE_SHIFT and passed to pin_user_pages_fast() which takes int nr_pages, causing implicit truncation if length is very large. Validate map.length at the ioctl entry point before calling afu_dma_map_region(), rejecting values whose page count exceeds INT_MAX.	8.8	More Details
	The MDJM Event Management plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 1.7.8.4. This is due to missing capability checks and nonce verification in the		

CVE-2026-15017	`MDJM_Permissions::set_permissions()` and `MDJM_Employee_Manager::init()` functions, combined with the absence of server-side allow-list validation on the `employee_roles[]` and `new_role` POST parameters before they are passed to `mdjm_set_employee_role()` and `WP_User::set_role()`. This makes it possible for unauthenticated attackers to grant arbitrary MDJM capabilities — including `mdjm_employee` and `mdjm_employee_edit` — to any registered WordPress role, and subsequently leverage a subscriber-level account to escalate privileges to Administrator. `MDJM_Permissions::init()` is registered on the public WordPress `init` hook without any authentication gate, meaning the role-manipulation endpoint is reachable without any prior login.	8.8	More Details
CVE-2026-14167	A low privileged remote attacker can perform privileged configuration changes reserved for the administrator level including permission management due to incorrect authorization.	8.8	More Details
CVE-2026-14168	A low privileged remote attacker can gain administrator privileges due to missing authorization at the insert path of the configuration table resulting in gaining full system access.	8.8	More Details
CVE-2026-66039	FFmpeg through 8.1.2, fixed in commit aafb5c6, contains a signed integer overflow vulnerability in the MACE6 audio decoder that allows attackers to corrupt heap memory by supplying a crafted CAF file with a malicious bytes_per_packet value. Attackers can craft a CAF file with oversized bytes_per_packet and frames_per_packet values in the desc chunk to trigger an integer overflow in mace_decode_frame() during output sample count computation, resulting in an undersized buffer allocation and heap out-of-bounds write that could enable code execution.	8.8	More Details
CVE-2026-14328	The Eazy Plugin Manager – Powerful Plugin Management Solution for WordPress plugin for WordPress is vulnerable to Privilege Escalation in all versions up to, and including, 4.4.1. This is due to insufficient authorization on the `wp_ajax_pos_get_option` AJAX handler, which verifies only a nonce that is localized to every logged-in admin-area user via `admin_enqueue_scripts` — without any capability check — before returning the value of any arbitrary WordPress option via `get_option()`, combined with the `admin_login_endpoint_handler` REST endpoint (`GET /wp-json/epm/v1/admin/login`) being registered as publicly accessible and authenticating callers solely by a whirlpool hash of values stored in those same options. This makes it possible for authenticated attackers, with Subscriber-level access and above, to read the `site_url`, `connection_key`, and `remote_user_id` values stored in the `eazywp_connecting_info` and `eazywp_connection` options, compute the required `auth_key`, call the `admin/login` REST endpoint to obtain Administrator authentication cookies, and fully take over the site. Exploitation requires the plugin's remote connection feature to have been configured, as the `eazywp_connecting_info` and `eazywp_connection` options must be populated with valid credentials.	8.8	More Details
CVE-2026-66040	FFmpeg through 8.1.2, fixed in commit b506faf, contains a heap out-of-bounds write vulnerability in the native PNG and APNG encoders that allows remote attackers to corrupt heap memory by supplying a crafted PNG image with a malicious eXIf chunk. Attackers can craft an eXIf chunk where multiple IFD entries reference the same large value payload, causing canonical serialization to expand the output far beyond the undersized allocation estimated by add_exif_profile_size(), resulting in png_write_chunk() writing tens of thousands of bytes past the buffer boundary, leading to deterministic heap corruption, process crash, and potentially arbitrary code execution.	8.8	More Details
CVE-2026-28931	A buffer overflow was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, watchOS 26.6. Connecting to a malicious NFS server may lead to kernel memory corruption.	8.8	More Details
CVE-2026-64876	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in GeoIP extension - Database-update requests lacked consistent token and Super User checks, this could cause unauthorized updates.	8.8	More Details
CVE-2021-32087	An issue was discovered in Quest KACE Systems Deployment Appliance (SMA) 11.0.273. It installs with default user credentials. The kbftp account has a password of getbxf, which is publicly known and documented. This allows remote attackers to trivially gain privileged access to the FTP service interface, which contains MySQL backups. Sensitive information is stored in the database, such as privileged credentials for other systems.	8.8	More Details
CVE-2026-16745	A flaw was found in odh-dashboard, the web console component of Red Hat OpenShift AI (RHOAI). Due to incorrect network binding, a malicious actor within the cluster can bypass authentication and impersonate any user by providing an arbitrary access token. This allows an attacker to gain unauthorized access to the Kubernetes API, potentially leading to arbitrary code execution, privilege escalation, or information disclosure.	8.8	More Details
CVE-2026-62426	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] To manage the system, sysctl and platform operations are used by the control domain or a possible Xenstore domain. Some of these operations may not be executed in parallel, so a system-wide lock each is used. The way those locks are acquired is, however, not providing any fairness. Furthermore, with XSM/Flask in use, the lock acquire will, for some operations, occur ahead of any permission checking. The sysctl issue is CVE-2026-62426. The platform-op issue is CVE-2026-62427.	8.8	More Details
CVE-2021-32085	An issue was discovered in Quest KACE Systems Deployment Appliance (SMA) 11.0.273. It installs with default user credentials. The report and R1 MySQL accounts have a password of box747, which is publicly known and documented. This allows remote attackers to trivially gain privileged access to the MySQL databases. Sensitive information is stored in the database, such as privileged credentials for other systems.	8.8	More Details
CVE-2026-64783	A use-after-free issue was addressed with improved memory management. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, visionOS 26.6, watchOS 26.6. Processing maliciously crafted web content may lead to an unexpected Safari crash.	8.8	More Details

CVE-2026-49258	Nebula Mesh is a self-hosted control plane for the Slack Nebula mesh VPN. In versions 0.3.5 and below, the web UI (/ui/*) does not apply the per-operator CA scoping employed by the JSON API. This was partially addressed by GHSA-598g-h2vc-h5vg, but the changes were not implemented in the web read/mutation surface. Any authenticated non-admin operator (for example, one created via self-registration or OIDC) can access resources belonging to other operators. The host create/edit/mobile-bundle/network-create paths and all CA-management routes were already correctly scoped. A malicious operator could block or delete any other operator's host, or read any operator's hosts and networks. This issue has been fixed in version 0.3.6.	8.8	More Details
CVE-2026-64835	FFmpeg versions 4.4 through 8.1.2 contain an out-of-bounds memory access vulnerability in the ADX audio decoder within libavcodec/adxdec.c that allows attackers to trigger both out-of-bounds reads and writes by supplying a crafted ADX or AAX audio file with a mid-stream channel layout change. When AV_PKT_DATA_NEW_EXTRADATA side data is received mid-stream, the adx_decode_frame function re-parses the stream header but fails to update the internal channel state, causing subsequent decoding operations to access the prev[] state array using a stale channel count.	8.8	More Details
CVE-2026-62427	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] To manage the system, sysctl and platform operations are used by the control domain or a possible Xenstore domain. Some of these operations may not be executed in parallel, so a system-wide lock each is used. The way those locks are acquired is, however, not providing any fairness. Furthermore, with XSM/Flask in use, the lock acquire will, for some operations, occur ahead of any permission checking. The sysctl issue is CVE-2026-62426. The platform-op issue is CVE-2026-62427.	8.8	More Details
CVE-2026-64832	FFmpeg versions 4.4 through 8.1.2 contain a double-free vulnerability in the NVIDIA NVDEC hardware decoder within libavcodec/nvdec.c that allows attackers to trigger memory corruption by supplying a crafted video file. When no decoder surfaces remain, the ff_nvdec_start_frame_sep_ref error path frees memory via nvdec_fdd_priv_free while the calling layer subsequently frees the same frame description data, resulting in a double-free of the underlying decoder context in any FFmpeg-based application using NVDEC hardware-accelerated decoding.	8.8	More Details
CVE-2026-66014	JFrog Artifactory contains an authentication handling weakness in internal request processing that, under specific conditions, may allow an attacker to escalate privileges beyond the intended access level.	8.8	More Details
CVE-2026-16801	Improper control of generation of code ('Code Injection') in the variables feature in Devolutions PowerShell Universal 2026.2.2 and earlier allows an authenticated user with variable write permission to execute arbitrary PowerShell code via a crafted variable value that is not properly escaped when written to the variables configuration file.	8.8	More Details
CVE-2026-65013	Onlook through 0.2.32, fixed in commit 423e2e9, contains a broken object level authorization vulnerability that allows authenticated attackers to access and manipulate other users' resources by supplying arbitrary UUID values to tRPC API procedures including project.get, member.remove, and chat.conversation.delete. Attackers can provide arbitrary projectId or conversationId values without authorization validation to read, modify, and delete other users' project data, members, and conversation history.	8.8	More Details
CVE-2026-64313	In the Linux kernel, the following vulnerability has been resolved: crypto: ecc - Fix carry overflow in vli multiplication The carry flag calculation fails when r01.m_high is saturated (0xFFFFFFFFFFFFFFFF) and addition of lower bits overflows. The condition (r01.m_high < product.m_high) doesn't handle the case where r01.m_high == product.m_high and an additional carry exists from lower-bit overflow. When commit 3c4b23901a0c ("crypto: ecdh - Add ECDH software support") introduced crypto/ecc.c, it split the muladd() function in the micro-ecc library into separate mul_64_64() and add_128_128() helpers. It seems the check got lost in translation. Add proper handling for this boundary by accounting for the carry from the lower addition.	8.8	More Details
CVE-2026-16800	Improper control of generation of code ('Code Injection') in the schedule feature in Devolutions PowerShell Universal 2026.2.2 and earlier allows an authenticated user with schedule creation permission to execute arbitrary PowerShell code via crafted schedule parameter names concatenated into a script invocation.	8.8	More Details
CVE-2026-64831	FFmpeg versions 8.0 through 8.1.2 contains a stack buffer overflow vulnerability in the Vulkan HEVC hardware decoder that allows remote attackers to overwrite return addresses and adjacent stack frames by supplying a crafted HEVC/H.265 bitstream. Attackers can embed a malicious vps_num_hrd_parameters value exceeding HEVC_MAX_SUB_LAYERS in any supported container format to overflow stack-allocated arrays in the vk_hevc_end_frame function, potentially achieving arbitrary code execution.	8.8	More Details
CVE-2026-64830	FFmpeg versions 2.1 through 8.1.2 contains a heap buffer overflow vulnerability in the VobSub subtitle demuxer that allows attackers to corrupt adjacent heap memory by supplying a malicious .sub/.idx subtitle file declaring more distinct stream IDs than the fixed-size array bounds in libavformat/mpeg.c. Attackers can craft a subtitle file with excessive distinct stream IDs to trigger unbounded writes beyond the vobsub->q[] array boundary via ff_subtitles_queue_insert(), potentially achieving arbitrary code execution in any application using FFmpeg's VobSub demuxer.	8.8	More Details
CVE-2026-45813	Out-of-bounds Write, Integer Underflow (Wrap or Wraparound) vulnerability in Apache NimBLE BASS service. Improper validation when parsing BASS service "Add Source" and "Modify Source" operation PDU could results in stack buffer overflow or arbitrary out-of-bound read. This can be triggered by nearby devices over Bluetooth connection, however pairing is required prior to accessing BASS service, which depending on device configuration may or may not require user action. This issue affects Apache NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.	8.8	More Details
CVE-2026-	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. From 0.17.0 until 0.60.2, datamodel-code-generator preserves attacker-controlled default_factory values in src/datamodel_code_generator/parser/jsonschema.py through JsonSchemaObject.init and get_field_extras and emits	8.8	More Details

54653	them into Field(default_factory=...) or field(default_factory=...), allowing Python expression execution when the generated model is imported. This issue is fixed in version 0.60.2.		
CVE-2026-16771	In firmware versions 2.7.7 and earlier, the Arris BGW210-700 gateway fails to enforce any server-side authentication on its /cgi-bin/*.ha management endpoints, relying solely on client-side CSS/JavaScript gating that can be bypassed by any HTTP client. This allows unauthenticated attackers on the LAN to read sensitive configuration data, modify persistent device settings, or trigger backend diagnostic operations. The issue appears systemic across the CGI handler chain.	8.8	More Details
CVE-2026-14551	The servereye client (also known as sensorhub, technically ClientAgentContainerService) versions 20.15 and earlier are vulnerable to Local Privilege Escalation. The high-privileged service SE3Recovery (EmergencyRecoveryService.exe), running as SYSTEM, periodically monitors the directory %ProgramData%\ServerEye3\update\ for a trigger file named "update_available". Due to insufficient access restrictions on this directory, a local standard user can create the trigger file and provide a path to a directory containing malicious JSON instructions. The service subsequently executes the utility UpdaterAction.exe with SYSTEM privileges, which parses the instructions and performs an unvalidated file copy from a user-controlled source to a protected system destination (e.g., overwriting a service binary). This leads to full system compromise as the service automatically restarts the overwritten binary with SYSTEM privileges.	8.8	More Details
CVE-2026-64394	In the Linux kernel, the following vulnerability has been resolved: ksmbd: add a WRITE_DAC/WRITE_OWNER check to SMB2 SET_INFO SECURITY commit cc57232cae23 ("ksmbd: fix FSCTL permission bypass by adding a permission check for FSCTL_SET_SPARSE") added a fp->daccess gate to fsctl_set_sparse and noted that "similar handle-level checks exist in other functions but are missing here." The SMB2 SET_INFO SECURITY arm is one of the missing ones, and the most security-relevant: smb2_set_info_sec() calls set_info_sec() with no per-handle access check. set_info_sec() (fs/smb/server/smbacl.c) re-permissions the file: it rewrites owner/group/mode via notify_change(), rewrites the POSIX ACL via set_posix_acl(), and on KSMBD_SHARE_FLAG_ACL_XATTR shares removes and rewrites the Windows security descriptor via ksmbd_vfs_set_sd_xattr(). Every other persistent-mutation arm of the sibling handler smb2_set_info_file() checks fp->daccess first (FILE_WRITE_DATA / FILE_DELETE / FILE_WRITE_EA / FILE_WRITE_ATTRIBUTES); the SECURITY arm — which mutates the access control itself — is the only one with no gate. A client can therefore open a handle with FILE_WRITE_ATTRIBUTES only (no FILE_WRITE_DAC / FILE_WRITE_OWNER) and use SMB2_SET_INFO with InfoType SMB2_O_INFO_SECURITY to rewrite the file's DACL and owner, granting itself access the handle's daccess never carried. Unlike the FSCTL data arms this is a metadata/xattr operation, so there is no FMODE_WRITE VFS backstop — the missing fp->daccess check is the entire gate. Setting a security descriptor is the WRITE_DAC / WRITE_OWNER operation, so require at least one of those on the handle before re-permissioning the file. -EACCES is mapped to STATUS_ACCESS_DENIED by smb2_set_info().	8.8	More Details
CVE-2026-60373	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	8.8	More Details
CVE-2026-64396	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix UAF of struct file_lock in SMB2_LOCK deferred-lock cancellation When a blocking byte-range lock request is deferred in the FILE_LOCK_DEFERRED path, ksmbd registers the asynchronous work into the connection's async_requests list via setup_async_work(). The cancel callback smb2_remove_blocked_lock() holds a reference to the flock. If the lock waiter is subsequently woken up but the work state is no longer KSMBD_WORK_ACTIVE (e.g., due to a concurrent cancellation), the cleanup path calls locks_free_lock(flock) without dequeuing the work from the async_requests list. Concurrently, smb2_cancel() walks the list under conn->request_lock and invokes the cancel callback, which then dereferences the already freed 'flock'. This leads to a slab-use-after-free inside __wake_up_common. Fix this by restructuring the cleanup logic after the worker returns from ksmbd_vfs_posix_lock_wait(). Move list_del(&smb_lock->llist) and release_async_work(work) to the top of the cleanup block. This guarantees that the async work is completely dequeued and serialized under conn->request_lock before locks_free_lock(flock) is called, rendering the flock unreachable for any concurrent smb2_cancel().	8.8	More Details
CVE-2026-65591	n8n contains a sanitizer bypass vulnerability in the legacy expression evaluator's computed-member handler. An authenticated user with workflow create or modify permissions can craft a malicious expression to bypass the sanitizer and achieve host-level code execution as the n8n process. The legacy expression engine is the default in affected versions. Fixed in n8n 1.123.64, 2.29.8, and 2.30.1.	8.8	More Details
CVE-2026-65015	n8n versions before 2.30.1 contain a privilege escalation vulnerability in the AI Agents feature where the node-execution tool lacks proper authorization checks. A Project Viewer user can escalate privileges by chatting with an agent that has node tools enabled, executing arbitrary nodes and accessing credential secrets without proper authorization verification.	8.8	More Details
CVE-2026-51235	LibRaw 0.21 is vulnerable to Buffer Overflow in the stretch() function (src/libraw_cxx.cpp) and fuji_rotate() function (src/decoders/fuji.cpp).	8.8	More Details
CVE-2026-51274	In schreibfaul1 ESP32-audioI2S 3.4.5, a heap-based buffer overflow in the ID3v2 SYLT synchronized lyrics parser in audiolib allows remote attackers to cause a denial of service (application crash), information disclosure, or potential arbitrary code execution via a crafted MP3 file. The vulnerability occurs due to missing bounds validation on attacker-controlled frame size and improper memory access during lyric parsing.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix WEP length underflow and		

CVE-2026-64445	OOB read in OnAuth() OnAuth() has two bugs in the shared-key authentication path. When the Privacy bit is set, rtw_wep_decrypt() is called without verifying that the frame is long enough to contain a valid WEP IV and ICV. Inside rtw_wep_decrypt(), length is computed as: length = len - WLAN_HDR_A3_LEN - iv_len and then passed as (length - 4) to crc32_le(). If len is less than WLAN_HDR_A3_LEN + iv_len + icv_len (32 bytes), length - 4 is negative and, after the implicit cast to size_t, causes crc32_le() to read far beyond the frame buffer. Add a minimum length check before accessing the IV field and calling the decryption path. When processing a seq=3 response, rtw_get_ie() stores the Challenge Text IE length in ie_len, but the subsequent memcmp() always reads 128 bytes regardless of ie_len. IEEE 802.11 mandates a challenge text of exactly 128 bytes; reject any IE whose length field differs, matching the check already applied to OnAuthClient().	8.8	More Details
CVE-2026-51275	In schreibfaul1 ESP32-audioI2S 3.4.5, a heap-based buffer overflow in the ID3v2 APIC frame parsing function in audiolib allows remote attackers to execute arbitrary code or cause a denial of service (crash) via a crafted MP3 file. The vulnerability exists due to missing length validation when processing the APIC frame size field, leading to an out-of-bounds memory write.	8.8	More Details
CVE-2026-15962	The Fluent Forms Pro Add On Pack plugin for WordPress is vulnerable to PHP Object Injection in all versions up to, and including, 6.2.6 via deserialization of untrusted input. This makes it possible for authenticated attackers, with Subscriber-level access and above, to inject a PHP Object. The additional presence of a POP chain allows attackers to change user passwords and potentially take over administrator accounts. Note: This can only be exploited if user update integration is enabled and a user meta field is mapped.	8.8	More Details
CVE-2026-64408	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: bnep: pin L2CAP connection during netdev registration bnep_add_connection() reads the L2CAP connection without holding the channel lock, then passes its HCI device to register_netdev(). Controller teardown can clear and release that connection concurrently, leaving the network device registration path to dereference a freed parent device. Take a reference to the L2CAP connection while holding the channel lock. Retain it until register_netdev() has taken the parent device reference.	8.8	More Details
CVE-2026-65608	Grav versions >= 1.7.0 and before 2.0.9 contain a remote code execution vulnerability. FlexDirectory::dynamicDataField() resolves blueprint data-*@: directives by calling call_user_func_array() on attacker-influenced input, validating only that the target is callable (is_callable()) without restricting dangerous functions such as exec, system, passthru, or shell_exec. Because FlexDirectory registers this handler for every Flex directory, it bypasses the validation added to Blueprint::dynamicData() in 2.0.7 (GHSA-fj2p-qj2f-74v5). Any authenticated user with create or update permission on any Flex-based directory (Flex Users, Flex Pages, Flex Objects, or custom Flex types) can execute arbitrary shell commands on the server.	8.8	More Details
CVE-2026-65897	Grav API Plugin versions before 1.0.10 fail to validate the groups field in InvitationsController::create(), allowing authenticated api.users.write callers to assign invited accounts to groups that grant api.super permissions. Attackers can create invitation records with elevated group membership, and when accepted, the new account gains full super-admin API access without the inviter holding those permissions.	8.8	More Details
CVE-2026-15212	The WPO365 | Login plugin for WordPress is vulnerable to Cross-Site Request Forgery in versions up to, and including, 43.2. This is due to the Ajax_Service::verify_ajax_request() helper gating its wp_verify_nonce() call behind the boolean option 'enable_nonce_check', which is absent from the default 'wpo365_options' array and therefore evaluates to false via get_global_boolean_var(); as a result, the wp_ajax_wpo365_update_settings handler (Ajax_Service::update_settings) accepts POSTs from cross-origin pages and forwards the attacker-supplied 'settings' payload (base64/JSON) to Options_Service::update_options(), which merges every key/value into wpo365_options without a key allowlist. This makes it possible for unauthenticated attackers to overwrite arbitrary plugin options — including enabling the SCIM REST endpoint (enable_scim), planting an attacker-known scim_secret_token, and setting new_usr_default_role to 'administrator' — via a forged request granted they can trick a site administrator into performing an action such as clicking on a link.	8.8	More Details
CVE-2026-65906	In JetBrains TeamCity before 2026.1.2, 2025.11.6 code execution via Kotlin DSL sandbox escape was possible	8.8	More Details
CVE-2026-64441	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB reads in rtw_get_sec_ie(), rtw_get_wapi_ie(), and rtw_get_wps_attr() Three IE/attribute parsing functions have missing bounds checks. rtw_get_sec_ie() and rtw_get_wapi_ie() iterate over a raw IE buffer without verifying that the header bytes (tag + length) are within the remaining buffer before reading them. Additionally, rtw_get_sec_ie() compares the 4-byte WPA OUI at cnt+2 without checking that at least 6 bytes remain, and rtw_get_wapi_ie() compares a 4-byte WAPI OUI at cnt+6 without checking that at least 10 bytes remain. rtw_get_wps_attr() reads wps_ie[0] and wps_ie+2 unconditionally at entry, before verifying that wps_ielen is large enough to contain the 6-byte WPS IE header (element_id + length + 4-byte OUI). Inside the attribute loop, get_unaligned_be16() is called on attr_ptr and attr_ptr+2 without checking that 4 bytes remain in the buffer. Add a cnt+2 bounds check before each loop body in rtw_get_sec_ie() and rtw_get_wapi_ie(), guard each multi-byte comparison with a minimum IE length requirement, add a wps_ielen < 6 early return in rtw_get_wps_attr(), and add a 4-byte bounds check in its inner loop.	8.8	More Details
CVE-2026-65917	CyberPanel through 1.9.1, fixed in commit b198460, contains an insecure direct object reference (IDOR) vulnerability in the IncBackups application's incremental-backup handlers (deleteBackup, fetchRestorePoints, and restorePoint) that allows authenticated panel users to access or manipulate other tenants' backup resources by supplying an attacker-controlled globally sequential IncJob integer ID that is never re-scoped to the authorized domain. Attackers can enumerate sequential backup IDs to read another tenant's backup metadata, irrecoverably delete another tenant's backup snapshots, or trigger unauthorized restoration of another tenant's backup job with root privileges.	8.8	More Details
CVE-2026-	Supermicro (SMC) SMASH services contain an Arbitrary code execution issue in X14DBG-DAP and X14DBI. An authorized attacker can exploit SMASH’s input capability to compromise data integrity or launch a Denial-of-Service	8.8	More

3821	(DoS) attack against the BMC.		Details
CVE-2026-65690	Bold Reports Standalone Report Designer before 14.1.12 contains a missing filepath validation vulnerability in its file upload functionality that allows authenticated attackers to traverse outside the intended directory by supplying a crafted filename. Attackers can exploit this path traversal weakness to execute arbitrary commands with high privileges on the server. The vulnerability is specific to the DataHub module, which was introduced in Bold Reports 6.3. Therefore, versions prior to 6.3 are not affected.	8.8	More Details
CVE-2026-55578	Pheditor is a single-file editor and file manager written in PHP. From version 2.0.1 to before version 2.0.6, the terminal feature in Pheditor uses an incomplete character blocklist to sanitize user-supplied commands before passing them to shell_exec(). After the fix for GHSA-9643-6xjp-vx57 (which added $ to the blocklist), the characters | (single pipe), ` (backtick), and the newline byte (0x0A) remain unblocked. An authenticated user with the terminal permission (enabled by default) can leverage any of these to bypass the TERMINAL_COMMANDS allowlist and execute arbitrary OS commands as the web server user. This issue has been patched in version 2.0.6.	8.8	More Details
CVE-2026-12968	The Product Addons and Product Options With Custom Fields WordPress plugin before 1.6.15 does not restrict an unauthenticated file-upload endpoint and accepts SVG files that are stored and served inline, allowing an unauthenticated attacker to upload a malicious SVG whose embedded script executes in the session of any user (such as an administrator) who later opens the file.	8.8	More Details
CVE-2026-60134	Weintek cMT3092X HMI allows a non-privileged user to modify cookies to gain elevated privileges.	8.8	More Details
CVE-2026-64438	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - fix VF2PF work teardown race in adf_disable_sriov() The VF2PF interrupt handler queues PF-side response work that stores a raw pointer to per-VF state (struct adf_accel_vf_info). Currently, adf_disable_sriov() destroys per-VF mutexes and frees vf_info without stopping new VF2PF work or waiting for in-flight workers to complete. A concurrently scheduled or already queued worker can then dereference freed memory. This manifests as a use-after-free when KASAN is enabled: BUG: KASAN: null-ptr-deref in mutex_lock+0x76/0xe0 Write of size 8 at addr 0000000000000260 by task kworker/24:2/... Workqueue: qat_pf2vf_resp_wq adf_iov_send_resp [intel_qat] Call Trace: kasan_report+0x119/0x140 mutex_lock+0x76/0xe0 adf_gen4_pfvf_send+0xd4/0x1f0 [intel_qat] adf_recv_and_handle_vf2pf_msg+0x290/0x360 [intel_qat] adf_iov_send_resp+0x8c/0xe0 [intel_qat] process_one_work+0x6ac/0xfd0 worker_thread+0x4dd/0xd30 kthread+0x326/0x410 ret_from_fork+0x33b/0x670 Add a PF-local flag, vf2pf_disabled, that gates work queueing, worker processing, and interrupt re-enabling during teardown. Set this flag atomically with the hardware interrupt mask inside adf_disable_all_vf2pf_interrupts(). After masking, synchronize the AE cluster MSI-X interrupt and flush the PF response workqueue before tearing down per-VF locks and state so all in-flight work completes before vf_info is destroyed. Introduce adf_enable_all_vf2pf_interrupts() to clear the flag and unmask all VF2PF interrupts under the same lock when SR-IOV is re-enabled. This ensures the software flag and hardware state transition atomically on both the enable and disable paths.	8.8	More Details
CVE-2026-64437	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix use-after-free of a deferred file_lock on SMB2_CLOSE then SMB2_CANCEL Commit f580d27e8928 ("ksmbd: fix use-after-free of a deferred file_lock on double SMB2_CANCEL") made smb2_cancel() skip a work whose state is KSMBD_WORK_CANCELLED, so its cancel_fn cannot be fired a second time. But KSMBD_WORK has three states (ACTIVE, CANCELLED, CLOSED), and the same freeing producer path is reached for CLOSED too: SMB2_CLOSE on the locking handle -> set_close_state_blocked_works() sets the deferred work's state to KSMBD_WORK_CLOSED and wakes the smb2_lock() worker. The worker takes the non-ACTIVE early-exit, locks_free_lock()s the file_lock and, because the state is not KSMBD_WORK_CANCELLED, takes the STATUS_RANGE_NOT_LOCKED branch with "goto out2" -- which, like the cancelled branch, skips release_async_work(). The work stays on conn->async_requests with a live cancel_fn = smb2_remove_blocked_lock pointing at the freed file_lock. A subsequent SMB2_CANCEL for the same AsyncId then passes the KSMBD_WORK_CANCELLED-only guard (its state is KSMBD_WORK_CLOSED), so smb2_cancel() fires cancel_fn again over the freed file_lock -- the same use-after-free fixed, via SMB2_CLOSE instead of a first SMB2_CANCEL: BUG: KASAN: slab-use-after-free in __locks_delete_block __locks_delete_block locks_delete_block ksmbd_vfs_posix_lock_unblock smb2_remove_blocked_lock smb2_cancel <- 2nd SMB2_CANCEL fires cancel_fn handle_ksmbd_work Allocated by ...: locks_alloc_lock <- smb2_lock Freed by ...: locks_free_lock <- smb2_lock (non-ACTIVE early-exit) ... cache file_lock_cache of size 192 Reproduced on mainline 7.1-rc7 (which already contains f580d27e8928) with KASAN by an authenticated SMB client; the double-SMB2_CANCEL control is silent on that kernel, so the splat is attributable to the CLOSE trigger. Only an ACTIVE deferred work may have its cancel_fn fired: both terminal states (CANCELLED and CLOSED) reach the smb2_lock() early-exit that frees the file_lock and skips release_async_work(). Guard on KSMBD_WORK_ACTIVE so any non-active work is skipped.	8.8	More Details
CVE-2026-54540	Pheditor is a single-file editor and file manager written in PHP. Prior to version 2.0.5, there is an authenticated terminal command whitelist bypass. The terminal feature checks whether the submitted command starts with one of the configured TERMINAL_COMMANDS values, then passes the full command string to shell_exec(). Shell command substitution such as $() is not blocked, so an authenticated user with the terminal permission can bypass a restricted command allowlist and execute arbitrary shell commands as the web server user. This issue has been patched in version 2.0.5.	8.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: Fix UAF in channel timeout by holding conn ref l2cap_chan_timeout() runs asynchronously and accesses chan->conn. If the connection is torn down while the timer is running or pending, chan->conn can be freed, leading to a use-after-free when the timer worker attempts to lock conn->lock: | BUG: KASAN: slab-use-after-free in instrument_atomic_read_write include/linux/instrumented.h:112 [inline] | BUG: KASAN: slab-use-after-free in atomic_long_try_cmpxchg_acquire include/linux/atomic/atomic-instrumented.h:4456 [inline] | BUG: KASAN: slab-use-after-free in __mutex_trylock_fast		

CVE-2026-64434	kernel/locking/mutex.c:161 [inline] | BUG: KASAN: slab-use-after-free in mutex_lock+0x4f/0xa0 kernel/locking/mutex.c:318 | Write of size 8 at addr ffff8881298d9550 by task kworker/2:1/83 | | CPU: 2 UID: 0 PID: 83 Comm: kworker/2:1 Not tainted 7.1.0-rc6-next-20260601-dirty #6 PREEMPT(full) | Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.17.0-debian-1.17.0-1 04/01/2014 | Workqueue: events l2cap_chan_timeout | Call Trace: | <TASK> | instrument_atomic_read_write include/linux/instrumented.h:112 [inline] | atomic_long_try_cmpxchg_acquire include/linux/atomic/atomic-instrumented.h:4456 [inline] | __mutex_trylock_fast kernel/locking/mutex.c:161 [inline] | mutex_lock+0x4f/0xa0 kernel/locking/mutex.c:318 | l2cap_chan_timeout+0x5d/0x1b0 net/bluetooth/l2cap_core.c:422 | process_one_work kernel/workqueue.c:3326 [inline] | process_scheduled_works+0x7c8/0xfb0 kernel/workqueue.c:3409 | worker_thread+0x8a9/0xcf0 kernel/workqueue.c:3490 | kthread+0x346/0x430 kernel/kthread.c:436 | ret_from_fork+0x1a3/0x470 arch/x86/kernel/process.c:158 | ret_from_fork_asm+0x1a/0x30 arch/x86/entry/entry_64.S:245 | </TASK> | | Allocated by task 320: | l2cap_conn_add+0xa7/0x820 net/bluetooth/l2cap_core.c:7075 | l2cap_connect_cfm+0xdb/0xd70 net/bluetooth/l2cap_core.c:7452 | hci_connect_cfm include/net/bluetooth/hci_core.h:2139 [inline] | hci_remote_features_evt+0x52f/0x9f0 net/bluetooth/hci_event.c:3760 | hci_event_func net/bluetooth/hci_event.c:7796 [inline] | hci_event_packet+0x561/0xa70 net/bluetooth/hci_event.c:7847 | hci_rx_work+0x370/0x890 net/bluetooth/hci_core.c:4040 | process_one_work kernel/workqueue.c:3326 [inline] | process_scheduled_works+0x7c8/0xfb0 kernel/workqueue.c:3409 | worker_thread+0x8a9/0xcf0 kernel/workqueue.c:3490 | kthread+0x346/0x430 kernel/kthread.c:436 | ret_from_fork+0x1a3/0x470 arch/x86/kernel/process.c:158 | ret_from_fork_asm+0x1a/0x30 arch/x86/entry/entry_64.S:245 | | Freed by task 322: | hci_disconn_cfm include/net/bluetooth/hci_core.h:2154 [inline] | hci_conn_hash_flush+0x101/0x1f0 net/bluetooth/hci_conn.c:2736 | hci_dev_close_sync+0x889/0xde0 net/bluetooth/hci_sync.c:5405 | hci_dev_do_close net/bluetooth/hci_core.c:502 [inline] | hci_unregister_dev+0x1f7/0x370 net/bluetooth/hci_core.c:2679 | vhci_release+0x12a/0x180 drivers/bluetooth/hci_vhci.c:690 | __fput+0x369/0x890 fs/file_table.c:510 | task_work_run+0x160/0x1d0 kernel/task_work.c:233 | get_signal+0xf5b/0x1120 kernel/signal.c:2810 | arch_do_signal_or_restart+0x4d/0x600 arch/x86/kernel/signal.c:337 | __exit_to_user_mode_loop kernel/entry/common.c:64 [inline] | exit_to_user_mode_loop+0x85/0x510 kernel/entry/common.c:98 | do_syscall_64+0x263/0x3d0 arch/x86/entry/syscall_64.c:100 | entry_SYSCALL_64_after_hwframe+0x77/0x7f | | The buggy address belongs to the object at ffff8881298d9400 | which belongs to the cache kmalloc-512 of size 512 | The buggy address is located 336 bytes inside of | freed 512-byte region [ffff8881298d9400, ffff8881298d9600) Fix it by having chan->conn hold a reference to l2cap_conn (via l2cap_conn_get) when the channel is added to the connection, and releasing it in the channel destructor. This ensures the l2cap_conn remains alive as long as the channel exists. A new FLAG_DEL channel flag is introduced to indicate that the ch ---truncated---	8.8	More Details
CVE-2026-65595	n8n before 2.30.1 and 2.29.8 assigns all Public API key scopes to JWTs issued through the Token Exchange module regardless of the acting user's role. On instances where the Token Exchange feature and Public API are enabled, a low-privileged user who can obtain a valid external JWT trusted by a configured issuer can use the resulting access token to invoke administrator-only Public API operations such as role escalation, user creation, and user deletion (role escalation requires an Advanced Permissions license), and, when unverified Community Package installation is enabled, achieve remote code execution.	8.8	More Details
CVE-2026-65016	n8n versions before 1.123.64, 2.29.8, and 2.30.1 contain a privilege escalation vulnerability in Enterprise SSO instance-role provisioning. The provisioning path maps an IdP-asserted role claim to an n8n global role but does not prevent assignment of the global:owner role (unlike the token-exchange identity path, which rejects it). An SSO-authenticated user whose instance-role claim resolves to global:owner is provisioned as instance owner, gaining full administrative control over workflows, credentials, users, and instance configuration. Exploitation requires that Enterprise SSO is configured, instance-role provisioning is enabled via N8N_SSO_SCOPES_PROVISION_INSTANCE_ROLE (disabled by default), and the attacker controls the instance-role claim value issued by the IdP.	8.8	More Details
CVE-2026-64390	In the Linux kernel, the following vulnerability has been resolved: ksmbd: track the connection owning a byte-range lock SMB2_LOCK adds each granted byte-range lock to both the file lock list and the lock list of the connection which handled the request. The final close and durable handle paths, however, remove the connection list entry while holding fp->conn->llist_lock. With SMB3 multichannel, the connection handling the LOCK request can be different from the connection which opened the file. The entry can therefore be removed under a different spinlock from the one protecting the list it belongs to. A concurrent traversal can then access freed struct ksmbd_lock and struct file_lock objects. Record the connection owning each lock's clist entry and hold a reference to it while the entry is linked. Use that connection and its llist_lock for unlock, rollback, close, and durable preserve. Durable reconnect assigns the new connection as the owner when publishing the locks again.	8.8	More Details
CVE-2026-65603	The Grav Login plugin (grav-plugin-login) versions <= 3.8.11 contain a privilege escalation flaw in the authenticated profile self-update handler (processUserProfile(), the update_user task). Unlike the registration handler, this handler does not strip privilege fields ('groups','access') from user-submitted form data before persisting them. When an administrator has added 'groups' and/or 'access' to plugins.login.user_registration.fields and the default 'regular'/DataUser account backend is in use, a low-privilege authenticated user can POST crafted profile form data (e.g. access[admin][super]=true) to escalate to super-admin, enabling admin panel access, scheduler abuse (RCE), and Twig evaluation. Fixed in 3.8.12.	8.8	More Details
CVE-2026-16807	Out of bounds write in Codecs in Google Chrome prior to 150.0.7871.186 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-16805	Use after free in Blink in Google Chrome prior to 150.0.7871.186 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized		

CVE-2026-60439	Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	8.8	More Details
CVE-2026-17568	Improper access control in the role membership management endpoint in Devolutions Server allows an authenticated non-administrative user holding the user-group membership management permission to escalate privileges to administrator via a crafted API request. This issue affects : * Devolutions Server 2026.2.4.0 through 2026.2.12.0 * Devolutions Server 2026.1.23.0 and earlier	8.8	More Details
CVE-2026-63684	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in various admin/import/export actions of multiple Regular Labs extension - Administrator actions, editor popups and import/export requests lacked consistent token, item-permission and input-validation checks. Unauthorized backend users or CSRF attacks could expose, create or modify extension configuration and items.	8.8	More Details
CVE-2026-16806	Use after free in WebMCP in Google Chrome prior to 150.0.7871.186 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	8.8	More Details
CVE-2026-60455	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	8.8	More Details
CVE-2026-63685	Joomla Extension - regularlabs.com - Authorization bypass in DB Replacer extension - Administrator routes and replacement requests did not consistently require Super User permission and a valid token. An unauthorized backend user or CSRF attack could perform database replacements, potentially causing major data corruption or site compromise.	8.8	More Details
CVE-2026-63280	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in Regular Labs conditions manager - Conditions administration did not consistently enforce tokens and component/mapped-item permissions.	8.8	More Details
CVE-2026-15992	The WP Password Policy plugin for WordPress is vulnerable to Privilege Escalation in all versions up to and including 3.7.1. This is due to missing authorization checks and nonce verification in the `get_user()` function of the `Module_Password_Hint` class, which unconditionally calls `WP_User::set_role()` with the attacker-supplied `role` parameter on any account resolved via `$_POST['user_login']`, without confirming the requesting user holds the capability to assign roles. This makes it possible for authenticated attackers, with subscriber-level access and above, to escalate their own privileges to Administrator by submitting a crafted POST request — with `action` set to `createuser` and `role` set to `administrator` — to the password-reset form endpoint. The vulnerable code path is reachable via the `password_hint` filter hooked during the WordPress password-reset form render, meaning an attacker need only possess a valid password-reset cookie to reach the sink.	8.8	More Details
CVE-2026-66748	Camaleon CMS versions 2.1.1 through 2.9.1 contains an authenticated remote code execution vulnerability that allows users with custom_fields manage permission to execute arbitrary Ruby code by supplying a malicious expression through the select_eval custom field type. Attackers can store an attacker-controlled Ruby expression in the field options command parameter, which is evaluated via instance_eval within an ERB view whenever a post edit page is rendered, achieving server-side code execution with web server process privileges.	8.8	More Details
CVE-2026-64382	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix double-free in SMB2_open() replay A response-bearing attempt can return a replayable error and free its response buffer. If SMB2_open_init() fails before the next send, cleanup retains the previous buffer type and frees that response again. Reset response bookkeeping before each attempt to prevent the stale free.	8.8	More Details
CVE-2026-64791	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in Regular Labs Extension Manager - Administrator routes and install/update/uninstall processing did not consistently enforce component-management and installation permissions. An unauthorized backend user or CSRF attack could install, update or remove extensions.	8.8	More Details
CVE-2026-64522	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Fix eswitch mode block underflow on IPsec acquire SA mlx5e_xfrm_add_state() handles acquire-flow temporary SAs by allocating software state and skipping hardware offload setup. That path jumps to the common success label before taking the eswitch mode block. After tunnel-mode validation was moved earlier, the common success label unconditionally calls mlx5_eswitch_unblock_mode(). For acquire SAs, this decrements esw->offloads.num_block_mode without a matching increment. Return directly after installing the acquire SA offload handle, so only the paths that successfully called mlx5_eswitch_block_mode() call the matching unblock.	8.8	More Details
CVE-2026-15064	IBM WebSphere Application Server 9.0, and 8.5 and IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.7 is vulnerable to HTTP Response Smuggling due to improper handling of non-standard HTTP version tokens.	8.7	More Details
CVE-2026-	SiYuan before v3.7.3 contains stored and reflected cross-site scripting vulnerabilities in SVG sanitization that allows authenticated attackers to execute scripts by bypassing the HTML parser-based cleaner. Attackers can hide script	8.7	More

66394	tags within desc, style, or noscript elements which the HTML parser treats as raw text but browsers interpret as executable SVG content when served as image/svg+xml, enabling script execution in the application origin.		Details
CVE-2026-15325	IBM WebSphere Application Server 9.0, and 8.5 and IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.7 is vulnerable to HTTP request smuggling due to improper handling of TRACE requests.	8.7	More Details
CVE-2026-47743	Shopper is a Headless e-commerce Admin Panel. Prior to 2.8.0, three related defects on admin Livewire components allowed data tampering, sensitive data disclosure, and stored XSS. First, several Livewire components in the admin panel exposed Eloquent model identifiers as public properties without the `#[Locked]` attribute. An authenticated user could rewrite the wire payload from the browser to target any record id, bypassing the implicit scoping enforced by the page routing. Second, `Customers/Create::store()` re-passed a `Hidden` `_password` form field straight into the create payload. The plaintext password was rendered into the HTML and transported through the Livewire snapshot in clear text, exposing credentials in the page DOM and in any logging that captures Livewire payloads. Finally, the product barcode field was rendered through `DNS1DFacade::getBarcodeHTML()` with `{!! !!}`. An attacker with `edit_products` permission could persist malicious payload in the barcode field that would execute in the browser of any admin user viewing that product, enabling session theft and privileged-action chaining. Starting in v2.8.0, all vulnerable Livewire model identifiers are now marked `#[Locked]`; `Customers/Create` no longer round-trips the password through a Hidden form field; the plaintext password is hashed at action boundary and never returned to the client; and the product barcode rendering now escapes the value before passing it to the barcode generator and the output is wrapped in an `<svg>` context that does not interpret event handlers. No known workarounds are available.	8.7	More Details
CVE-2026-28698	Pronetiqs IntraVUE versions 3.2.1a14 and prior have an exposure of sensitive system information to an unauthorized control sphere vulnerability which could expose the underlying host/share filesystem.	8.6	More Details
CVE-2026-28973	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, watchOS 26.6. A malicious app may be able to break out of its sandbox.	8.6	More Details
CVE-2025-15662	The Printcart Web to Print Product Designer for WooCommerce WordPress plugin before 2.5.3 does not restrict a user-supplied URL before fetching it server-side and does not enforce a valid authorization check, allowing unauthenticated attackers to read arbitrary local files (including configuration files containing database credentials and secret keys) and to make server-side requests to internal resources.	8.6	More Details
CVE-2026-64814	In JetBrains IntelliJ IDEA before 2026.2 unauthorized file access was possible in a Remote Development session	8.6	More Details
CVE-2026-54650	openhole exposes localhost to the internet in one command. In 0.1.1 and earlier, openhole-server in internal/server/public_proxy.go forwarded r.URL.Path instead of preserving the original request target with r.URL.EscapedPath(), allowing percent encoded dot segments %2e and separators %2f to reach tunneled local services as ../ and / for path traversal. This issue is fixed in version 0.1.2.	8.6	More Details
CVE-2026-65908	In JetBrains PyCharm before 2026.1.4, 2026.2 arbitrary code execution via malicious Python executable was possible on untrusted project open	8.6	More Details
CVE-2026-54609	QTI Neon is a minimal, game-agnostic, relay-based UDP multiplayer protocol library. In version 1.0.0, the relay's handleReconnectRequest forwards RECONNECT_REQUEST packets to the host without bounding them, so an unauthenticated client can drive relay-to-host amplification and cause a denial of service on the host. No fixed version is available as of this review.	8.6	More Details
CVE-2026-13321	The BIND resolver accepts validly-signed NSEC records where the "Next Domain Name" field points outside the signer's zone. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	8.6	More Details
CVE-2026-45293	WordPress Coding Standards is a set of PHP_CodeSniffer rules (sniffs) that enforce WordPress coding conventions. From 0.14.1 until 3.4.1, the WordPress.WP.EnqueuedResourceParameters sniff (active in the WordPress and WordPress-Extra rulesets) reconstructed the $ver argument passed to functions such as wp_enqueue_script() and ran it through eval() inside its is_falsy() method, so a maliciously crafted argument such as 'system'('id') would execute during a scan; as a result, running PHPCS with WordPressCS over untrusted PHP (for example a CI pipeline that lints pull requests, or a developer reviewing third-party code) could lead to arbitrary command execution on the scanning host. The WordPress-Core and WordPress-Docs rulesets are not affected. This issue is fixed in version 3.4.1.	8.6	More Details
CVE-2026-48396	Bridge is affected by an Incorrect Authorization vulnerability that could result in arbitrary code execution in the context of the current user. An attacker could exploit this vulnerability to execute arbitrary code. Exploitation of this issue requires user interaction in that a victim must open a malicious file. Scope is changed.	8.6	More Details
CVE-2026-48388	Adobe Photoshop Installer was affected by an Uncontrolled Search Path Element vulnerability that could have resulted in arbitrary code execution in the context of the current user. An attacker could have exploited this vulnerability by placing a malicious library in a directory searched by the installer. Exploitation of this issue required user interaction in that a victim must have been running the installer. Scope is changed.	8.6	More Details
CVE-	The terraform-mcp-server before version 1.1.0 is vulnerable to a server-side request forgery issue in the streamable-HTTP transport that may allow an unauthenticated remote client to redirect the server's Terraform API requests, and		More

2026-14869	the server-side authorization token, to an attacker-controlled endpoint. This vulnerability, CVE-2026-14869, is fixed in terraform-mcp-server 1.1.0.	8.6	Details
CVE-2026-48395	Bridge is affected by an Untrusted Search Path vulnerability that could result in arbitrary code execution in the context of the current user. An attacker could exploit this vulnerability to execute arbitrary code. Exploitation of this issue requires user interaction in that a victim must open a malicious file. Scope is changed.	8.6	More Details
CVE-2026-54603	OAuth2 is a Ruby wrapper for the OAuth 2.0 and 2.1 authorization frameworks, including OpenID Connect (OIDC). From 0.4.0 to 2.0.21, a protocol-relative redirect Location returned to OAuth2::Client#request overrides the request authority, so the bearer Authorization header is sent to an attacker-controlled host, leaking the credential. This issue is fixed in version 2.0.22.	8.6	More Details
CVE-2026-65702	Vanna through 2.0.2 contains a path traversal vulnerability in the FileSystemConversationStore persistence integration that allows unauthenticated remote attackers to write attacker-controlled JSON files to arbitrary filesystem locations and read conversation metadata from outside the intended store base directory. Attackers can supply path traversal sequences in the conversation_id parameter submitted to the unauthenticated chat API endpoints to escape the base directory during both write and read operations, enabling arbitrary file write with attacker-controlled content and unauthorized file read on the server filesystem.	8.6	More Details
CVE-2026-64400	In the Linux kernel, the following vulnerability has been resolved: ksmbd: prevent path traversal bypass by restricting caseless retry ksmbd_vfs_path_lookup() enforces LOOKUP_BENEATH to restrict path resolution within the share root. When a crafted path attempts to escape the share boundary using parent-directory components ('..'), vfs_path_parent_lookup() detects this and immediately fails, returning -EXDEV. However, a bug exists in __ksmbd_vfs_kern_path() under caseless mode. The function fails to intercept the -EXDEV error and erroneously falls through to the caseless retry logic, which is intended only for genuinely missing files. During this retry process, the path is reconstructed, leading to an unintended LOOKUP_BENEATH bypass that allows write-capable users to create zero-length files or directories outside the exported share. Fix this by ensuring that the execution only proceeds to the caseless lookup retry when the error is specifically -ENOENT. Any other errors, such as -EXDEV from a path traversal attempt, must be returned immediately.	8.6	More Details
CVE-2026-25405	Contributor SQL Injection in eRoom <= 1.7.1 versions.	8.5	More Details
CVE-2026-65451	Contributor SQL Injection in MapSVG <= 8.14.0 versions.	8.5	More Details
CVE-2026-65454	Contributor SQL Injection in Quiz And Survey Master <= 11.2.0 versions.	8.5	More Details
CVE-2026-59551	Subscriber SQL Injection in rtMedia for WordPress, BuddyPress and bbPress <= 4.7.10 versions.	8.5	More Details
CVE-2026-17192	A VCO feature does not sufficiently validate caller-supplied input, allowing requests to be made on behalf of authenticated tenant accounts to internal services that are not otherwise accessible. This vulnerability requires a minimum role of Enterprise Standard Admin. This issue was discovered internally by Arista and the company is not aware of any malicious uses of this issue in customer networks.	8.5	More Details
CVE-2026-56167	Server-side request forgery (ssrf) in Azure AI Search allows an authorized attacker to elevate privileges over a network.	8.5	More Details
CVE-2026-65450	Contributor SQL Injection in MapSVG <= 8.14.0 versions.	8.5	More Details
CVE-2026-65526	Contributor SQL Injection in Visualizer <= 4.0.6 versions.	8.5	More Details
CVE-2026-49332	A flaw was found in openshift/oauth-proxy. The proxy sets authenticated identity headers using only dash-variant keys (X-Forwarded-User) but does not strip underscore-variant keys (X_Forwarded_User) from incoming requests. WSGI and PHP frameworks normalize both variants to the same variable, allowing an authenticated low-privilege user to smuggle a forged identity that may override the legitimate authenticated identity in the upstream application.	8.5	More Details
CVE-2026-65895	Grav API Plugin versions before 1.0.10 fail to restrict write access to security-critical plugin configuration scopes, allowing authenticated users with api.config.write privilege to modify rate limiting and CORS settings. Attackers can disable rate limiting site-wide to enable credential brute-forcing attacks and reconfigure CORS policies to include attacker-controlled origins with credentials enabled.	8.5	More Details
CVE-2026-24552	Contributor SQL Injection in Create by Mediavine <= 2.5.3 versions.	8.5	More Details

CVE-2026-17107	A flaw was found in the cluster-proxy service-proxy component used in Red Hat Advanced Cluster Management for Kubernetes (RHACM) and multicluster-engine (MCE). The service-proxy appends impersonation group headers to proxied requests without first removing caller-supplied values, and the spoke ServiceAccount holds unrestricted impersonation permissions. An authenticated hub principal can inject an Impersonate-Group header to escalate to cluster-admin on every managed cluster.	8.5	More Details
CVE-2026-59686	An OS Command Injection vulnerability in Progress Software LoadMaster, ECS Connection Manager, Object Scale Connection Manager, and MOVEit WAF allows an authenticated attacker with high privileges to execute arbitrary operating system commands on the affected appliance via the management interface, potentially resulting in complete system compromise.	8.4	More Details
CVE-2026-59688	An OS Command Injection vulnerability in Progress Software LoadMaster, ECS Connection Manager, Object Scale Connection Manager, and MOVEit WAF allows an authenticated attacker with high privileges to execute arbitrary operating system commands on the affected appliance via the backup restore functionality, potentially resulting in complete system compromise.	8.4	More Details
CVE-2026-59687	An OS Command Injection vulnerability in Progress Software LoadMaster, ECS Connection Manager, Object Scale Connection Manager, and MOVEit WAF allows an authenticated attacker with high privileges to execute arbitrary operating system commands on the affected appliance via the Geo Location management interface, potentially resulting in complete system compromise.	8.4	More Details
CVE-2026-64804	In JetBrains WebStorm before 2026.2 arbitrary code execution was possible before granting project trust via project-local linter tooling	8.4	More Details
CVE-2024-58023	Information disclosure in Bosch Configuration Manager in Version 7.72.0106 allows an attacker to access sensitive information.	8.4	More Details
CVE-2026-64805	In JetBrains WebStorm before 2026.2 arbitrary code execution was possible before granting project trust via project-local package-manager tooling	8.4	More Details
CVE-2026-64520	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_ffa: Bound PARTITION_INFO_GET_REGS copies The register-based PARTITION_INFO_GET path trusted the firmware-provided indices when copying partition descriptors into the caller buffer. Reject inconsistent counts or index progressions so the copy loop cannot write past the allocated array. (fixed cur_idx when exactly one descriptor in the first fragment)	8.4	More Details
CVE-2026-66396	SiYuan before v3.7.2 fails to escape the title-img Individual Attribute List value when rendering Gallery and Kanban cover images, allowing stored cross-site scripting via unescaped style attribute interpolation. Attackers with editor permissions can inject onload handlers that execute arbitrary code in the Electron renderer with full Node.js access when victims open affected documents.	8.4	More Details
CVE-2026-64808	In JetBrains PhpStorm before 2026.2 arbitrary code execution was possible before granting project trust via project tooling	8.4	More Details
CVE-2026-64806	In JetBrains WebStorm before 2026.2 arbitrary code execution was possible before granting project trust via the configured Node.js interpreter	8.4	More Details
CVE-2026-64809	In JetBrains PhpStorm before 2026.2 arbitrary code execution was possible before granting project trust via the configured interpreter	8.4	More Details
CVE-2026-66140	Exim before 4.99.5 allows directory traversal to access files outside of the spool area, and consequently gain privileges, because arguments related to queue-name are mishandled.	8.4	More Details
CVE-2026-64490	In the Linux kernel, the following vulnerability has been resolved: ALSA: virtio: Validate control metadata from the device virtio-snd control handling trusts the device-provided control type and value count returned by the device. That metadata is then used directly to index g_v2a_type_map[] in virtsnd_kctl_info(), and to size loops and memcpy() operations in virtsnd_kctl_get() and virtsnd_kctl_put() against fixed-size virtio_snd_ctl_value and snd_ctl_elem_value arrays. A buggy or malicious device can therefore trigger out-of-bounds access by advertising an invalid control type or an oversized value count. Validate control type and count once in virtsnd_kctl_parse_cfg(), before querying enumerated items or exposing the control to ALSA.	8.4	More Details
	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: hyper-v: Bound the bank index when querying sparse banks When checking if a VP ID is included in a sparse bank set, explicitly check that the ID can actually be contained in a sparse bank (the TLFS allows for a maximum of 64 banks of 64 vCPUs each). When handling a paravirtual TLB flush for L2, the VP ID is copied verbatim from the enlightened VMCS, without any bounds check, i.e. isn't guaranteed to be under the limit of 4096. Failure to check the bounds of the VP ID leads to an out-of-bounds read when testing the sparse bank, and super strictly speaking could lead to KVM performing an unnecessary TLB flush for an L2 vCPU. == BUG: KASAN: use-after-free in hv_is_vp_in_sparse_set+0x85/0x100 [kvm] Read of size 8 at addr ffff88811ba5f598 by task hyperv_evmcs/2802 CPU: 12 UID: 1000 PID: 2802 Comm: hyperv_evmcs Not tainted 7.1.0-rc2 #7 PREEMPT Hardware		

CVE-2026-64247	name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 0.0.0 02/06/2015 Call Trace: <TASK> dump_stack_lvl+0x51/0x60 print_report+0xcb/0x5d0 kasan_report+0xb4/0xe0 kasan_check_range+0x35/0x1b0 hv_is_vp_in_sparse_set+0x85/0x100 [kvm] kvm_hv_flush_tlb+0xe9e/0x16c0 [kvm] kvm_hv_hypercall+0xe6b/0x1e60 [kvm] vmx_handle_exit+0x485/0x1b60 [kvm_intel] kvm_arch_vcpu_ioctl_run+0x22e3/0x5070 [kvm] kvm_vcpu_ioctl+0x5d0/0x10c0 [kvm] __x64_sys_ioctl+0x129/0x1a0 do_syscall_64+0xb9/0xcf0 entry_SYSCALL_64_after_hwframe+0x4b/0x53 RIP: 0033:0x7f0e62d1a9bf </TASK> The buggy address belongs to the physical page: page: refcount:0 mapcount:0 mapping:0000000000000000 index:0xffffffffffffffff pfn:0x11ba5f flags: 0x4000000000000000(zone=1) raw: 4000000000000000 0000000000000000 00000000ffffffff 0000000000000000 raw: ffffffffffffffff 0000000000000000 00000000ffffffff 0000000000000000 page dumped because: kasan: bad access detected Memory state around the buggy address: ffff88811ba5f480: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ffff88811ba5f500: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff >ffff88811ba5f580: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ^ ffff88811ba5f600: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ffff88811ba5f680: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ========== Disabling lock debugging due to kernel taint Opportunistically add a compile time assertion to ensure the maximum number of sparse banks exactly matches the number of possible bits in the passed in mask. [sean: add KASAN splat, drop comment, add assert, massage changelog]	8.4	More Details
CVE-2026-64515	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fix MLE defragmentation If either reconf or EPCS multi-link element (MLE) is contained in a non-transmitted profile, the defragmentation routine is called with a pointer to the defragmented copy, but the original elements. This is incorrect for two reasons: - if the original defragmentation was needed, it will not find the correct data - if the original frame is at a higher address, the parsing will potentially overrun the heap data (though given the layout of the buffers, only into the new defragmentation buffer, and then it has to stop and fail once that's filled with copied data. Fix it by tracking the container along with the pointer and in doing so also unify the two almost identical defragmentation routines.	8.3	More Details
CVE-2026-66027	Suna before 0.9.102 contains a broken access control vulnerability in the message queue API that allows authenticated attackers to access and manipulate queue resources belonging to other users by exploiting missing ownership and account isolation checks. Attackers can read pending prompt queues of all users, read or delete individual sessions, and inject arbitrary prompts into another user's session queue, causing the background drainer to forward malicious messages to the victim's running AI agent with the victim's credentials and permissions.	8.3	More Details
CVE-2026-16804	Use after free in Input in Google Chrome prior to 150.0.7871.186 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	8.3	More Details
CVE-2026-17497	NoteGen before 0.32.0 grants the Tauri shell plugin shell:allow-execute capability for bash, python, and python3 with arbitrary arguments in the default desktop capabilities. JavaScript running in the application webview can therefore invoke plugin:shell|execute to run attacker-controlled operating system commands with the privileges of the NoteGen process. In combination with script execution in the webview (for example via chat XSS), this enables full remote code execution on the user's machine.	8.3	More Details
CVE-2026-65709	sysPass through version 3.2.11 contains a missing object-level authorization vulnerability in the JSON-RPC API that allows API token holders to enumerate account metadata, overwrite passwords, and delete accounts across the entire vault without per-account access control. Attackers can invoke AccountController methods such as viewAction, editAction, deleteAction, and editPassAction without AccountFilterUser checks to modify or delete accounts beyond the scope of their assigned token permissions.	8.3	More Details
CVE-2026-47483	NVIDIA DCGM Exporter for all platforms contains a vulnerability in the /debug/pprof endpoints, where an attacker could cause uncontrolled resource consumption by submitting concurrent unauthenticated profiling requests. A successful exploit of this vulnerability might lead to denial of service and information disclosure.	8.2	More Details
CVE-2026-43910	Appium Java Client is the Java language binding for writing Appium tests that conform to the W3C WebDriver protocol. From 8.2.1 until 10.1.1, when directConnect(true) is enabled, AppiumCommandExecutor.setDirectConnect() reads the directConnectHost, directConnectPort, and directConnectPath fields from the server's NEW_SESSION response and rebuilds the client's server URL from them, validating only that the protocol is https, with no host allowlist or IP validation; a rogue or compromised server can therefore redirect all subsequent session traffic to an arbitrary destination, enabling full interception of session traffic and a server-side request forgery pivot to internal hosts, including cloud metadata (IMDS) credential theft. This vulnerability is fixed in 10.1.1.	8.2	More Details
CVE-2026-48391	Bridge is affected by an Untrusted Search Path vulnerability that could result in arbitrary code execution in the context of the current user. A low-privileged attacker could exploit this vulnerability to execute arbitrary code. Exploitation of this issue requires user interaction in that a victim must open a malicious file. Scope is changed.	8.2	More Details
CVE-2026-48390	Bridge is affected by an Incorrect Authorization vulnerability that could result in privilege escalation. An attacker could leverage this vulnerability to gain unauthorized read and write access. Exploitation of this issue requires user interaction in that a victim must open a malicious file. Scope is changed.	8.2	More Details
CVE-2026-65604	Skipper contains an incomplete fix for CVE-2026-50197 in which oversized request bodies bypass Open Policy Agent (OPA) deny-on-presence Rego policies. When a request body exceeds the configured maxBodyBytes limit, Skipper forwards the full payload to the upstream service while OPA evaluates against an empty parsed_body, so policies that deny requests based on body content are not enforced and forbidden actions proceed. No fixed version is available; v0.27.26 adds documentation guidance only.	8.2	More Details
CVE-2026-9830	The bookingpress-appointment-booking-pro WordPress plugin before 5.7.3 does not correctly invoke its REST permission callback, leaving every route in one of its API namespaces reachable without authentication and allowing unauthenticated attackers to read customer booking data and modify other users' bookings.	8.2	More Details

CVE-2026-16002	The affected product is vulnerable to an Out-of-bounds read, which may allow an attacker to crash the parsing process and cause a denial of service.	8.2	More Details
CVE-2026-63765	Chatwoot before 4.16.0 contains an authentication bypass vulnerability in the direct uploads controller that allows unauthenticated attackers to create arbitrary ActiveStorage blobs in any tenant account. Attackers can exploit missing authentication checks to resolve any account and conversation, then obtain signed PUT URLs to write arbitrary data to the application's storage backend.	8.2	More Details
CVE-2026-64435	In the Linux kernel, the following vulnerability has been resolved: audit: Fix data races of skb_queue_len() readers on audit_queue Multiple readers access audit_queue.qlen via skb_queue_len() without holding the queue lock or using READ_ONCE(), while kauditd writes to this field via the skb_dequeue() → __skb_unlink() path with WRITE_ONCE() protected by a spinlock. This constitutes data races. All affected skb_queue_len(&audit_queue) call sites: - kauditd_thread() wait_event_freezable() condition - audit_receive_msg() AUDIT_GET handler (s.backlog assignment) - audit_receive() backlog check - audit_log_start() backlog check and pr_warn() KCSAN reports the following conflicting access pattern (one example): ========== BUG: KCSAN: data-race in audit_log_start / skb_dequeue write (marked) to 0xffffffff8512ee20 of 4 bytes by task 661 on cpu 57: skb_dequeue+0x70/0xf0 kauditd_send_queue+0x71/0x220 kauditd_thread+0x1cb/0x430 kthread+0x1c2/0x210 ret_from_fork+0x162/0x1a0 ret_from_fork_asm+0x1a/0x30 read to 0xffffffff8512ee20 of 4 bytes by task 36586 on cpu 1: audit_log_start+0x2a0/0x6b0 audit_core_dumps+0x64/0xa0 do_coredump+0x14b/0x1260 get_signal+0xeb2/0xf70 arch_do_signal_or_restart+0x41/0x170 exit_to_user_mode_loop+0xa2/0x1c0 do_syscall_64+0x1a3/0x1c0 entry_SYSCALL_64_after_hwframe+0x76/0xe0 value changed: 0x00000001 -> 0x00000000 ========== Resolve the race by switching to lockless helper skb_queue_len_lockless(), which internally uses READ_ONCE() and properly pairs with the WRITE_ONCE() write accesses already present on the writer side. [PM: line length tweak]	8.2	More Details
CVE-2026-64389	In the Linux kernel, the following vulnerability has been resolved: ksmbd: validate NTLMv2 response before updating session key ksmbd_auth_ntlmv2() derives the NTLMv2 session key into sess->sess_key before it verifies the NTLMv2 response. ksmbd_decode_ntlmssp_auth_blob() then continues into KEY_XCH even when ksmbd_auth_ntlmv2() failed. With SMB3 multichannel binding, the failed authentication operates on an existing session and the session setup error path does not expire binding sessions. A client can send a binding session setup with a bad NT proof and KEY_XCH and still modify sess->sess_key before STATUS_LOGON_FAILURE is returned. Relevant path: smb2_sess_setup() -> conn->binding = true -> ntlm_authenticate() -> session_user() -> ksmbd_decode_ntlmssp_auth_blob() -> ksmbd_auth_ntlmv2() -> calc_ntlmv2_hash() -> hmac_md5_usingrawkey(..., sess->sess_key) -> crypto_memneq() returns mismatch -> KEY_XCH arc4_crypt(..., sess->sess_key, ...) -> out_err without expiring the binding session Derive the base session key into a local buffer and copy it to sess->sess_key only after the proof matches. Return immediately on authentication failure so KEY_XCH is only processed after successful authentication.	8.2	More Details
CVE-2026-54690	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. From 0.9.1 until 0.61.0, datamodel-code-generator silently dereferences attacker-controlled JSON Schema $ref HTTP or HTTPS URLs in src/datamodel_code_generator/parser/jsonschema.py through _get_ref_body, and the --allow-remote-refs gate can warn instead of blocking, allowing server-side request forgery through src/datamodel_code_generator/http.py. This issue is fixed in version 0.61.0.	8.2	More Details
CVE-2026-64380	In the Linux kernel, the following vulnerability has been resolved: smb: client: harden POSIX SID length parsing posix_info_sid_size() reads sid[1] to obtain the subauthority count, but its existing boundary check still accepts buffers with only one remaining byte. Require two bytes before reading sid[1] so all client paths that reuse the helper reject truncated POSIX SIDs safely.	8.2	More Details
CVE-2026-64737	An authorization issue was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious app may be able to break out of its sandbox.	8.2	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: smb: client: restrict implied bcc[0] exemption to responses without data area smb2_check_message() has a long-standing quirk that accepts a response whose calculated length is one byte larger than the bytes actually received ("server can return one byte more due to implied bcc[0]"). This was introduced to accommodate servers that omit the trailing bcc[0] overlap byte when no data area is present. However, the exemption is applied unconditionally, regardless of whether the command actually carries a data area (has_smb2_data_area[]). When a response with a data area is subject to the +1 exemption, the reported data can extend one byte beyond the bytes actually received, yet smb2_check_message() still accepts it. The subsequent decoder then reads past the end of the receive buffer. This is reachable during NEGOTIATE and SESSION_SETUP, before the session is established. The resulting out-of-bounds reads are visible under KASAN when mounting against a non-conforming server; both the SPNEGO/negTokenInit and the NTLMSSP challenge decoders are affected: BUG: KASAN: slab-out-of-bounds in asn1_ber_decoder+0x16a7/0x1b00 Read of size 1 at addr ffff8880084d67c0 by task mount.cifs/81 CPU: 1 UID: 0 PID: 81 Comm: mount.cifs Not tainted 7.1.0-rc6 #1 Call Trace: <TASK> dump_stack_lvl+0x4e/0x70 print_report+0x157/0x4c9 kasan_report+0xce/0x100 asn1_ber_decoder+0x16a7/0x1b00 decode_negTokenInit+0x19/0x30 SMB2_negotiate+0x31d9/0x4c90 cifs_negotiate_protocol+0x1f2/0x3f0 cifs_get_smb_ses+0x93f/0x17e0 cifs_mount_get_session+0x7f/0x3a0 cifs_mount+0xb4/0xcf0 cifs_smb3_do_mount+0x23a/0x1500 smb3_get_tree+0x3b0/0x630 vfs_get_tree+0x82/0x2d0 fc_mount+0x10/0x1b0 path_mount+0x50d/0x1de0 __x64_sys_mount+0x20b/0x270 do_syscall_64+0xee/0x590 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 85: kmem_cache_alloc_noprof+0x106/0x380 mempool_alloc_noprof+0x116/0x1e0 cifs_small_buf_get+0x31/0x80		More

2026-64448	allocate_buffers+0x10d/0x2b0 cifs_demultiplex_thread+0x1d5/0x1d50 kthread+0x2c6/0x390 ret_from_fork+0x36e/0x5a0 ret_from_fork_asm+0x1a/0x30 The buggy address is located 0 bytes to the right of allocated 448-byte region [ffff8880084d6600, ffff8880084d67c0) which belongs to the cache cifs_small_rq of size 448 BUG: KASAN: slab-out-of-bounds in kmemdup_noprof+0x36/0x50 Read of size 329 at addr ffff88800726c678 by task mount.cifs/89 CPU: 0 UID: 0 PID: 89 Comm: mount.cifs Tainted: G B 7.1.0-rc6 #1 Call Trace: <TASK> dump_stack_lvl+0x4e/0x70 print_report+0x157/0x4c9 kasan_report+0xce/0x100 kasan_check_range+0x10f/0x1e0 __asan_memcpy+0x23/0x60 kmemdup_noprof+0x36/0x50 decode_ntlmssp_challenge+0x457/0x680 SMB2_sess_auth_rawntlmssp_negotiate+0x6f0/0xcb0 SMB2_sess_setup+0x219/0x4f0 cifs_setup_session+0x248/0xaf0 cifs_get_smb_ses+0xf79/0x17e0 cifs_mount_get_session+0x7f/0x3a0 cifs_mount+0xb4/0xcf0 cifs_smb3_do_mount+0x23a/0x1500 smb3_get_tree+0x3b0/0x630 vfs_get_tree+0x82/0x2d0 fc_mount+0x10/0x1b0 path_mount+0x50d/0x1de0 __x64_sys_mount+0x20b/0x270 do_syscall_64+0xee/0x590 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 93: kmem_cache_alloc_noprof+0x106/0x380 mempool_alloc_noprof+0x116/0x1e0 cifs_small_buf_get+0x31/0x80 allocate_buffers+0x10d/0x2b0 cifs_demultiplex_thread+0x1d5/0x1d50 kthread+0x2c6/0x390 ret_from_fork+0x36e/0x5a0 ret_from_fork_asm+0x1a/0x30 The buggy address is located 120 bytes inside of allocated 448-byte region [ffff88800726c600, ffff88800726c7c0) which belongs to the cache cifs_small_rq of size 448 Restrict the +1 exemption to responses that have no data area, so that it still covers the bcc[0] omission it was meant for. When a data area is present, the +1 discrepancy instead means the reported data length overruns the --- truncated---	8.2	Details
CVE-2026-43772	A path traversal issue was addressed with improved input validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to break out of its sandbox.	8.2	More Details
CVE-2026-64286	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Clear __hyp_running_vcpu when flushing the pKVM hyp vCPU flush_hyp_vcpu() copies the host vCPU context into the hyp's private vCPU on every run. ctxt_to_vcpu() expects a guest context to have a NULL __hyp_running_vcpu, which is only ever set on the host context, so that it resolves the vCPU via container_of(). While this is generally the case, flush_hyp_vcpu() copies the context verbatim and does not enforce this, so a value provided by the host is dereferenced at EL2 (host -> EL2). Fix by clearing __hyp_running_vcpu after the copy.	8.2	More Details
CVE-2026-64287	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Bound used_lrs when flushing the pKVM hyp vCPU flush_hyp_vcpu() copies the host vGIC state into the hyp's private vCPU on every run. The vGIC list register save and restore use used_lrs as their loop bound and expect it to stay within the number of implemented list registers. While this is generally the case, flush_hyp_vcpu() copies vgic_v3 verbatim and does not enforce this, so a value provided by the host is used at EL2 to index vgic_lr[] and access ICH_LR<n>_EL2 (host -> EL2). Fix by clamping used_lrs to the number of implemented list registers after the copy, as the trusted path already does in vgic_flush_lr_state(). The number of implemented list registers is constant after init, so it is replicated once from kvm_vgic_global_state.nr_lr into hyp_gicv3_nr_lr rather than read on every entry.	8.2	More Details
CVE-2026-14996	IBM Aspera Faspex 5 5.0.0 through 5.0.15.4 has addressed a vulnerability related to session management.	8.2	More Details
CVE-2026-54691	datamodel-code-generator generates Python data models from schema definitions. From 0.9.1 until 0.61.0, src/datamodel_code_generator/http.py http.get_body accepts --url targets and redirect chain targets without host/IP validation, allowing server-side request forgery against loopback, private, link-local, metadata, and other network-accessible resources. This issue is fixed in version 0.61.0.	8.2	More Details
CVE-2026-64815	In JetBrains IntelliJ IDEA before 2026.2 arbitrary code injection was possible via UI Designer form files	8.1	More Details
CVE-2026-14169	Due to incorrect behavior order a low privileged remote attacker could trigger account inconsistent state via crafted input and overwrites existing user passwords which could result in complete administrative unavailability of the device.	8.1	More Details
CVE-2026-7769	IBM Sterling B2B Integrator 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 is vulnerable to SQL injection. A remote attacker could send specially crafted SQL statements, which could allow the attacker to view, add, modify, or delete information in the back-end database.	8.1	More Details
CVE-2026-13072	When compute mode is enabled on a standalone mongod instance, insufficient validation of externally sourced BSON data during aggregation pipeline processing can result in memory corruption, potentially leading to process termination or other unintended behavior. This configuration is non-default and requires explicit enablement at startup.	8.1	More Details
CVE-2026-42016	JFrog Artifactory (Self Hosted) versions before 7.133.11 are vulnerable to a privilege escalation attack due to a validation check of the token signature/issuer and not the token’s scope.	8.1	More Details
CVE-2026-65708	sysPass through version 3.2.11 contains an insecure direct object reference vulnerability that allows any authenticated attacker to access account file attachments belonging to accounts they do not have ACL permissions for by exploiting missing authorization checks in AccountFileController. Attackers can supply arbitrary numeric file IDs through the download, view, delete, upload, and list actions to enumerate and manipulate any attachment in the	8.1	More Details

	vault, bypassing account-level access controls entirely.		
CVE-2026-65596	n8n before 1.123.64, 2.29.8, and 2.30.1 fails to enforce the "Allowed HTTP Request Domains" restriction on HTTP-based credentials (Header Auth, Basic Auth, Query Auth, OAuth) in the GraphQL node, unlike the HTTP Request node. An authenticated user able to create or edit workflows can point the node's endpoint at a server they control and exfiltrate restricted credentials. Only instances where a credential has "Allowed HTTP Request Domains" configured and is usable by non-owner users are affected.	8.1	More Details
CVE-2026-64235	In the Linux kernel, the following vulnerability has been resolved: x86/ftrace: Relocate %rip-relative percpu refs in dynamic trampolines With CONFIG_CALL_DEPTH_TRACKING enabled on an x86 retbleed-affected platform (eg: Skylake), with retbleed=stuff, registering a dynamic ftrace trampoline crashes on the first call into the traced function: BUG: unable to handle page fault for address: ffff88817ae18880 #PF: supervisor write access in kernel mode #PF: error_code(0x0002) - not-present page PGD 4b53067 P4D 4b53067 PUD 0 Oops: Oops: 0002 [#1] SMP PTI CPU: 3 UID: 0 PID: 187 Comm: usleep Not tainted 7.0.10 #243 PREEMPT(full) Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS Arch Linux 1.17.0-2-2 04/01/2014 Code: 24 78 00 00 00 00 48 89 ea 48 89 54 24 20 48 8b b4 24 b8 00 00 00 48 8b bc 24 b0 00 00 00 48 89 bc 24 80 00 00 00 48 83 ef 05 <65> 48 c1 3d 1f a8 b6 02 05 48 8b 15 f6 00 00 00 4c 89 3c 24 4c 89 Call Trace: <TASK> ? find_held_lock ? exc_page_fault ? lock_release ? __x64_sys_clock_nanosleep ? lockdep_hardirqs_on_prepare ? trace_hardirqs_on __x64_sys_clock_nanosleep do_syscall_64 ? exc_page_fault ? call_depth_return_thunk entry_SYSCALL_64_after_hwframe ... Kernel panic - not syncing: Fatal exception This small reproducer allows to easily trigger the crash: # echo 'p __x64_sys_clock_nanosleep' > /sys/kernel/tracing/kprobe_events # echo 1 > /sys/kernel/tracing/events/kprobes/p___x64_sys_clock_nanosleep_0/enable # usleep 1 Monitoring the crash under GDB points to the exact instruction in charge of incrementing the call depth: sarq $5, %gs:__x86_call_depth(%rip) This instruction matches the one inserted by the ftrace_regs_caller from ftrace_64.S. This emitted code was likely working fine until the introduction of 59bec00ace28 ("x86/percpu: Introduce %rip-relative addressing to PER_CPU_VAR()"): it has made the call depth accounting addressing relative to $rip, instead of being based on an absolute address. As this code exact location depends on where the trampoline lives in memory, the corresponding displacement needs to be adjusted at runtime to actually correctly find the per-cpu __x86_call_depth value, otherwise the targeted address is wrong, leading to the page fault seen above. Fix the %rip-relative displacement of the copied CALL_DEPTH_ACCOUNT instruction (from ftrace_regs_caller) by calling text_poke_apply_relocation(), as it is done for example by the x86 BPF JIT compiler through x86_call_depth_emit_accounting(). This corrects both CALL_DEPTH_ACCOUNT slots, in ftrace_caller and ftrace_regs_caller. [bp: Massage.]	8.1	More Details
CVE-2026-13181	In Progress® Telerik® UI for AJAX prior to v2026.2.708, forged upload metadata can influence AsyncUploadTypeName processing and trigger unsafe attacker-controlled type resolution, enabling remote code execution in affected deployments.	8.1	More Details
CVE-2026-49035	The affected product is vulnerable to a heap-based buffer overflow via a crafted MMS Initiate request. Remote code execution (RCE) has been demonstrated when ASLR is disabled; memory corruption or denial of service may occur in configurations where ASLR is enabled.	8.1	More Details
CVE-2026-64713	This issue was addressed with improved checks. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Websites may know if the user has visited a given link.	8.1	More Details
CVE-2026-64223	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: consume only present negotiated TTLM maps ieee80211_tid_to_link_map_size_ok() validates negotiated TTLM elements against the number of link-map entries indicated by link_map_presence. ieee80211_parse_neg_ttlm() must consume the same layout. The parser advanced its cursor for every TID, including TIDs whose presence bit is clear and therefore have no map bytes in the element. A sparse map can then make a later present TID read past the validated element. The bad bytes land in neg_ttlm->{up,down}link[tid] but are gated by valid_links before being applied to driver state, so a peer cannot turn the read into a policy change. Under KUnit + KASAN with an exact-sized element allocation the OOB read is reported as a slab-out-of-bounds; whether the same trigger fires under the production RX path depends on surrounding allocator state. Advance the cursor only when the current TID has a map present.	8.1	More Details
CVE-2026-13185	In Progress® Telerik® UI for AJAX prior to v2026.2.708, applications using cookie-based storage in RadPersistenceManager or RadDockLayout deserialize attacker-controlled cookie content, allowing unauthenticated remote code execution.	8.1	More Details
CVE-2026-13186	In Progress® Telerik® UI for AJAX prior to v2026.2.708, a path traversal vulnerability in the file-based persistence storage provider can be exploited when the storage key is derived from user-controlled input, enabling attacker-controlled deserialization and remote code execution.	8.1	More Details
CVE-2026-65757	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in Modules Anywhere extension - The editor popup could expose restricted module data to authenticated users without the required module permissions or valid request tokens.	8.1	More Details
CVE-2026-13187	In Progress® Telerik® UI for AJAX prior to v2026.2.708, DialogHandler provider type input may be tampered with, potentially altering dialog processing and enabling chained exploitation.	8.1	More Details
CVE-2026-13059	An authenticated user with low privileges may be able to perform unauthorized reads and writes on data protected by role-based query-level access controls, due to insufficient validation of certain client-supplied command parameters. The issue affects find, update, delete, and aggregate commands in non-apiStrict configurations.	8.1	More Details
CVE-	IBM WebSphere Application Server 8.5, and 9.0 traditional could allow a remote attacker to execute arbitrary code		More

2026-14974	caused by unsafe deserialization of untrusted data.	8.1	Details
CVE-2026-13190	In Progress® Telerik® UI for AJAX prior to v2026.2.708, a deserialization vulnerability in the persistence utilities allows unsafe type instantiation from attacker-influenced persisted state, which can lead to remote code execution.	8.1	More Details
CVE-2026-64768	An out-of-bounds read issue was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. A remote attacker may cause an unexpected app termination.	8.1	More Details
CVE-2026-64368	In the Linux kernel, the following vulnerability has been resolved: mm/slab: do not limit zeroing to orig_size when only red zoning is enabled When init (zeroing) on allocation is requested, for kmalloc() we generally have to zero the full object size even if a smaller size is requested, in order to provide krealloc()'s __GFP_ZERO guarantees. But if we track the requested size, krealloc() uses that information to do the right thing, so we can zero only the requested size. With red zoning also enabled, any extra size became part of the red zone, so it must not be zeroed and thus we must zero only the requested size. However the current check is imprecise, and will trigger also when only SLAB_RED_ZONE is enabled without SLAB_STORE_USER (which enables tracking the requested size). This means enabling red zoning alone can compromise krealloc()'s __GFP_ZERO contract. Fix this by using slub_debug_orig_size() instead, which is the exact check for whether the requested size is tracked. We don't need to care if red zoning is also enabled or not. Also update and expand the comment accordingly.	8.1	More Details
CVE-2026-64398	In the Linux kernel, the following vulnerability has been resolved: ksmbd: add a permission check for FSCTL_SET_ZERO_DATA FSCTL_SET_ZERO_DATA in smb2_ioctl() destroys file data via ksmbd_vfs_zero_data() -> vfs_fallocate(PUNCH_HOLE/ZERO_RANGE) after checking only the share-level KSMBD_TREE_CONN_FLAG_WRITABLE, with no per-handle access check. A handle opened with only FILE_WRITE_ATTRIBUTES still yields an FMODE_WRITE filp (FILE_WRITE_ATTRIBUTES is part of FILE_WRITE_DESIRE_ACCESS_LE, so smb2_create_open_flags() opens it O_WRONLY), so the vfs_fallocate FMODE_WRITE check does not stop it; only the missing fp->daccess gate would. Reproduced on mainline 7.1-rc7 with KASAN by an authenticated SMB client: a FILE_WRITE_ATTRIBUTES-only handle zeroed 4096 bytes of file data it had no FILE_WRITE_DATA right to (6/6; a FILE_READ_DATA-only handle was correctly denied). This is the unfixed sibling of commit cc57232cae23 ("ksmbd: fix FSCTL permission bypass by adding a permission check for FSCTL_SET_SPARSE"). Because SET_ZERO_DATA writes data (not an attribute), require FILE_WRITE_DATA.	8.1	More Details
CVE-2026-8789	The Easy Appointments plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check and missing nonce verification on the `ea_delete_multiple_connections` AJAX action in all versions up to, and including, 3.12.27. This makes it possible for authenticated attackers, with Contributor-level access and above, to delete arbitrary connection records from the `wp_ea_connections` table, disrupting the plugin's core booking functionality.	8.1	More Details
CVE-2026-54342	In epa4all, prior to version 2026-05-20, an attacker on the network path between epa4all and any backend (ePA Aktensystem, Konnektor, IDP, TSS) can present a self-signed TLS certificate and intercept the connection. For non-VAU connections (Konnektor, IDP), this allows direct read and modification of the inner traffic, including smartcard operations and OIDC authentication exchanges. For the ePA backend, the disabled TLS verification is the transport-level enabler for the VAU MITM described in GHSA-vvh7-x6c7-46gh. This issue has been patched in version 2026-05-20.	8.1	More Details
CVE-2026-54593	Pterodactyl is a free, open-source game server management panel. Prior to Panel version 1.12.3 and Wings version 1.12.2, the Wings /upload/file endpoint accepted any valid panel-signed JWT that contained server_uuid, user_uuid, and unique_id claims without checking the token's intended purpose; because the Panel issues JWTs carrying those same claims for lower-privilege operations such as WebSocket authentication and file-download links, an authenticated subuser could reuse one of those tokens (for example a WebSocket token obtained with only the websocket.connect permission) by replaying it against /upload/file to write arbitrary files to the same server, despite never being granted the file.create permission. This issue is fixed in Panel version 1.12.3 and Wings version 1.12.2.	8.1	More Details
CVE-2026-59545	Unauthenticated Broken Authentication in miniOrange Discord Integration <= 2.2.4 versions.	8.1	More Details
CVE-2026-64442	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB reads in IE loops in issue_assocreq() and join_cmd_hdl() Two IE parsing loops are missing the header bounds checks before they dereference pIE->length: - issue_assocreq() walks pmlmeinfo->network.ies to build the association request. If the stored IE data ends with only an element_id byte and no length byte, pIE->length is read one byte past the end of the buffer. - join_cmd_hdl() walks pnetwork->ies during station join and has the same problem under the same conditions. Both buffers are filled from AP beacon and probe-response frames, so a malicious AP that sends a truncated final IE can trigger the issue. Apply the two-guard pattern established in update_beacon_info(): 1. Break if fewer than sizeof(*pIE) bytes remain. 2. Break if the IE's declared data extends past the buffer end.	8.1	More Details
CVE-2026-12255	The MainWP Child WordPress plugin before 6.1.2 does not verify the requester's identity in its site-registration request handler when password authentication has been disabled for the targeted account, allowing an unauthenticated attacker to obtain a valid authentication session as that account, including an administrator, by naming its login in a single registration request.	8.1	More Details
CVE-2026-	An out-of-bounds access issue was addressed with improved bounds checking. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing maliciously crafted	8.1	More Details

64719	web content may lead to an unexpected Safari crash.		
CVE-2026-64440	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB write in HT_caps_handler() HT_caps_handler() iterates pIE->length bytes and writes into HT_caps.u.HT_cap[], which is a fixed 26-byte array (sizeof struct HT_caps_element). Because pIE->length is a raw u8 from an over-the-air 802.11 AssocResponse frame and is never validated, a malicious AP can set it up to 255, causing up to 229 bytes of out-of-bounds writes into adjacent fields of struct mlme_ext_info. Truncate the iteration count to the size of HT_caps.u.HT_cap using umin() so that data from a longer-than-expected IE is silently ignored rather than written out of bounds, preserving interoperability with APs that pad the element. An early return on oversized IEs was considered but rejected: it would bypass the pmlmeinfo->HT_caps_enable = 1 assignment that precedes the loop, silently disabling HT mode for APs that append extra bytes to the HT Capabilities IE.	8.1	More Details
CVE-2026-17496	NoteGen before 0.32.0 renders AI chat responses with markdown-it configured with html:true and injects the result into the DOM via dangerouslySetInnerHTML in chat-preview, without HTML sanitization and with CSP set to null. Attacker-controlled content that reaches the model prompt (for example a malicious skill REFERENCE.md that instructs the model to emit HTML) can cause the model response to include executable markup such as an img onerror handler. When the user views the chat response, that markup runs as JavaScript in the privileged Tauri webview, enabling arbitrary script execution in the application context (cross-site scripting).	8.1	More Details
CVE-2026-10818	The WPForms Pro plugin for WordPress is vulnerable to Arbitrary File Upload in all versions up to, and including, 1.10.1.1 via the ajax_chunk_upload_finalize function. This is due to the file type validation occurring after chunk metadata and file contents have already been written to disk, and the assembled file not being deleted upon validation failure. This makes it possible for unauthenticated attackers to upload files that may be executable, which makes remote code execution possible.	8.1	More Details
CVE-2026-13152	The Custom Fields Account Registration For Woocommerce WordPress plugin before 1.4 does not prevent its custom registration fields from writing to the user capabilities meta key on sites that use a non-default database table prefix, so an unauthenticated user who registers an account can be granted the administrator role when a correspondingly named field has been configured.	8.1	More Details
CVE-2026-66374	Knot Resolver before 6.4.1 allows remote code execution via a heap-based buffer overflow in the DoQ (DNS-over-QUIC) receive path.	8.1	More Details
CVE-2026-64443	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB read in update_beacon_info() IE loop The IE parsing loop in update_beacon_info() advances by (pIE->length + 2) each iteration but only guards on i < len. When a malicious AP sends a Beacon whose last IE has only one byte remaining in the frame (the element_id byte lands at len-1), the loop reads pIE->length from one byte past the allocated receive buffer. Additionally, even when the header bytes are in bounds, pIE->length itself can extend the data window beyond len, passing a truncated IE to the handler functions. Add two guards at the top of the loop body: 1. Break if fewer than sizeof(*pIE) bytes remain (can't read header). 2. Break if the IE's declared data extends past len. Also replace i += (pIE->length + 2) with i += sizeof(*pIE) + pIE->length for consistency with the sizeof(*pIE) guards added above.	8.1	More Details
CVE-2026-65916	CyberPanel through 1.9.1, fixed in commit b198460, contains a missing authorization vulnerability in the cancelBackupCreation handler that allows authenticated users to kill, delete, and corrupt other tenants' backups. Attackers can send crafted POST requests with arbitrary backupCancellationDomain and fileName parameters to terminate backup processes, delete backup archives, corrupt backup status files, and remove database records belonging to other tenants.	8.1	More Details
CVE-2026-15802	The WP Foodbakery plugin for WordPress is vulnerable to arbitrary file deletion due to insufficient file path validation in the 'delete_locations_backup_file_callback' function in all versions up to, and including, 4.9. This makes it possible for authenticated attackers, with subscriber-level access and above, to delete arbitrary files on the server, which can easily lead to remote code execution when the right file is deleted (such as wp-config.php).	8.1	More Details
CVE-2026-48060	Litestar is an Asynchronous Server Gateway Interface (ASGI) framework. Prior to version 2.20.0, Litestar instances which use a template engine in conjunction with CSRF protection are vulnerable to HTML Injection which can be escalated to Cross Site Scripting due to the contents of the CSRF cookie being excluded from automatic escaping by the template engine when configured inline with documentation recommendations. This issue has been patched in version 2.20.0.	8.1	More Details
CVE-2026-64444	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB read in OnAssocRsp() IE loop The IE parsing loop in OnAssocRsp() advances by (pIE->length + 2) each iteration but only guards on i < pkt_len. When a malicious AP sends an AssocResponse whose last IE has only one byte remaining in the frame (the element_id byte lands at pkt_len-1), the loop reads pIE->length from pframe[pkt_len], which is one byte past the allocated receive buffer. Additionally, even when the header bytes are in bounds, pIE->length itself can extend the data window beyond pkt_len, silently passing a truncated IE to the handler functions. Add two guards at the top of the loop body: 1. Break if fewer than sizeof(*pIE) bytes remain (can't read header). 2. Break if the IE's declared data extends past pkt_len.	8.1	More Details
CVE-2026-4773	Improper validation of specified type of input vulnerability in Magarsus Consulting Ltd. Co. IDM-MFA allows Authentication Bypass. This issue affects IDM-MFA: from 2025.11.27 before 2026.03.10.	8.1	More Details
CVE-2026-	A Missing Authorization vulnerability in Progress Software LoadMaster, ECS Connection Manager, Object Scale Connection Manager, MOVEit WAF, and Multi Tenant allows an authenticated attacker with low privileges to perform privileged administrative operations via the REST API that should not be accessible to their permission level,	8.0	More Details

59690	potentially resulting in a system compromise.		
CVE-2026-63265	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in various Regular Labs extension AJAX endpoints - Privileged Regular Labs AJAX endpoints did not consistently require valid CSRF tokens, matching component/item permissions and trusted server-generated form configuration. Authenticated lower-privileged users or CSRF attacks could invoke lookups or mutations outside their authorization.	8.0	More Details
CVE-2026-64406	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: fix UAF in bt_accept_dequeue() bt_accept_get() takes a temporary reference before dropping the accept queue lock. bt_accept_dequeue() currently drops that reference before bt_accept_unlink(), leaving only the queue reference. bt_accept_unlink() drops the queue reference. The subsequent sock_hold() therefore accesses freed memory if it was the final reference, as observed by KASAN during listening L2CAP socket cleanup. Retain the temporary queue-walk reference through unlink and hand it to the caller on success. Drop it explicitly on the closed and not-yet-connected paths.	8.0	More Details
CVE-2026-59689	An Incorrect Authorization vulnerability in Progress Software LoadMaster, ECS Connection Manager, Object Scale Connection Manager, and MOVEit WAF allows an authenticated attacker with low privileges to escalate privileges to root on the affected appliance, potentially resulting in full system compromise.	8.0	More Details
CVE-2026-12736	The Wpify Woo plugin for WordPress is vulnerable to Privilege Escalation in versions up to, and including, 5.4.16. This is due to the SettingsApi::save_option() REST route (POST /wp-json/wpify-woo/v1/option) passing the request-supplied 'option' and 'data' parameters directly to update_option() without any option-name allowlist or value sanitization, while the permission_callback only verifies the manage_woocommerce capability. This makes it possible for authenticated attackers, with Shop Manager-level access and above, to elevate their privileges to Administrator by overwriting arbitrary WordPress options (for example setting default_role to administrator and users_can_register to 1, or disabling security plugins via active_plugins).	8.0	More Details
CVE-2026-35425	Improper access control in Azure API Management (APIM) allows an authorized attacker to execute code over a network.	8.0	More Details
CVE-2026-60371	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Difficult to exploit vulnerability allows low privileged attacker with access to the physical communication segment attached to the hardware where the Oracle Platform Security for Java executes to compromise Oracle Platform Security for Java. While the vulnerability is in Oracle Platform Security for Java, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 8.0 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:A/AC:H/PR:L/UI:N/S:C/C:H/I:H/A:H).	8.0	More Details
CVE-2026-43698	An injection issue was addressed with improved validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to gain root privileges.	7.8	More Details
CVE-2026-64367	In the Linux kernel, the following vulnerability has been resolved: HID: hid-goodix-spi: validate report size to prevent stack buffer overflow goodix_hid_set_raw_report() builds a protocol frame in a 128-byte stack buffer (tmp_buf), writing an 11-12 byte header followed by the caller-supplied report data. The HID core caps report size at HID_MAX_BUFFER_SIZE (16384) by default, while the driver does not set hid_ll_driver.max_buffer_size and performs no bounds checking before copying the payload: memcpy(tmp_buf + tx_len, buf, len); A hidraw SET_REPORT ioctl with a report larger than ~116 bytes overflows the stack buffer. Add a size check after constructing the header, rejecting reports that would exceed the buffer capacity. Discovered by Atuin - Automated Vulnerability Discovery Engine.	7.8	More Details
CVE-2026-48374	Bridge is affected by an Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability that could lead to arbitrary file system read. An attacker could exploit this vulnerability to access sensitive files and directories outside the intended access scope. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-54655	datamodel-code-generator generates Python data models from schema definitions. From 0.51.0 until 0.60.2, x-python-type values parsed by src/datamodel_code_generator/parser/jsonschema.py in _get_python_type_override are inserted into generated field annotations without sufficient validation, allowing attacker-controlled JSON Schema content to execute Python code when the generated module is imported. This issue is fixed in version 0.60.2.	7.8	More Details
CVE-2026-54654	datamodel-code-generator generates Python data models from schema definitions. From 0.14.1 until 0.60.2, the --extra-template-data comment field is rendered into Python comments in src/datamodel_code_generator/model/template/TypeAliasAnnotation.jinja2, src/datamodel_code_generator/model/template/TypedDict.jinja2, src/datamodel_code_generator/model/template/dataclass.jinja2, src/datamodel_code_generator/model/template/msgspec.Struct.jinja2, src/datamodel_code_generator/model/template/pydantic/BaseModel.jinja2, and src/datamodel_code_generator/model/template/pydantic_v2/BaseModel.jinja2 without neutralizing carriage returns in Python # comments, allowing an attacker-controlled comment value to inject Python code into generated models that runs when imported. This issue is fixed in version 0.60.2.	7.8	More Details
CVE-2026-48392	Bridge is affected by an out-of-bounds write vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details

CVE-2026-64361	In the Linux kernel, the following vulnerability has been resolved: hfs/hfsplus: fix u32 overflow in check_and_correct_requested_length check_and_correct_requested_length() compares (off + len) against node_size using u32 arithmetic. When the caller passes a large len value (e.g. from an underflowed subtraction in hfs_brec_remove()), off + len can wrap past 2^32 and produce a small result, causing the bounds check to pass when it should fail. For example, with off=14 and len=0xFFFFFFF2 (underflowed from data_off - keyoffset - size in hfs_brec_remove), off + len wraps to 6, which is less than a typical node_size of 512, so the check passes and the subsequent memmove reads ~4GB past the node buffer. Fix this by widening the addition to u64 before comparing against node_size. This prevents the u32 wrap while keeping the logic straightforward.	7.8	More Details
CVE-2026-43711	A memory corruption issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted video file may lead to unexpected app termination.	7.8	More Details
CVE-2026-54621	datamodel-code-generator generates Python data models from schema definitions. Prior to 0.60.1, GraphQL Union description values in src/datamodel_code_generator/model/template/UnionTypeStatement.jinja2 and src/datamodel_code_generator/model/template/UnionTypeStatement.py312.jinja2 are rendered into Python comments without neutralizing carriage returns in Python # comments, allowing attacker-controlled GraphQL schema content to inject Python code into generated models that runs when imported. This issue is fixed in version 0.60.1.	7.8	More Details
CVE-2026-64802	In JetBrains GoLand before 2026.2 arbitrary code execution was possible before granting project trust in the Go Modules integration	7.8	More Details
CVE-2026-64469	In the Linux kernel, the following vulnerability has been resolved: binder: fix UAF in binder_thread_release() When a thread exits, binder_thread_release() walks its transaction stack to clear the t->from and t->to_proc that correspond with the exiting thread. However, a process dying in parallel might attempt to kfree some of these transactions. And if one of them has no associated t->to_proc, the t->to_proc->inner_lock will not be acquired. This means that transaction accesses in binder_thread_release() after t->to_proc has been cleared might race with binder_free_transaction() and cause a use-after-free error as reported by KASAN: == BUG: KASAN: slab-use-after-free in binder_thread_release+0x5d0/0x798 Write of size 8 at addr ffff000016627500 by task X/715 CPU: 17 UID: 0 PID: 715 Comm: X Not tainted 7.1.0-rc5-00149-g8fde5d1d47f6 #30 PREEMPT Hardware name: linux,dummy-virt (DT) Call trace: binder_thread_release+0x5d0/0x798 binder_ioctl+0x12c0/0x299c [...] Allocated by task 717 on cpu 18 at 67.267803s: __kasan_kmalloc+0xa0/0xbc __kmalloc_cache_noprof+0x174/0x444 binder_transaction+0x554/0x8150 binder_thread_write+0xa30/0x4354 binder_ioctl+0x20f0/0x299c [...] Freed by task 202 on cpu 18 at 90.416221s: __kasan_slab_free+0x58/0x80 kfree+0x1a0/0x4a4 binder_free_transaction+0x150/0x294 binder_send_failed_reply+0x398/0x6d8 binder_release_work+0x3e4/0x4ec binder_deferred_func+0xbd8/0x104c [...] == In order to avoid this, make sure that binder_free_transaction() reads the t->to_proc under the transaction lock. This will serialize the transaction release with the accesses in binder_thread_release(). Plus, it matches the documented locking rules for @to_proc.	7.8	More Details
CVE-2026-64716	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted image may corrupt process memory.	7.8	More Details
CVE-2026-43749	A parsing issue in the handling of directory paths was addressed with improved path validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to gain root privileges.	7.8	More Details
CVE-2026-64811	In JetBrains IntelliJ IDEA before 2026.2 arbitrary code execution was possible before granting project trust via development container configuration	7.8	More Details
CVE-2026-64463	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpci_rt1711h: unregister TCPCI port with devres rt1711h_probe() registers the TCPCI port before requesting the interrupt and enabling alert interrupts. If either of those later steps fails, the probe function returns without unregistering the TCPCI port. The explicit unregister currently only happens from the remove callback. Register a devres action immediately after tcpci_register_port() succeeds, so tcpci_unregister_port() runs on later probe failures and on driver detach. Drop the remove callback to avoid unregistering the same port twice. This issue was identified during our ongoing static-analysis research while reviewing kernel code.	7.8	More Details
CVE-2026-64354	In the Linux kernel, the following vulnerability has been resolved: bpf: Validate BTF repeated field counts before expansion btf_parse_struct_metas() walks user-supplied BTF during BPF_BTF_LOAD, and btf_repeat_fields() expands repeatable fields from array elements into the fixed BTF_FIELDS_MAX scratch array used by btf_parse_fields(). The remaining-capacity check performs the expanded field count calculation in u32. A malformed BTF can wrap that calculation, causing the check to pass even when the expanded field count exceeds the scratch array capacity. The following memcpy() can then write past the end of the array. Use checked addition and multiplication before copying repeated fields and reject impossible counts.	7.8	More Details
CVE-2026-64807	In JetBrains WebStorm before 2026.2 arbitrary code execution was possible via a project-supplied linter configuration	7.8	More Details

CVE-2026-43733	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6. Processing a maliciously crafted image may corrupt process memory.	7.8	More Details
CVE-2026-48393	Bridge is affected by an out-of-bounds write vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-43729	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6. Processing a maliciously crafted image may corrupt process memory.	7.8	More Details
CVE-2026-64468	In the Linux kernel, the following vulnerability has been resolved: binder: fix UAF in binder_free_transaction() In binder_free_transaction(), the t->to_proc is read under the t->lock. However, once the t->lock is dropped, the to_proc can die in parallel. This leads to a use-after-free error when we attempt to acquire its inner lock right afterwards: == BUG: KASAN: slab-use-after-free in _raw_spin_lock+0xe4/0x1a0 Write of size 4 at addr ffff00001125da70 by task B/672 CPU: 20 UID: 0 PID: 672 Comm: B Not tainted 7.1.0-rc6-00284-g8e65320d91cd #4 PREEMPT Hardware name: linux,dummy-virt (DT) Call trace: _raw_spin_lock+0xe4/0x1a0 binder_free_transaction+0x8c/0x320 binder_send_failed_reply+0x21c/0x2f8 binder_thread_release+0x488/0x7e0 binder_ioctl+0x12c0/0x29a0 [...] Allocated by task 675: __kmalloc_cache_noprof+0x174/0x444 binder_open+0x118/0xb70 do_dentry_open+0x374/0x1040 vfs_open+0x58/0x3bc [...] Freed by task 212: __kasan_slab_free+0x58/0x80 kfree+0x1a0/0x4a4 binder_proc_dec_tmpref+0x32c/0x5e0 binder_deferred_func+0xc48/0x104c process_one_work+0x53c/0xbc0 [...] == To prevent this, pin the target thread (t->to_thread) to guarantee the target process remains alive. Undelivered transactions without a target thread are already safe, as the target process can only be the current context in those paths.	7.8	More Details
CVE-2026-64803	In JetBrains GoLand before 2026.2 arbitrary code execution was possible before granting project trust via the configured Go SDK	7.8	More Details
CVE-2026-39875	A permissions issue was addressed with additional restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious app may be able to gain root privileges.	7.8	More Details
CVE-2026-44191	A flaw was found in the Visual Studio Code Ansible Lightspeed extension. This command injection vulnerability (CWE-78) arises from improper handling of the ansible.executionEnvironment.containerOptions and ansible.executionEnvironment.volumeMounts settings, allowing an attacker to inject shell separators. This can be triggered automatically during Language Server initialization or manually when executing a playbook. Successful exploitation leads to remote code execution (RCE) on the victim's machine with the privileges of the Visual Studio Code user, potentially resulting in a complete system compromise.	7.8	More Details
CVE-2026-18107	A flaw was found in CRIU's handling of restartable sequences (rseq) during checkpoint/restore. A malicious process inside a container can register an rseq critical section that hijacks CRIU's parasite code injection during checkpoint, allowing it to spoof the process credentials saved in the checkpoint image. On restore, the container process gains elevated capabilities and zeroed UIDs/GIDs. The practical impact on Red Hat products is limited by several factors: checkpoint/restore requires root privileges (podman) or cluster-admin RBAC (OpenShift) to trigger and cannot be initiated from within the container itself; on OpenShift prior to 4.17 the feature required explicit opt-in, and on 4.17+ the kubelet checkpoint API RBAC is not configured by default; OpenShift enforces user namespaces by default for regular workloads (hostUsers is gated behind admin-only SCCs), which makes the spoofed capabilities namespace-scoped and ineffective for privilege escalation; SELinux type enforcement (container_t) blocks privilege transitions independently of capabilities; seccomp filters persist through checkpoint/restore and cannot be corrupted via the parasite; and kernel mount namespace ownership checks on RHEL 9/10 kernels prevent mount-based container escape even with spoofed capabilities.	7.8	More Details
CVE-2026-43780	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted texture may lead to unexpected app termination.	7.8	More Details
CVE-2026-64529	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - remove unused character device and IOCTLs The QAT driver exposes a character device (qat_adf_ctl) with IOCTLs for device configuration, start, stop, status query and enumeration. These IOCTLs are not part of any public uAPI header and have no known in-tree or out-of-tree users. Device lifecycle is already managed via sysfs. The ioctl interface also increases the attack surface and is the subject of a number of bug reports. Remove the character device, the IOCTL definitions, and the related data structures (adf_dev_status_info, adf_user_cfg_key_val, adf_user_cfg_section, adf_user_cfg_ctl_data). Drop the now-unused adf_cfg_user.h header and strip adf_ctl_drv.c down to the minimal module_init/module_exit hooks for workqueue, AER, and crypto/compression algorithm registration. Clean up leftover dead code that was only reachable from the removed IOCTL paths: adf_cfg_del_all(), adf_devmgr_verify_id(), adf_devmgr_get_num_dev(), adf_devmgr_get_dev_by_id(), adf_get_vf_real_id() and the unused ADF_CFG macros. Additionally, drop the entry associated to QAT IOCTLs in ioctl-number.rst.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: smb: client: resolve SWN tcon from live registrations cifs_swn_notify() looks up a witness registration by id under cifs_swnreg_idr_mutex, drops the mutex, and then uses the registration's cached tcon pointer. That pointer is not a lifetime reference, and it is not a stable		

CVE-2026-64401	representative once cifs_get_swn_reg() lets multiple tcons for the same net/share name share one registration id. A same-share second mount can keep the cifs_swn_reg alive after the first tcon unregisters and is freed. The registration then still points at the freed first tcon, so taking tc_lock or incrementing tc_count through swnreg->tcon only moves the use-after-free earlier. Taking tc_lock while holding cifs_swnreg_idr_mutex also violates the documented CIFS lock order. Fix this by making the registration store only the stable witness identity: id, net name, share name, and notify flags. When a notify arrives, copy that identity under cifs_swnreg_idr_mutex, drop the mutex, then find and pin a live witness tcon that currently matches the net/share pair under the normal cifs_tcp_ses_lock -> tc_lock order. The notification path uses that pinned tcon directly and drops the reference when done. Registration and unregister messages now use the live tcon passed by the caller instead of a cached tcon in the registration. The final unregister send is folded into cifs_swn_unregister() while the registration is still protected by cifs_swnreg_idr_mutex. This removes the previous find/drop/reacquire raw-pointer window. The release path only removes the idr entry and frees the stable identity strings. This preserves the intended one-registration/many-tcon behavior: a registration id represents a net/share pair, and notify handling acts on a live representative selected at use time. It also preserves CLIENT_MOVE ordering for the representative tcon because the old-IP unregister is sent before cifs_swn_register() sends the new-IP register.	7.8	More Details
CVE-2026-64402	In the Linux kernel, the following vulnerability has been resolved: coresight: ultrasoc-smb: Fix OOB write in smb_sync_perf_buffer() When the SMB sink is used as a perf AUX sink, smb_update_buffer() calls smb_sync_perf_buffer() to copy hardware trace data into the perf AUX ring buffer pages. It derives pg_idx = head >> PAGE_SHIFT from @head, which is handle->head, and indexes dst_pages[pg_idx]. The pg_idx %= nr_pages normalization is only applied after the first loop iteration. This leaves the initial page index underived from the buffer size, which can result in an out-of-bounds write past dst_pages[] when head exceeds the AUX buffer size. Normalize head modulo the AUX buffer size before deriving the page index and offset, mirroring tmc_etr_sync_perf_buffer().	7.8	More Details
CVE-2026-51273	In schreibfaul1 ESP32-audioI2S 3.4.5, a heap-based buffer overflow vulnerability exists in the ID3 tag parsing function showID3Tag() of the embedded audio streaming library. The program reads untrusted long ID3 tag value from malicious audio files and uses unbounded appendf() to write formatted strings into ps_ptr heap buffer without length validation. Successful exploitation allows attackers to execute arbitrary code, leak sensitive memory data, cause device crash, or escalate privileges via a crafted malicious audio file.	7.8	More Details
CVE-2026-44190	A flaw was found in the Ansible Lightspeed Visual Studio Code extension. This Command Injection vulnerability (CWE-78) allows a remote attacker to execute unauthorized commands on a user's system. The issue occurs because the `ansible.python.activationScript` setting, intended for a virtual environment activation script, does not properly validate user input as a file path. If a user opens or executes a specially crafted project, an attacker could exploit this to gain complete control over the user's system with the privileges of the Visual Studio Code application.	7.8	More Details
CVE-2026-44189	A flaw was found in the Visual Studio Code Ansible Lightspeed extension's AnsiblePlaybookRunProvider. This command injection vulnerability allows an attacker to craft a malicious playbook filename containing special characters. When a victim runs the playbook, these characters are not properly sanitized, leading to the execution of arbitrary code with the privileges of the user running VS Code. This could result in a full system compromise, including the exfiltration of sensitive data, modification of project files, and permanent data loss.	7.8	More Details
CVE-2024-14040	In the Linux kernel, the following vulnerability has been resolved: net: nexthop: Increase weight to u16 In CLOS networks, as link failures occur at various points in the network, ECMP weights of the involved nodes are adjusted to compensate. With high fan-out of the involved nodes, and overall high number of nodes, a (non-)ECMP weight ratio that we would like to configure does not fit into 8 bits. Instead of, say, 255:254, we might like to configure something like 1000:999. For these deployments, the 8-bit weight may not be enough. To that end, in this patch increase the next hop weight from u8 to u16. Increasing the width of an integral type can be tricky, because while the code still compiles, the types may not check out anymore, and numerical errors come up. To prevent this, the conversion was done in two steps. First the type was changed from u8 to a single-member structure, which invalidated all uses of the field. This allowed going through them one by one and audit for type correctness. Then the structure was replaced with a vanilla u16 again. This should ensure that no place was missed. The UAPI for configuring nexthop group members is that an attribute NHA_GROUP carries an array of struct nexthop_grp entries: struct nexthop_grp { __u32 id; /* nexthop id - must exist */ __u8 weight; /* weight of this nexthop */ __u8 resvd1; __u16 resvd2; }; The field resvd1 is currently validated and required to be zero. We can lift this requirement and carry high-order bits of the weight in the reserved field: struct nexthop_grp { __u32 id; /* nexthop id - must exist */ __u8 weight; /* weight of this nexthop */ __u8 weight_high; __u16 resvd2; }; Keeping the fields split this way was chosen in case an existing userspace makes assumptions about the width of the weight field, and to sidestep any endianness issues. The weight field is currently encoded as the weight value minus one, because weight of 0 is invalid. This same trick is impossible for the new weight_high field, because zero must mean actual zero. With this in place: - Old userspace is guaranteed to carry weight_high of 0, therefore configuring 8-bit weights as appropriate. When dumping nexthops with 16-bit weight, it would only show the lower 8 bits. But configuring such nexthops implies existence of userspace aware of the extension in the first place. - New userspace talking to an old kernel will work as long as it only attempts to configure 8-bit weights, where the high-order bits are zero. Old kernel will bounce attempts at configuring >8-bit weights. Renaming reserved fields as they are allocated for some purpose is commonly done in Linux. Whoever touches a reserved field is doing so at their own risk. nexthop_grp::resvd1 in particular is currently used by at least strace, however they carry an own copy of UAPI headers, and the conversion should be trivial. A helper is provided for decoding the weight out of the two fields. Forcing a conversion seems preferable to bending backwards and introducing anonymous unions or whatever.	7.8	More Details
CVE-2026-17523	A flaw was found in the kernel. An unprivileged local user can exploit this vulnerability to execute arbitrary code within the kernel, which leads to a local privilege escalation (LPE). This allows the attacker to gain root privileges and take full control of the affected system.	7.8	More Details
CVE-	A buffer overflow was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6,		

2026-43776	macOS Sequoia 15.7.8, macOS Tahoe 26.6. Processing a maliciously crafted file may lead to unexpected app termination or arbitrary code execution.	7.8	More Details
CVE-2026-48372	Format Plugins is affected by a Heap-based Buffer Overflow vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-64418	In the Linux kernel, the following vulnerability has been resolved: mm: shrinker: fix shrinker_info teardown race with expansion expand_shrinker_info() iterates all visible memcgs under shrinker_mutex, including memcgs that have not finished ->css_online() yet. Once pn->shrinker_info has been published, teardown must stay serialized with expand_shrinker_info() until that memcg is either fully online or no longer visible to iteration. Today alloc_shrinker_info() breaks that rule by dropping shrinker_mutex before freeing a partially initialized shrinker_info array, which may cause the following race: CPU0 CPU1 ==== ==== css_create --> list_add_tail_rcu(&css->sibling, &parent_css->children); online_css --> mem_cgroup_css_online --> alloc_shrinker_info --> alloc node0 info rcu_assign_pointer(C->node0->shrinker_info, old0) alloc node1 info -> FAIL -> goto err mutex_unlock(shrinker_mutex) shrinker_alloc() --> shrinker_memcg_alloc --> mutex_lock(shrinker_mutex) expand_shrinker_info --> mem_cgroup_iter see the memcg expand_one_shrinker_info --> old0 = C->node0->shrinker_info memcpy(new->unit, old0->unit, ...); free_shrinker_info --> kvfree(old0); /* double free !! */ kvfree_rcu(old0, rcu); The same problem exists later in mem_cgroup_css_online(). If alloc_shrinker_info() succeeds but a subsequent objcg allocation fails, the free_objcg -> free_shrinker_info() unwind path tears down the already published pn->shrinker_info arrays without shrinker_mutex. The expand_one_shrinker_info() can race with that teardown in the same way, leading to use-after-free or double-free of the old shrinker_info. Fix this by serializing shrinker_info teardown with shrinker_mutex, and by keeping alloc_shrinker_info() error cleanup inside the locked section.	7.8	More Details
CVE-2026-64423	In the Linux kernel, the following vulnerability has been resolved: ipv4: igmp: remove multicast group from hash table on device destruction When a device is destroyed under RTNL, ip_mc_destroy_dev() iterates through the multicast list and calls ip_ma_put() on each membership, scheduling them for RCU reclamation. However, they are not unlinked from the device's multicast hash table (mc_hash). Since the device remains published in dev->ip_ptr until after ip_mc_destroy_dev() completes, concurrent RCU readers traversing mc_hash can still locate and access the multicast group after its refcount is decremented. If the RCU callback runs and frees the group while a reader is accessing it, a use-after-free occurs. Fix this by unlinking the multicast group from mc_hash using ip_mc_hash_remove() before scheduling it for reclamation. BUG: KASAN: slab-use-after-free in ip_check_mc_rcu+0x149/0x3f0 Read of size 4 at addr ffff888009bf1408 by task mausezahn/2276 Call Trace: <IRQ> dump_stack_lvl+0x67/0x90 print_report+0x175/0x7c0 kasan_report+0x147/0x180 ip_check_mc_rcu+0x149/0x3f0 udp_v4_early_demux+0x36d/0x12d0 ip_rcv_finish_core+0xb8b/0x1390 ip_rcv_finish+0x54/0x120 NF_HOOK+0x213/0x2b0 __netif_receive_skb+0x126/0x340 process_backlog+0x4f2/0xf00 __napi_poll+0x92/0x2c0 net_rx_action+0x583/0xc60 handle_softirqs+0x236/0x7f0 do_softirq+0x57/0x80 </IRQ> Allocated by task 2239: kasan_save_track+0x3e/0x80 __kasan_kmalloc+0x72/0x90 ___ip_mc_inc_group+0x31a/0xa40 __ip_mc_join_group+0x334/0x3f0 do_ip_setsockopt+0x16fa/0x2010 ip_setsockopt+0x3f/0x90 do_sock_setsockopt+0x1ad/0x300 Freed by task 0: kasan_save_track+0x3e/0x80 kasan_save_free_info+0x40/0x50 __kasan_slab_free+0x3a/0x60 __rcu_free_sheaf_prepare+0xd4/0x220 rcu_free_sheaf+0x36/0x190 rcu_core+0x8d9/0x12f0 handle_softirqs+0x236/0x7f0	7.8	More Details
CVE-2026-64431	In the Linux kernel, the following vulnerability has been resolved: ntfs: avoid calling post_write_mst_fixup() for invalid index_block ntfs_icx_ib_sync_write() calls post_write_mst_fixup() when ntfs_ib_write() returns an error, intending to restore the buffer after a failed write. However, ntfs_ib_write() returns an error immediately if pre_write_mst_fixup() validation fails. The caller, ntfs_icx_ib_sync_write(), interprets any error as a write failure requiring rollback. It does not differentiate between I/O errors and validation failures, and calls post_write_mst_fixup() anyway. Since post_write_mst_fixup() assumes that the index_block contents is correct, it doesn't perform the boundary checks, which results in out-of-bounds memory access. An attacker can craft a malicious NTFS image with: - large index_block.usa_ofs offset, pointing outside the ntfs_record - index_block.usa_count = 0, causing integer underflow - or index_block.usa_count larger than actual number of sectors in the ntfs_record, causing out-of-bounds access KASAN reports describing the memory corruption: == BUG: KASAN: slab-out-of-bounds in post_write_mst_fixup+0x19c/0x1d0 Read of size 2 at addr ffff8881586c9018 by task p/9428 Call Trace: <TASK> dump_stack_lvl+0x100/0x190 print_report+0x139/0x4ad ? post_write_mst_fixup+0x19c/0x1d0 ? __virt_addr_valid+0x262/0x500 ? post_write_mst_fixup+0x19c/0x1d0 kasan_report+0xe4/0x1d0 ? post_write_mst_fixup+0x19c/0x1d0 post_write_mst_fixup+0x19c/0x1d0 ntfs_icx_ib_sync_write+0x179/0x220 ntfs_inode_sync_filename+0x83d/0x1080 __ntfs_write_inode+0x1049/0x1480 ntfs_file_fsync+0x131/0x9b0 == BUG: KASAN: slab-out-of-bounds in post_write_mst_fixup+0x1aa/0x1d0 Write of size 2 at addr ffff8881586c91fe by task p/9428 Call Trace: <TASK> dump_stack_lvl+0x100/0x190 print_report+0x139/0x4ad ? post_write_mst_fixup+0x1aa/0x1d0 ? __virt_addr_valid+0x262/0x500 ? post_write_mst_fixup+0x1aa/0x1d0 kasan_report+0xe4/0x1d0 ? post_write_mst_fixup+0x1aa/0x1d0 post_write_mst_fixup+0x1aa/0x1d0 ntfs_icx_ib_sync_write+0x179/0x220 ntfs_inode_sync_filename+0x83d/0x1080 __ntfs_write_inode+0x1049/0x1480 ntfs_file_fsync+0x131/0x9b0 == Let's move the post_write_mst_fixup() call to ntfs_ib_write(). The ntfs_ib_write() function calls pre_write_mst_fixup() at the beginning. If the index_block contents is invalid, pre_write_mst_fixup() fails and ntfs_ib_write() returns early without calling post_write_mst_fixup() on bad index_block.	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: validate Dirty Page Table capacity in		

CVE-2026-64432	log_replay copy_lcns In the analysis pass of $LogFile journal replay, log_replay() copies LCNs from each action log record into an existing Dirty Page Table (DPT) entry without bounding the destination index. A crafted NTFS image with DPT entry lcns_follow=1 and an action log record with lcns_follow=2 produces a kernel slab out-of-bounds write at mount time: BUG: KASAN: slab-out-of-bounds in log_replay+0x654c/0xdb60 Write of size 8 at addr ffff8880095e1040 by task mount Two attacker-controlled fields can drive j+i past the allocated page_lcns[] array: 1. dp->lcns_follow (capacity) can be smaller than lrh->lcns_follow. 2. lrh->target_vcn may be smaller than dp->vcn, making the u64 subtraction wrap to a huge size_t. Validate target VCN delta and per-record LCN count against the DPT entry capacity, bail via the existing out: cleanup label with -EINVAL. This mirrors the bounds-check pattern added in commit b2bc7c44ed17 ("fs/ntfs3: Fix slab-out-of-bounds read in DeleteIndexEntryRoot") and commit 0ca0485e4b2e ("fs/ntfs3: validate rec->used in journal-replay file record check").	7.8	More Details
CVE-2026-60122	gpsd through release-3.27.5, fixed at commit 4c06658, contains a code injection vulnerability in the gpsprof utility that allows an attacker who controls GPS input data to execute arbitrary OS commands by injecting malicious content into the SKY.satellites[].used field, which is inserted unsanitized into a gnuplot heredoc data block. Attackers can supply a used value containing the string EOD to terminate the heredoc early and append gnuplot system() calls, achieving OS command execution as the user running gpsprof when the generated plot script is processed by gnuplot in polar mode.	7.8	More Details
CVE-2026-65703	FFmpeg versions 2.7 through 8.1.2 contain an out-of-bounds write vulnerability in the TDSC video decoder that allows remote attackers to cause heap corruption by supplying a crafted AVI file that changes frame dimensions across TDSF frames. The tdsc_parse_tdsf() function fails to unreference the existing reference frame before calling av_frame_get_buffer(), causing tdsc_blit() and tdsc_yuv2rgb() to write attacker-controlled pixel data beyond the end of the undersized reference frame buffer, resulting in a process crash and potential code execution.	7.8	More Details
CVE-2026-65704	FFmpeg through 8.1.2 contains an out-of-bounds write vulnerability that allows attackers to cause heap corruption by supplying a crafted ffconcat file processed with the -safe 0 flag. The TY demuxer's demux_audio() function decrements packet size without bounds checking, producing a negative size value that is passed to memcpy() in shorten_decode_frame(), where conversion to size_t wraps the value to near SIZE_MAX and triggers reads beyond the source allocation and writes far beyond the Shorten decoder's bitstream buffer.	7.8	More Details
CVE-2026-65705	FFmpeg versions 3.4 through 8.1.2 contain an out-of-bounds write vulnerability in the vf_floodfill video filter that allows attackers to corrupt heap memory by supplying a dynamically sized video stream with filtergraph reinitialization disabled via -reinit_filter 0. When config_input() allocates the points traversal stack based on initial frame dimensions and a subsequent larger frame is processed, filter_frame() performs flood-fill neighbor pushes beyond the original allocation boundary, resulting in heap corruption and process crash with potential for code execution depending on heap layout and process hardening.	7.8	More Details
CVE-2026-65706	FFmpeg versions 3.0 through 8.1.2 contain an out-of-bounds write vulnerability in the vf_swaprect video filter that allows attackers to corrupt heap memory by supplying a crafted NV12 video frame with odd width dimensions. The filter_frame() function reuses a temporary row buffer sized for plane 0's single-byte pixel step across all planes, causing an 18-byte memcpy into a 17-byte heap allocation when processing the two-byte-per-sample interleaved chroma plane of a 17x16 NV12 frame, resulting in heap corruption and process crash with potential for code execution.	7.8	More Details
CVE-2026-64388	In the Linux kernel, the following vulnerability has been resolved: smb/client: fix chown/chgrp with SMB3 POSIX Extensions Ownership (chown) and group (chgrp) modifications were being ignored when mounting with SMB3 POSIX Extensions unless CIFS_MOUNT_CIFS_ACL or CIFS_MOUNT_MODE_FROM_SID were also explicitly set. Fix this by checking for posix_extensions in cifs_setattr_nounix() when updating UID and GID, ensuring that id_mode_to_cifs_acl() is called to map and set the ownership/group information on the server.	7.8	More Details
CVE-2026-64447	In the Linux kernel, the following vulnerability has been resolved: staging: media: ipu7: fix double-free and use-after-free in error paths In both ipu7_isys_init() and ipu7_psys_init(), pdata is allocated and then passed to ipu7_bus_initialize_device(), which stores it in adev->pdata. The ipu7_bus_release() function frees adev->pdata when the device's reference count drops to zero. Two error paths incorrectly call kfree(pdata) after the device teardown has already freed it: 1. When ipu7_mmu_init() fails: put_device() is called, which drops the reference count to zero and triggers ipu7_bus_release() -> kfree(pdata). The subsequent kfree(pdata) is a double-free. 2. When ipu7_bus_add_device() fails: it calls auxiliary_device_uninit() internally, which calls put_device() -> ipu7_bus_release() -> kfree(pdata). The subsequent kfree(pdata) is again a double-free. Note that the kfree(pdata) when ipu7_bus_initialize_device() itself fails is correct, because in that case auxiliary_device_init() failed and the release function was never set up, so pdata must be freed manually. Additionally, the error code was not saved before calling put_device(), causing ERR_CAST() to dereference the already-freed adev pointer when constructing the return value. Fix this by saving the error from dev_err_probe() before put_device() and returning ERR_PTR() instead. Remove the redundant kfree(pdata) calls and fix the use-after-free in the return values of the two affected error paths.	7.8	More Details
CVE-2026-66758	A flaw was found in the file-fits plugin in GIMP. When processing a FITS image file, the plugin calculates memory allocation sizes using signed 32-bit integers for width and height. If a crafted file sets both values to large values, their product exceeds 2^31 and overflows, resulting in an undersized heap-based buffer allocation. This integer overflow issue results in a heap-based buffer overflow when cfitsio subsequently writes a full row of pixels in the buffer, causing memory corruption, potentially leading to arbitrary code execution or a denial of service.	7.8	More Details
CVE-2026-38764	An issue in Unistal Systems Pvt. Ltd.Protegent 360 v2.0.0.4 allows a local attacker to escalate privileges via the kernel driver pgsecdl.sys	7.8	More Details
CVE-	Multiple Lenze products are affected by an improper signature verification vulnerability in the SSH enablement		

2026-14837	mechanism. A low-privileged local attacker can bypass verification of the SSH enable file signature and enable SSH access on the device. Successful exploitation may result in unauthorized administrative access and complete system compromise.	7.8	[More Details]()
CVE-2026-14172	Rapid7 InsightVM, Nexpose, and the Insight Agent execute discovered executables during authenticated assessment without validating file ownership, allowing a local low-privileged user to run code as the scan credential (Scan Engine) or as root/SYSTEM (Insight Agent). Fixed in Scan Engine content 1.1.3935 and Insight Agent content component 0.0.245.0.	7.8	[More Details]()
CVE-2026-54656	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. From 0.52.1 until 0.60.2, datamodel-code-generator interpolates validators from --extra-template-data in src/datamodel_code_generator/model/pydantic_v2/base_model.py through _process_validators into @field_validator decorators without safe validation, allowing Python code execution when the generated Pydantic v2 model is imported. This issue is fixed in version 0.60.2.	7.8	[More Details]()
CVE-2026-64378	In the Linux kernel, the following vulnerability has been resolved: writeback: fix race between cgroup_writeback_umount() and inode_switch_wbs() When a container exits, the following BUG_ON() is occasionally triggered: == VFS: Busy inodes after unmount of sdb (ext4) ------------[cut here]------------ kernel BUG at fs/super.c:695! CPU: 3 PID: 6 Comm: containerd-shim Tainted: G OE K 6.6 #1 pstate: 63400009 (nZCv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=--) pc : generic_shutdown_super+0xf0/0x100 lr : generic_shutdown_super+0xf0/0x100 Call trace: generic_shutdown_super+0xf0/0x100 kill_block_super+0x20/0x48 ext4_kill_sb+0x28/0x60 deactivate_locked_super+0x54/0x130 deactivate_super+0x84/0xa0 cleanup_mnt+0xa4/0x140 __cleanup_mnt+0x18/0x28 task_work_run+0x78/0xe0 do_notify_resume+0x204/0x240 == The root cause is a race between cgroup_writeback_umount() and inode_switch_wbs()/cleanup_offline_cgwb(). There is a window between inode_prepare_wbs_switch() returning true and the subsequent wb_queue_isw() call. Following is the process that triggers the issue: CPU A (umount) | CPU B (writeback) ~~~~~~ inode_switch_wbs/cleanup_offline_cgwb atomic_inc(&isw_nr_in_flight) inode_prepare_wbs_switch -> passes SB_ACTIVE check __iget(inode) generic_shutdown_super sb->s_flags &= ~SB_ACTIVE cgroup_writeback_umount(sb) smp_mb() atomic_read(&isw_nr_in_flight) rcu_barrier() -> no pending RCU callbacks flush_workqueue(isw_wq) -> nothing queued, returns evict_inodes(sb) -> Inode skipped as isw still holds a ref. sop->put_super(sb) /* destroys percpu counters */ -> VFS: Busy inodes after unmount! wb_queue_isw() queue_work(isw_wq, ...) /* later in work function */ inode_switch_wbs_work_fn process_inode_switch_wbs iput() -> evict percpu_counter_dec() // UAF! Fix this by extending the RCU read-side critical section in inode_switch_wbs() and cleanup_offline_cgwb() to cover from inode_prepare_wbs_switch() through wb_queue_isw(). Since there is no sleep in this window, rcu_read_lock() can be used. Then add a synchronize_rcu() in cgroup_writeback_umount() before the existing rcu_barrier(), so that all in-flight switchers that have passed the SB_ACTIVE check have completed queue_work() before flush_workqueue() is called. The existing rcu_barrier() is intentionally retained so this fix can be backported unchanged to stable kernels (5.10.y, 6.6.y, ...) that still queue switches via queue_rcu_work(). It is a no-op on current mainline (since commit e1b849cfa6b6 ("writeback: Avoid contention on wb->list_lock when switching inodes")) and is removed in a follow-up patch.	7.8	[More Details]()
CVE-2026-24252	NVIDIA NeMo for Linux contains a vulnerability where an attacker may cause OS command injection. A successful exploit of this vulnerability may lead to code execution, data tampering, escalation of privileges and information disclosure.	7.8	[More Details]()
CVE-2026-64375	In the Linux kernel, the following vulnerability has been resolved: proc: protect ptrace_may_access() with exec_update_lock (FD links) proc_pid_get_link() and proc_pid_readlink() currently look up the task from the pid once, then do the ptrace access check on that task, then look up the task from the pid a second time to do the actual access. That's racy in several ways. To fix it, pass the task to the ->proc_get_link() handler, and instead of proc_fd_access_allowed(), introduce a new helper call_proc_get_link() that looks up and locks the task, does the access check, and calls ->proc_get_link().	7.8	[More Details]()
CVE-2026-64449	In the Linux kernel, the following vulnerability has been resolved: staging: vme_user: bound slave read/write to the kern_buf size The SLAVE-path helpers buffer_to_user() and buffer_from_user() copy 'count' bytes into/out of the fixed-size kern_buf (size_buf == PCI_BUF_SIZE == 0x20000, 128 KiB) using *ppos as the offset, without bounding *ppos + count against size_buf. vme_user_write()/vme_user_read() only clamp count to the VME window size (image_size = vme_get_size(resource)), which VME_SET_SLAVE sets from the user-supplied slave.size -- validated against the VME address space (up to VME_A32_MAX = 4 GiB), not against PCI_BUF_SIZE. When the window exceeds 128 KiB, a write()/read() copies past the kern_buf allocation. Clamp count against size_buf in both helpers, with an early return when *ppos is already at/after the buffer end. *ppos is >= 0 here (the caller rejects negative offsets), so size_buf - *ppos cannot wrap. This mirrors the existing clamp in the MASTER-path helpers resource_to_user() / resource_from_user(), and matches the read()/write() convention of a short transfer at end-of-buffer. Found by static analysis (CodeQL taint tracking + CBMC bounded model checking) and confirmed dynamically under KASAN with the vme_fake bridge: BUG: KASAN: slab-out-of-bounds in _copy_from_user+0x2d/0x80 Write of size 262144 at addr ffff888004100000 by task trigger/68 _copy_from_user+0x2d/0x80 vme_user_write+0x13e/0x240 [vme_user] vfs_write+0x1b8/0x7a0 ksys_write+0xb8/0x150	7.8	[More Details]()
CVE-	In the Linux kernel, the following vulnerability has been resolved: cpufreq: pcc: fix use-after-free and double free in _OSC evaluation pcc_cpufreq_do_osc() calls acpi_evaluate_object() twice for the two-phase _OSC negotiation. Between the two calls it freed output.pointer but left output.length unchanged. Since acpi_evaluate_object() treats a		

2026-64372	non-zero length with a non-NULL pointer as an existing buffer to write into, the second call wrote into freed memory (use-after-free). The subsequent kfree(output.pointer) at out_free then freed the same pointer a second time (double free). Reset output.pointer to NULL and output.length to ACPI_ALLOCATE_BUFFER after freeing the first result, so ACPICA allocates a fresh buffer for each phase independently.	7.8	More Details
CVE-2026-64333	In the Linux kernel, the following vulnerability has been resolved: USB: serial: digi_acceleport: fix write buffer corruption The digi_write_inb_command() is supposed to wait for the write urb to become available or return an error, but instead it updates the transfer buffer and tries to resubmit the urb on timeout. To make things worse, for commands like break control where no timeout is used, the driver would corrupt the urb immediately due to a broken jiffies comparison (on 32-bit machines this takes five minutes of uptime to trigger due to INITIAL_JIFFIES). Fix this by adding the missing return on timeout and waiting indefinitely when no timeout has been specified as intended. This issue was (sort of) flagged by Sashiko when reviewing an unrelated change to the driver.	7.8	More Details
CVE-2026-43723	A path handling issue was addressed with improved validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to gain root privileges.	7.8	More Details
CVE-2026-64600	In the Linux kernel, the following vulnerability has been resolved: xfs: resample the data fork mapping after cycling ILOCK xfs_reflink_fill_{cow_hole,delalloc} are both presented with an inode, a data fork mapping, and a cow fork mapping. Unfortunately, these two helpers cycle the ILOCK to grab a transaction, which means that the mappings are stale as soon as we reacquire the ILOCK. Currently we refresh the cow fork mapping by re-calling xfs_find_trim_cow_extent, but we don't refresh the data fork mapping beforehand, which means that the xfs_bmap_trim_cow in that function queries the refcount btree about the wrong physical blocks and returns an inaccurate value in *shared. If *shared is now false, the directio write proceeds with a stale data fork mapping. Fix this by querying the data fork mapping if the sequence counter changes across the ILOCK cycle.	7.8	More Details
CVE-2026-28981	A buffer overflow was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. Processing a maliciously crafted image may lead to arbitrary code execution.	7.8	More Details
CVE-2026-64260	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: Avoid queue->stopped races and set/read that value under lock There are several readers of queue->stopped that check the value under lock, but fuse_uring_commit_fetch() did not and actually the value was not set under the lock in fuse_uring_abort_end_requests() either. Especially in fuse_uring_commit_fetch it is important to check under a lock, because due to races 'struct fuse_req' might be freed with fuse_request_end, but another thread/cpu might already do teardown work.	7.8	More Details
CVE-2026-48394	Bridge is affected by an out-of-bounds write vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue requires user interaction in that a victim must open a malicious file.	7.8	More Details
CVE-2026-64764	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to unexpected app termination or arbitrary code execution.	7.8	More Details
CVE-2026-64265	In the Linux kernel, the following vulnerability has been resolved: fuse: clear intr_entry in fuse_resend and fuse_remove_pending_req When fuse_resend() moves a request from fpq->processing back to fiq->pending, it sets FR_PENDING and clears FR_SENT but does not remove the requests intr_entry from fiq->interrupts. If the request had FR_INTERRUPTED set from a prior signal, intr_entry remains dangling on fiq->interrupts. When the requesting task then receives a fatal signal, fuse_remove_pending_req() sees FR_PENDING=1, removes the request from fiq->pending and frees it via the refcount path, also without cleaning intr_entry. The stale intr_entry causes use-after-free when fuse_read_interrupt() iterates fiq->interrupts: - list_del_init(&req->intr_entry) -> UAF write on freed slab - req->in.h.unique -> UAF read, data leaked to userspace Remove intr_entry from fiq->interrupts in fuse_resend() for interrupted requests before they are placed back on fiq->pending. Add a WARN_ON if the intr_entry is not empty on request destruction.	7.8	More Details
CVE-2026-16157	Duplicati v2.3.0.1 backup software gives Authenticated Users MODIFY permissions that propagate to all subdirectories. Installing the software outside of the Program Files directory, or on a custom path, creates a LocalSystem service running from a directory that any standard local user can write to. A standard local user can overwrite any DLL in the service directory. On service restart, the OS loads the attacker's DLL before any managed code runs, executing arbitrary code as SYSTEM.	7.8	More Details
CVE-2026-64261	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: Avoid use-after-free in fuse_uring_async_stop_queues fuse_uring_async_stop_queues() might run when the last reference on ring->queue_refs was already dropped. In order to avoid an early destruction a reference on struct fuse_conn is now taken before starting fuse_uring_async_stop_queues() and that reference is only released when that delayed work queue terminates.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: exfat: bound uniname advance in exfat_find_dir_entry() In exfat_find_dir_entry(), each TYPE_EXTEND (file name) entry advances the output pointer by a fixed amount while the loop guard only tracks the accumulated name length: if (++order == 2) uniname = p_uniname->name; else uniname += EXFAT_FILE_NAME_LEN; len = exfat_extract_uni_name(ep, entry_uniname); name_len += len; unichar = *(uniname+len); *(uniname+len) = 0x0; uniname grows by EXFAT_FILE_NAME_LEN (15) per name entry, but name_len grows only by the actual extracted length, which is shorter when a name fragment contains an early NUL. The only guard is `name_len >= MAX_NAME_LENGTH`, so a crafted directory with many short		More

2026-64296	name fragments lets uniname run far past the p_uniname->name[MAX_NAME_LENGTH + 3] buffer while name_len stays small, causing an out-of-bounds read and write at *(uniname+len). The sibling extractor exfat_get_uniname_from_ext_entry() already stops on a short fragment (the lockstep `len != EXFAT_FILE_NAME_LEN` guard added in commit d42334578eba ("exfat: check if filename entries exceeds max filename length")); exfat_find_dir_entry() never got the equivalent. Track the per-entry write offset as a count and reject a fragment once the offset, or the offset plus the extracted length, would exceed MAX_NAME_LENGTH, before forming the output pointer.	7.8	Details
CVE-2026-64747	A buffer overflow was addressed with improved size validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to execute arbitrary code with kernel privileges.	7.8	More Details
CVE-2026-64502	In the Linux kernel, the following vulnerability has been resolved: iio: adc: ad_sigma_delta: fix clear_pending_event for registerless devices ad_sigma_delta_clear_pending_event() falls through to the status register read path for devices with has_registers = false and no rdy_gpiod. For such devices, ad_sd_read_reg() skips the address byte entirely and clocks raw MISO bytes with no address phase — making it byte-for-byte identical to reading conversion data. If a pending conversion result is present, this partially consumes it and corrupts the data stream for the subsequent ad_sd_read_reg() call in ad_sigma_delta_single_conversion(). Furthermore, with num_resetclks = 0 on these devices, data_read_len evaluates to 0. If the clocked byte has bit 7 clear, pending_event is set and the code attempts memset(data + 2, 0xff, 0 - 1), overflowing to SIZE_MAX and corrupting the heap. Fix by returning 0 immediately when neither rdy_gpiod nor has_registers is set. This is safe for all current registerless devices: ad7191 and ad7780 (with powerdown GPIO) are reset between conversions by CS deassertion, so there is no stale result to drain; ad7780 (without powerdown GPIO) and max11205 are continuously-converting and cycle ~DRDY at the output data rate regardless of whether the previous result was read, so the next falling edge fires naturally. A future registerless device that holds ~DRDY asserted until data is read would be broken by this early return and would require either num_resetclks set or a rdy-gpio. The same heap corruption is reachable on any device with rdy_gpiod set but num_resetclks = 0: if the GPIO indicates a pending event, the drain path executes memset(data + 2, 0xff, 0 - 1) regardless of has_registers. Add an explicit data_read_len == 0 guard after the pending event check; the stale result is then consumed by the first ad_sd_read_reg() call in ad_sigma_delta_single_conversion().	7.8	More Details
CVE-2026-62428	When grant-copy operations are processed, the respective grant may or may not already be in use by another operation (a mapping or another copy). For all copy operations the referenced guest frame is looked up. When another operation is already active for the grant (the grant is "pinned"), what is being supplied back to actually carry out permission checks and copy operation may not be consistent: The permission check may be carried out on a page different from the one involved in the copy.	7.8	More Details
CVE-2026-14881	When importing connections in Compass it is possible to override some connection options that are otherwise can't be changed via connection form. In particular it is possible to provide a custom browser open command for OIDC auth flow that is usually can be set only globally via Compass settings.	7.8	More Details
CVE-2026-64300	In the Linux kernel, the following vulnerability has been resolved: perf/aux: Fix page UAF in map_range() map_range() reads rb->aux_pages[], rb->aux_nr_pages and rb->aux_pgoff via perf_mmap_to_page() while holding only event->mmap_mutex. Those fields are serialized by rb->aux_mutex, and mmap_mutex is per event. Thus, two events sharing one rb via PERF_EVENT_IOC_SET_OUTPUT can race rb_alloc_aux() with map_range(), leading to a page-UAF scenario as follows: CPU 0 CPU 1 ===== ===== rb_alloc_aux() map_range() [1]: allocate rb->aux_pages[0] [2]: rb->aux_nr_pages++ [3]: perf_mmap_to_page() returns rb->aux_pages[0] [4]: map it as VM_PFNMAP [5]: rb->aux_pgoff = 1 munmap the page [6]: free rb->aux_pages[0] Pages mapped as VM_PFNMAP have no refcount protection, so CPU 1 holds a mapping to a freed physical frame. Fix this by taking rb->aux_mutex across the page walk in map_range().	7.8	More Details
CVE-2026-64221	In the Linux kernel, the following vulnerability has been resolved: spi: ti-qspi: fix use-after-free after DMA setup failure The driver falls back to PIO mode if DMA setup fails during probe. Make sure to clear the DMA channel pointer also if buffer allocation fails to avoid passing a pointer to the released channel to the DMA engine (or trying to free the channel a second time on late probe errors or driver unbind). This issue was flagged by Sashiko when reviewing a devres allocation conversion patch.	7.8	More Details
CVE-2026-64311	In the Linux kernel, the following vulnerability has been resolved: crypto: loongson - Remove broken and unused loongson-rng The loongson-rng rng_alg has several vulnerabilities, including not providing forward security, and a use-after-free bug due to the use of wait_for_completion_interruptible(). Meanwhile, the rng_alg framework doesn't really have any purpose in the first place other than to access the software algorithms crypto/drbg.c and crypto/jitterentropy.c. Hardware-specific rng_algs have no in-kernel user, and unlike hwrng there's no feed into the actual Linux RNG. As such, there's really no point to this code. There are of course other rng_alg drivers that are similarly unused, but they're similarly in the process of being phased out, e.g. https://lore.kernel.org/r/20260529193648.18172-1-ebiggers@kernel.org and https://lore.kernel.org/r/20260529220430.34135-1-ebiggers@kernel.org Given that, there's no point in fixing forward these vulnerabilities, and it makes much more sense to simply roll back the addition of this driver. If this platform provides TRNG (not PRNG) functionality, it could make sense to add a hwrng driver, but it would be quite different.	7.8	More Details
CVE-2026-64277	In the Linux kernel, the following vulnerability has been resolved: Input: synaptics-rmi4 - bound the F3A keymap to the GPIO count rmi_f3a_initialize() takes the GPIO count from the device query register (f3a->gpio_count = buf & RMI_F3A_GPIO_COUNT, range 0..127). rmi_f3a_map_gpios() then allocates gpio_key_map with min(gpio_count, TRACKSTICK_RANGE_END) == at most 6 entries, but rmi_f3a_attention() iterates the full gpio_count and dereferences gpio_key_map[i], and input->keycodemax is set to the full gpio_count while input->keycode points at the 6-entry allocation. A device that reports gpio_count > 6 therefore causes an out-of-bounds read of gpio_key_map[] on every attention interrupt, and out-of-bounds accesses through the input core's default keymap ioctls: EVIOCGKEYCODE reads past the buffer (leaking adjacent slab memory to user space) and EVIOCSKEYCODE writes a caller-controlled	7.8	More Details

	value past it, for any process able to open the evdev node, since input_default_getkeycode() and input_default_setkeycode() only bound the index against keycodemax. Size the keymap for the full gpio_count. The mapping loop is unchanged: it still assigns only the first min(gpio_count, TRACKSTICK_RANGE_END) entries; the remaining slots stay KEY_RESERVED (devm_kcalloc zero-fills) and are skipped when reporting.		
CVE-2026-64765	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to unexpected app termination or arbitrary code execution.	7.8	More Details
CVE-2026-64766	An integer overflow was addressed with improved input validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to unexpected app termination or arbitrary code execution.	7.8	More Details
CVE-2026-28912	A logic issue was addressed with improved restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Tahoe 26.6. A user may be able to elevate privileges.	7.8	More Details
CVE-2026-64481	In the Linux kernel, the following vulnerability has been resolved: ALSA: hda/cs35l41: Fix firmware load work teardown cs35l41_hda creates ALSA controls whose private data points at the cs35l41_hda object. The firmware load control can also queue fw_load_work. Those controls are not removed on component unbind, and device remove only cancels fw_load_work through cs35l41_remove_dsp(). That helper is skipped when halo_initialized is false. With firmware_autostart disabled, a firmware load can be requested before the DSP has been initialized. If the component or device is removed before the queued work runs, the worker can run after teardown and dereference driver state that is no longer valid. Track the created controls and remove them on unbind so no new control callback can reach the driver data or queue more work. Then cancel fw_load_work to drain any request that was already queued. Also cancel the work unconditionally during device remove before runtime PM teardown.	7.8	More Details
CVE-2026-39874	A permissions issue was addressed with additional restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious app may be able to gain root privileges.	7.8	More Details
CVE-2026-64217	In the Linux kernel, the following vulnerability has been resolved: netfs: Fix overrun check in netfs_extract_user_iter() Fix netfs_extract_user_iter() so that if iov_iter_extract_pages() overfills pages[], then those pages don't get included in the iterator constructed at the end of the function. If there was an overfill, memory corruption has already happened.	7.8	More Details
CVE-2026-64304	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - validate RSA CRT component lengths The generic RSA key parser (rsa_helper.c) bounds each CRT component (p, q, dp, dq, qinv) by the modulus size n_sz, but qat_rsa_setkey_crt() allocates half-size DMA buffers (key_sz / 2) and right-aligns each component with: memcpy(dst + half_key_sz - len, src, len) When a CRT component is larger than half_key_sz the subtraction underflows and memcpy writes past the DMA buffer, causing memory corruption. Add a len > half_key_sz check next to the existing !len check for each of the five CRT components so the driver falls back to the non-CRT path instead of writing out of bounds.	7.8	More Details
CVE-2026-64485	In the Linux kernel, the following vulnerability has been resolved: ALSA: compress: Fix task creation error unwind snd_compr_task_new() allocates the driver task before validating the returned DMA buffers and reserving file descriptors. When either of those later steps fails, the core frees its task wrapper and DMA-buffer references without calling the driver's task_free() callback. Any driver resources allocated by task_create() are therefore leaked. The dual-fd allocation path also jumps to cleanup without storing the negative get_unused_fd_flags() result in retval. Since retval still contains the successful task_create() return value, TASK_CREATE can incorrectly report success although the task was discarded. Preserve the fd allocation errors and call task_free() when failure occurs after a successful task_create() callback.	7.8	More Details
CVE-2026-64218	In the Linux kernel, the following vulnerability has been resolved: batman-adv: bla: fix report_work leak on backbone_gw purge batadv_bla_purge_backbone_gw() removes stale backbone gateway entries, but fails to properly handle their associated report_work: - If report_work is running, the purge must wait for it to finish before freeing the backbone_gw, otherwise the worker may access freed memory (e.g. bat_priv). - If report_work is pending, the purge must cancel it and release the reference held for that pending work item. The previous implementation called hlist_for_each_entry_safe() inside a spin_lock_bh() section, but cancel_work_sync() may sleep and therefore cannot be called from within a spinlock-protected region. Restructure the loop to handle one entry per spinlock critical section: acquire the lock, find the next entry to purge, remove it from the hash list, then release the lock before calling cancel_work_sync() and dropping the hash_entry reference. Repeat until no more entries require purging.	7.8	More Details
CVE-2026-64763	An out-of-bounds write issue was addressed by removing the vulnerable code. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to unexpected app termination or arbitrary code execution.	7.8	More Details
CVE-2025-60835	An issue in the unrar.dll component of IZArc v4.6 allows attackers to execute a path traversal.	7.8	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: iommufd: Use sizeof(*hdr) instead of sizeof(hdr) in veventq read The bound-check in iommufd_veventq_fops_read() for the normal vEVENT path uses sizeof(hdr) where the surrounding code uses sizeof(*hdr): if (!vevent_for_lost_events_header(cur) && sizeof(hdr) + cur->data_len > count - done) { hdr is declared as struct iommufd_vevent_header *, so sizeof(hdr) evaluates to the size of the pointer. Surrounding code uses sizeof(*hdr) consistently: if (done >= count || sizeof(*hdr) > count - done) { ... if (copy_to_user(buf + done, hdr, sizeof(*hdr))) { ... done += sizeof(*hdr); struct iommufd_vevent_header is currently 8		

2026-64293	bytes (two __u32 fields, flags and sequence), so on 64-bit (sizeof(void *) == 8) the two expressions happen to be equal and the check works as intended. On 32-bit (sizeof(void *) == 4) the check under-counts the header by 4 bytes: a vEVENT whose data_len causes 8 + cur->data_len to exceed count - done while 4 + cur->data_len does not will pass the check, then the loop will copy_to_user 8 bytes of header followed by data_len bytes of payload, writing past the user-supplied buffer. It is also a latent bug for any future expansion of struct iommufd_vevent_header beyond sizeof(void *) on 64-bit; the check should not depend on the type happening to match the host pointer width. Use sizeof(*hdr) to match the rest of the function and the actual amount that will be copied.	7.8	More Details
CVE-2025-71408	NLTK (Natural Language Toolkit) before version 3.9.3 contains an eval injection vulnerability in the nltk.collocations module that allows an attacker who controls command-line arguments to execute arbitrary Python code. When collocations.py is invoked directly, the __main__ block passes command-line arguments directly to eval() as suffixes of BigramAssocMeasures without allowlist validation or sanitization, enabling an attacker to supply a Python expression that escapes the intended attribute lookup and executes arbitrary code including OS commands via the os module.	7.8	More Details
CVE-2026-64749	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6, visionOS 26.6. An app may be able to cause unexpected system termination or corrupt kernel memory.	7.8	More Details
CVE-2026-16287	Improper neutralization of special elements used in an OS command ('OS command injection') vulnerability in TUBITAK BILGEM Software Technologies Research Institute pardus-update allows OS Command Injection. This issue affects pardus-update: from 0.6.6 before 0.7.0.	7.8	More Details
CVE-2026-14985	The Analog Way Picturall Quad Compact Mark II version 3.5.8, contains a local privilege escalation vulnerability in the core firmware. This is due to improper privilege delegation and insufficient input validation in a maintenance script.	7.8	More Details
CVE-2026-43673	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted audio file may corrupt process memory.	7.8	More Details
CVE-2026-64276	In the Linux kernel, the following vulnerability has been resolved: Input: synaptics-rmi4 - bound the F30 keymap to the GPIO/LED count rmi_f30_map_gpios() allocates gpioled_key_map with min(gpioled_count, TRACKSTICK_RANGE_END) == at most 6 entries, but rmi_f30_attention() iterates the full f30->gpioled_count (device query register, range 0..31) and dereferences gpioled_key_map[i], and input->keycodemax is set to the full gpioled_count while input->keycode points at the 6-entry allocation. A device that reports gpioled_count > 6 with GPIO support enabled therefore causes an out-of-bounds read on the attention interrupt and out-of-bounds read/write through the EVIOCGKEYCODE/EVIOCSKEYCODE ioctls, which bound the index only against keycodemax. This is the same defect as the F3A handler, which was copied from F30. Size the keymap for the full gpioled_count; the mapping loop still assigns only the first min(gpioled_count, TRACKSTICK_RANGE_END) entries.	7.8	More Details
CVE-2026-64279	In the Linux kernel, the following vulnerability has been resolved: i2c: core: fix adapter deregistration race Adapters can be looked up by their id using i2c_get_adapter() which takes a reference to the embedded struct device. Remove the adapter from the IDR before tearing it down during deregistration (and on registration failure) to make sure its resources are not accessed after having been freed (e.g. the device name).	7.8	More Details
CVE-2026-38765	An issue in Unistal Systems Pvt. Ltd.Protegent 360 v2.0.0.4 allows a local attacker to escalate privileges via the kernel driver pgsecdl.sys	7.8	More Details
CVE-2026-49743	Software installed and run as a non-privileged user may conduct improper GPU system calls to manipulate the lifetimes of synchronisation objects in the kernel, leading to read/write UAFs. During workload submission involving a fence exported by the GPU driver, the reference count of the underlying synchronisation primitive is not properly incremented. This can be exploited, by destroying the exported fence and prematurely release the underlying primitive, resulting in a potential use-after-free condition.	7.8	More Details
CVE-2026-49744	Kernel software installed and running inside a Guest VM may post improper commands to the GPU Firmware to trigger a write of data outside the Guest's virtualised GPU memory. Out of bounds accesses triggered by malware introduced to a Guest KMD could allow privilege escalation which escapes virtualization boundaries.	7.8	More Details
CVE-2026-49745	Kernel software installed and running inside a Guest VM may post improper commands to the GPU Firmware to trigger a write of data outside the Guest's virtualised GPU memory. Software installed and run under a Guest VM can send commands to the GPU which result in out of bounds memory accesses. These can be used to escalate privileges.	7.8	More Details
CVE-2026-16607	A vulnerability in Fujitsu Software Linux openFT and Fujitsu Software Oracle Solaris openFT before version 12.1D00 allows for local privilege escalation to root of an already authenticated user on GNU/Linux or Oracle Solaris. The Fsas Technologies PSIRT obtained that intelligence internally and covers the CVE beyond its CNA scope under existing agreement with Fujitsu Germany.	7.8	More Details
CVE-2026-38766	An issue in Unistal Systems Pvt. Ltd.Protegent 360 v2.0.0.4 allows a local attacker to escalate privileges via the sub_186f4 function	7.8	More Details
	In the Linux kernel, the following vulnerability has been resolved: udf: validate free block extents against the partition length udf_free_blocks() checks the logical block number and count against the partition length, but drops the extent offset from that final bound. A crafted extent can pass the guard while logicalBlockNum + offset + count points past the partition, which later indexes past the space bitmap array. A single ftruncate(2) on a file backed by such an extent reliably panics the kernel. This is a local availability issue. On desktop systems where UDisks/polkit allows the		

CVE-2026-64324	active user to mount removable UDF media without CAP_SYS_ADMIN, an unprivileged local user can supply the crafted filesystem and trigger the panic by truncating a writable file on it. Systems that require root or CAP_SYS_ADMIN to mount the image have a higher prerequisite. No confidentiality or integrity impact is claimed: the reproduced primitive is an out-of-bounds read of a bitmap pointer slot followed by a kernel panic. Use the already computed logicalBlockNum + offset + count value for the partition length check. Also make load_block_bitmap() reject an out-of-range block group before indexing s_block_bitmap[], so corrupted callers cannot walk past the flexible array.	7.8	More Details
CVE-2026-51254	schreibfaul1 ESP32-audioI2S v3.4.5 has an integer underflow vulnerability in the MP3Decoder::GetBits() function of the MP3 decoder due to unchecked bit reading operations. The lack of validation on the nBits parameter causes the cachedBits counter to underflow to negative values, leading to invalid bit manipulation, incorrect bitstream parsing, application crash, or arbitrary code execution via a specially crafted MP3 file.	7.8	More Details
CVE-2026-64251	In the Linux kernel, the following vulnerability has been resolved: pwrseq: core: fix use-after-free in pwrseq_debugfs_seq_next() pwrseq_debugfs_seq_next() declares 'next' with __free(put_device), which causes put_device() to be called on the returned pointer when the variable goes out of scope. This results in a use-after-free since the seq_file framework receives a pointer whose reference has already been dropped. Simply removing __free(put_device) would fix the UAF but would leak the reference acquired by bus_find_next_device(), as stop() only calls up_read(&pwrseq_sem) and never releases the device reference. Fix this by making the reference counting consistent across all seq_file callbacks, matching the standard pattern used by PCI and SCSI: - start(): use get_device() so it returns a referenced pointer. - next(): explicitly put_device(curr) to release the previous device's reference (no NULL check needed - the seq_file framework only calls next() while the previous return was non-NULL). - stop(): put_device(data) to release the last iterated device's reference, with a NULL guard since stop() may be called with NULL when start() returned NULL or next() reached end-of-sequence.	7.8	More Details
CVE-2026-64322	In the Linux kernel, the following vulnerability has been resolved: udf: validate sparing table length as an entry count, not a byte count udf_load_sparable_map() accepts a sparing table when sizeof(*st) + le16_to_cpu(st->reallocationTableLen) > sb->s_blocksize is false, i.e. it treats reallocationTableLen as a number of BYTES that must fit in the block. But the table is walked as an array of 8-byte sparingEntry elements: for (i = 0; i < le16_to_cpu(st->reallocationTableLen); i++) { struct sparingEntry *entry = &st->mapEntry[i]; ... entry->origLocation ... } in udf_get_pblock_spar15() and udf_relocate_blocks(). A reallocationTableLen of N therefore passes the check whenever sizeof(*st) + N <= blocksize, yet the consumers index sizeof(*st) + N * sizeof(struct sparingEntry) bytes -- up to ~8x the block. On a crafted UDF image this is an out-of-bounds read in udf_get_pblock_spar15(); udf_relocate_blocks() additionally feeds the same length to udf_update_tag(), whose crc_itu_t() reads far past the block, and its memmove() through st->mapEntry[] is an out-of-bounds write. Validate reallocationTableLen as the entry count it is, with struct_size().	7.8	More Details
CVE-2026-64758	The issue was addressed with improved bounds checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to unexpected app termination.	7.8	More Details
CVE-2026-64266	In the Linux kernel, the following vulnerability has been resolved: fuse: re-lock request before returning from fuse_ref_folio() fuse_ref_folio() unlocks the request but does not re-lock it before returning. fuse_chan_abort() can end the request and the async end callback (eg fuse_writepage_free()) can free the args while the subsequent copy chain logic after fuse_ref_folio() accesses them, leading to use-after-free issues. Fix this by locking the request in fuse_ref_folio() before returning.	7.8	More Details
CVE-2026-39877	A memory corruption issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. An app may be able to disclose kernel memory.	7.8	More Details
CVE-2026-64259	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: make a fuse_req on SQE commit only findable after memcpy Bad userspace might try to trick us and send commit SQEs request unique / commit-id of requests that are not even send to fuse-server (io_uring_cmd_done() not called) yet. fuse_uring_commit_fetch() ends the fuse request when the ring entry has a wrong state, but that could have caused a use-after-free with the memcpy operations in fuse_uring_send_in_task(). In order to avoid such races the call of fuse_uring_add_to_pq() is moved after the copy operations and just before completing the io-uring request - malicious userspace cannot find the request anymore until all prepration work in fuse-client/kernel is completed. This also moves fuse_uring_add_to_pq() a bit up in the code to avoid a forward declaration. Also not with a preparation commit, to make it easier to back port to older kernels.	7.8	More Details
CVE-2026-64226	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Avoid UAF in scx_root_enable_workfn() init failure path In scx_root_enable_workfn(), put_task_struct(p) is called before scx_error() dereferences p->comm and p->pid. If the iterator's reference is the last drop, the task is freed synchronously and the deref becomes a UAF. Move put_task_struct() past scx_error().	7.8	More Details
CVE-2026-61390	There is a heap buffer overflow vulnerability in some Hikvision cameras, which may allow unauthenticated attackers to cause device malfunction by sending specially crafted packets.	7.7	More Details
CVE-2026-13078	A vulnerability was discovered in MongoDB Server where the server-side MozJS scripting engine unconditionally registered a module loading hook that enables JavaScript calls to read arbitrary files from the host filesystem using the mongod process's privileges. An authenticated user could exploit this through crafted aggregation pipeline commands to read sensitive files accessible to the MongoDB server process.	7.7	More Details
	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. In versions 4.0.0 through 5.8.0, 3.3.0		

CVE-2026-59931	through 3.10.6, 2.2.0 through 2.4.6, 2.0.0 through 2.1.17, and all releases up to and including 1.30.5, the WEBSERVICE() domain whitelist can be bypassed via an HTTP redirect (SSRF). In Calculation/Web/Service.php, the webService() method validates a URL's host against the whitelist set via Spreadsheet::setDomainWhiteList(), then fetches content with file_get_contents($url, false, $ctx); because PHP's HTTP stream wrapper follows 301/302 redirects automatically (up to 20 hops) and the redirect target is never re-validated, an attacker who can trigger a redirect from a whitelisted domain can reach arbitrary URLs, including internal addresses. An attacker able to upload XLSX files to an application that uses setDomainWhiteList() and getCalculatedValue() can achieve a full-read SSRF, returning up to 32,767 bytes of the response body as a cell's calculated value, which enables exfiltration of cloud metadata (AWS/GCP/Azure credentials via http://169.254.169.254/), access to internal-only services, and internal port scanning (the port is not validated). This issue has been fixed in versions 5.8.1, 3.10.7, 2.4.7, 2.1.18, and 1.30.6.	7.7	More Details
CVE-2026-64524	In the Linux kernel, the following vulnerability has been resolved: drm/hyperv: validate resolution_count and fix WIN8 fallback A SYNTHVID_RESOLUTION_RESPONSE with resolution_count > 64 walks past the supported_resolution[SYNTHVID_MAX_RESOLUTION_COUNT] array in the parse loop. Bound resolution_count against the array size, folded into the existing zero-check. When the WIN10 resolution probe fails, the caller in hyperv_connect_vsp() left hv->screen_*_max / preferred_* unpopulated, which sets mode_config.max_width / max_height to 0 and makes drm_internal_framebuffer_create() reject every userspace framebuffer with -EINVAL. The pre-WIN10 branch had the same gap for preferred_width / preferred_height. Use a single post-probe fallback guarded by screen_width_max == 0 so both paths converge on the WIN8 defaults.	7.7	More Details
CVE-2026-43820	NIOSSLCertificate._subjectAlternativeNames provides access to the raw bytes for a cert's SANs. NIOSSL provides access to a buffer assumed to be backed by an ASN1_STRING, but not all SANs are backed by ASN1_STRING, so accessing the buffer for such a type can lead to out-of-bounds memory access. This vulnerability is addressed in swift-nio-ssl version 2.37.2.	7.7	More Details
CVE-2026-17527	In containerized-data-importer (CDI), the aggregated cdi.kubevirt.io:view ClusterRole, intended to provide read-only access to CDI resources, includes a rule granting create on the datavolumes/source subresource. CDI's DataVolume clone authorization accepts this permission as sufficient to authorize cloning the contents of any PVC the caller can name, without requiring write access to the source namespace. A user or service account bound to the view role, commonly granted cluster-wide via ClusterRoleBinding, who also has ordinary write access (edit/admin) to any single namespace, can use this to exfiltrate the contents of any PVC in the cluster into a namespace they control, bypassing namespace isolation and the read-only guarantee of the view role.	7.7	More Details
CVE-2026-28896	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. An attacker may be able to cause unexpected system termination or read kernel memory.	7.7	More Details
CVE-2026-63313	9Router before 0.4.72 contains a server-side request forgery (SSRF) vulnerability in the /v1/web/fetch endpoint. The endpoint accepts a user-controlled url parameter and passes it to a configured external scraping provider (Firecrawl, Jina Reader, Tavily, or Exa) to fetch content. The URL is only validated as syntactically valid via new URL() with no blocklist for private IP ranges, cloud metadata endpoints (e.g., 169.254.169.254), link-local addresses, or internal hostnames. An authenticated or locally-connected user can cause the server to fetch arbitrary internal URLs and have the response content returned, enabling read-access SSRF that can expose cloud metadata credentials, reach internal services, and bypass authentication on localhost endpoints.	7.7	More Details
CVE-2026-59542	Subscriber Arbitrary File Deletion in Kali Forms <= 2.4.18 versions.	7.7	More Details
CVE-2026-64456	In the Linux kernel, the following vulnerability has been resolved: hwrng: virtio: clamp device-reported used.len at copy_data() random_recv_done() stores the device-reported used.len directly into vi->data_avail. copy_data() then indexes vi->data[] using vi->data_idx (advanced by previous copy_data() calls) and issues a memcpy() without re-validating either value against the posted buffer size sizeof(vi->data) (SMP_CACHE_BYTES bytes, typically 32 or 64). A malicious or buggy virtio-rng backend can set used.len beyond sizeof(vi->data), steering the memcpy() past the end of the inline array into adjacent kmalloc-1k slab bytes. hwrng_fillfn() mixes those bytes into the guest RNG, and guest root can also observe them directly via /dev/hwrng. Concrete impact is inside the guest: - Memory-safety / hardening: any virtio-rng backend that over-reports used.len causes the driver to read past vi->data into unrelated slab contents. hwrng_fillfn() is a kernel thread that runs as soon as the device is probed; no guest userspace interaction is required to first-trigger the OOB. - Cross-boundary leak (confidential-compute threat model): a malicious hypervisor cooperating with a malicious or compromised guest root userspace can use /dev/hwrng as a leak channel for guest-kernel heap data. The host sets a large used.len, guest root reads /dev/hwrng, and the returned bytes contain guest kernel slab contents that were adjacent to vi->data. In practice, confidential-compute guests (SEV-SNP, TDX) usually disable virtio-rng entirely, so this path is narrow, but the fix is still worth carrying because the underlying memory-safety bug contaminates the guest RNG on any host. KASAN confirms the OOB on a 7.1-rc4 guest whose virtio-rng backend has been patched to report used.len = 0x10000: BUG: KASAN: slab-out-of-bounds in virtio_read+0x394/0x5d0 Read of size 64 at addr ffff88800ae0ba20 by task hwrng/52 Call Trace: __asan_memcpy+0x23/0x60 virtio_read+0x394/0x5d0 hwrng_fillfn+0xb2/0x470 kthread+0x2cc/0x3a0 Allocated by task 1: probe_common+0xa5/0x660 virtio_dev_probe+0x549/0xbc0 The buggy address belongs to the object at ffff88800ae0b800 which belongs to the cache kmalloc-1k of size 1024 The buggy address is located 0 bytes to the right of allocated 544-byte region [ffff88800ae0b800, ffff88800ae0ba20) Same class of bug as commit c04db81cd028 ("net/9p: Fix buffer overflow in USB transport layer"), which hardened usb9pfs_rx_complete() against unchecked device-reported length in the USB 9p transport. With the clamp at point of use and array_index_nospec() in place, the same harness boots cleanly: copy_data() returns zero for the bogus report, the device-supplied bytes after data_idx are discarded, and the driver issues a fresh request.	7.7	More Details
CVE-			

2026-65532	Shop manager SQL Injection in Persian Woocommerce SMS <= 7.2.2 versions.	7.6	More Details
CVE-2026-66427	Administrator SQL Injection in WP Google Review Slider <= 18.4 versions.	7.6	More Details
CVE-2026-59537	Administrator SQL Injection in Sender – Newsletter, SMS and Email Marketing Automation for WooCommerce <= 2.10.22 versions.	7.6	More Details
CVE-2026-65462	Administrator SQL Injection in Uncanny Automator <= 7.3.2 versions.	7.6	More Details
CVE-2026-16313	A flaw was found in sg3_utils. The sg_inq command, when invoked with the --export option, outputs device identification data without sanitizing control characters in SCSI name string fields. A newline character embedded in a device-supplied name string can inject arbitrary properties into the udev device database. This could allow an attacker who can present a crafted SCSI device to execute arbitrary commands as root when the device is disconnected.	7.6	More Details
CVE-2026-66035	libssh2 through 1.11.1, fixed in commit 42e33d8, contains a pre-authentication heap buffer overflow vulnerability that allows a malicious SSH server to corrupt heap metadata in any connecting client by sending a packet with a packet_length smaller than the cipher's block size during Encrypt-then-MAC cipher negotiation. In the fullpacket() function in src/transport.c, the ETM path allocates a buffer of packet_length bytes but copies blocksize minus one bytes via memcpy, causing an overflow that on 32-bit glibc writes attacker-controlled bytes into an adjacent chunk's SIZE field, enabling tcache bin confusion, overlapping live objects, and function pointer overwrite during the session handshake before authentication.	7.5	More Details
CVE-2026-66033	libssh2 through 1.11.1, fixed in commit a2ed82d, contains a pre-authentication integer underflow vulnerability in the ssh2_cipher_crypt() function in src/openssl.c that allows a malicious SSH server to crash any connecting client by negotiating AES-GCM ciphers during handshake. Attackers can exploit the underflow in the expression computing blocksize minus aadlen minus authentication tag length to trigger an out-of-bounds read and a memcpy call with a near-SIZE_MAX length argument, causing immediate process crash before any authentication occurs.	7.5	More Details
CVE-2026-15615	Logto omits validation of the SAML <Conditions> element, enabling attackers to strip time and audience restrictions and replay assertions indefinitely.	7.5	More Details
CVE-2026-66034	libssh2 through 1.11.1, fixed in commit a13bb6c, contains a missing bounds check vulnerability that allows a malicious SSH server to trigger an arbitrary-length heap out-of-bounds read and a free of an uninitialized pointer via the publickey subsystem. In libssh2_publickey_list_fetch(), the version 1 response parser reads a server-controlled comment_len value and advances the parse pointer without verifying sufficient bytes remain in the buffer, causing the out-of-bounds read to leak heap pointers from adjacent allocations defeating ASLR, followed by heap allocator state corruption when the error cleanup path frees an uninitialized pointer from a non-zeroed realloc() region.	7.5	More Details
CVE-2026-15614	Logto silently fails to delete IdP-initiated SAML sessions, enabling session replay and reuse within the session’s validity window.	7.5	More Details
CVE-2026-16756	Missing connection and header-read timeouts and the absence of a concurrent-connection cap in the default serve() path of Amazon aws-smithy-http-server might allow remote attackers to cause a denial of service by opening many connections and sending partial requests that are never completed, exhausting server sockets and tasks. To mitigate this issue, users should upgrade to aws-smithy-http-server 0.66.5 or later.	7.5	More Details
CVE-2026-65919	Meshery before 1.0.57 contains an unauthenticated arbitrary file read vulnerability in the /api/system/fileView and /api/system/fileDownload endpoints that pass user-supplied file parameters directly to os.Open without path validation. Attackers can supply absolute paths or traversal sequences in the file parameter to read arbitrary files from the host filesystem without authentication.	7.5	More Details
CVE-2026-43823	When initializing an RSA public key from DER or PEM bytes throws an error, the EVP_PKEY* is double-freed: first in the catch block, then in the deinit. This can lead to a crash on future memory allocations. This double-free manifests when BoringSSL cannot decode the public key from the bytes provided. This vulnerability is addressed in swift-crypto version 4.5.1.	7.5	More Details
CVE-2026-64430	In the Linux kernel, the following vulnerability has been resolved: NTB: epf: Avoid calling pci_irq_vector() from hardirq context ntb_epf_vec_isr() calls pci_irq_vector() in hardirq context to derive the vector number. pci_irq_vector() calls msi_get_virq() that takes a mutex and can therefore trigger "scheduling while atomic" splats: BUG: scheduling while atomic: kworker/u33:0/55/0x00010001 ... Call trace: ... schedule+0x38/0x110 schedule_preempt_disabled+0x28/0x50 __mutex_lock.constprop.0+0x848/0x908 __mutex_lock_slowpath+0x18/0x30 mutex_lock+0x4c/0x60 msi_domain_get_virq+0xe8/0x138 pci_irq_vector+0x2c/0x60 ntb_epf_vec_isr+0x28/0x120 [ntb_hw_epf] __handle_irq_event_percpu+0x70/0x3a8 handle_irq_event+0x48/0x100 handle_edge_irq+0x100/0x1c8 ... Cache the Linux IRQ number for vector 0 when vectors are allocated and use it as a base in the ISR. Running the ISR in a threaded IRQ handler would also avoid the problem, but that would be unnecessary here.	7.5	More Details
	datamodel-code-generator prior to version 0.70.0 contains a code injection vulnerability that allows attackers who		

CVE-2026-63720	control input schemas to achieve remote code execution by supplying a malicious customBasePath value containing embedded newlines and a dot-free Python expression. The crafted value is emitted verbatim into a generated 'from ... import ...' statement without identifier validation, causing arbitrary Python code to execute when the generated module is imported.	7.5	More Details
CVE-2026-67183	TinyWeb through 0.0.8 contains a memory leak vulnerability that allows unauthenticated attackers to exhaust available memory by sending ordinary well-formed HTTP requests. Each request causes HttpParser::execute() to allocate Url objects, HttpHeaders objects, and HttpHeader instances via raw new expressions that are never freed due to missing destructors and unreachable delete calls, causing worker resident memory to grow monotonically by approximately 20 to 28 kB per request until the worker process is killed.	7.5	More Details
CVE-2026-14490	The Demi – One Click Demo Import, WP Backup & Site Migration plugin for WordPress is vulnerable to Arbitrary Directory Deletion in all versions up to, and including, 0.0.7. The vulnerability exists because the plugin stores its HMAC signing key and per-step restore token as dotfiles inside a publicly accessible subdirectory of the WordPress uploads folder — without any `.htaccess` or index file protection — and the `demi_restore_step` AJAX handler, registered for unauthenticated callers, explicitly accepts possession of the on-disk signing key as a standalone alternative to WordPress capability and nonce checks; an unauthenticated attacker who retrieves the exposed key can forge a valid signed state envelope to invoke `CleanDir::execute()` with a caller-supplied absolute path that is subject to no allow-list or path-canonicalization check. This makes it possible for unauthenticated attackers to recursively delete arbitrary directories on the server.	7.5	More Details
CVE-2026-17524	Versions of the package zip-lib before 1.1.0 are vulnerable to Directory Traversal via the caching mechanism for path validation during the extraction process. An attacker can bypass security checks designed to prevent directory traversal. The intended security function, isOutsideTargetFolder, only checks and caches the path status when the initial directory symlink is created during the first extraction.	7.5	More Details
CVE-2026-12741	The WP Fast Total Search – The Power of Indexed Search plugin for WordPress is vulnerable to generic SQL Injection via the 'form_data[s]' parameter in all versions up to, and including, 1.80.280 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-66473	Unauthenticated Broken Access Control in Xendit Payment <= 7.1.0 versions.	7.5	More Details
CVE-2026-43777	This issue was addressed with improved input validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A remote attacker may be able to cause a denial of service.	7.5	More Details
CVE-2026-44909	Proxygen lacked a generalized slow-consumer detection mechanism in its core HTTP session layer. A remote, unauthenticated attacker could exploit HTTP/2 flow-control by setting SETTINGS_INITIAL_WINDOW_SIZE to 0 or withholding WINDOW_UPDATE frames, causing the server to buffer complete response bodies in memory indefinitely for stalled streams. By opening many simultaneous streams requesting large resources while preventing the server from transmitting responses, an attacker could induce unbounded memory growth leading to service degradation, resource exhaustion, or denial of service. Versions v2017.01.16.00 through v2026.07.20.00 are affected.	7.5	More Details
CVE-2026-14924	The Tablesome Table WordPress plugin before 1.1.31 does not perform any authentication, capability, or nonce checks in one of its AJAX actions, allowing unauthenticated users to create new published posts and to overwrite arbitrary existing posts and pages.	7.5	More Details
CVE-2026-64414	In the Linux kernel, the following vulnerability has been resolved: netfilter: handle unreadable frags sashiko reports: When an skb with unreadable fragments (such as from devmem TCP, where skb_frags_readable(skb) returns false) is processed by the u32 module, skb_copy_bits() will safely return a negative error code [..] xt_u32: bail out with hotdrop in this case. gather_frags: return -1, just as if we had no fragment header. nfnetlink_queue: restrict to the linear part. nfnetlink_log: restrict to the linear part. v2: - skb_zerocopy helpers don't copy readable flag, i.e. nfnetlink_queue is broken too xt_u32 shouldn't return true if hotdrop was set.	7.5	More Details
CVE-2026-67182	Rouille 0.3.3 through 3.6.2 contains an HTTP request smuggling vulnerability that allows remote attackers to bypass access controls by injecting bare line feed characters (0x0A) into client-supplied request header values that are copied verbatim to upstream connections without validation. Attackers can craft a header value containing a complete additional HTTP request that is interpreted as a separate request by backends such as Go net/http and Python http.server, causing the backend to process a smuggled request with attacker-chosen method, path, and headers that bypasses the rouille handler's access control logic.	7.5	More Details
CVE-2026-64395	In the Linux kernel, the following vulnerability has been resolved: ksmbd: require source read access for duplicate extents FSCTL_DUPLICATE_EXTENTS_TO_FILE passes the source file directly to vfs_clone_file_range() or vfs_copy_file_range() without checking the SMB access mask granted to the source handle. A handle opened with attribute access can consequently be used to copy file contents into an attacker-readable destination. Require FILE_READ_DATA on the source handle before either VFS operation, matching other ksmbd data-copy paths.	7.5	More Details
CVE-2026-14603	The WowOptin: Next-Gen Popup Maker WordPress plugin before 1.4.38 does not have proper authorization on a REST endpoint, allowing unauthenticated users to disable all of the site's opt-in forms and insert new template-based opt-in rows into the database.	7.5	More Details
	Quinn is a pure-Rust, async-compatible implementation of the IETF QUIC transport protocol. Starting in version 0.1.0		

CVE-2026-25800	and prior to version 0.11.15, the `Assembler` component that assembles unordered stream fragments into consecutive chunks of the stream incurs some overhead for non-contiguous fragments. Readers that read from a `RecvStream` in order (through an `AsyncRead` impl for example) will be sensitive to peers that send fragments while leaving out early parts of the stream, and in particular, fragments with many gaps (because these cannot be defragmented). In such a scenario, the receiving connection suffers from high buffer overhead, enabling memory exhaustion. Version 0.11.15 fixes the issue.	7.5	More Details
CVE-2026-12383	A flaw was found in the Event-Driven Ansible (EDA) server. The ExternalEventStreamViewSet uses permissive access controls (permission_classes=[AllowAny], authentication_classes=[]) and relies solely on the Subject HTTP header value for mTLS authentication without verifying that the header originated from a trusted proxy. Additionally, the expected certificate Distinguished Name is leaked in the 403 error response body. An attacker who can reach the EDA API endpoint with a spoofed Subject header can inject arbitrary events into mTLS-protected event streams, triggering downstream automation actions.	7.5	More Details
CVE-2026-40430	Pronetiqs IntraVUE Versions 3.2.1a14 and prior have a plaintext storage of a password vulnerability that could expose cleartext credentials through the API.	7.5	More Details
CVE-2026-64208	In the Linux kernel, the following vulnerability has been resolved: crypto/krb5, rxrpc: Fix lack of pre-decrypt/pre-verify length checks Change the krb5 crypto library to provide facilities to precheck the length of the message about to be decrypted or verified. Fix AF_RXRPC to make use of this to validate DATA packets secured with RxGK.	7.5	More Details
CVE-2026-62430	Accesses to the CMOS memory contents are done using an indirect IO port pair. Therefore Xen needs to cache the guest chosen index, and one of the usages of the index didn't take the necessary locking to avoid concurrent changes. As a result, a guest could change the index after it being checked, causing a subsequent out-of-bound read access to the contents of an array.	7.5	More Details
CVE-2026-64312	In the Linux kernel, the following vulnerability has been resolved: crypto: pcrypt - restore callback for non-parallel fallback pcrypt installs pcrypt_aead_done() on the child AEAD request before trying to submit it through padata. If padata_do_parallel() returns -EBUSY, pcrypt falls back to calling the child AEAD directly. That fallback must not keep the padata completion callback. Otherwise an asynchronous completion runs pcrypt_aead_done() even though the request was never enrolled in padata. Restore the original request callback and callback data before calling the child AEAD directly. This keeps the fallback path aligned with a direct AEAD request while leaving the parallel path unchanged.	7.5	More Details
CVE-2026-62431	The logic to handle periodic Viridian STIMERs performs a division with an unchecked user-controlled divisor value, that can be set to zero to cause a #DE fault.	7.5	More Details
CVE-2026-65881	Joomla Extension - joomdle.com - Insecure default configuration allows read/write user account access in Joomdle < 3.1.1 - The default configuration of the extension allowed read access and password reset of CMS accounts.	7.5	More Details
CVE-2026-66144	Although remote policy references are not retrieved during policy normalization, if they are manually retrieved via the API it can cause a denial of service attack if a huge policy is retrieved. Users are recommended to upgrade to version 3.2.3, which fixes this issue by imposing a default maximum size on data read from remote policy references.	7.5	More Details
CVE-2026-66143	It is possible to bypass the maximum number of normalized policy alternatives that was introduced in Apache Neethi 3.2.2 via certain crafted policies, which may lead to a denial of service attack via resource consumption. Users are recommended to upgrade to version 3.2.3, which fixes this issue.	7.5	More Details
CVE-2026-66142	Apache Neethi is vulnerable to uncontrolled recursion when parsing policies that lack policy Ids or with deeply nested structures, which may lead to a denial of service attack when parsing policies due to runtime memory exhaustion. Users are recommended to upgrade to version 3.2.3, which fixes this issue.	7.5	More Details
CVE-2026-59878	Improper Input Validation vulnerability in Apache ActiveMQ AMQP, Apache ActiveMQ, Apache ActiveMQ All. A remote unauthenticated peer that can reach an exposed AMQP NIO connector can trigger denial-of-service behavior by sending a frame size value. This cause the NIO threads to die and if done rapidly enough can lead to exhaustion of the NIO thread pool denying service to other connections. This issue affects Apache ActiveMQ AMQP: before 5.19.9, from 6.0.0 before 6.2.8; Apache ActiveMQ: before 5.19.9, from 6.0.0 before 6.2.8; Apache ActiveMQ All: before 5.19.9, from 6.0.0 before 6.2.8. Users are recommended to upgrade to version 5.19.9, 6.2.8, or 6.3.0 which fixes the issue.	7.5	More Details
CVE-2026-45816	NULL Pointer Dereference vulnerability in Apache NimBLE in LE Long Term Key Request event. This requires disabled asserts (otherwise assert would trigger before NULL dereference) and bogus (or misbehaving) controller, thus severity is low. This issue affects Apache NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.	7.5	More Details
CVE-2026-45815	Reachable Assertion vulnerability in Apache NimBLE. A specially crafted ATT Read Multiple Variable Response (BLE_ATT_OP_READ_MULT_VAR_RSP) may trigger assert in ATT parser. Severity is medium as this requires DUT to first send ATT Read Multiple Variable Request. This issue affects Apache NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.	7.5	More Details
CVE-2026-45811	Buffer Copy without Checking Size of Input ('Classic Buffer Overflow') vulnerability in Apache NimBLE. The HCI socket transport did not check whether a received HCI event would fit the configured event pool before copying it, allowing a buffer overflow. Severity is low: exploitation requires either a misconfigured pool size or a malicious/compromised controller on the other end of the HCI socket link, not over-the-air Bluetooth access. This issue affects Apache	7.5	More Details

	NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.		
CVE-2026-51251	Schreibfaul1 ESP32-audioI2S 3.4.5 has a buffer overflow vulnerability in the MP3Decoder::decode() function of the MP3 decoder due to missing size validation on untrusted mainDataBegin and nSlots values.	7.5	More Details
CVE-2026-66299	Uncontrolled Resource Consumption vulnerability in Apache Tomcat's WebSocket chat example. This issue affects Apache Tomcat: from 11.0.0-M20 through 11.0.24, from 10.1.24 through 10.1.57, from 9.0.89 through 9.0.120. Users who have followed the security guidance to remove the examples web application are not affected by this issue. Users are recommended to remove the examples web application or to upgrade to version 11.0.25, 10.1.58 or 9.0.121 (when released), which fix the issue.	7.5	More Details
CVE-2026-43728	This issue was addressed through improved state management. This issue is fixed in macOS Tahoe 26.6. An attacker may be able to modify the state of the Keychain.	7.5	More Details
CVE-2026-66373	Redis before 8.8.0, in the unusual case where an authenticated attacker can execute RESTORE, allows remote code execution via a RESTORE payload where the same NACK (pending entry) is referenced by more than one consumer, because deleting both consumers via XGROUP DELCONSUMER leads to a double free. NOTE: this issue exists because of an incomplete fix for CVE-2026-25243.	7.5	More Details
CVE-2026-12497	The Paid Membership Plugin, Ecommerce, User Registration Form, Login Form, User Profile & Restrict Content WordPress plugin before 4.16.18 does not consistently enforce the role restriction configured on its front-end registration role-selection field. The set of roles offered to the visitor and the set of roles the registration handler accepts are derived by two different parsers, and for some valid ways of configuring the offered roles the handler ignores the restriction and falls back to accepting any non-administrator role. Combined with the absence of a nonce on the public registration handler, this allows an unauthenticated visitor to register an account with a higher role, such as Editor or Author, than the form was configured to offer.	7.5	More Details
CVE-2026-47427	GitHub MCP Server is GitHub's official MCP Server. Prior to 1.1.0, the CompletionsHandler function in pkg/github/server.go accesses params.Ref without first checking whether it is nil, so a completion/complete request with a missing or empty ref field triggers a nil pointer dereference and a Go runtime panic; because the crash occurs before any authentication or token validation, any unauthenticated client able to send JSON-RPC messages can crash the server, resulting in a complete denial of service. This issue is fixed in version 1.1.0.	7.5	More Details
CVE-2026-64210	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: xsk: Fix unlocked writing to ICOSQ During napi poll, when the affinity changes and there's still XSK work to be done, we trigger an ICOSQ interrupt on the new CPU. However, this triggering on the ICOSQ is done unprotected. There are 2 such races: A) mlx5e_trigger_irq() is called while mlx5e_xsk_alloc_rx_mpwqe() is running from a different CPU due to affinity change. This can happen because IRQ triggering is done after napi_complete_done(). At this point the NAPI can be scheduled on a different CPU. Like this: CPU A (old affinity, NAPI tail) CPU B (new affinity, fresh NAPI) ---------- ---------- napi_complete_done() clears SCHED mlx5e_cq_arm(...) napi_schedule_prep() sets SCHED mlx5e_napi_poll() mlx5e_xsk_alloc_rx_mpwqe() mlx5e_icosq_sync_lock() // noop memcpy 640 B UMR body advance sq->pc by 10 mlx5e_trigger_irq(&c->icosq) wqe_info[pi] = {NOP, 1} mlx5e_post_nop() advances sq->pc B) mlx5e_trigger_irq() is called on the ICOSQ when mlx5e_trigger_napi_icosq() is running. The obvious fix would be to lock the ICOSQ. But ICOSQ has an optimized locking scheme that doesn't work for this scenario. Kick the async ICOSQ instead which is always locked. This issue was noticed in the wild with the following splat: netdevice: ge-0-0-1: Bad OP in ICOSQ CQE: 0xd WARNING: drivers/net/ethernet/mellanox/mlx5/core/en_rx.c:826 [...] [...] Call Trace: <IRQ> mlx5e_napi_poll+0x11d/0x7f0 [mlx5_core] __napi_poll+0x30/0x200 ? skb_defer_free_flush+0x9c/0xc0 net_rx_action+0x2fe/0x3f0 handle_softirqs+0xd8/0x340 __irq_exit_rcu+0xbc/0xe0 common_interrupt+0x85/0xa0 </IRQ> <TASK> asm_common_interrupt+0x26/0x40 [...] ---[end trace 0000000000000000]--- mlx5_core 0000:08:00.0 ge-0-0-1: Error cqe on cqn 0x548, ci 0x2022, qn 0x8f4, opcode 0xd, syndrome 0x2, vendor syndrome 0x68 00000000: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000010: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000020: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000030: 00 00 00 00 01 00 68 02 01 00 08 f4 de 14 59 d2 WQE DUMP: WQ size 16384 WQ cur size 0, WQE index 0x1e14, len: 64 00000000: 00 00 00 01 d9 ed 80 02 00 00 00 01 d9 ed 90 02 00000010: 00 00 00 01 d9 ed a0 02 00 00 00 01 d9 ed b0 02 00000020: 00 00 00 01 d9 ed c0 02 00 00 00 01 d9 ed d0 02 00000030: 00 00 00 01 d9 ed e0 02 00 00 00 01 d9 ed f0 02 mlx5_core 0000:08:00.0 ge-0-0-1: Error cqe on cqn 0x548, ci 0x2023, qn 0x8f4, opcode 0xd, syndrome 0x5, vendor syndrome 0xf9 00000000: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000010: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000020: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00000030: 00 00 00 00 01 00 f9 05 01 00 08 f4 de 15 cf d2	7.5	More Details
	In the Linux kernel, the following vulnerability has been resolved: sched/rt: Have RT_PUSH_IPI be default off for non PREEMPT_RT RT migration is done aggressively. When a CPU schedules out a high priority RT task for a lower priority task, it will look to see if there's any RT tasks that are waiting to run on another CPU that is of higher priority than the task this CPU is about to run. If it finds one, it will pull that task over to the CPU and allow it to run there instead. Normally, this pulling is done by looking at the RT overloaded mask (rto) which contains all the CPUs in the scheduler domain with RT tasks that are waiting to run due to a higher priority RT task currently running on their CPU. The CPU that is about to schedule a lower priority task will grab the rq lock of the overloaded CPU and move the RT task from that CPU's runqueue to the local one and schedule the higher priority RT task. This caused issues when a lot of CPUs would schedule a lower priority task at the same time. They would all try to grab the same runqueue lock of the CPU with the overloaded RT tasks. Only the first CPU that got in will get that task. All the others would wait until they got the runqueue lock and see there's nothing to pull and do nothing. On systems with lots of CPUs, this caused a large latency (up to 500us) which is beyond what PREEMPT_RT is to allow. The solution to that was to create an RT_PUSH_IPI logic. When any CPU wanted to pull a task, instead of grabbing the runqueue lock of the overloaded CPU, it would start by sending an IPI to the overloaded CPU, and that IPI handler would have the CPU with the waiting RT		

CVE-2026-64374	task do a push instead. Then that handler would send an IPI to the next CPU with overloaded RT tasks, and so on. Note, after the first CPU starts this process, if another CPU wanted to do a pull, it would see that the process has already begun and would only increment a counter to have the IPIs continue again. The RT_PUSH_IPI solved the latency problem with PREEMPT_RT but could cause a new issue with non PREEMPT_RT. Namely, softirqs run in a threaded context on PREEMPT_RT but they can run in an interrupt context in non-RT. If an IPI lands on a CPU that has just woken up multiple RT tasks and the current CPU is running a non RT or a low priority RT task, instead of doing a push, it would simply do a schedule on that CPU. But if a softirq was also executing on this CPU, the schedule would need to wait until the softirq finished. Until then, the CPU would still be considered overloaded as there are RT tasks still waiting to run on it. A live lock occurred on a workload that was doing heavy networking traffic on a large machine where the softirqs would run 500us out of 750us. And it would also be waking up RT tasks, causing the RT pull logic to be constantly executed. When a softirq triggered on a CPU with RT tasks queued but not running yet, and the other CPUs would see this CPU as being overloaded, they would send an IPI over to it. The CPU would notice that the waiting RT tasks are of higher priority than the currently running task and simply schedule that CPU instead. But because the softirq was executing, before it could schedule, it would receive another IPI to do the same. The amount of IPIs would slow down the currently running softirq so much that before it could return back to task context, it would execute another softirq never allowing the CPU to schedule. This live locked that CPU. As RT_PUSH_IPI was created to help PREEMPT_RT, make it default off if PREEMPT_RT is not enabled.	7.5	More Details
CVE-2026-65694	Microweber CMS through 2.0.20 contains a path traversal vulnerability in the static file controller that allows unauthenticated remote attackers to read arbitrary files by supplying directory traversal sequences in the path query parameter. Attackers can send a single unauthenticated HTTP GET request exploiting the failure of normalize_path() to strip traversal sequences, disclosing sensitive files such as environment configuration files containing credentials and system files.	7.5	More Details
CVE-2026-14516	The Online Scheduling and Appointment Booking System – Bookly plugin for WordPress is vulnerable to time-based SQL Injection via the 'staff_ids' parameter in all versions up to, and including, 27.5 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. Exploitation requires a two-request chain: an attacker first calls the unauthenticated bookly_get_form_id action to seed a booking session carrying malicious staff_ids values, then triggers bookly_render_time to cause the tainted array to reach the vulnerable query; CSRF/nonce validation is absent on both endpoints, meaning this chain can be initiated cross-site.	7.5	More Details
CVE-2026-51244	schreibfaul1 ESP32-audioI2S 3.4.5 has a buffer overflow vulnerability in UnpackFrameHeader(). Multiple attacker-controlled index parameters are used to access static and heap table arrays without range limitation. Invalid index values lead to out-of-bounds memory writing and heap buffer overflow.	7.5	More Details
CVE-2026-67184	TinyWeb through 0.0.8 contains a null pointer dereference vulnerability that allows unauthenticated remote attackers to crash worker processes by sending a malformed HTTP request line with an invalid version string. The HttpParser::execute() function fails to allocate the Url object when version parsing fails, leaving the url pointer NULL, and buildResponse() subsequently dereferences this NULL pointer without checking the valid_requ flag, producing a SIGSEGV that terminates the worker process and, when repeated across all workers, takes the server permanently offline until manually restarted.	7.5	More Details
CVE-2026-12800	The Premium Packages – Sell Digital Products Securely plugin for WordPress is vulnerable to SQL Injection via the 'code' parameter of the POST /wp-json/wpdmpp/v1/cart/coupon REST API endpoint in versions up to, and including, 6.2.0. This is due to insufficient escaping on the user-supplied parameter, which is interpolated directly into a raw SQL query string in the CouponCodes::find() method without use of $wpdb->prepare() or esc_sql(). This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-13161	The TrueBooker – Appointment Booking and Scheduler System plugin for WordPress is vulnerable to generic SQL Injection via the 'alldata[truebooker_user]' parameter in all versions up to, and including, 1.2.2 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. The check_ajax_referer() nonce guard does not constitute an authentication or authorization barrier because the nonce is exposed to unauthenticated visitors on TrueBooker front-end booking pages; exploitation additionally requires that the required booking fields (category, service, person, date, and time slot) be present in the alldata POST parameter so that execution reaches the vulnerable SQL query branch.	7.5	More Details
CVE-2026-64281	In the Linux kernel, the following vulnerability has been resolved: svcrdma: wake sq waiters when the transport closes Threads parked in svc_rdma_sq_wait() on sc_sq_ticket_wait or sc_send_wait can hang indefinitely in TASK_UNINTERRUPTIBLE state across transport teardown, pinning svc_xprt references and blocking svc_rdma_free(). The close path sets XPT_CLOSE before invoking xpo_detach and both wait_event predicates include an XPT_CLOSE term, but the predicates are re-evaluated only on wakeup. sc_sq_ticket_wait has no completion-driven wake path; it is advanced solely by the chained ticket handoff inside svc_rdma_sq_wait() itself. Without an explicit wake at close, parked threads never observe XPT_CLOSE, hold their svc_xprt_get reference forever, and svc_rdma_free() blocks on xpt_ref dropping to zero. Two close entry points reach this transport. Local teardown runs svc_rdma_detach() from svc_handle_xprt() -> svc_delete_xprt() -> xpo_detach() on a worker thread. A remote disconnect arrives at svc_rdma_cma_handler(), which calls svc_xprt_deferred_close(): that sets XPT_CLOSE and enqueues the transport but does not access either RDMA waitqueue, so a worker already parked in svc_rdma_sq_wait() never re-evaluates its predicate. With every worker parked on this transport, no thread is available to run the local teardown either, and the wake site there is unreachable. Introduce svc_rdma_xprt_deferred_close(), a thin svcrdma wrapper that calls svc_xprt_deferred_close() and then wakes both sc_sq_ticket_wait and sc_send_wait. Convert the svcrdma producers that called svc_xprt_deferred_close() directly: svc_rdma_cma_handler(), qp_event_handler(),	7.5	More Details

	svc_rdma_post_send_err(), svc_rdma_wc_send(), the sendto drop path, the rw completion error paths, and the recvfrom flush and read-list error paths. Wake both waitqueues from svc_rdma_detach() as well. The synchronous svc_xprt_close() path (backchannel ENOTCONN, device removal via svc_rdma_xprt_done) reaches detach without flowing through svc_xprt_deferred_close() and therefore does not invoke the new helper. [cel: add svc_rdma_xprt_deferred_close() to complete the fix]		
CVE-2026-10697	Improper Authentication vulnerability in Progress MOVEit Transfer. This issue affects MOVEit Transfer: before 2025.1.5, from 2026.0.0 before 2026.0.3.	7.5	More Details
CVE-2026-15966	Permissive cross-domain security policy with untrusted domains vulnerability in Progress MOVEit Transfer. This issue affects MOVEit Transfer: before 2025.1.5, from 2026.0.0 before 2026.0.3.	7.5	More Details
CVE-2026-15967	Insufficient session expiration vulnerability in Progress MOVEit Transfer. This issue affects MOVEit Transfer: before 2025.1.5, from 2026.0.0 before 2026.0.3.	7.5	More Details
CVE-2025-63913	An issue was discovered in OpenSBI 1.3 allowing attackers to cause a denial of service via crafted request to the SBI function #2 or the 'Find and configure a matching counter' function of SBI PMU extension.	7.5	More Details
CVE-2026-42493	Addressing certain issues, in particular related to operations which may take excessively long and therefore would need preemption, has turned out overly costly. Since alternatives (HVM/PVH: HAP, PV: shim) are commonly available, the decision was to deprecate the functionality, while still retaining it for people to use at their own (security) risk. Memory-wise small enough guests may still be okay to run.	7.5	More Details
CVE-2026-10207	The PickPlugins Question Answer plugin for WordPress is vulnerable to SQL Injection in versions up to and including 1.2.73. This is due to insufficient sanitization of user-supplied input via the 'id' GET parameter in the user profile template combined with the use of wp_unslash() which removes WordPress's magic quotes protection, followed by direct concatenation into a SQL query without proper escaping or prepared statements in the qa_user_profile_card() function. This makes it possible for unauthenticated attackers to append additional SQL queries into existing queries, which can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-14785	The Web Directory Free plugin for WordPress is vulnerable to generic SQL Injection via the 'levels' parameter in all versions up to, and including, 1.7.13 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-15025	The Uncanny Automator – Easy Automation, Integration, Webhooks & Workflow Builder plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 7.3.2 via the automator_google_contacts_fetch_labels, automator_mautic_segment_fetch, automator_mautic_tags_fetch, and automator_mautic_render_contact_fields AJAX actions due to a missing capability check and missing nonce verification in the corresponding handlers (ajax_fetch_labels, segments_fetch, tags_fetch, and render_contact_fields). This makes it possible for authenticated attackers, with Subscriber-level access and above, to enumerate sensitive Google Contacts groups/labels and Mautic segments, tags, and contact-field definitions retrieved via integration credentials configured by an administrator, and to consume third-party API quota.	7.5	More Details
CVE-2026-50032	A NULL pointer dereference in the MMS Write Named Variable List handler, which may allow a network adjacent attacker to crash the server by sending a WriteRequest with an empty listOfData field.	7.5	More Details
CVE-2026-50039	The affected product is vulnerable to a stack-based buffer overflow, which may allow an attacker to cause a memory corruption via a Read Request.	7.5	More Details
CVE-2026-61609	Pterodactyl is a free, open-source game server management panel. From 1.7.0 until 1.13.0, the authentication rate limiter defined in RouteServiceProvider::configureRateLimiting() applied a single global bucket to the login and two-factor checkpoint endpoints instead of keying by IP or account: the fall-through Limit::perMinute(10) covering POST /auth/login and POST /auth/login/checkpoint omitted ->by(), so Laravel derived a constant cache key (md5('authentication')) shared by every request. An unauthenticated attacker sending roughly ten requests per minute from a single IP, most cheaply against the checkpoint endpoint (which has no reCAPTCHA), exhausts the shared counter and causes HTTP 429 for every user attempting to log in or complete two-factor authentication, a panel-wide authentication denial of service that also locks out administrators. This issue is fixed in version 1.13.0.	7.5	More Details
CVE-2026-51077	SQL injection vulnerability in Dede CMS v.5.7.118 allows a remote attacker to obtain sensitive information via the sqlquery parameter of the sys_sql_query.php component	7.5	More Details
CVE-2026-42492	Xenstore, to have an up-to-date picture of the entire system, wants to know of domains appearing and disappearing. To make this more robust, a new XEN_DOMCTL_get_domain_state was introduced. The management of the bitmap underlying that operation is tied into the binding of the VIRQ_DOM_EXC virtual IRQ. Unfortunately an error path there would tear down the bitmap even in cases when it wasn't set up. Unprivileged domains can trigger that error path.	7.5	More Details
CVE-2026-	An issue in Dede CMS v.5.7.118 allows a remote attacker to obtain sensitive information via the str parameter of the file_manage_control.php component	7.5	More Details

51078			
CVE-2026-12981	The CAFEHAUS API WordPress plugin through 1.0.0 does not have any authentication or authorisation when updating user passwords, allowing unauthenticated attackers to set the password of any user, including administrators, and fully take over their accounts.	7.5	More Details
CVE-2026-13183	In Progress® Telerik® UI for AJAX prior to v2026.2.708, RadAsyncUpload upload metadata processing may leak cryptographic validity through measurable timing differences, enabling remote attackers to recover protected metadata values.	7.5	More Details
CVE-2026-63683	Joomla Extension - regularlabs.com - Client IP spoofing vulnerability in Regular Labs conditions manager - IP and GeoIP conditions trusted spoofable forwarded headers, allowing remote clients to bypass location-based rules.	7.5	More Details
CVE-2026-45112	Allocation of Resources Without Limits or Throttling vulnerability in Apache Thrift Java bindings. This issue affects Apache Thrift: from 0.19.0 before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-65493	Subscriber PHP Object Injection in Dokan Pro <= 5.0.2 versions.	7.5	More Details
CVE-2026-64797	Joomla Extension - regularlabs.com - IP spoofing vulnerability in IP login extension - IP Login trusted forwarded client-IP headers without requiring a configured trusted proxy. Attackers could spoof the IP used for automatic login and potentially impersonate mapped accounts.	7.5	More Details
CVE-2026-60370	Vulnerability in the Oracle Platform Security for Java product of Oracle Fusion Middleware (component: Centralized Thirdparty Jars). Supported versions that are affected are 12.2.1.4.0 and 14.1.2.0.0. Difficult to exploit vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Platform Security for Java. Successful attacks of this vulnerability can result in takeover of Oracle Platform Security for Java. CVSS 3.1 Base Score 7.5 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:L/UI:N/S:U/C:H/I:H/A:H).	7.5	More Details
CVE-2026-15074	@fastify/static up to and including version 10.1.0 fails to reject dot-dot path segments in request pathnames before the file-resolution stage. This is a bypass of the earlier fix for CVE-2026-6414, which only covered encoded forward slashes. Because the underlying send library normalizes dot segments before applying its own path-traversal guard, an unauthenticated attacker can bypass any route-scoped middleware and read files inside the static root that live under the guarded URL prefix. The bypass does not allow access outside the configured static root by itself, it defeats route-guard filtering only. The issue is patched in @fastify/static 10.1.1.	7.5	More Details
CVE-2026-65495	Unauthenticated Broken Access Control in Dokan Pro <= 5.0.3 versions.	7.5	More Details
CVE-2026-12082	The Praison AI SEO WordPress plugin before 5.0.7 does not perform authorization checks on several of its REST API routes, allowing unauthenticated users to modify the permalink of any published post and to read Praison AI SEO WordPress plugin before 5.0.7 configuration data.	7.5	More Details
CVE-2026-14981	IBM WebSphere Application Server 9.0, and 8.5 and IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.7 are affected by a denial of service vulnerability in the HTTP channel due to unbounded allocation of resources without limits.	7.5	More Details
CVE-2026-14291	The security-ninja-premium WordPress plugin before 5.290 does not verify the second authentication factor in one of its two-factor authentication code paths, allowing an unauthenticated attacker who knows a user's password to complete authentication without the one-time code and bypass enforced two-factor authentication for any account, including administrators. The affected two-factor module ships only in the premium build.	7.5	More Details
CVE-2026-64834	FFmpeg versions 0.6.3 through 8.1.2 contain an infinite loop vulnerability in the RTP/ASF demuxer within libavformat/rtpdec_asf.c that allows remote attackers to cause denial of service by sending a crafted RTP/ASF stream. The rtp_asf_fix_header function fails to validate a minimum chunksize when iterating over ASF objects, causing the loop pointer to never advance when a chunksize is smaller than the 24-byte minimum ASF object header size, resulting in CPU exhaustion that denies service to legitimate users.	7.5	More Details
CVE-2026-44690	In NLnet Labs Unbound 1.7.0 up to and including 1.25.1, insufficient validation of the RRSIG.Labels field combined with premature cache writes during RFC 8198 aggressive NSEC processing leads to cache poisoning that permits a malicious actor controlling a single delegated zone to poison arbitrary sibling zones under NSEC-signed parent domains. A malicious actor with one registered domain under an NSEC-signed TLD can serve malicious insecure DNS responses for unrelated sibling domains (sharing the same parent zone). Arbitrary delegations that do not exist under the parent domain and are covered by the parent's NSEC chain can be brought into insecure existence by fraudulent wildcard DS records (less labels than expected, unknown algorithm) from the malicious sibling domain. This allows the malicious actor to inject insecure wildcard records for those delegations.	7.5	More Details
CVE-2026-65500	Unauthenticated Broken Access Control in Manual - Documentation, Knowledge Base & Education WordPress Theme <= 7.5.4 versions.	7.5	More Details
CVE-			More

2026-59528	Subscriber Sensitive Data Exposure in ShipTime: Discounted Shipping Rates <= 1.1.1 versions.	7.5	Details
CVE-2026-15057	IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.7 is vulnerable to a denial of service due to uncontrolled heap allocation.	7.5	More Details
CVE-2026-59548	Unauthenticated Sensitive Data Exposure in Byteflows Travel & Hotel Booking <= 1.0.0 versions.	7.5	More Details
CVE-2026-45623	PostCSS takes a CSS file and provides an API to analyze and modify its rules by transforming the rules into an Abstract Syntax Tree. In versions 8.5.11 and prior, the PreviousMap parses the /*# sourceMappingURL=PATH */ comment from any CSS string passed to process() and dereferences PATH against the local filesystem with no scheme, allowlist, or traversal check. An attacker who controls the CSS input can cause the host process to read any file readable by Node and leak the first ~10 bytes of its content through the resulting JSON.parse SyntaxError message. The bug also yields a precise file-existence oracle and a controllable-read primitive that may be combined with large-file targets for DoS. The behaviour is triggered with PostCSS's default options — no from, no map, no plugins required — and is therefore reachable from any pipeline that runs untrusted CSS through PostCSS (CMS themes, user-uploaded styles, browser-extension/userstyle processors, build pipelines for third-party packages, blog comment renderers, etc.). This issue has been fixed in version 8.5.12.	7.5	More Details
CVE-2026-40691	In Unbound 1.9.0 up to and including 1.25.1, when a DNSCrypt query is received over TCP, the routine that encrypts the reply in place fails to bound the reply length against the destination buffer size. The size clamp that protects the UDP path is not applied on the TCP path, so a reply larger than 65504 bytes is shifted forward by 48 bytes inside a buffer of capacity equal to 'msg-buffer-size', writing past the end of the heap allocation. A single malicious encrypted query crashes the resolver and lead to denial of service. This vulnerability needs Unbound to be compiled with DNSCrypt support ('--enable-dnscrypt') and the 'dnscrypt:' clause to be configured and enabled for the listening interfaces.	7.5	More Details
CVE-2026-48145	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift C++ bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-14235	The Download Manager WordPress plugin before 3.3.62 does not bind its temporary download token to the requesting session nor expire it promptly, making the token a long-lived, multi-use, portable bearer token, so that an attacker who obtains one leaked download key can repeatedly download a role- or password-protected package file without authorization.	7.5	More Details
CVE-2026-64792	Joomla Extension - regularlabs.com - disclosure of restricted content via search index in various Regular Labs extensions - Smart Search indexing could render generated content using the indexing administrator’s identity instead of a public guest. Restricted or administrator-only content could consequently be stored in the public search index and disclosed to visitors.	7.5	More Details
CVE-2026-66731	facil.io 0.7.5 through 0.7.6 contains a denial-of-service vulnerability in the HTTP/1.1 chunked transfer encoding parser that allows unauthenticated remote attackers to crash the server by sending a negative chunk size value. Attackers can send a single POST request with a Transfer-Encoding: chunked header containing a leading minus sign in the chunk size field, causing the parser in http1_parser.h to compute a large positive integer from the negated value, corrupting internal state and moving the read pointer into unmapped memory resulting in a fault.	7.5	More Details
CVE-2026-9713	The Lumise Product Designer for WooCommerce plugin for WordPress is vulnerable to SQL Injection via the 'id' and 'table' parameters in the uploaded cart JSON file processed by the checkout AJAX action in versions up to, and including, 2.1.1. This is due to insufficient escaping on the user-supplied parameters before they are appended directly to a raw SQL query in the find_resource() function — the 'id' field is interpolated without quotes into a WHERE clause (numeric context) and 'table' is interpolated into the FROM clause, neither of which is protected by wp_magic_quotes or passed through $wpdb->prepare(). This makes it possible for unauthenticated attackers to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	7.5	More Details
CVE-2026-66730	facil.io 0.6.0 through 0.7.6 contains a denial-of-service vulnerability in the multipart body parser that allows an unauthenticated remote attacker to permanently freeze worker processes at 100% CPU by sending a multipart/form-data request with a partial closing boundary. The missing progress guard in the parser loop causes http_mime_parse to return 0 bytes consumed without setting done or error flags, causing the calling loop to re-invoke the parser on the same buffer indefinitely, exhausting all workers and permanently disabling the server until manually restarted.	7.5	More Details
CVE-2026-61954	Unauthenticated Broken Access Control in PayU India <= 3.8.9 versions.	7.5	More Details
CVE-2026-13463	IBM Cloud Pak System 2.3.5.0 could allow a local attacker to obtain sensitive information due to the insertion of credentials into log files.	7.5	More Details
CVE-2026-	Joomla Extension - regularlabs.com - SSRF via remote image downloads in Articles Anywhere and Users Anywhere extensions - Content-controlled image URLs could request private or reserved network services, follow unsafe redirects and save responses without validating that they were images. This could result in SSRF, internal-data access	7.5	More Details

64799	or writing attacker-controlled files into a web-accessible folder.		
CVE-2026-51304	sqlite 3.41 has a use-after-free (UAF) vulnerability in the ORDER BY clause parsing routine. The affected code first releases the memory of an ExprList object via sqlite3ExprListDelete(), then attempts to access the nExpr member of the already freed object. This dangling pointer access causes invalid memory read operations. By constructing a malicious SQL statement containing an ORDER BY clause with a large number of items, a remote adversary can trigger this vulnerability. Successful exploitation can result in application crash (denial of service), leakage of sensitive memory contents, and under certain memory layout conditions, arbitrary code execution on the affected system.	7.5	More Details
CVE-2026-61943	Unauthenticated Broken Access Control in WPDM – Premium Packages <= 6.2.0 versions.	7.5	More Details
CVE-2026-59554	Unauthenticated Broken Authentication in Ziina <= 1.2.21 versions.	7.5	More Details
CVE-2026-59539	Subscriber Insecure Direct Object References (IDOR) in Paid Member Subscriptions <= 3.0.7 versions.	7.5	More Details
CVE-2026-59547	Unauthenticated Broken Access Control in Payment Gateway for PayPal on WooCommerce <= 9.1.4 versions.	7.5	More Details
CVE-2026-14899	The code to parse MIME headers for display when forwarding a message (if the setting to view all headers was enabled) had an off-by-one error, allowing a single byte to be read from the memory after the buffer for the headers, and potentially crashing Thunderbird. This vulnerability was fixed in Thunderbird 153 and Thunderbird 140.13.	7.5	More Details
CVE-2026-65477	Contributor Local File Inclusion in Tonda Core <= 2.1.2 versions.	7.5	More Details
CVE-2026-55415	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. From 0.11.6 until 0.64.0, datamodel-code-generator allows attacker-controlled x-python-import or customTypePath schema extensions to reach src/datamodel_code_generator/parser/jsonschema.py and generated import handling through Import.from_full_path and Imports.create_line in src/datamodel_code_generator/imports.py, allowing a newline to break out of an import statement and execute Python code when the generated model is imported. This issue is fixed in version 0.64.0.	7.5	More Details
CVE-2026-55391	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. Prior to 0.63.0, datamodel-code-generator validates a URL host once in src/datamodel_code_generator/http.py through get_body, _validate_url_for_fetch, and _get_ips_from_host, but then lets httpx resolve the host again for the connection, allowing DNS rebinding to bypass allow_private_network=False and reach internal services. This issue is fixed in version 0.63.0.	7.5	More Details
CVE-2026-47219	find-my-way is a framework-independent HTTP router that internally uses a Radix Tree and supports route parameters and wildcards. Versions prior to 9.7.0 are vulnerable to remotely triggerable DoS in find-my-way when it is used with Node's HTTP/2 server. The lookup() function passes req.method into find(), and find() indexes this.trees[method]. Since this.trees is a normal object, HTTP/2 method values like constructor, toString, or __proto__ can resolve inherited object properties instead of returning undefined. The code then treats that value like a router node and crashes when it reaches currentNode.prefix.length. This issue has been fixed in version 9.0.7.	7.5	More Details
CVE-2026-66050	NitroShare Desktop through 0.3.4 contains a path traversal vulnerability in its LAN file transfer server that allows unauthenticated attackers on the same network to write arbitrary files by sending a crafted filename containing directory traversal sequences in the JSON item header name field. Attackers can exploit the lack of path validation to write files outside the transfer root directory to arbitrary locations the current user has write access, including the Windows Startup folder, enabling persistent code execution on the next user login.	7.5	More Details
CVE-2026-65481	Contributor Local File Inclusion in Vino <= 1.9 versions.	7.5	More Details
CVE-2026-57600	Insufficient validation of input parameters in the firmware of some Hikvision cameras allows unauthenticated attackers to retrieve partial sensitive data.	7.5	More Details
CVE-2026-66729	facil.io 0.6.0 through 0.7.6 contains an integer underflow vulnerability in the multipart MIME body parser that allows unauthenticated remote attackers to crash the server process by sending a crafted Content-Disposition header with an empty field name. Attackers can trigger a uint32_t wraparound in http_mime_parser.h causing an out-of-bounds memory read past the name pointer, resulting in a bus fault that crashes the handling worker with a single POST request.	7.5	More Details
	OIDC::Lite versions through 0.12.1 for Perl allow ID Token signature verification bypass via a token-controlled algorithm allowlist in verify. When the caller does not pin an algorithm, OIDC::Lite::Model::IDToken::verify sets $self->alg($self->header->{alg}) from the token's own header and then calls decode_jwt(token, key, 1, [$self->alg]),		

CVE-2026-13089	handing JSON::WebToken an accepted-algorithm allowlist taken from the untrusted token. A token with alg=none yields ['none'], so decode_jwt returns the claims with no signature check, and a token with alg=HS256 is verified with the RP's RSA public key as the HMAC secret (RS to HS confusion). The ID Token is the OpenID Connect authentication assertion delivered to the Relying Party. Any caller that verifies an ID Token through the unpinned load(token)->verify path, or load(token, key) with only the key pinned, accepts a forged token carrying attacker-chosen claims such as sub and is authenticated as any user. Passing an explicit algorithm so $self->alg is already set bypasses the header-derived allowlist and is not affected. Note that the latest version uploaded to CPAN is 0.10. Later versions are available in the git repository.	7.5	More Details
CVE-2026-54638	gotd/td is a T Telegram MTProto API client in Go. Prior to 0.145.1, proto.UnencryptedMessage.Decode in proto/unencrypted_message.go read attacker controlled dataLen from an unauthenticated MTProto unencrypted packet and allocated make([]byte, dataLen) before checking the remaining buffer, allowing remote unauthenticated denial of service through excessive memory allocation and CPU or garbage collection pressure. This issue is fixed in version 0.145.1.	7.5	More Details
CVE-2026-41608	Improper Handling of Highly Compressed Data (Data Amplification) vulnerability in Apache Thrift Python bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-48586	Improper Handling of Highly Compressed Data (Data Amplification) vulnerability in Apache Thrift C++, Java, Python, Go, D, C/GLib bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-32665	In NLnet Labs Unbound 1.22.0 up to and including 1.25.1, when downstream DNS-over-QUIC (DoQ) is enabled, the first two bidirectional streams on a new QUIC connection (stream_id 0 and 4) bypass the per-stream 'quic-size' gate entirely, and large input buffers are allocated later, after only the 2-byte length prefix has been received from the initial streams. As a result, a remote client can make Unbound exceed the configured 'quic-size' limit with low-cost input. Using only one connection and two streams, each sending a declared 65535-byte length prefix and then holding the streams open, a client can already trivially make Unbound roughly allocate double that amount. This is a remote availability issue / memory-accounting bypass in the downstream DoQ implementation that leads to denial of service for new DoQ clients. This vulnerability needs Unbound to be compiled with DoQ support ('--with-libngtcp2') and the 'quic-port' to be configured for the listening interfaces.	7.5	More Details
CVE-2026-63047	Joomla Extension - joomdonation.com - Invoice data exfiltration via incorrect ACL check in Events Booking 5.0.0-5.8.1 - The Joomla extension Events Booking prior version 5.0-5.8.1 did not properly verify that an actor is allowed to download invoice information.	7.5	More Details
CVE-2026-13189	In Progress® Telerik® UI for AJAX prior to v2026.2.708, insufficient validation of the language parameter in the spell check handler may allow an attacker to influence server-side file path resolution and trigger unintended server-side requests.	7.5	More Details
CVE-2026-62145	A vulnerability in Check Point Gaia Portal allows an authenticated attacker with read-only Gaia Portal privileges to execute commands with root privileges.	7.5	More Details
CVE-2026-65754	Joomla Extension - regularlabs.com - Insecure path handling in ReReplacer Pro extension - ReReplacer XML include paths could read files outside the site directory.	7.5	More Details
CVE-2026-14257	brace-expansion through 5.0.7 is vulnerable to denial of service via memory exhaustion. The expand() function limits the number of results with a max option (default 100,000) but does not bound the length of each result string. By chaining multiple brace groups, an attacker keeps the result count under the limit while making each result progressively longer, so total memory scales with both count and string length until the process hits a fatal, uncatchable out-of-memory error. About 7.5 KB of input ('{a,b}'.repeat(1500)) crashes a default Node.js process. Any application that passes attacker-influenced strings to brace-expansion.expand() - directly or transitively via minimatch / glob brace patterns - can be crashed by a small request. Fixed in 5.0.8 by adding a maxLength option (default 4,000,000) that bounds accumulated output and intermediate arrays.	7.5	More Details
CVE-2026-12493	The Clover Payment Gateway by Zaytech for WooCommerce WordPress plugin before 1.3.6 does not verify that an approved external payment record actually belongs to the WooCommerce order being completed, nor that the paid amount matches the order total, allowing unauthenticated users to mark arbitrary orders as paid by replaying a single genuinely-approved payment reference (for example one obtained from their own minimal purchase).	7.5	More Details
CVE-2026-59529	Unauthenticated Sensitive Data Exposure in Ebook Store <= 6.19 versions.	7.5	More Details
CVE-2026-54719	goshs is a feature-rich single-binary file server for red teamers and developers. Prior to 2.1.1, the httpserver/updown.go bulkDownload handler for ?bulk&file= ZIP downloads did not call findEffectiveACL or applyCustomAuth, allowing unauthenticated reads of files protected only by .goshs folder ACLs and block lists. This issue is fixed in version 2.1.1. This vulnerability exists due to an incomplete fix for CVE-2026-40189.	7.5	More Details
CVE-2026-55390	datamodel-code-generator generates Python data models from schema definitions. From 0.59.0 until 0.62.0, XML Schema parsing in src/datamodel_code_generator/parser/xmlschema.py for --input-file-type xmlschema resolves xs:include, xs:import, xs:redefine, and xs:override schemaLocation values outside the input base path, allowing arbitrary local files to be read and reflected into generated models. This issue is fixed in version 0.62.0.	7.5	More Details

CVE-2026-15280	IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.8 ND Collective Controller is affected by a path-segment injection vulnerability in the collective routing mechanism.	7.5	More Details
CVE-2026-55389	datamodel-code-generator generates Pydantic v2 models, dataclasses, TypedDict, and msgspec.Struct from OpenAPI, JSON Schema, GraphQL, Avro, Protobuf, and raw JSON, YAML, or CSV. Prior to 0.62.0, datamodel-code-generator resolves JSON Schema $ref targets in src/datamodel_code_generator/parser/jsonschema.py through is_url and _get_ref_body without containing file:// or ../ traversal references to the input directory and without honoring --no-allow-remote-refs, allowing arbitrary local file reads. This issue is fixed in version 0.62.0.	7.5	More Details
CVE-2026-65430	Joomla Extension - regularlabs.com - MaxMind Credential leakage in GeoIP extension - MaxMind credentials where leaked in request URLs, causing a credential leakage vulnerability.	7.5	More Details
CVE-2026-12987	The Events Manager WordPress plugin before 7.3.7 does not safely handle booking-registration data on sites using No-User-Account Booking Mode: a booker-supplied registration field is stored as booking meta and later deserialized without restricting allowed classes, enabling PHP object injection. The resulting gadget chain reaches a database query that is built without parameterisation, so an unauthenticated attacker can read arbitrary database data (e.g. user password hashes, secret keys) when the booking is later loaded.	7.5	More Details
CVE-2026-13182	In Progress® Telerik® UI for AJAX prior to v2026.2.708, RadAsyncUpload client-state processing can distinguish decrypt failures from invalid-JSON parse failures, creating an oracle that reveals protected metadata values to remote attackers.	7.5	More Details
CVE-2026-59530	Unauthenticated Broken Access Control in Stripe For WooCommerce <= 4.0.7 versions.	7.5	More Details
CVE-2026-59531	Unauthenticated Unknown in Falcon – WordPress Optimizations & Tweaks <= 2.10.0 versions.	7.5	More Details
CVE-2026-67185	TinyWeb through 0.0.8 contains a path traversal vulnerability that allows unauthenticated attackers to read arbitrary files by submitting ../ sequences in the URL path, which are concatenated directly to the configured web root in HttpBuilder::buildResponse() without normalization, dot-segment removal, or boundary checks. Attackers can craft a single request with ../ sequences that pass through the URL parser unchanged and reach the filesystem call via HttpFile::setFile(), exposing sensitive files such as credential stores and private keys when the server process runs as root.	7.5	More Details
CVE-2026-13184	In Progress® Telerik® UI for AJAX prior to v2026.2.708, when Telerik.Upload.ConfigurationHashKey is absent and machineKey is not explicitly configured, upload metadata integrity protection may fall back to a predictable default key, enabling attackers to forge protected upload metadata and unlock further exploit chains.	7.5	More Details
CVE-2026-55973	In NLnet Labs Unbound 1.23.0 up to and including 1.25.1, when 'dns-error-reporting: yes' is set, the EDNS Report-Channel option (code 18) from the last upstream response is read and uses the option's length as the length of the agent domain. When a domain name check is performed on the agent domain, the returned lenght is not used and if the agent domain is followed by garbage, those bytes are moved onto the tail of the synthetic '_er.' report query name. That query name is later used in the iterator via a subquery to send out the DNS Error Report and when Unbound tries to walk that query name during 'find_closest_of_type()', it strips labels using the query name length rather than stopping at the embedded root, walks one byte past it, and feeds the first garbage byte to 'dname_query_hash()' as a label length writing over the stack variable 'labuf'. One ordinary upstream response from a delegated zone the attacker controls is sufficient to terminate the daemon.	7.5	More Details
CVE-2026-59532	Unauthenticated Other Vulnerability Type in Booking and Rental Manager <= 2.7.2 versions.	7.5	More Details
CVE-2026-59536	Unauthenticated Broken Access Control in CoCart – Headless ecommerce <= 4.8.4 versions.	7.5	More Details
CVE-2026-65755	Joomla Extension - regularlabs.com - Date-sensitive query-cache leakage in Articles Anywhere and Users Anywhere extension - Date-sensitive query cache keys did not retain a bounded time component. Cached results could remain active across future publication or expiry boundaries, potentially exposing content after it should become unavailable.	7.5	More Details
CVE-2026-11331	An attacker who knows (or guesses) that a resolver uses RPZ with wildcard CNAME policies can craft query names long enough to trigger a NAMETOOLONG error condition during RPZ processing. This is not handled correctly and may lead to defeating the RPZ rule. It also may lead to an unexpected exit of the BIND 9 software. This issue affects BIND 9 versions 9.16.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.16.8-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details
CVE-2026-66745	Artica Proxy before 4.50.000000 Service Pack 7 (fixed in hotfix 20260724-02) contains a session fixation vulnerability that allows unauthenticated attackers to hijack administrative sessions by setting a known PHPSESSID on a victim's browser prior to authentication. Attackers can pre-set a controlled session identifier and wait for a victim to authenticate through fw.login.php, after which the attacker gains a fully authenticated administrative session on port 9000.	7.5	More Details

CVE-2024-58330	A missing authentication check in Bosch IP cameras of families CPP13 and CPP14 allows an unauthenticated attacker to retrieve video analytics event data.	7.5	More Details
CVE-2026-49158	Improper Handling of Highly Compressed Data (Data Amplification) vulnerability in Apache Thrift Ruby bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-55968	Inefficient Algorithmic Complexity, Allocation of Resources Without Limits or Throttling vulnerability in Apache Thrift Node.js bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-52688	RRSIGs with too few labels can lead to bypass of DNSSEC wildcard validation	7.5	More Details
CVE-2026-59534	Unauthenticated Broken Access Control in Post My CF7 Form <= 6.2.0 versions.	7.5	More Details
CVE-2026-13204	If a provably insecure domain is covered by both an NSEC and NSEC3 record at the parent, and there exist an RRSIG for only one of these types, then BIND may exit unexpectedly with an assertion while validating this proof. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details
CVE-2026-12617	The issue is unexpected program termination based on ordering and/or specific content in responses to queries for CNAME or DNAME, and A records. Specifically, if a client queries for a DNAME and A record below the DNAME to the resolver, and the authoritative server responds positively to the A query but delays the DNAME response and later responds negatively, `named` may quit unexpectedly. Or, if a client queries for a CNAME and A record for the same name to the resolver, and the authoritative server responds positively to the A query but delays the CNAME response and later responds with a self-referential CNAME, the same failure may occur. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.18.11-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details
CVE-2026-55969	Integer Overflow or Wraparound vulnerability in Apache Thrift C++, c_glib, Go, netstd, Delphi and Haxe bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-51296	SQLite 3.41 has a use-after-free vulnerability in jsonRemoveFunc of SQLite JSON module. The parsed JSON object is freed at line 3555, while line 3575 still calls jsonLookupStep with the released pointer. Remote attackers can exploit this flaw to crash the service and leak heap memory information.	7.5	More Details
CVE-2026-65598	n8n before 1.123.64, 2.29.8, and 2.30.1 contains a TOCTOU race condition in the Git node's clone operation that allows authenticated users to bypass path restrictions by swapping a directory for a symlink after the path is validated but before the clone runs. This lets an attacker plant a crafted repository in the community node directory, which n8n loads as a custom node on the next restart, executing arbitrary JavaScript on the server. Both self-hosted and cloud instances are affected.	7.5	More Details
CVE-2026-59933	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. In versions 4.0.0 through 5.8.0, 3.3.0 through 3.10.6, 2.2.0 through 2.4.6, 2.0.0 through 2.1.17, and all releases up to and including 1.30.5, the OLE reader follows sector chains from attacker-controlled XLS/OLE metadata without detecting cycles or enforcing a maximum chain length. A tiny malformed .xls/OLE file can set the small-block depot sector chain to point back to itself. During normal XLS detection, OLERead::read() appends the same sector data repeatedly until the PHP process exhausts memory. This is reachable from Reader\Xls::canRead() and therefore from automatic spreadsheet type detection. Applications that accept attacker-controlled spreadsheet uploads can suffer denial of service from a very small file. This issue has been fixed in versions 5.8.1, 3.10.7, 2.4.7, 2.1.18 and 1.30.6.	7.5	More Details
CVE-2026-11721	It is possible for an attacker's zone to respond to a query with an RRSIG that has a smaller number of labels than the zone in which the RRSIG is contained. This causes `named` to produce a wildcard name for a zone that is shorter than the attacker's zone, which can result in cache poisoning. For this attack to have any effect, the resolver under attack must have set `synth-from-dnssec yes;` (which is the default). This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details
CVE-2026-59932	PhpSpreadsheet is a pure PHP library for reading and writing spreadsheet files. In versions 4.0.0 through 5.8.0, 3.3.0 through 3.10.6, 2.2.0 through 2.4.6, 2.0.0 through 2.1.17, and all releases up to and including 1.30.5, the Gnumeric reader reads attacker-supplied .gnumeric files into memory and, when the file starts with gzip magic bytes, calls gzdecode() on the full compressed contents without enforcing a decompressed-size limit. A very small compressed .gnumeric file can expand to data larger than the PHP memory limit and crash the process during Gnumeric::canRead() before the file is rejected or fully parsed. This is reachable through normal file-type detection and Gnumeric loading paths, so applications that accept attacker-controlled spreadsheet uploads can suffer denial of service. This issue has been fixed in versions 5.8.1, 3.10.7, 2.4.7, 2.1.18 and 1.30.6.	7.5	More Details
CVE-2026-11622	A DNSSEC validating resolver that is under a random subdomain attack against a DNSSEC-signed zone can suffer from runaway memory usage. The attacker needs to be able to send queries faster than the resolver can perform validation. The increased memory usage can be orders of magnitude beyond the limit configured in the `max-cache-size` parameter. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details

CVE-2026-54635	pytonapi is a Python SDK for TONAPI that provides REST API, streaming, and webhook access to the TON blockchain. From 2.0.0 to 2.2.0, TonapiWebhookDispatcher fails to validate the Authorization header when a webhook handler is registered with the documented path argument, because setup() stores bearer tokens only under the default suffix paths and never adds the custom path to the token map, so self._tokens.get(path) returns None and the authentication guard is skipped. An unauthenticated remote attacker can POST forged payloads to the custom webhook endpoint and trigger victim-defined handlers. This issue is fixed in version 2.2.1.	7.5	More Details
CVE-2026-11605	The issue is a resource exhaustion vulnerability associated with DNSSEC validation. BIND always validates all RRSIG records in an answer, even if they are not strictly needed. A query to an authoritative server/zone which returns many valid but superfluous RRSIG records causes the validator to waste disproportionate CPU time. This issue affects BIND 9 versions 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, and 9.20.9-S1 through 9.20.24-S1.	7.5	More Details
CVE-2026-58389	Allocation of Resources Without Limits or Throttling vulnerability in Apache Thrift Rust bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-64611	A flaw was found in libcupsfilters. The cfIEEE1284NormalizeMakeModel() function enters an infinite loop when processing a printer-advertised IEEE-1284 device ID with an empty model field, causing sustained CPU consumption. A network-adjacent attacker could exploit this by broadcasting a specially crafted printer advertisement, leading to denial of service.	7.5	More Details
CVE-2026-43871	Loop with Unreachable Exit Condition ('Infinite Loop') vulnerability in Apache Thrift Python, Go, PHP and Java bindings.This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	7.5	More Details
CVE-2026-64829	Question2Answer through 1.8.8 contains a session invalidation vulnerability that allows attackers with a previously obtained remember-me cookie to retain authenticated access by exploiting the forgot-password reset flow's failure to clear the sessioncode field in qa-include/app/users-edit.php. While the normal password-change flow in qa-include/pages/account.php explicitly clears the sessioncode to invalidate persistent qa_session cookies, the forgot-password handler qa_finish_reset_user() omits this step, allowing any valid persistent cookie issued before the reset to continue authenticating the account after the password reset completes.	7.4	More Details
CVE-2026-59546	Subscriber Broken Authentication in Hide My WP Ghost <= 7.0.06 versions.	7.4	More Details
CVE-2026-66141	Exim before 4.99.5 allows .forward privilege escalation because force_command for a pipe transport is mishandled.	7.4	More Details
CVE-2026-15328	IBM WebSphere Application Server 9.0, and 8.5 and IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.7 is vulnerable to HTTP request smuggling.	7.4	More Details
CVE-2026-57990	Files or directories accessible to external parties in Microsoft Edge (Chromium-based) allows an unauthorized attacker to disclose information over a network.	7.4	More Details
CVE-2026-57989	Origin validation error in Microsoft Edge (Chromium-based) allows an unauthorized attacker to disclose information over a network.	7.4	More Details
CVE-2026-14528	IBM WebSphere Application Server 9.0, and 8.5 traditional could allow a remote attacker to obtain sensitive information.	7.4	More Details
CVE-2026-16632	A flaw has been found in boazsegev facil.io up to 0.7.4. Affected is the function websocket_on_protocol_error in the library lib/facil/http/parsers/websocket_parser.h of the component WebSocket Frame Parser. This manipulation of the argument on_message 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.	7.3	More Details
CVE-2026-16519	A DLL hijacking vulnerability exists in the GeoVision GV-IP Device Utility desktop application. The application loads one or more dynamic-link libraries (DLLs) from an unsafe search path, allowing a local attacker to place a malicious DLL in a location searched before the legitimate library location.	7.3	More Details
CVE-2026-8164	Uncontrolled Search Path Element vulnerability in ArkSigner Software and Hardware Industry and Trade Inc. ArkSigner Desktop Client allows Search Order Hijacking. This issue affects ArkSigner Desktop Client: from v2.2.16.10 through 17062026.	7.3	More Details
CVE-2026-10033	The EventON Action User plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.5.14. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for unauthenticated attackers to grant EventON management capabilities and the upload_files capability to any non-administrator WordPress role or user, escalating their privileges within the site. The administrator role is protected by an early-return guard in update_role_caps(), so only non-administrator roles and individual users can be targeted; however, the same unauthenticated exposure also allows attackers to enumerate all	7.3	More Details

	WordPress users with their IDs and display names, disclose role and user capability state along with nonce values, and tamper with event-to-user term assignments.		
CVE-2026-14893	IBM Observability with Instana (Agent) Build 1.0.303 through 1.0.320 IBM Instana Node.js tracer component @instana/core version 6.2.1 is vulnerable to prototype pollution through its configuration normalization API.	7.3	More Details
CVE-2026-59535	Unauthenticated Broken Access Control in Thrive Product Manager <= 10.9.2 versions.	7.3	More Details
CVE-2026-16765	A vulnerability was determined in CodeAstro Online Classroom 1.0. Affected by this issue is some unknown functionality of the file /OnlineClassroom/loginlinkadmin.php. Executing a manipulation of the argument aid can lead to sql injection. The attack can be executed remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2026-62432	The EVTCHNOP_expand_array hypercall checks for whether FIFO event channels are enabled, but without holding the correct lock. It can race with EVTCHNOP_reset, resulting in dereferencing a NULL pointer.	7.3	More Details
CVE-2026-16796	Improper neutralization of argument delimiters in the install_packages() method in AWS Bedrock AgentCore Python SDK before 1.18.1 might allow a remote authenticated user to execute arbitrary commands within the Code Interpreter sandbox via crafted package name arguments. To mitigate this issue, users should upgrade to the patched version 1.18.1.	7.3	More Details
CVE-2026-62433	Parts of the DM_OP handling code assumes the caller has provided the required number of buffers for the given operation without any checking being done. As a result, certain operations might access stack rubble as structures are possibly uninitialized.	7.3	More Details
CVE-2026-66015	An authenticated privilege-escalation vulnerability in JFrog Platform may be exploited under admin-provisioned account conditions. Successful exploitation may grant temporary platform administrator access.	7.2	More Details
CVE-2026-65711	sysPass through version 3.2.11 contains an OS command injection vulnerability that allows authenticated administrators to execute arbitrary commands as the web server process user by setting a malicious backup path and triggering a backup. The FileBackupService builds a tar shell command via string concatenation, inserting the admin-configurable siteBackupPath setting without escapeshellarg() or equivalent sanitization before passing it to exec(), causing injected commands to persist and execute on every subsequent backup trigger.	7.2	More Details
CVE-2026-12421	The ARforms plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'password' Field Values in all versions up to, and including, 7.2.1 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	7.2	More Details
CVE-2026-61391	There is a stack-based buffer overflow vulnerability in some Hikvision cameras, which may allow authenticated attackers to cause device malfunction by sending specially crafted packets.	7.2	More Details
CVE-2026-54605	OAuth is a Ruby wrapper for the OAuth 1.0 and 1.0a protocols, providing clients and servers. From 0.5.5 to 1.1.5, OAuth::Consumer#token_request parses the raw Location header of a 300 to 399 redirect returned by the OAuth server and follows the redirect recursively, which can mutate the consumer's configuration and expose signed OAuth request metadata, including the Authorization header, to a cross-origin host. This issue is fixed in version 1.1.6.	7.2	More Details
CVE-2026-59552	Unauthenticated Server Side Request Forgery (SSRF) in 3D Flipbook PDF Viewer & Embedder <= 1.4.2 versions.	7.2	More Details
CVE-2026-65693	Microweber CMS through 2.0.20 contains a server-side template injection vulnerability that allows authenticated administrators to achieve arbitrary OS command execution by injecting Twig expressions into mail templates. Attackers can exploit the unsandboxed Twig environment in TwigView::render(), which lacks SandboxExtension or a SecurityPolicy, to inject malicious expressions such as filter('system') into mail template bodies stored unsanitized in the database, causing automatic payload execution on each subsequent application event that triggers a mail dispatch.	7.2	More Details
CVE-2026-40714	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) an Improper Input Validation vulnerability. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Elevation of privileges.	7.2	More Details
CVE-2026-65516	Unauthenticated Server Side Request Forgery (SSRF) in PeproDev Ultimate Invoice <= 2.2.6 versions.	7.2	More Details
CVE-2026-15401	The VikBooking Hotel Booking Engine & PMS plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'vbfX' parameter in all versions up to, and including, 1.8.13 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The vbfX custom-field value is stored via the public-facing saveorder task, which has no capability or authentication check enforced by default, enabling fully unauthenticated submission of malicious payloads.	7.2	More Details

CVE-2026-65442	Unauthenticated Server Side Request Forgery (SSRF) in FormCraft <= 3.9.15 versions.	7.2	More Details
CVE-2026-13440	The StoreGrowth: Smart Sales Booster for WooCommerce | BOGO, Upsells, Direct Checkout, Quick View, Side Cart plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'message_popup' parameter in all versions up to, and including, 2.1.0 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The exploit is possible because the 'ajd_protected' nonce required by the create_popup handler is exposed to all unauthenticated frontend visitors via wp_localize_script under bogo_save_url.ajd_nonce, effectively bypassing the nonce-only access control.	7.2	More Details
CVE-2026-16585	The Better Messages – Chat Rooms, Group Chat, Private Messages & AI Chat Bots plugin for WordPress is vulnerable to arbitrary file deletion due to insufficient file path validation in the delete_sticker function in all versions up to, and including, 2.15.19. This makes it possible for authenticated attackers, with administrator-level access and above, to delete arbitrary files on the server, which can easily lead to remote code execution when the right file is deleted (such as wp-config.php). The prefix check intended to restrict deletion to the uploads directory can be bypassed by crafting a URL that begins with the legitimate uploads base URL but embeds ../ traversal sequences in the path portion, as the normalize_sticker function only applies esc_url_raw(), which does not strip ../ sequences, allowing the traversal payload to be stored verbatim in WordPress options.	7.2	More Details
CVE-2026-7232	The FormCraft plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the '[parameter name]' parameter in all versions up to, and including, 3.9.14 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The exploit chain combines a server-side gap — where composite matrix sub-field keys such as field2_0 and field2_1 are never passed through the sanitization loop and are stored raw via $wpdb->insert() — with a client-side gap where DOMPurify is only invoked when typeof field.value === 'string', but matrix values arrive from the server as arrays, bypassing the check before being mapped to strings and injected into the DOM. Additionally, the same sink is reachable via a second attack vector: array-typed field values are passed through htmlentities() on submission but later reversed by html_entity_decode() at formcraft-main.php:2608 and :2122, restoring the malicious payload before storage and rendering.	7.2	More Details
CVE-2026-7534	The SUMO Reward Points plugin for WordPress is vulnerable to Unauthenticated Stored Cross-Site Scripting via the REST API endpoint `/wp-json/wc-srp/v1/earning` in versions up to, and including, 32.7.0. This is due to the `user_has_cap` filter in the `SRP_REST_Earning_Controller` class unconditionally granting the custom `rs_earning_read` capability to all users — including unauthenticated visitors — combined with missing sanitization of the `reason` parameter in the `create_items()` function and missing output escaping in the `column_default()` method of `SRP_Master_Log`. This makes it possible for unauthenticated attackers to inject arbitrary web scripts into the reward points log that will execute whenever an administrator accesses the Master Log or User Reward Points admin pages.	7.2	More Details
CVE-2026-65898	DOMPurify before 3.4.11 fails to clone the ALLOWED_ATTR allowlist when setConfig() is used with an uponSanitizeAttribute hook, allowing the hook to permanently mutate the shared allowlist. Attackers can register a hook that conditionally allows dangerous attributes like onerror for trusted elements, then submit untrusted content that inherits the polluted allowlist and executes event handlers as stored XSS.	7.2	More Details
CVE-2026-61953	Unauthenticated Server Side Request Forgery (SSRF) in Simple Link Directory Pro <= 15.0.6 versions.	7.2	More Details
CVE-2026-65497	Administrator PHP Object Injection in Complianz <= 7.5.0 versions.	7.2	More Details
CVE-2026-66138	In OpenStack Ironic Python Agent through 11.6.0, a project-scoped user with the manager role can achieve arbitrary code execution on a running Ironic-Python-Agent via a maliciously constructed configuration, because the value of ntp_server is passed to a shell.	7.2	More Details
CVE-2026-64318	In the Linux kernel, the following vulnerability has been resolved: partitions: aix: bound the pp_count scan to the ppe array aix_partition() reads the physical volume descriptor into a fixed-size struct pvd and then scans its physical-partition-extent array: int numpps = be16_to_cpu(pvd->pp_count); ... for (i = 0; i < numpps; i += 1) { struct ppe *p = pvd->ppe + i; ... lp_ix = be16_to_cpu(p->lp_ix); pvd points at a single kmalloc()'d struct pvd whose ppe[] member holds a fixed ARRAY_SIZE(pvd->ppe) (1016) entries, but the loop runs up to the on-disk pp_count. pp_count is an unvalidated __be16 read straight from the descriptor, so a crafted AIX image with pp_count larger than 1016 drives the loop to read pvd->ppe[i] past the end of the allocation (up to 65535 entries, ~2 MB out of bounds). The partition scan runs without mounting anything, when a block device with a crafted AIX/IBM partition table appears (an attacker-supplied image attached with losetup -P, or a device auto-scanned by udev), via msdos_partition() -> aix_partition(). Clamp the scan to the number of entries the ppe[] array can hold.	7.1	More Details
CVE-2026-59558	Unauthenticated Cross Site Scripting (XSS) in Booking Calendar <= 11.4.2 versions.	7.1	More Details
CVE-2026-57367	Subscriber Broken Access Control in WP Booking System < 5.12.8.1 versions.	7.1	More Details

CVE-2026-16192	IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.8 is affected by a denial of service vulnerability when the restConnector-2.0 feature is enabled.	7.1	More Details
CVE-2026-65510	Unauthenticated Cross Site Scripting (XSS) in PeproDev Ultimate Invoice <= 2.2.6 versions.	7.1	More Details
CVE-2026-65511	Unauthenticated Cross Site Scripting (XSS) in Manual - Documentation, Knowledge Base & Education WordPress Theme <= 7.5.4 versions.	7.1	More Details
CVE-2026-64317	In the Linux kernel, the following vulnerability has been resolved: isofs: bound Rock Ridge symlink components to the SL record get_symlink_chunk() and the SL handling in parse_rock_ridge_inode_internal() walk the variable-length components of a Rock Ridge "SL" (symbolic link) record. Each component is a two-byte header (flags, len) followed by len bytes of text, so it occupies slp->len + 2 bytes. Both loops read slp->len and advance to the next component, and get_symlink_chunk() additionally does memcpy(rpnt, slp->text, slp->len), but neither checks that the component lies within the SL record before dereferencing it. A crafted SL record whose component declares a len that runs past the record (rr->len) therefore triggers an out-of-bounds read of up to 255 bytes. When the record sits at the tail of its backing buffer - for example a small kmalloc()ed continuation block reached through a CE record - the read crosses the allocation; get_symlink_chunk() then copies the out-of-bounds bytes into the symlink body returned to user space by readlink(), disclosing adjacent kernel memory. ISO 9660 images are routinely mounted from untrusted removable media - desktop environments auto-mount them (e.g. via udisks2) without CAP_SYS_ADMIN - so the record contents are attacker-controlled. Reject any component that does not fit in the remaining record bytes before using it. In get_symlink_chunk() return NULL, like the existing output-buffer (plimit) checks, so a malformed record makes readlink() fail with -EIO rather than silently returning a truncated target; in parse_rock_ridge_inode_internal() stop the inode-size walk.	7.1	More Details
CVE-2026-64323	In the Linux kernel, the following vulnerability has been resolved: udf: validate VAT header length against the VAT inode size udf_load_vat() takes the virtual partition's start offset straight from the on-disk VAT 2.0 header without checking it against the VAT inode size: map->s_type_specific.s_virtual.s_start_offset = le16_to_cpu(vat20->lengthHeader); map->s_type_specific.s_virtual.s_num_entries = (sbi->s_vat_inode->i_size - map->s_type_specific.s_virtual.s_start_offset) >> 2; lengthHeader is a fully attacker-controlled 16-bit value. If it exceeds the VAT inode size, the s_num_entries subtraction underflows to a huge count, which defeats the "block > s_num_entries" bound in udf_get_pblock_virt15(); and on the ICB-inline path that function reads ((__le32 *)(iinfo->i_data + s_start_offset))[block] so a large s_start_offset indexes past the inode's in-ICB data. Mounting a crafted UDF image with a virtual (VAT) partition then triggers an out-of-bounds read. Reject a VAT whose header length does not leave room for at least one entry within the VAT inode.	7.1	More Details
CVE-2026-65441	Unauthenticated Cross Site Scripting (XSS) in GiveWP <= 4.16.3 versions.	7.1	More Details
CVE-2026-65918	PyTorch torchvision through 0.28.0, fixed in commit 4e05dc2, contains an out-of-bounds heap read vulnerability in the GIF decoder's read_from_tensor callback that passes unclamped length to memcpy. Attackers can supply malicious or truncated GIF files to cause denial of service via segmentation fault or disclose adjacent heap memory contents.	7.1	More Details
CVE-2026-65710	sysPass through version 3.2.11 contains a missing authorization vulnerability that allows authenticated users with the PUBLICLINK_CREATE profile flag to trigger unauthorized decryption and persistent storage of any vault account's password by exploiting the absence of AccountAcl checks in the public link creation flow. Attackers can invoke the saveCreateFromAccountAction endpoint to cause AccountService::getDataForLink to load arbitrary target accounts without AccountFilterUser restrictions, decrypt credentials using the session master key, and serialize cleartext passwords into Vault storage on the PublicLink database row, enabling subsequent unauthenticated retrieval if the generated link hash is recovered.	7.1	More Details
CVE-2026-64501	In the Linux kernel, the following vulnerability has been resolved: iio: adc: ad_sigma_delta: fix CS held asserted and state leaks In ad_sigma_delta_single_conversion(), set_mode(AD_SD_MODE_IDLE) and disable_one() were called from the out: block while keep_cs_asserted was still true. This caused any SPI transfer issued by those callbacks to carry cs_change=1, leaving CS permanently asserted after the conversion. Fix by moving both calls into the out_unlock: block, after keep_cs_asserted is cleared, matching the pattern already used in ad_sd_calibrate(). In the error path of ad_sd_buffer_postenable(), if an operation fails after set_mode(AD_SD_MODE_CONTINUOUS) has already succeeded (e.g. spi_offload_trigger_enable()), the device is left in continuous conversion mode with CS physically asserted. Additionally, bus_locked remaining true after spi_bus_unlock() causes subsequent SPI operations to call spi_sync_locked() without the bus lock actually held, allowing concurrent SPI access. Fix the error path by clearing keep_cs_asserted first, then calling set_mode(AD_SD_MODE_IDLE) to revert the device mode and deassert CS, then clearing bus_locked before releasing the bus. For devices that implement neither set_mode nor disable_one (such as MAX11205, which has no physical CS pin), no SPI transfer is issued during cleanup and the cs_change flag has no effect on any physical line.	7.1	More Details
	In the Linux kernel, the following vulnerability has been resolved: net: af_key: initialize alg_key_len for IPComp states pfkey_msg2xfrm_state() handles the IPComp (SADB_X_SATYPE_IPCOMP) case by allocating x->calg and copying only the algorithm name: x->calg = kmalloc_obj(*x->calg); if (!x->calg) { err = -ENOMEM; goto out; } strcpy(x->calg->alg_name, a->name); x->props.calgo = sa->sadb_sa_encrypt; Unlike the authentication (x->aalg) and encryption (x->ealg) branches of the same function, the compression branch never initializes calg->alg_key_len. IPComp carries		

CVE-2026-64436	no key and the allocation only reserves sizeof(struct xfrm_algo) (i.e. no room for a key), so the field is left containing uninitialized slab data. calg->alg_key_len is later used as a length by xfrm_algo_clone() when an IPComp state is cloned during XFRM_MSG_MIGRATE: xfrm_state_migrate() xfrm_state_clone_and_setup() x->calg = xfrm_algo_clone(orig->calg); kmemdup(orig, xfrm_alg_len(orig)); where xfrm_alg_len() returns sizeof(*alg) + (alg_key_len + 7) / 8. With a non-zero garbage alg_key_len, kmemdup() reads past the end of the 68-byte calg object. Adding an IPComp SA via PF_KEY and then migrating it triggers (net-next, KASAN, init_on_alloc=0): BUG: KASAN: slab-out-of-bounds in kmemdup_noprof+0x44/0x60 Read of size 4164 at addr ff11000025a74980 by task diag2/9287 CPU: 3 UID: 0 PID: 9287 Comm: diag2 7.1.0-rc6-g903db046d557 #1 Call Trace: <TASK> dump_stack_lvl+0x10e/0x1f0 print_report+0xf7/0x600 kasan_report+0xe4/0x120 kasan_check_range+0x105/0x1b0 __asan_memcpy+0x23/0x60 kmemdup_noprof+0x44/0x60 xfrm_state_migrate+0x70a/0x1da0 xfrm_migrate+0x753/0x18a0 xfrm_do_migrate+0xb47/0xf10 xfrm_user_rcv_msg+0x411/0xb50 netlink_rcv_skb+0x158/0x420 xfrm_netlink_rcv+0x71/0x90 netlink_unicast+0x584/0x850 netlink_sendmsg+0x8b0/0xdc0 ____sys_sendmsg+0x9f7/0xb90 ___sys_sendmsg+0x134/0x1d0 __sys_sendmsg+0x16d/0x220 do_syscall_64+0x116/0x7d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f </TASK> Allocated by task 9287: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0xaa/0xb0 pfkey_add+0x2652/0x2ea0 pfkey_process+0x6d0/0x830 pfkey_sendmsg+0x42c/0x850 __sys_sendto+0x461/0x4b0 __x64_sys_sendto+0xe0/0x1c0 do_syscall_64+0x116/0x7d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f The buggy address belongs to the object at ff11000025a74980 which belongs to the cache kmalloc-96 of size 96 The buggy address is located 0 bytes inside of allocated 68-byte region [ff11000025a74980, ff11000025a749c4) Depending on the uninitialized value the same field can instead request an oversized kmemdup() allocation and make the migration clone fail. The XFRM netlink path is not affected: verify_one_alg() rejects an XFRMA_ALG_COMP attribute shorter than xfrm_alg_len(), so a calg added via XFRM_MSG_NEWSA is always self-consistent. Initialize calg->alg_key_len to 0, matching the aalg/ealg branches.	7.1	More Details
CVE-2026-28945	A permissions issue was addressed with additional sandbox restrictions. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to bypass network restrictions.	7.1	More Details
CVE-2026-64379	In the Linux kernel, the following vulnerability has been resolved: smb: client: mask server-provided mode to 07777 in modefromsid When modefromsid is active, parse_dacl() applies the server-provided sub_auth[2] value from the NFS mode SID to cf_mode without masking to 07777. Apply the correct masking, same as in the read path.	7.1	More Details
CVE-2026-64243	In the Linux kernel, the following vulnerability has been resolved: ASoC: codecs: simple-mux: Fix enum control bounds check simple_mux_control_put() rejects values greater than e->items, but enum control values are zero based. For the two-entry mux used by this driver, valid values are 0 and 1, so value 2 must be rejected as well. Accepting e->items can store an invalid mux state, pass it to the GPIO setter, and pass it on to the DAPM mux update path where it is used as an index into the enum text array. Use the same >= e->items check used by the ASoC enum helpers.	7.1	More Details
CVE-2026-43771	A stack overflow was addressed with improved input validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause a denial-of-service.	7.1	More Details
CVE-2026-59556	Unauthenticated Cross Site Scripting (XSS) in Dynamic Pricing With Discount Rules for WooCommerce <= 4.5.11 versions.	7.1	More Details
CVE-2026-65896	Grav API Plugin (Composer package getgrav/grav-plugin-api) before 1.0.10 fails to properly validate the slug field in the POST /pages/{route}/move endpoint. PagesController::move() sanitizes the slug only with ltrim($body['slug'], '.'), which strips leading periods but does not neutralize '/' or '..' segments. An authenticated API caller with the api.pages.write permission can supply path traversal sequences (e.g., 01.home/../../../pwned) to move an entire page directory (content and media) to an arbitrary writable location outside user/pages/, including outside the Grav installation.	7.1	More Details
CVE-2026-64725	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause a denial-of-service.	7.1	More Details
CVE-2026-64422	In the Linux kernel, the following vulnerability has been resolved: net: ipv4: bound TCP reordering sysctl writes and MTU probe sizes Reject invalid `net.ipv4.tcp_reordering` values before they reach TCP socket state. The sysctl is stored as an `int` but copied into the `u32` `tp->reordering` field for new sockets, so negative writes wrap to large values. With `tcp_mtu_probing=2`, the wrapped value can overflow the `tcp_mtu_probe()` size calculation and drive the MTU probing path into an out-of-bounds read. Route `tcp_reordering` writes through `proc_dointvec_minmax()` and require it to be at least 1. Also require `tcp_max_reordering` to be at least 1 so the configured maximum cannot become negative either. When registering the table for a non-init network namespace, relocate `extra2` pointers that refer into `init_net.ipv4` so the `tcp_reordering` upper bound follows that namespace's `tcp_max_reordering`. Harden `tcp_mtu_probe()` itself by computing `size_needed` as `u64`. This keeps the send queue and window checks from being bypassed through signed integer overflow.	7.1	More Details
CVE-2026-59553	Unauthenticated Cross Site Scripting (XSS) in Product Feed Manager <= 7.6.1 versions.	7.1	More Details
CVE-2026-43813	A validation issue was addressed with improved input sanitization. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. A maliciously crafted app may be able to bypass code signing enforcement.	7.1	More Details

CVE-2026-65540	Unauthenticated Cross Site Request Forgery (CSRF) in Popup for CF7 with Sweet Alert <= 1.6.5 versions.	7.1	More Details
CVE-2026-65539	Unauthenticated Cross Site Request Forgery (CSRF) in Kwayy HTML Sitemap <= 4.0 versions.	7.1	More Details
CVE-2026-64496	In the Linux kernel, the following vulnerability has been resolved: iio: event: Fix event FIFO reset race `iio_event_getfd()` creates the event file descriptor with `anon_inode_getfd()`, which allocates a new fd, creates the anonymous file and installs it in the process fd table before returning to the caller. The IIO code resets the event FIFO after `anon_inode_getfd()` has returned, but before `IIO_GET_EVENT_FD_IOCTL` has copied the fd number to userspace. But since fd tables are shared between threads, another thread can guess the newly allocated fd number and issue a `read()` on it as soon as the fd has been installed. This means the `kfifo_to_user()` in `iio_event_chrdev_read()` can run in parallel with the `kfifo_reset_out()` in `iio_event_getfd()`. The kfifo documentation says that `kfifo_reset_out()` is only safe when it is called from the reader thread and there is only one concurrent reader. Otherwise it is dangerous and must be handled in the same way as `kfifo_reset()`. If that happens, `kfifo_to_user()` can advance the FIFO `out` index based on state from before the reset, after the reset has already moved the `out` index to the current `in` index. That can leave the FIFO with an `out` index past the `in` index. A later `read()` can then see an underflowed FIFO length and copy more data than the event FIFO buffer contains. This can result in an out-of-bounds read and leak adjacent kernel memory to userspace. Move the FIFO reset before `anon_inode_getfd()`. At that point the event fd is marked busy, but the new fd has not been installed yet, so userspace cannot access it while the FIFO is reset.	7.1	More Details
CVE-2026-64452	In the Linux kernel, the following vulnerability has been resolved: 6lowpan: fix NHC entry use-after-free on error path lowpan_nhc_do_uncompression() looks up an NHC descriptor while holding lowpan_nhc_lock. If the descriptor has no uncompress callback, the error path drops the lock before printing nhc->name. lowpan_nhc_del() removes descriptors under the same lock and then relies on synchronize_net() before the owning module can be unloaded. That only waits for net RX RCU readers. lowpan_header_decompress() is also exported and can be reached from callers that are not necessarily covered by the net core RX critical section, for example the Bluetooth 6LoWPAN L2CAP receive path. This leaves a race where one task drops lowpan_nhc_lock in the error path, another task unregisters and frees the matching descriptor after synchronize_net() returns, and the first task then dereferences nhc->name for the warning. With the post-unlock window widened, KASAN reports: BUG: KASAN: slab-use-after-free in lowpan_nhc_do_uncompression+0x1f4/0x220 Read of size 8 lowpan_nhc_do_uncompression lowpan_header_decompress Fix this by printing the warning before dropping lowpan_nhc_lock, so the descriptor name is read while unregister is still excluded. The malformed packet is still rejected with -ENOTSUPP.	7.1	More Details
CVE-2026-13726	The MPG WordPress plugin before 4.1.8 does not sanitise and escape a parameter before reflecting it back in the response, allowing unauthenticated attackers to perform Reflected Cross-Site Scripting against a victim who is induced to send a crafted request.	7.1	More Details
CVE-2026-43681	A buffer overflow was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A local user may be able to read kernel memory.	7.1	More Details
CVE-2026-48029	libheif is a HEIF and AVIF file format decoder and encoder. Versions 1.19.0 through 1.21.2 have a heap OOB read in ImageItem_Grid::decode_grid_tile via irot-induced tile-coordinate underflow. Version 1.22.0 fixes the issue.	7.1	More Details
CVE-2026-43672	An authorization issue was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. A malicious application may be able to bypass Privacy preferences.	7.1	More Details
CVE-2026-57370	Unauthenticated Cross Site Scripting (XSS) in Visitor Traffic Real Time Statistics Pro <= 11.9.1 versions.	7.1	More Details
CVE-2026-57397	Unauthenticated Cross Site Scripting (XSS) in Coaching <= 3.9.2 versions.	7.1	More Details
CVE-2026-9765	Note: The CVE and blog post don't exist because we determined this is actually a cloud-only issue. Access Controls are “Broken” when a user can access resources they are not authorized to access. An attacker can bypass any access control mechanisms in a web application, and gain unauthorized access to resources that are not available with their permissions. Broken access control can allow attackers to: Access resources only accessible to certain users, thus allowing unauthorized access to data Perform operations on behalf of other users, leading to account takeovers in the worst cases Attempt privilege escalation Attempt to take over an account	7.1	More Details
CVE-2026-65446	Unauthenticated Cross Site Scripting (XSS) in Kali Forms <= 2.4.18 versions.	7.1	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: tracing: Prevent out-of-bounds read in glob matching String event fields are not necessarily NUL-terminated, so the filter predicate functions (filter_pred_string(), filter_pred_strloc() and filter_pred_strrelloc()) pass the field length to the regex match callbacks, and the length-aware matchers honour it. regex_match_glob() was the exception: it ignored the length and called glob_match(), which		

2026-64299	scans the string until it hits a NUL byte. Some string fields are not NUL-terminated. One example is the dynamic char array of the xfs_* namespace tracepoints, which is copied without a trailing NUL. For such a field, glob matching reads past the end of the event field, causing a KASAN slab-out-of-bounds read in glob_match(), reached via regex_match_glob() and filter_match_preds() from the xfs_lookup tracepoint. Add a length-bounded glob_match_len() and use it from regex_match_glob() so glob matching always stops at the field boundary. The matching loop is factored into a shared helper so glob_match() keeps its behaviour.	7.1	More Details
CVE-2026-66759	A flaw was found in the file-icns plugin in GIMP. When applying a decompressed mask during ICNS image processing, the plugin reads from the mask data buffer without verifying if the cursor exceeds the allocated resource size. If a crafted file contains a truncated mask resource, the icns_decompress function continues reading past the bounds of the buffer. This out-of-bounds read vulnerability results in information disclosure of heap contents, where memory contents are leaked as alpha channel pixel values, or a crash leading to a denial of service if unmapped memory is accessed.	7.1	More Details
CVE-2026-64298	In the Linux kernel, the following vulnerability has been resolved: NFSv4: include MAY_WRITE in open permission mask for O_TRUNC POSIX requires write permission to truncate a file, so an open() that specifies O_TRUNC must be authorized for write access regardless of the O_ACCMODE access mode. nfs_open_permission_mask() builds the access mask passed to nfs_may_open(), which is the local authorization gate for OPENs the client serves itself from a cached write delegation via the can_open_delegated() path in nfs4_try_open_cached(). The mask is derived from O_ACCMODE alone, so an open(O_RDONLY | O_TRUNC) against a file the caller cannot write requests only MAY_READ and passes the local check. The OPEN is then satisfied locally and the truncation is issued to the server as a SETATTR(size=0) over the delegation stateid, which the server accepts under standard write-delegation semantics. POSIX requires that this open fail with EACCES. Include MAY_WRITE in the mask whenever O_TRUNC is set so the local check matches the access the server would have enforced.	7.1	More Details
CVE-2026-43747	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. Parsing a maliciously crafted file may lead to an unexpected app termination.	7.1	More Details
CVE-2026-59512	Unauthenticated Cross Site Scripting (XSS) in Product Enquiry for WooCommerce <= 2.2.34.43 versions.	7.1	More Details
CVE-2026-59517	Unauthenticated Cross Site Scripting (XSS) in Easy Form Builder <= 4.0.12 versions.	7.1	More Details
CVE-2026-13077	A missing bounds check in the BSON CodeWScope element accessors allows an attacker to trigger an out-of-bounds heap read via a crafted aggregation pipeline. The vulnerability can be exploited by an authenticated user by generating a malformed BSONColumn data containing a CodeWScope element, bypassing wire-level BSON validation. When the forged element is decompressed, the unchecked size value is used in pointer arithmetic, causing either a server crash or disclosure of adjacent heap memory contents.	7.1	More Details
CVE-2026-65447	Unauthenticated Cross Site Scripting (XSS) in Contest Gallery <= 30.0.6 versions.	7.1	More Details
CVE-2026-61957	Unauthenticated Cross Site Scripting (XSS) in miniorange otp verification <= 5.5.1 versions.	7.1	More Details
CVE-2026-65437	Unauthenticated Cross Site Scripting (XSS) in Spam protection, AntiSpam, FireWall by CleanTalk <= 6.82 versions.	7.1	More Details
CVE-2026-64833	FFmpeg versions 0.7.1 through 8.1.2 contain an out-of-bounds read vulnerability in the S/PDIF muxer that allows attackers to access memory beyond buffer boundaries by supplying a crafted DTS stream with a core_size value larger than the actual packet length. Attackers can exploit the missing bounds check in the spdif_header_dts4 function by providing a malicious DTS-HD audio stream during S/PDIF re-muxing to trigger unauthorized memory reads beyond the packet buffer.	7.1	More Details
CVE-2026-64284	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Ensure vendor's exit handler runs before fastpath userspace exits Move the handling of fastpath userspace exits into vendor code to ensure KVM runs vendor specific operations that need to run before userspace gains control of the vCPU. E.g. for VMX (and soon to be for SVM as well), KVM needs to flush the PML buffer prior to exiting to userspace, otherwise any memory written by the final KVM_RUN might never be flagged as dirty. Note, waiting to snapshot CR0 and CR3 until svm_handle_exit() is flawed in general, as that risks consuming stale state in a fastpath handler. That will be addressed in a future change.	7.1	More Details
CVE-2026-65438	Unauthenticated Cross Site Scripting (XSS) in Message Filter for Contact Form 7 <= 1.6.3.9 versions.	7.1	More Details
CVE-2026-65439	Unauthenticated Cross Site Scripting (XSS) in Ultimate Addons for Contact Form 7 <=3.5.45 versions.	7.1	More Details

CVE-2026-65443	Unauthenticated Cross Site Scripting (XSS) in BackWPup <= 5.7.4 versions.	7.1	More Details
CVE-2026-64411	In the Linux kernel, the following vulnerability has been resolved: netfilter: ebtables: terminate table name before find_table_lock() update_counters() and compat_update_counters() forward a user-supplied 32-byte table name to find_table_lock() without NUL-terminating it. On a lookup miss, find_inlist_lock() calls try_then_request_module(..., "%s%s", "ebtable_", name), and vsnprintf() reads past the name field and the stack object until it hits a zero byte. BUG: KASAN: stack-out-of-bounds in string (lib/vsprintf.c:648 lib/vsprintf.c:730) Read of size 1 at addr ffff8880119dfb20 by task exploit/147 Call Trace: ... string (lib/vsprintf.c:648 lib/vsprintf.c:730) vsnprintf (lib/vsprintf.c:2945) __request_module (kernel/module/kmod.c:150) do_update_counters.isra.0 (net/bridge/netfilter/ebtables.c:371 net/bridge/netfilter/ebtables.c:380) update_counters (net/bridge/netfilter/ebtables.c:1440) do_ebt_set_ctl (net/bridge/netfilter/ebtables.c:2573) nf_setsockopt (net/netfilter/nf_sockopt.c:101) ip_setsockopt (net/ipv4/ip_sockglue.c:1424) raw_setsockopt (net/ipv4/raw.c:847) __sys_setsockopt (net/socket.c:2393) ... compat_do_replace() shares the same unterminated name via compat_copy_ebt_replace_from_user(); terminate it there too so all find_table_lock() callers behave alike. The other callers already terminate the name after the copy.	7.1	More Details
CVE-2026-54545	wakaru is a JavaScript decompiler and unminifier toolkit. From 1.0.0 until 1.4.0, @wakaru/cli sanitizes bundle-controlled module filenames only once before writing extracted modules, so a crafted filename containing overlapping traversal sequences such as ....// collapses to ../ after sanitization and lets the final output path escape the selected output directory, allowing an attacker who can cause a user to run wakaru --unpack on a malicious bundle to write files outside that directory and, depending on the target path and environment, potentially achieve code execution. This issue is fixed in @wakaru/cli 1.4.0.	7.1	More Details
CVE-2026-64412	In the Linux kernel, the following vulnerability has been resolved: netfilter: ebtables: module names must be null-terminated We need to explicitly check the length, else we may pass non-null terminated string to request_module().	7.1	More Details
CVE-2026-61947	Unauthenticated Cross Site Scripting (XSS) in Form Vibes – Database Manager for Forms <= 1.5.2 versions.	7.1	More Details
CVE-2026-65440	Unauthenticated Cross Site Scripting (XSS) in GetGenie <= 4.4.3 versions.	7.1	More Details
CVE-2026-64403	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: validate option length before reading conf opt value l2cap_get_conf_opt() derives the option length from the attacker-controlled opt->len field and immediately dereferences opt->val (as u8, get_unaligned_le16() or get_unaligned_le32(), or a raw pointer for the default case) before any caller has confirmed that opt->len bytes are present in the buffer. The callers (l2cap_parse_conf_req(), l2cap_parse_conf_rsp() and l2cap_conf_rfc_get()) only detect a malformed option afterwards, once the running length has gone negative, by which point the out-of-bounds read has already executed. An existing post-hoc length check keeps the garbage value from being consumed, so this is not a data leak in the current control flow. It is still a validate-after-use ordering bug: up to 4 bytes are read past the end of the buffer before it is known to contain them, and it is fragile to future changes in the callers. Fix it at the source. Pass the end of the buffer into l2cap_get_conf_opt() and refuse to touch opt->val unless the full option (header + value) fits. Each caller computes an end pointer once before the loop and checks the return value directly instead of inferring the error from a negative length.	7.1	More Details
CVE-2026-57809	Unauthenticated Cross Site Scripting (XSS) in AffiliateWP <= 2.34.0 versions.	7.1	More Details
CVE-2026-14870	The Database for Contact Form 7, WPforms, Elementor forms WordPress plugin before 1.5.3 does not properly sanitise and escape a parameter before reflecting it back in an admin page, leading to a Reflected Cross-Site Scripting which could be used against high privilege users such as admin.	7.1	More Details
CVE-2026-13442	IBM Langflow OSS 1.0.0 through 1.10.1 can allow an attacker to reuse another user's FAISS namespace to access owner-only vector content and influence later query results. This causes cross-user information disclosure and limited integrity impact through persistent poisoning of returned results.	7.1	More Details
CVE-2026-57374	Unauthenticated Cross Site Scripting (XSS) in Funnel Kit Funnel Builder PRO <= 3.15.0.7 versions.	7.1	More Details
CVE-2026-15968	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Progress MOVEit Transfer. This issue affects MOVEit Transfer: before 2025.1.5, from 2026.0.0 before 2026.0.3.	7.1	More Details
CVE-2026-64692	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause a denial-of-service.	7.1	More Details
CVE-2026-	Unauthenticated Cross Site Scripting (XSS) in Download Monitor - WPForms Lock <= 1.0.4 versions.	7.1	More Details

57427			
CVE-2026-57428	Unauthenticated Cross Site Scripting (XSS) in Sprout Clients <= 3.2.3 versions.	7.1	More Details
CVE-2026-57626	Cross-Site Request Forgery (CSRF) vulnerability in MailPoet allows Cross Site Request Forgery. This issue affects MailPoet: from 5.30.0 through 5.33.0.	7.1	More Details
CVE-2026-65922	An authorization weakness in JFrog Artifactory internal metadata handling could allow a user with limited repository access to write to restricted internal metadata areas under specific conditions. Successful abuse is limited to integrity and availability impact at a low level; confidentiality is not affected.	7.1	More Details
CVE-2026-57696	Contributor Arbitrary File Deletion in Picture Gallery <= 1.6.5 versions.	7.1	More Details
CVE-2026-57699	Subscriber Cross Site Scripting (XSS) in Slider Pro <= 4.8.13 versions.	7.1	More Details
CVE-2026-57701	Unauthenticated Cross Site Scripting (XSS) in Real Estate Manager Pro <= 12.8.5 versions.	7.1	More Details
CVE-2026-61944	Unauthenticated Cross Site Scripting (XSS) in Bookly <= 27.7 versions.	7.1	More Details
CVE-2026-57704	Unauthenticated Cross Site Scripting (XSS) in Smart Manager <= 8.90.0 versions.	7.1	More Details
CVE-2026-14976	IBM WebSphere Application Server - Liberty 17.0.0.3 through 26.0.0.8 is affected by remote code execution with the collectiveController-1.0 feature enabled.	7.1	More Details
CVE-2026-65494	Subscriber SQL Injection in Dokan Pro <= 5.0.2 versions.	7.1	More Details
CVE-2026-65492	Unauthenticated Cross Site Scripting (XSS) in Dokan Pro <= 5.0.0 versions.	7.1	More Details
CVE-2026-57735	Unauthenticated Cross Site Scripting (XSS) in Breakdance <= 2.7.1 versions.	7.1	More Details
CVE-2026-57767	Unauthenticated Cross Site Scripting (XSS) in WP Google Maps Pro <= 10.1.02 versions.	7.1	More Details
CVE-2026-57769	Unauthenticated Cross Site Scripting (XSS) in Grand Photography <= 5.7.8 versions.	7.1	More Details
CVE-2026-65488	Unauthenticated Cross Site Request Forgery (CSRF) in LA-Studio Element Kit for Elementor <= 1.6.2 versions.	7.1	More Details
CVE-2026-16184	IBM WebSphere Application Server 9.0, and 8.5 could allow a remote attacker to bypass authentication by sending a crafted unauthenticated request.	7.0	More Details
CVE-2026-64460	In the Linux kernel, the following vulnerability has been resolved: PCI/IOV: Skip VF Resizable BAR restore on read error sriov_restore_vf_rebar_state() uses the VF Resizable BAR Control register to decide how many VF BARs to restore (nbars) and which VF BAR each iteration addresses (bar_idx). bar_idx indexes into dev->sriov->barsz[], which has only PCI_SRIOV_NUM_BARS (6) entries. When a device does not respond, config reads typically return PCI_ERROR_RESPONSE (~0). Both fields are 3 bits wide, so nbars and bar_idx both evaluate to 7. The barsz[] access then goes out of bounds. UBSAN reports this as: UBSAN: array-index-out-of-bounds in drivers/pci/iov.c:948:51 index 7 is out of range for type 'resource_size_t [6]' Observed on an NVIDIA RTX PRO 1000 GPU (GB207GLM) that stopped responding during a failed GC6 power state exit. The subsequent pci_restore_state() invoked sriov_restore_vf_rebar_state() while config reads returned 0xffffffff, triggering the splat. Bail out if any VF Resizable BAR Control read returns PCI_ERROR_RESPONSE. No further VF BARs are touched, which is safe because a config read that returns PCI_ERROR_RESPONSE indicates the device is unreachable and restoration is pointless. This mirrors the guard in pci_restore_rebar_state().	7.0	More Details

CVE-2026-43693	A race condition was addressed with improved state handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to gain root privileges.	7.0	More Details
CVE-2026-64315	In the Linux kernel, the following vulnerability has been resolved: crypto: caam - use print_hex_dump_devel to guard key hex dumps Use print_hex_dump_devel() for dumping sensitive key material in *_setkey() to avoid leaking secrets at runtime when CONFIG_DYNAMIC_DEBUG is enabled.	7.0	More Details
CVE-2026-43755	A race condition was addressed with improved state management. This issue is fixed in macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to gain root privileges.	7.0	More Details
CVE-2026-16584	Improper handling of an initialization failure in AWS API MCP Server from 0.2.13 through 1.3.46 might allow an actor to bypass the user-configured security policy and execute AWS API operations that the policy was set to deny or gate. When initialization of the security policy enforcement data fails at server startup, the policy check is skipped for the lifetime of the process. IAM permissions on the configured credentials remain in effect and are unaffected. To remediate this issue, users should upgrade to version 1.3.47.	7.0	More Details
CVE-2026-28926	A race condition was addressed with improved state handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. An app may be able to elevate privileges.	7.0	More Details
CVE-2026-64413	In the Linux kernel, the following vulnerability has been resolved: netfilter: ebtables: zero chainstack array sashiko reports: looking at ebtables table translation, could a sparse cpu_possible_mask lead to an uninitialized pointer free? If cpu_possible_mask is sparse (for example, CPU 0 and CPU 2 are possible, but CPU 1 is not), the allocation loop skips CPU 1. If vmalloc_node() fails at CPU 2, the cleanup loop will blindly decrement and call vfree() on newinfo->chainstack[1]. Not a real-world bug, such allocation isn't expected to fail in the first place.	7.0	More Details
CVE-2026-64510	In the Linux kernel, the following vulnerability has been resolved: ACPI: NFIT: core: Fix acpi_nfit_init() error cleanup If acpi_nfit_init() fails after adding the acpi_desc object to the acpi_descs list, that object is never removed from that list because the acpi_nfit_shutdown() devm action is not added for the NFIT device in that case. Next, the acpi_nfit_init() failure causes acpi_nfit_probe() to fail, the acpi_desc object is freed, and a dangling pointer is left behind in the acpi_descs. Any subsequent ACPI Machine Check Exception will trigger nfit_handle_mce() which iterates over acpi_descs and so a use-after-free will occur. Moreover, if acpi_nfit_probe() returns 0 after installing a notify handler for the NFIT device and without allocating the acpi_desc object and setting the NFIT device's driver data pointer, the acpi_desc object will be allocated by acpi_nfit_update_notify() and acpi_nfit_init() will be called to initialize it. Regardless of whether or not acpi_nfit_init() fails in that case, the acpi_nfit_shutdown() devm action is not added for the NFIT device and acpi_desc is never removed from the acpi_descs list. If the acpi_desc object is freed subsequently on driver removal, any subsequent ACPI MCE will lead to a use-after-free like in the previous case. To address the first issue mentioned above, make acpi_nfit_probe() call acpi_nfit_shutdown() directly on acpi_nfit_init() failures and to address the other one, add a remove callback to the driver and make it call acpi_nfit_shutdown(). Also, since it is now possible to pass NULL to acpi_nfit_shutdown() or the acpi_desc object passed to it may not have been initialized, add checks against NULL for acpi_desc and its nvdimm_bus field to that function and make acpi_nfit_unregister() clear the latter after unregistering the NVDIMM bus.	7.0	More Details
CVE-2026-64420	In the Linux kernel, the following vulnerability has been resolved: mfd: cros_ec: Delay dev_set_drvdata() until probe success If ec_device_probe() fails, cros_ec_class_release releases memory for the cros_ec_dev structure. However, because the drvdata was already set, sub-drivers like cros_ec_typec can still retrieve the stale pointer via the platform device. This leads to a use-after-free when cros_ec_typec attempts to access &typec->ec->ec->dev on a device that has already been released. Move dev_set_drvdata() to ensure that the pointer is only made available once all initialization steps have succeeded. sysfs: cannot create duplicate filename '/class/chromeos/cros_ec' Call trace: sysfs_do_create_link_sd+0x94/0xdc sysfs_create_link+0x30/0x44 device_add_class_symlinks+0x90/0x13c device_add+0xf0/0x50c ec_device_probe+0x150/0x4f0 platform_probe+0xa0/0xe0 ... BUG: KASAN: invalid-access in __memcpy+0x44/0x230 Write at addr f5ffff809e2d33ac by task kworker/u32:5/125 Pointer tag: [f5], memory tag: [fe] Tainted : [W]=WARN, [O]=OOT_MODULE Hardware name: Google Navi unprovisioned 0x7FFFFFFF/sku0 board/sku3 Workqueue: events_unbound deferred_probe_work_func Call trace: __memcpy+0x44/0x230 cros_ec_check_features+0x60/0xcc [cros_ec_proto] cros_typec_probe+0xe8/0x6e0 [cros_ec_typec] platform_probe+0xa0/0xe0	7.0	More Details
CVE-2026-64219	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Validate payload length and link_index in dc_process_dmub_aux_transfer_async [Why&How] dc_process_dmub_aux_transfer_async() copies payload->length bytes into a 16-byte stack buffer (dpaux.data[16]) guarded only by an ASSERT(), which is a no-op in release builds. If a caller ever passes length > 16 this results in a stack buffer overflow via memcpy. Additionally, link_index is used to dereference dc->links[] without bounds checking against dc->link_count, risking an out-of-bounds access. Replace the ASSERT with a hard runtime check that returns false when payload->length exceeds the destination buffer size, and add a bounds check for link_index before it is used. (cherry picked from commit ba4caa9fecdf7a38f98c878ad05a8a64148b6881)	7.0	More Details
CVE-2026-64222	In the Linux kernel, the following vulnerability has been resolved: octeontx2-pf: avoid double free of pool->stack on AQ init failure otx2_pool_aq_init() frees pool->stack when mailbox sync or retry allocation fails, but leaves the pointer unchanged. Later, otx2_sq_aura_pool_init() unwinds the partial setup through otx2_aura_pool_free(), which frees pool->stack again. The CN20K-specific cn20k_pool_aq_init() implementation has the same bug in its corresponding error path. Set pool->stack to NULL immediately after the local free so the shared cleanup path does not free the same stack again while cleaning up partially initialized pool state. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under	7.0	More Details

	development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1-rc3. Runtime validation was not performed because reproducing this path requires OcteonTX2/CN20K hardware.		
CVE-2026-53667	React Router is a router for React. In versions 7.11.0 through 7.17.0, the RSCErrorHandler is missing protocol validation, allowing for redirects from untrusted sources. This issue is a follow up to CVE-2026-53667, and only affects consuming applications if they are using the unstable RSC APIs. This issue has been fixed in version 7.18.0.	6.9	More Details
CVE-2026-53668	React Router is a router for React. In versions 6.30.2 through 6.30.4 and 7.9.6 through 7.12.0, applications that allow open redirects are vulnerable to XSS. An attacker could craft a malicious link that redirects users to an unexpected external site or that exploits an XSS vector.This issue has been fixed in version 7.13.0.	6.9	More Details
CVE-2026-65923	A URL validation weakness in JFrog Artifactory Ansible repository handling could allow a user, under specific repository access conditions, to cause unintended server-side requests. The issue primarily affects confidentiality and integrity and has been addressed in fixed Artifactory versions.	6.8	More Details
CVE-2026-65695	Office-Word-MCP-Server through 1.1.11 contains a path traversal vulnerability in its document tools that allows attackers who can influence the filename argument to read arbitrary .docx files or create and overwrite .docx files outside the intended working directory. Attackers can supply absolute paths or ../ traversal sequences directly to document open and save operations, bypassing the check_file_writeable and ensure_docx_extension helpers which perform no base-directory confinement or realpath validation.	6.8	More Details
CVE-2026-16615	A flaw was found in librest. The PKCE implementation for OAuth authorization uses the GRand function from the GLib API, a cryptographically insecure pseudo-random number generator. Because the generated "code verifier" lacks sufficient cryptographic entropy, a malicious actor can reverse-engineer the pseudo-random number generator (PRNG) seed to predict or reconstruct the code verifier string, allowing an attacker to bypass PKCE protections and successfully impersonate the client during the OAuth 2.0 authorization flow.	6.8	More Details
CVE-2026-14827	The Calendar WordPress plugin before 1.3.18 does not properly escape a user-supplied event field before outputting it inside an HTML attribute on a public-facing page, allowing users with the Contributor role to inject arbitrary JavaScript that executes in the browser of anyone viewing the calendar.	6.8	More Details
CVE-2026-6390	A flaw was found in GNU nano's multi-buffer error message handling. When a user opens multiple files at startup and one triggers an ALERT-level error, a specially crafted filename containing printf format specifiers can be reinterpreted. This format string vulnerability may allow an attacker to achieve stack information disclosure, cause a denial of service (crash), or potentially perform arbitrary memory writes.	6.8	More Details
CVE-2026-65436	Editor Arbitrary File Deletion in Kirki <= 6.0.13 versions.	6.8	More Details
CVE-2026-50044	Pronetiqs IntraVUE versions 3.2.1a14 and prior have an inadequate encryption strength vulnerability which could allow an attacker to steal admin credentials via weak hash or a pass-the-hash attack.	6.8	More Details
CVE-2026-10723	BIND may accept incorrect child-zone NSEC3 records as valid, which could allow an attacker to forge authenticated NXDOMAIN responses. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	6.8	More Details
CVE-2026-27377	Booking Agent Broken Access Control in QuickCal - Appointment Booking Calendar for WordPress <= 1.0.16 versions.	6.7	More Details
CVE-2026-46737	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) an Improper Input Validation vulnerability in the REST API. A high privileged attacker with remote access could potentially exploit this vulnerability, leading to Remote execution.	6.7	More Details
CVE-2026-66028	Ekushey Project Manager CRM through version 5.0 contains a missing uniqueness constraint vulnerability that allows authenticated administrators to create duplicate client accounts with identical email and password credentials. Attackers can exploit the lack of email field uniqueness enforcement to create conflicting account states where multiple accounts share the same email address with different passwords, resulting in unpredictable authentication behavior and unauthorized account access.	6.7	More Details
CVE-2026-65010	Datasets through 5.00, fixed in commit ad2d853, contains a symlink-following vulnerability in Extractor.extract() that allows local attackers to write arbitrary files by pre-planting symlinks at predictable output paths. Attackers can redirect archive extraction to arbitrary filesystem locations in shared-cache environments, enabling overwrite of sensitive files and potential privilege escalation or code execution.	6.6	More Details
CVE-2026-44192	A flaw was found in the Ansible Lightspeed Model Context Protocol (MCP) server. This vulnerability, known as path traversal, allows an attacker to manipulate an AI agent through indirect prompt injection. By doing so, the attacker can cause the server to write files to unauthorized locations on the user's system. This can result in the exposure of sensitive host information and enable the attacker to execute malicious commands, potentially leading to a full system compromise.	6.6	More Details
	The userspace verifier z_vrfy_log_filter_set() for the log_filter_set syscall in subsys/logging/log_mgmt.c performed a signed comparison against the int16_t src_id parameter: src_id < (int16_t)log_src_cnt_get(domain_id). Any negative value for src_id (e.g. -1) trivially satisfied this check and was forwarded into z_impl_log_filter_set, where it propagated to filter_set() and ultimately to get_dynamic_filter(), which uses source_id as an unsigned index into the linker-section		

CVE-2026-10682	array &TYPE_SECTION_START(log_dynamic)[source_id].filters. After implicit conversion through uint32_t, an int16_t -1 becomes 0xFFFFFFFF, indexing log_dynamic far out of bounds and causing the kernel to perform an OOB read and an OOB read-modify-write (LOG_FILTER_SLOT_GET/SET) against memory adjacent to the log_dynamic section. The written value is a constrained 3-bit log level slot within the targeted 32-bit word, but the target address is attacker-chosen (a small negative offset from log_dynamic) and the write occurs in supervisor mode following a syscall from an unprivileged user thread, providing a kernel memory-corruption / privilege-escalation primitive. The defect is reachable on any build with CONFIG_USERSPACE=y and CONFIG_LOG_RUNTIME_FILTERING=y. Present from Zephyr v3.3.0 through v4.4.1. The fix replaces the signed bound check with an unsigned comparison: (uint32_t)src_id < log_src_cnt_get(domain_id), which correctly rejects negative inputs.	6.6	More Details
CVE-2026-57599	There is a privilege escalation vulnerability in some Hikvision cameras. Due to incorrect permission allocation in the device program, attackers can escalate privileges and gain full control of the device after authenticating via SSH.	6.6	More Details
CVE-2026-65561	Contributor Cross Site Scripting (XSS) in WordPress Social Login and Register <= 7.8.0 versions.	6.5	More Details
CVE-2026-65480	Contributor Cross Site Scripting (XSS) in TheGem <= 5.11.1 versions.	6.5	More Details
CVE-2026-39155	Knot DNS before 3.4.10 and 3.5.x before 3.5.4 contains a vulnerability in mod-onlinesign where the next NSEC owner name can be computed incorrectly. This can create an overly broad authenticated denial interval, allowing downstream validating resolvers using aggressive negative caching to synthesize negative answers for legitimate names and causing resolver-side denial of service.	6.5	More Details
CVE-2026-65482	Contributor Cross Site Scripting (XSS) in LA-Studio Element Kit for Elementor <= 1.6.2 versions.	6.5	More Details
CVE-2026-50103	A NULL pointer dereference in the L2 GOOSE and R-GOOSE shared parser, which may allow a network-adjacent attacker to crash a subscribing application by sending a crafted GOOSE frame containing a malformed TLV value.	6.5	More Details
CVE-2026-13062	An authenticated user with write privileges on a Queryable Encryption-enabled collection may be able to modify internal encryption metadata fields that are intended to be server-controlled, by sending crafted write commands through the mongos router on a sharded cluster. This can result in corruption of encrypted query correctness.	6.5	More Details
CVE-2026-13060	An authenticated user with limited read privileges may be able to access documents from collections they are not authorized to read, due to an inconsistency in how the $graphLookup aggregation stage is evaluated during authorization and during execution. Affected scenarios involve collections referenced within existing view pipeline definitions.	6.5	More Details
CVE-2026-13064	Certain query operations involving deeply nested $jsonSchema constructs can trigger disproportionate CPU consumption in affected MongoDB deployments, potentially leading to resource exhaustion. The resulting CPU-bound operation cannot be interrupted through standard administrative controls.	6.5	More Details
CVE-2026-66391	Use of Insufficiently Random Values, Protection Mechanism Failure vulnerability in Apache Wicket. This issue affects Apache Wicket: from 9.0.0 through 9.23.0, from 10.0.0 through 10.9.0. Users are recommended to upgrade to version 10.10.0, which fixes the issue.	6.5	More Details
CVE-2026-13065	A user with read-only privileges is able to craft an aggregation pipeline using the $linearFill window function operator with a specific sortBy expression type to cause the mongod process to terminate abnormally, resulting in denial of service. The issue stems from insufficient validation of sort specifications during execution.	6.5	More Details
CVE-2026-50248	In NLnet Labs Unbound 1.7.0 up to and including 1.25.1, when an auth/rpz zone has a configured primary hostname that resolves to BOGUS A/AAAA, it is still considered as a possible XFR endpoint. A malicious actor that can spoof the hostname's A/AAAA record (no valid RRSIG required) becomes the zone's XFR primary and can replaces the entire zone/the resolver's entire response policy.	6.5	More Details
CVE-2026-13066	Improper handling of DBPointer objects during BSON serialization in MongoDB's server-side JavaScript engine can result in internal process memory contents being included in data returned to the client. This constitutes an unintended information disclosure affecting deployments that use server-side JavaScript.	6.5	More Details
CVE-2026-13055	The `$_internalIndexKey` aggregation expression can be used by any authenticated user to crash a MongoDB server (mongod). The expression fails to handle compound wildcard index specifications, triggering an internal consistency check that aborts the server process. The user must be able to run an aggregation pipeline.	6.5	More Details
CVE-2025-13146	The The Contact Form 7 – Dynamic Text Extension plugin for WordPress is vulnerable to arbitrary shortcode execution in all versions up to, and including, 5.0.6. This is due to the software allowing users to execute an action that does not properly validate a value before running do_shortcode. This makes it possible for unauthenticated attackers to execute arbitrary shortcodes. The vulnerability was partially patched in version 5.0.4.	6.5	More Details
CVE-2026-65475	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WP Chill Modula Image Gallery allows Stored XSS. This issue affects Modula Image Gallery: from 2.14.25 through 2.14.30.	6.5	More Details

CVE-2026-65473	Contributor Cross Site Scripting (XSS) in Virtue/Ascend/Pinnacle Toolkit <= 4.9.12 versions.	6.5	More Details
CVE-2026-66399	phpMyFAQ before 4.1.6 contains a privilege escalation vulnerability in GroupController::updateMembers() that allows administrators with only group-management permissions to join privileged groups without verification of required rights. Attackers can add themselves to pre-existing groups holding user-management rights and immediately inherit those permissions to modify or delete user accounts.	6.5	More Details
CVE-2026-13069	An authenticated user can cause excessive CPU consumption or out-of-memory conditions on a MongoDB server by sending a crafted Queryable Encryption find payload containing an unvalidated field used to control an internal computation loop. The resulting resource exhaustion degrades availability for other operations.	6.5	More Details
CVE-2026-16767	A vulnerability was detected in Ne-Lexa php-zip up to 4.0.2. This affects the function ZipFile::extractTo of the file src/ZipFile.php of the component ZIP Handler. Performing a manipulation of the argument entryName results in path traversal. It is possible to initiate the attack remotely. The exploit is now public and may be used. The project was informed of the problem early through an issue report but has not responded yet.	6.5	More Details
CVE-2026-65470	Contributor Cross Site Scripting (XSS) in Fluent Support <= 2.3.0 versions.	6.5	More Details
CVE-2026-65465	Contributor Cross Site Scripting (XSS) in JetElements For Elementor <= 2.9.1.1 versions.	6.5	More Details
CVE-2026-65562	Contributor Cross Site Scripting (XSS) in BetterDocs <= 4.6.2 versions.	6.5	More Details
CVE-2026-13071	An authenticated user with read access can cause the mongod process to be terminated through certain aggregation expressions that execute server-side JavaScript. The issue involves improper memory handling during document processing.	6.5	More Details
CVE-2026-13056	Using expressions that generate large arrays it is possible to craft a query that creates very large intermediate objects in memory, causing the server to crash with OOM error.	6.5	More Details
CVE-2026-64872	Joomla Extension - regularlabs.com - Path traversal in Cache Cleaner Pro extension - Custom purge and log paths could escape the site webroot directory.	6.5	More Details
CVE-2026-10819	Mattermost versions 11.6.x <= 11.6.5, 10.11.x <= 10.11.20, 11.8.x <= 11.8.1, 11.7.x <= 11.7.4 fail to limit the number of frames and enforce the file size cap on animated GIF uploads, which allows an authenticated attacker to cause a denial of service via a crafted animated GIF uploaded as a custom emoji.. Mattermost Advisory ID: MMSA-2026-00695	6.5	More Details
CVE-2026-66749	Let's Chat 0.4.0 through 0.4.8 contains a null dereference vulnerability that allows authenticated attackers to crash the server by supplying a valid 24-character hex string room parameter that matches no document in the database. Attackers can send a crafted GET /messages request causing an uncaught TypeError in an asynchronous Mongoose callback that terminates the Node.js server process, with the same defect reachable through multiple code paths including the socket.io interface.	6.5	More Details
CVE-2026-66063	goshs is a feature-rich single-binary file server for red teamers and developers. Prior to 2.1.5, the httpserver/updown.go multipart upload handler split part.FileName() on / but did not reject .., allowing an unauthenticated upload with filename .. to create a file outside the served tree. This issue is fixed in version 2.1.5.	6.5	More Details
CVE-2026-63263	Uncontrolled Resource Consumption (CWE-400) in Elasticsearch can lead to denial of service via Exponential Data Expansion (CAPEC-197). An authenticated user may submit a specially crafted query to the ES|QL engine that causes exponential CPU consumption during query evaluation. Because the resource exhaustion persists beyond query completion, repeated requests can fully exhaust the available query worker resources, rendering ES|QL queries unavailable until the node is restarted.	6.5	More Details
CVE-2026-65435	Unauthenticated Broken Access Control in Thrive Leads Version <= 10.9.2 versions.	6.5	More Details
CVE-2026-65434	Subscriber Sensitive Data Exposure in ЮKassa для WooCommerce <= 2.16.1 versions.	6.5	More Details
CVE-2026-65607	SiYuan before v3.7.2 contains a path traversal vulnerability in the /export/temp/ short-circuit branch of the serveExport handler (kernel/server/serve.go). Unlike the main export branch, this branch joins the raw, percent-decoded request path with util.TempDir and serves the file without the IsSubPath or IsSensitivePath checks added in the earlier export-disclosure hardening (GHSA-6865-qjcf-286f). An authenticated attacker can send percent-encoded traversal sequences (e.g. /export/temp/%2e%2e/.../etc/passwd, where %2e%2e is decoded to '..') to read arbitrary	6.5	More Details

	files outside TempDir, including /etc/passwd, SSH keys (~/.ssh/*), and SiYuan workspace *.db and *.log files, bypassing the sensitive-file protection.		
CVE-2026-10822	If BIND encounters a particular invalid data structure in a DNS record, it will accept the invalid data, and may subsequently abort and exit. BIND will first need to store a DNS record for a key (KEY, DNSKEY, etc.). That key must specify a PRIVATEDNS algorithm (253), and in the algorithm identifier, improperly give a length longer than the actual identifier data. The invalid identifier will be stored. If BIND later needs to render that record to text, it will use the invalid length during processing, leading to a consistency check failing. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.18.11-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.	6.5	More Details
CVE-2026-64728	A permissions issue was addressed with improved validation. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Maliciously crafted web content may violate iframe sandboxing policy.	6.5	More Details
CVE-2026-64730	The issue was addressed with improved UI. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Visiting a website that frames malicious content may lead to UI spoofing.	6.5	More Details
CVE-2026-64735	An inconsistent user interface issue was addressed with improved state management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. A remote attacker may be able to bypass network filters.	6.5	More Details
CVE-2026-65536	Unauthenticated Cross Site Request Forgery (CSRF) in => (افزونه حمل و نقل ووکامرس (پست پیشتاز و سفارشی، پیک موتوری 4.4.5 versions.	6.5	More Details
CVE-2026-47755	ITFlow provides an IT documentation, ticketing and accounting system for small managed service providers. Prior to version 26.05, low-privileged authenticated agent can retrieve plaintext credentials and TOTP secrets belonging to another client by directly requesting the credential edit modal with an arbitrary `credential_id`. The endpoint does not enforce client scoping or object-level authorization before loading and decrypting the credential record. Version 26.05 fixes the issue.	6.5	More Details
CVE-2026-65533	Contributor Cross Site Scripting (XSS) in Smart SEO Tool <= 4.1.2 versions.	6.5	More Details
CVE-2026-64742	This issue was addressed by using HTTPS when sending information over the network. This issue is fixed in iOS 26.6 and iPadOS 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	6.5	More Details
CVE-2026-64743	An authorization issue was addressed with improved state management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	6.5	More Details
CVE-2026-65528	Contributor Cross Site Scripting (XSS) in BSK PDF Manager <= 3.8 versions.	6.5	More Details
CVE-2026-65527	Contributor Cross Site Scripting (XSS) in LIQUID SPEECH BALLOON <= 1.2.5 versions.	6.5	More Details
CVE-2026-65522	Contributor Cross Site Scripting (XSS) in Manual - Documentation, Knowledge Base & Education WordPress Theme <= 7.5.4 versions.	6.5	More Details
CVE-2026-7868	IBM OPENBMC FW1110.00 through FW1110.20, and FW1060.00 through FW1060.71 allows ReadOnly users to escalate privileges and give themselves administrator privileges.	6.5	More Details
CVE-2026-2406	Authorization bypass through User-Controlled key vulnerability in Universe Software Computer Marketing Trade and Industry Inc. Online Registration and Workflow Management System allows Exploiting Trust in Client. This issue affects Online Registration and Workflow Management System: through 12022026.	6.5	More Details
CVE-2026-65519	Author Cross Site Scripting (XSS) in Photo Gallery <= 2.7.7.29 versions.	6.5	More Details
CVE-2026-10681	In Zephyr's userspace dynamic-objects subsystem, thread_idx_alloc() in kernel/userspace/userspace.c allocated a new thread permission index from the global _thread_idx_map[] bitmap without holding lists_lock. On SMP systems, two user-mode threads invoking the k_object_alloc(K_OBJ_THREAD) syscall concurrently can both observe the same low free bit, perform the same non-atomic RMW to clear it, and return the identical tidx. The two newly created K_OBJ_THREAD objects are then assigned the same thread_id, so the two user threads alias a single bit position in every kernel object's perms[] bitfield: any subsequent grant of access on a kernel object to one thread is implicitly a grant to the other, defeating userspace ACL isolation. A secondary lost-update window between the unlocked &=~BIT() in alloc and the locked |= BIT() in thread_idx_free() can also leak entries from the thread-index pool. The	6.5	More Details

	defect is reachable from any user-mode thread via the unrestricted __syscall k_object_alloc and is gated on CONFIG_USERSPACE, CONFIG_DYNAMIC_OBJECTS, and CONFIG_SMP. The flaw was introduced when the per-thread permission index was added in 2018 and is present in every release up to and including v4.4.0. Fixed by holding lists_lock across the bitmap RMW and the permissions clear (and inlining the obj_list traversal that previously took the lock itself).		
CVE-2026-65518	Contributor Cross Site Scripting (XSS) in Accept Donations with PayPal & Stripe <= 1.5.5 versions.	6.5	More Details
CVE-2026-65514	Contributor Cross Site Scripting (XSS) in Appointment Hour Booking <= 1.5.86 versions.	6.5	More Details
CVE-2026-14568	The User Frontend: AI Powered Frontend Post Submission, User Directory, User Profile, Membership & User Registration WordPress plugin before 4.3.8 does not correctly verify ownership before deleting an attachment, allowing unauthenticated attackers to permanently delete author-less attachments such as guest uploads and User Frontend: AI Powered Frontend Post Submission, User Directory, User Profile, Membership & User Registration WordPress plugin before 4.3.8-installed placeholder media.	6.5	More Details
CVE-2026-65503	Contributor Cross Site Scripting (XSS) in Ultimate Store Kit Elementor Addons <= 3.0.5 versions.	6.5	More Details
CVE-2026-49159	Exposure of sensitive information to an unauthorized actor in Microsoft Graph allows an authorized attacker to disclose information over a network.	6.5	More Details
CVE-2026-65499	Unauthenticated Broken Access Control in PeproDev Ultimate Invoice <= 2.2.6 versions.	6.5	More Details
CVE-2026-66412	Leantime 3.6.2 and prior contains a broken access control vulnerability that allows authenticated users to read milestone data from projects they are not assigned to by supplying arbitrary integer milestone IDs to the tickets.getMilestone JSON-RPC endpoint. Attackers can enumerate integer milestone IDs through the JSON-RPC API to access project planning information, milestone titles, descriptions, and timelines across all projects on the instance regardless of project membership.	6.5	More Details
CVE-2026-15304	The Plugin Organizer plugin for WordPress is vulnerable to SQL Injection via the 'PO_plugin_path' parameter in versions up to, and including, 10.2.4. This is due to insufficient escaping on the user-supplied parameter in the perform_plugin_search() function, where esc_sql() output is passed as the replacement string to preg_replace(), which collapses backslash escapes and defeats the quoting protection; additionally, the AJAX handler lacks both nonce verification and capability checks. This makes it possible for authenticated attackers, with subscriber-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-16544	A flaw was found in AWX. The websocket event consumer performs RBAC authorization checks only for event groups that are mapped in the consumer_access() function (job_events, workflow_events, ad_hoc_command_events). Three event groups - inventory_update_events, project_update_events, and system_job_events — are not mapped, causing the authorization check to be skipped. Any authenticated user can subscribe to these unmapped websocket event groups for any object ID and receive real-time stdout output from jobs belonging to organizations they have no access to. This is an incomplete remediation of CVE-2020-10698.	6.5	More Details
CVE-2026-61945	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in MultiVendorX WooCommerce Product Stock Alert allows Retrieve Embedded Sensitive Data. This issue affects WooCommerce Product Stock Alert: from n/a through 3.0.6.	6.5	More Details
CVE-2026-65449	Contributor Cross Site Scripting (XSS) in MapSVG <= 8.14.0 versions.	6.5	More Details
CVE-2026-16802	Cleartext storage of sensitive information in the variables feature in Devolutions PowerShell Universal 2026.2.2 and earlier allows a local actor with file system access to read secret values via secret variables stored in cleartext on disk when no vault is selected.	6.5	More Details
CVE-2026-65707	Likeshop through 3.0.5 contains an authenticated SQL injection vulnerability that allows admin-level users to extract arbitrary database contents by submitting unsanitized POST parameters to the adjustAccount endpoint. The adjustAccount method in UserLogic.php concatenates the money, integral, growth, and earnings parameters directly into Db::raw() SQL fragments without type casting, numeric validation, or parameter binding, enabling boolean-based binary-search extraction of credentials, PII, and session tokens via distinct success and failure response messages.	6.5	More Details
CVE-2026-14932	In Progress® Telerik® UI for AJAX prior to v2026.2.708, the obsolete RadChart component's ChartImage.axd handler is vulnerable to unauthenticated file read and deletion of image-extension files within the application directory.	6.5	More Details
	n8n before 2.29.8 and 2.30.x before 2.30.1 (affected from 2.27.0, when the OAuth 2.1 consent and token-issuance flow was introduced) does not verify that the authenticated user has access to the workflow referenced as the OAuth		

CVE-2026-65594	resource. On instances with at least one active MCP Server Trigger workflow configured with n8n OAuth2 authentication, a member-level user can register an OAuth client, self-approve consent for another user's workflow, and obtain a valid token. The workflow then runs in the owner's project context with the owner's stored credentials, and the attacker can set tool inputs and read outputs (potentially including data from the owner's connected integrations), breaking user and project isolation.	6.5	More Details
CVE-2026-27403	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in NerdPress Hubbub Lite allows Stored XSS. This issue affects Hubbub Lite: from n/a through 1.36.3.	6.5	More Details
CVE-2026-49326	Missing Authorization vulnerability in Apache HBase thrift and rest delegation service. A scan operation in thrift/rest service has 3 steps, open, fetch(possible multiple times), close. The open step will return an id which will be passed back to server for identifying the scanner instances stored at server side. We missed the owner check in fetch and close steps which means a user can fetch rows from the scanner which is opened by other users, and close scanners which belongs to other users. This issue affects Apache HBase:from 3.0.0-alpha-1 through 3.0.0-beta-1, from 2.6.0 through 2.6.5, from 2.5.0 through 2.5.14, through 2.4.*. Users are recommended to upgrade to version 3.0.0-beta-2, 2.6.6 and 2.5.15, which fixes the issue.	6.5	More Details
CVE-2026-17059	A flaw was found in the role-users endpoint of the keycloak-services library, which is the core component of the Keycloak identity and access management solution. The issue occurs because the system fails to check if an administrator has permission to view individual users when listing members of a role. This allows a restricted administrator to see private information, such as names and email addresses, for users they should not be able to access.	6.5	More Details
CVE-2026-62435	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] With the introduction of Grant Table v2 came the requirement to be able to switch between versions. Switching from v1 to v2 reduces the number of valid grant references, as a bigger shared entry structure is then needed while the shared table doesn't change size. Switching from v2 back to v1 the status frames, which are separate in v2, go away. Code holding, but intermediately dropping and then re-acquiring the grant table lock, sometimes wrongly assumes that said properties wouldn't change across the window in time where the lock is not being held. The v1 -> v2 issue is CVE-2026-62435. The v2 -> v1 issue is CVE-2026-62436.	6.5	More Details
CVE-2026-13075	An authenticated user can cause the mongod process to be terminated by the operating system under memory pressure via the $rankFusion and $scoreFusion aggregation stages. The issue originates in the server's error-handling path and requires the ability to run aggregation queries.	6.5	More Details
CVE-2026-16798	Insertion of sensitive information into sent data in the automation jobs API in Devolutions PowerShell Universal 2026.2.2 and earlier allows an authenticated user with scoped job or script read permission to obtain another user's stored OAuth refresh token via job read responses that fail to strip the refresh token.	6.5	More Details
CVE-2026-62436	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] With the introduction of Grant Table v2 came the requirement to be able to switch between versions. Switching from v1 to v2 reduces the number of valid grant references, as a bigger shared entry structure is then needed while the shared table doesn't change size. Switching from v2 back to v1 the status frames, which are separate in v2, go away. Code holding, but intermediately dropping and then re-acquiring the grant table lock, sometimes wrongly assumes that said properties wouldn't change across the window in time where the lock is not being held. The v1 -> v2 issue is CVE-2026-62435. The v2 -> v1 issue is CVE-2026-62436.	6.5	More Details
CVE-2026-65925	A user with JFrog Artifactory Cargo remote repository read access could make Artifactory request unintended URLs and return the response.	6.5	More Details
CVE-2026-6251	The Chaty Pro plugin for WordPress is vulnerable to Authenticated Time-Based Blind SQL Injection in versions up to and including 3.5.5. This is due to the fetch_custom_field() function in admin/class-admin-base.php retrieving the widget_id POST parameter via filter_input(INPUT_POST, ...) and directly concatenating the value into a raw SQL query in a numeric context without using $wpdb->prepare() or any integer casting. Additionally, the nonce verification check is performed after the SQL query has already executed, providing no protection against the injection. This makes it possible for authenticated attackers with subscriber-level access and above to inject arbitrary SQL commands, potentially leading to unauthorized extraction of sensitive database contents including user credentials and configuration data.	6.5	More Details
CVE-2026-65924	JFrog Artifactory support for Terraform remote repositories was found to be susceptible to Server-Side Request Forgery (SSRF). An authenticated user - or, if anonymous access is enabled on the repository, an unauthenticated user - could cause Artifactory to issue outbound HTTP requests to arbitrary destinations and receive the response content.	6.5	More Details
CVE-2026-57373	Customer Cross Site Scripting (XSS) in Funnel Kit Funnel Builder PRO <= 3.15.0.4 versions.	6.5	More Details
CVE-2026-66007	Datasets through 5.0.0, fixed in commit f989ef9, contains a path traversal vulnerability in folder-based dataset builders where the file_name metadata field is not properly validated before being joined to the dataset directory. Attackers can supply crafted file_name values with directory traversal sequences to read arbitrary local files, which are then embedded into output when save_to_disk or push_to_hub is called.	6.5	More Details
	The AI Copilot – Content Generator plugin for WordPress is vulnerable to generic SQL Injection via 'order[0][dir]'		

CVE-2026-13009	Parameter in all versions up to, and including, 1.5.4 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with subscriber-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. The required waic-nonce is emitted on the front-end whenever the [waic_form] or [aiwu-form] shortcode is rendered, enabling contributor-level users who can publish shortcodes to obtain a valid nonce and reach the vulnerable AJAX handler, which performs no capability check beyond nonce verification when the shortcodes are not already embedded in a page.	6.5	More Details
CVE-2026-62429	Accessing the vNUMA configuration data of a guest is still possible when domain destruction has already started. The cleaning up of that configuration information is not synchronized with its retrieval by a device model controlling the guest.	6.5	More Details
CVE-2026-13119	The Registrations For The Events Calendar plugin for WordPress is vulnerable to SQL Injection via JSON keys in the 'standard' parameter handled by the rtec_records_edit AJAX action in versions up to and including 3.2. The handler decodes attacker-controlled JSON from $_POST['standard'] and uses the JSON array keys directly as column identifiers in the SET clause of an UPDATE statement built inside RTEC_Db_Admin::update_entry(). Only esc_sql() (mysqli_real_escape_string) is applied to the identifier; that function escapes quotes, backslashes, and a few control characters but does not escape spaces, equals signs, parentheses, or hyphens, so an attacker can break out of the identifier context and inject subqueries (terminated with a SQL comment). This makes it possible for authenticated attackers, with Contributor-level access and above who can edit the targeted event, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-27372	Unauthenticated Sensitive Data Exposure in PeproDev Ultimate Invoice <= 2.2.6 versions.	6.5	More Details
CVE-2026-55970	Buffer Over-read vulnerability in Apache Thrift C++ bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.	6.5	More Details
CVE-2026-66018	Build readers can access another repository's environment properties. A caller with read access to an ordinary repository can select a readable repository parameter while retrieving environment properties for a protected build, exposing build environment secrets (confidentiality impact; no integrity or availability impact demonstrated).	6.5	More Details
CVE-2026-15448	The Tickera – Sell Tickets & Manage Events plugin for WordPress is vulnerable to generic SQL Injection via the 'tc_order_status_filter' parameter in all versions up to, and including, 3.6.0.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with staff-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-65599	n8n versions before 1.123.64, 2.29.8, and 2.30.1 contain a credential exposure vulnerability: when configured with a Google Service Account key, the full PEM private key was mistakenly placed in the JWT header's kid field (intended only for a key identifier). Because JWT headers are Base64-encoded rather than encrypted, the private key could be recovered by anything that logged or inspected the JWT. An attacker who obtained the key could impersonate the service account and access or modify any Google Cloud resource it was authorized to use. Only instances using Google Service Account credentials are affected.	6.5	More Details
CVE-2026-13192	In Progress® Telerik® UI for AJAX prior to v2026.2.708, insufficient validation of content submitted to the RadEditor PDF export feature may allow an authenticated attacker to trigger server-side requests to arbitrary hosts, resulting in outbound network connections and potential exposure of Windows authentication credentials.	6.5	More Details
CVE-2026-15267	The Taskbuilder – Project Management & Task Management Tool With Kanban Board plugin for WordPress is vulnerable to SQL Injection via the 'wppm_proj_filter' parameter in versions up to, and including, 5.0.9. This is due to insufficient escaping on the user-supplied parameter and the lack of sufficient preparation on the existing SQL query — the value is re-read at line 144 using only sanitize_text_field() (overwriting the earlier absint() result), then concatenated into the SQL WHERE clause as an unquoted numeric operand using only esc_sql(), which does not protect against injection in that context, and finally string-interpolated into the $wpdb->prepare() format string, bypassing parameterization entirely. This makes it possible for authenticated attackers, with subscriber-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-65713	Joomla Extension - regularlabs.com - Insecure path handling in Modals Pro extension - Modals gallery paths could enumerate unintended directories.	6.5	More Details
CVE-2026-15761	The Tickera – Sell Tickets & Manage Events plugin for WordPress is vulnerable to generic SQL Injection via the 'tc_event_filter' parameter in all versions up to, and including, 3.6.0.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with staff-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. The vulnerable code path is reachable by users holding the plugin's custom Staff role, as the plugin's add_required_capabilities() function grants that role the edit_tc_tickets_instances capability, providing access to the tc_tickets_instances admin list screen where the filter is applied.	6.5	More Details
CVE-2026-	A flaw was found in Dogtag PKI's ACME responder where the web.xml security constraints use exact URL pattern matching for admin-only enable/disable endpoints. By appending a trailing slash to the URL, an unauthenticated attacker can bypass the Tomcat authentication constraint while RESTEasy still routes the request to the handler,	6.5	More Details

18047	allowing unauthorized toggling of the ACME service state including persistent denial of service.		
CVE-2026-43804	This issue was addressed through improved state management. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, visionOS 26.6. Visiting a website may lead to an app denial-of-service.	6.5	More Details
CVE-2026-64875	Joomla Extension - regularlabs.com - IP spoofing vulnerability in GeoIP extension - GeoIP lookups trusted spoofable forwarded client-IP headers, this could cause GeoIP-rule bypass.	6.5	More Details
CVE-2026-15906	The Premium Packages – Sell Digital Products Securely plugin for WordPress is vulnerable to generic SQL Injection via the 'orderby' parameter in all versions up to, and including, 7.0.4 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with admin-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	6.5	More Details
CVE-2026-57384	Subscriber Cross Site Scripting (XSS) in WishList Member X <= 3.32.0 versions.	6.5	More Details
CVE-2026-59559	Subscriber Cross Site Scripting (XSS) in RT Mega Menu – Mega Menu Builder for Elementor & Gutenberg <= 1.5.1 versions.	6.5	More Details
CVE-2026-61487	Improper Authorization vulnerability in Apache ActiveMQ Broker, Apache ActiveMQ All, Apache ActiveMQ. An authenticated low-privilege user can bypass a per-destination write ACL by sending to an ActiveMQ temporary composite destination whose physical name is a comma-separated composite of real queues. This allows publishing messages to any of the destinations in the list without proper write ACL permissions because the authorization check is bypassed due to the composite destination being marked as temporary. This issue affects Apache ActiveMQ Broker: before 5.19.9, from 6.0.0 before 6.2.8; Apache ActiveMQ All: before 5.19.9, from 6.0.0 before 6.2.8; Apache ActiveMQ: before 5.19.9, from 6.0.0 before 6.2.8. Users are recommended to upgrade to version 5.19.9, 6.2.8 or 6.3.0, which fixes the issue.	6.5	More Details
CVE-2026-59522	Subscriber Broken Access Control in WP ERP <= 1.17.5 versions.	6.5	More Details
CVE-2026-59560	Subscriber Broken Access Control in FundEngine <= 1.7.8 versions.	6.5	More Details
CVE-2026-65433	Subscriber Broken Access Control in RT Mega Menu – Mega Menu Builder for Elementor & Gutenberg <= 1.5.1 versions.	6.5	More Details
CVE-2026-57808	Subscriber Arbitrary Content Deletion in WP EasyPay <= 4.5.0 versions.	6.5	More Details
CVE-2026-9737	During query planning when reading the sort pattern in raw BSONObj form, in some places we don’t explicitly handle the meta expression case. This may lead to incorrect transformations leading to invariant failure.	6.5	More Details
CVE-2026-43792	An authorization issue was addressed with improved state management. This issue is fixed in Safari 26.6, macOS Tahoe 26.6. An app may be able to access sensitive user data.	6.5	More Details
CVE-2026-59513	Subscriber Cross Site Scripting (XSS) in Masteriyo - LMS <= 2.3.0 versions.	6.5	More Details
CVE-2026-13076	An authenticated user can cause a {{mongod}} process to be terminated by the operating system under memory pressure by performing a specific data type conversion operation within MongoDB's aggregation framework. The behavior stems from disproportionate memory consumption during this operation, and requires both write access to the database and the ability to run aggregation queries.	6.5	More Details
CVE-2026-14955	The Checkout Field Editor for WooCommerce (Pro) plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 3.7.7 via the 'thwcfe_legacy_file' parameter. This makes it possible for authenticated attackers, with subscriber-level access and above, to read the contents of arbitrary files on the server, which can contain sensitive information.	6.5	More Details
CVE-2026-65448	Unauthenticated Cross Site Scripting (XSS) in Anti Spam and list cleaner – AcyChecker <= 1.8.1 versions.	6.5	More Details
CVE-2026-	Unauthenticated Broken Authentication in Easy Digital Downloads <= 3.6.7 versions.	6.5	More

59524			Details
CVE-2026-12688	The ProfileGrid WordPress plugin before 5.9.9.7 does not verify PayPal IPN notifications before granting paid group membership, allowing unauthenticated attackers to forge a payment notification and mark any user as a paid member of any group without any payment being made.	6.5	More Details
CVE-2026-65445	Unauthenticated Broken Access Control in Ad Invalid Click Protector (AICP) <= 1.3.0 versions.	6.5	More Details
CVE-2026-16078	The WCPOS – Point of Sale (POS) plugin for WooCommerce plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 1.9.8 via the 'type' parameter parameter. This makes it possible for authenticated attackers, with shop manager-level access and above, to read the contents of arbitrary files on the server, which can contain sensitive information. Successful exploitation requires supplying context=edit in the request, which bypasses the content-stripping logic in prepare_item_for_response() and returns the traversed file verbatim in the REST API response.	6.5	More Details
CVE-2026-61946	Unauthenticated Insecure Direct Object References (IDOR) in Easy Appointments <= 3.12.27 versions.	6.5	More Details
CVE-2026-43821	An access issue was addressed with improved access restrictions. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to read files outside of its sandbox.	6.5	More Details
CVE-2026-59557	Unauthenticated Broken Access Control in Events Made Easy <= 3.1.3 versions.	6.5	More Details
CVE-2026-66448	Contributor Cross Site Scripting (XSS) in Gallery PhotoBlocks <= 1.3.3 versions.	6.5	More Details
CVE-2026-65618	Improper URL validation when handling specific URLs, allows an attacker, under certain conditions, to make unauthorized requests from JFrog Artifactory, potentially exposing internal services and cached response data.	6.5	More Details
CVE-2026-66037	FFmpeg through 8.1.2, fixed in commit 5d7112c, contains an uncontrolled resource consumption vulnerability in the IAMF demuxer that allows an unauthenticated attacker to cause multi-gigabyte memory allocation from a 17-byte input file by supplying a crafted count_label field. The mix_presentation_obu() function in libavformat/iamf_parse.c calls av_calloc(count_label, sizeof(*language_label)) with an attacker-controlled value before validating available OBU data, enabling an allocation amplification of approximately 126 million bytes per input byte that exhausts process memory or triggers an OOM-kill during format probing.	6.5	More Details
CVE-2026-66038	FFmpeg through 8.1.2, fixed in commit 8670835, contains an information disclosure vulnerability in the LCL/ZLIB video decoder that allows attackers to expose uninitialized heap memory by supplying a valid zlib stream that inflates to fewer bytes than the expected frame size. The zlib_decomp() function in lcldec.c treats short decompression as non-fatal and continues to the RGB24 conversion path, which copies a full frame's worth of rows from the allocation buffer using original frame dimensions, causing uninitialized heap contents including pointer-derived allocator bytes to be copied into the attacker-observable AVFrame output and potentially defeating ASLR in long-lived media processing services.	6.5	More Details
CVE-2026-66445	Contributor Cross Site Scripting (XSS) in Open User Map <= 1.4.46 versions.	6.5	More Details
CVE-2026-57425	Unauthenticated Broken Access Control in Autopay dla WooCommerce <= 2.2.27 versions.	6.5	More Details
CVE-2026-45812	Incorrect Calculation of Buffer Size vulnerability in Apache NimBLE when processing Legacy Advertising Report HCI event. When a single HCI advertising report event bundles multiple reports, NimBLE miscalculated the offset to the next report. This can cause the host to read past the end of the buffer and deliver a GAP event with bogus data to the application. Severity is low: NimBLE's own controller never batches multiple reports into one event, so this only matters when NimBLE's host is paired with a third-party controller that does. This issue affects Apache NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.	6.5	More Details
CVE-2026-65589	n8n versions before 1.123.64 fail to properly mask custom HTTP header credentials in LLM sub-node execution data, writing plaintext API keys and secrets to workflow execution records. Authenticated users with access to execution data can read exposed header values and credentials that persist in the database and can be exported.	6.5	More Details
CVE-2026-60135	An attacker can modify data that should be restricted to read-only access.	6.5	More Details
CVE-2026-	Unauthenticated Broken Access Control in Knit Pay <= 9.6.0.0 versions.	6.5	More

57717			Details
CVE-2026-61886	Weintek cMT3092X HMI stores user account passwords in plaintext.	6.5	More Details
CVE-2026-66434	Contributor Cross Site Scripting (XSS) in Photonic Gallery & Lightbox for Flickr, SmugMug & Others <= 3.33 versions.	6.5	More Details
CVE-2026-66337	A flaw was found in libsoup. An unsigned integer underflow in the soup_filter_input_stream_read_until() function causes a heap buffer over-read when parsing multipart HTTP responses. A malicious HTTP server can exploit this by sending a crafted multipart response, potentially causing the client application to crash or disclose sensitive heap memory.	6.5	More Details
CVE-2026-66433	Contributor Cross Site Scripting (XSS) in Location Weather <= 3.0.6 versions.	6.5	More Details
CVE-2026-66339	A flaw was found in libsoup. After a CONNECT tunnel is established through an HTTP proxy, libsoup incorrectly attaches the Proxy-Authorization header to subsequent HTTPS requests sent through that tunnel to the destination server. This allows the destination server to capture proxy credentials, leading to information disclosure.	6.5	More Details
CVE-2026-64794	Joomla Extension - regularlabs.com - restricted user-data exposure in Users Anywhere and Articles Anywhere extensions - User tags, filters and conditions allowed access to insufficiently restricted user fields. Crafted content could expose authentication-related data, raw user parameters or restricted contact details.	6.5	More Details
CVE-2026-15787	The Ultimate Addons for Elementor plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Navigation Menu Widget data-toggle-icon/data-close-icon Attributes in all versions up to, and including, 2.9.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. wp_kses_post, applied on save for users without unfiltered_html, does not neutralize HTML-entity-encoded payloads stored inside data-* attributes on kses-allowed elements, as the browser decodes these values client-side before jQuery .html() renders them as markup.	6.4	More Details
CVE-2026-15794	The Grid/List View for WooCommerce plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'position' Shortcode Attribute in all versions up to, and including, 3.0.9 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The shortcode's all_page="1" attribute can be used to force the widget to render on any page, expanding the attack surface beyond shop and category pages.	6.4	More Details
CVE-2026-15016	The Paid Memberships Pro – Content Restriction, User Registration, & Paid Subscriptions plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Readonly User Field via [pmpro_member_profile_edit] Shortcode in all versions up to, and including, 3.8.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with subscriber-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15464	The WP Hotel Booking plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'widget_search' Shortcode Attribute in all versions up to, and including, 2.3.2 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This is only exploitable on browsers where access keys can be used as the payload is stored in a hidden attribute.	6.4	More Details
CVE-2026-15821	The SureDash – Community, Courses & Member Dashboard plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Shortcode Attributes in all versions up to, and including, 1.10.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15739	The Rich Showcase for Google Reviews plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'pagination' Shortcode Attribute in all versions up to, and including, 6.9.9 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15393	The Cozy Blocks – Page Builder for Gutenberg Editor & FSE with 600+ Patterns, 58 Blocks & Templates plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'postMeta.font.size' Block Attribute in all versions up to, and including, 2.2.11 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15755	The Open User Map – Interactive Leaflet Maps plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Shortcode Attributes in all versions up to, and including, 1.4.45 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The attack does not require post publication; a Contributor submitting a post for pending review is sufficient, as the payload executes when an Administrator opens the pending-review preview.	6.4	More Details

CVE-2026-15665	The Fluent Support – Helpdesk & Customer Support Ticket System plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'redirect-to' Shortcode Attribute in all versions up to, and including, 2.3.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The XSS payload is in a hidden attribute so it only fires in specific browsers when specific access keys are used making exploitation unlikely.	6.4	More Details
CVE-2026-15653	The Visualizer – Tables & Charts Manager with Built-in AI Generator plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'backend-title' parameter in all versions up to, and including, 4.0.5 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15646	The Brands for WooCommerce plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'style' Shortcode Attribute in all versions up to, and including, 3.8.8 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15648	The Brands for WooCommerce plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'width' Shortcode Attribute in all versions up to, and including, 3.8.8 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15394	The Header Footer Script Adder – Insert Code in Header, Body & Footer plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'asm_code' Snippet Meta in all versions up to, and including, 2.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with author-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-14481	The Equalize Digital Accessibility Checker – WCAG, ADA, EAA and Section 508 compliance plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'html' parameter in all versions up to, and including, 1.46.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This requires the attacker to have the ability to edit a post, as the REST endpoint /accessibility-checker/v1/post-scan-results/{id} is guarded only by the edit_post capability on the target post.	6.4	More Details
CVE-2026-15404	The Lpagery plugin for WordPress is vulnerable to Stored Cross-Site Scripting via post titles in versions up to, and including, 2.5.7. This is due to insufficient input sanitization and output escaping in the lpagery_add_filter_text_template_post() function, which is hooked to admin_footer and echoes the raw post_title of the post referenced by the ?lpagery_template query parameter directly inside a JavaScript single-quoted string literal, without esc_js(), esc_html(), or any other encoding. 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 (such as an administrator) accesses an admin page with the ?lpagery_template=<post_id> parameter pointing at the attacker's post.	6.4	More Details
CVE-2026-9729	The Webpushr Push Notifications plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'webpushr_notification_title' and 'webpushr_notification_body' parameters in versions up to, and including, 4.39.0. This is due to insufficient input sanitization in the save_send_notification_flag() function and missing output escaping in the wpp_notification_box() function, which concatenates raw post meta values directly into HTML attribute and textarea contexts. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15730	The GamiPress – Gamification plugin to reward points, achievements, badges & ranks in WordPress plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'heading_size' Shortcode Attribute in all versions up to, and including, 7.9.9.1 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. The wp_kses_post filter applied at post save does not neutralize this payload because only the inert shortcode text is stored in post_content; the dangerous HTML is synthesized at render time by the shortcode handler, entirely bypassing save-time sanitization.	6.4	More Details
CVE-2025-9205	The MapSVG plugin for WordPress is vulnerable to Stored Cross-Site Scripting in all versions up to, and including, 8.14.0. This is due to insufficient input sanitization and output escaping on user supplied attributes within the map options. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15334	The Cozy Blocks – Page Builder for Gutenberg Editor & FSE with 600+ Patterns, 58 Blocks & Templates plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'icon.view' Block Attribute in all versions up to, and including, 2.2.11 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-15425	The Yoast SEO – Advanced SEO with real-time guidance and built-in AI plugin for WordPress is vulnerable to Stored Cross-Site Scripting via Post Slug (post_name) in all versions up to, and including, 28.0 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with author-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This requires pretty permalinks to be enabled, as the exploit chain depends on get_permalink() embedding the stored percent-encoded post_name in the generated URL.	6.4	More Details

CVE-2026-15333	The Cozy Blocks – Page Builder for Gutenberg Editor & FSE with 600+ Patterns, 58 Blocks & Templates plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'cozyCustomFont' Block Attribute in all versions up to, and including, 2.2.11 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-9635	The WP Shortcode by MyThemeShop plugin for WordPress is vulnerable to Stored Cross-Site Scripting via the 'title' parameter of the [tab] shortcode in versions up to, and including, 1.4.17. This is due to insufficient input sanitization and output escaping in the mts_tabs() function, which outputs the title shortcode attribute directly into the HTML output between anchor tags without applying any escaping functions. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	6.4	More Details
CVE-2026-6454	The Firelight Lightbox plugin for WordPress is vulnerable to Stored DOM Cross-Site Scripting in versions up to and including 2.3.20. This is due to insufficient sanitization of the href attribute value within the FancyBox V2 PDF beforeLoad JavaScript callback generated in inc/fancybox-2.php, where this.href is string-concatenated directly into an HTML string without escaping, allowing a stored href containing entity-encoded double-quotes to break out of the data attribute and inject arbitrary event handlers into the DOM. This makes it possible for authenticated attackers with contributor-level access and above to inject arbitrary web scripts into pages that execute whenever a user clicks the malicious PDF link.	6.4	More Details
CVE-2026-15100	The Post Grid Gutenberg Blocks – PostX plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'searchnoresult' Block Attribute in all versions up to, and including, 5.0.32 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. This is particularly concerning in a cross-privilege scenario where a Contributor creates a pending or draft post containing the malicious block attribute, which then executes arbitrary JavaScript in the browser session of an Editor or Administrator who previews that post.	6.4	More Details
CVE-2026-16764	A vulnerability was identified in OWASP DefectDojo 2.59.0. This issue affects the function UserSerializer of the file dojo/api_v2/serializers.py of the component API/Web. Such manipulation of the argument is_staff leads to improper privilege management. The attack may be performed from remote. The exploit is publicly available and might be used. Upgrading to version 2.58.3 and 3.0.0 is capable of addressing this issue. The name of the patch is 68a272f299d096249fd3ba9c2676bf69012857bf. It is advisable to upgrade the affected component. 2.59.0 was not intended to be released and has been removed.	6.3	More Details
CVE-2026-66005	Jan through 0.8.4, fixed in commit 3e1c1e7, contains a CORS misconfiguration vulnerability in its local API server that allows network-adjacent attackers to bypass trusted host restrictions by exploiting the server's replacement of user-configured trusted hosts with a wildcard that reflects arbitrary origins with credentials. Attackers on the local network or using DNS rebinding can reach the unauthenticated OpenAI-compatible API to perform inference, enumerate models, invoke MCP tools, and read cross-origin responses.	6.3	More Details
CVE-2026-17530	A security flaw has been discovered in AstrBotDevs AstrBot up to 4.25.5. Affected by this vulnerability is the function _build_handoff_toolset of the file AstrBot/astrbot/core/astr_agent_tool_exec.py of the component Subagent. The manipulation results in incorrect authorization. The attack may be launched remotely. The exploit has been released to the public and may be used for attacks. The patch is identified as d23011262e8e75e1ec41b0f1f0091493a022327e. A patch should be applied to remediate this issue.	6.3	More Details
CVE-2026-17434	A flaw has been found in nanocoai NanoClaw up to 2.0.64. Affected is the function handleAddMcpServer of the file src/modules/self-mod/request.ts of the component add_mcp_server. Executing a manipulation can lead to improper authorization. The attack may be launched remotely. The exploit has been published and may be used. This patch is called e5b928783d5c485637565eb07d2967922dfbf8d8. A patch should be applied to remediate this issue.	6.3	More Details
CVE-2026-13067	When PROXY protocol v2 is used on the Unix domain socket path, roles derived from X.509 client certificates may not be validated against the configured tlsCATrusts allow-list. This can result in unintended role assignments following MONGODB-X509 authentication. Affected scenarios require local access to the proxy Unix domain socket and a valid X.509 certificate issued by a trusted certificate authority.	6.3	More Details
CVE-2026-17529	A vulnerability was identified in AstrBotDevs AstrBot up to 4.25.5. Affected is an unknown function of the file astrbot/core/astr_main_agent.py. The manipulation of the argument req.func_tool leads to incorrect authorization. The attack may be initiated remotely. The exploit is publicly available and might be used. The identifier of the patch is d23011262e8e75e1ec41b0f1f0091493a022327e. It is suggested to install a patch to address this issue.	6.3	More Details
CVE-2026-16490	A security flaw has been discovered in itsourcecode Hospital Management System 1.0. Impacted is an unknown function of the file /prescription.php. The manipulation of the argument editid results in sql injection. The attack can be executed remotely. The exploit has been released to the public and may be used for attacks.	6.3	More Details
CVE-2026-17458	A vulnerability was found in mf-yang openclaw-cn up to 0.2.1. This affects the function clickViaPlaywright of the file src/browser/routes/agent.act.ts of the component Browser Control HTTP API. Performing a manipulation results in server-side request forgery. It is possible to initiate the attack remotely. The exploit has been made public and could be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2026-11391	Tanium addressed a SQL injection vulnerability in Patch.	6.3	More Details

CVE-2026-65484	Contributor Broken Access Control in Style Kits <= 2.6.5 versions.	6.3	More Details
CVE-2026-15348	The Premium Packages – Sell Digital Products Securely plugin for WordPress is vulnerable to Authentication Bypass in all versions up to, and including, 7.0.4 via the `wpdmppdl` parameter. This is due to the `download()` function — hooked to the unauthenticated WordPress `wp` action — decoding the attacker-controlled `wpdmppdl` parameter using only `base64_decode()` and `json_decode()` with no HMAC, cryptographic signature, or nonce verification, and then issuing WordPress authentication cookies after a domain check that is trivially bypassed because both sides of the comparison are attacker-supplied values. This makes it possible for unauthenticated attackers to authenticate as any non-administrator WordPress user, including subscribers, customers, contributors, authors, editors, and shop managers, who owns an order, gaining full session-level access to that account.	6.3	More Details
CVE-2026-57703	Subscriber Broken Access Control in Sunshine Photo Cart <= 3.6.10.1 versions.	6.3	More Details
CVE-2026-51298	sqlite 3.41 is vulnerable to use after free in the JSON extraction function. After releasing JsonParse object memory via jsonParseFree(), the program still accesses internal member of the freed pointer, which can cause service crash and denial of service.	6.2	More Details
CVE-2026-65712	Joomla Extension - regularlabs.com - Insecure path handling in CDN for Joomla Pro extension - CDN versioning could check file paths outside the site directory, exposing local file existence and modification metadata.	6.2	More Details
CVE-2026-14190	The Sina Extension for Elementor WordPress plugin before 3.10.2 does not escape a value reconstructed from request input in one of its unauthenticated AJAX handlers before reflecting it into the HTML response, allowing unauthenticated attackers to execute arbitrary JavaScript in the browser of anyone who triggers a crafted request.	6.1	More Details
CVE-2026-13400	Simply Schedule Appointments is vulnerable to unauthenticated Stored Cross-Site Scripting in all versions up to and including 1.6.12.2. The root cause is a sanitization-ordering defect: the rendered notification content is decoded back into live HTML after it has already passed through the Simply Schedule Appointments WordPress plugin before 1.6.12.4's wp_kses_post() filter, so a double-encoded payload survives intake and is reintroduced as an executable element at render time.	6.1	More Details
CVE-2026-65697	Fathom Lite through 1.3.1 contains a stored cross-site scripting vulnerability in the analytics collection endpoint that allows unauthenticated attackers to inject a javascript: URI into the Top Pages dashboard by supplying a crafted hostname and pathname to the unauthenticated /collect endpoint. The parseHostname and parsePathname functions perform no URI scheme validation, allowing a javascript: hostname combined with a newline-prefixed pathname to be stored and later rendered as an anchor href in the authenticated dashboard without sanitization, enabling session hijacking and full account takeover when an operator clicks the poisoned entry.	6.1	More Details
CVE-2026-66390	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Apache Wicket. This issue affects Apache Wicket: from 9.0.0 through 9.23.0, from 10.0.0 through 10.9.0. Users are recommended to upgrade to version 10.10.0, which fixes the issue.	6.1	More Details
CVE-2026-12982	The Document Gallery WordPress plugin before 5.1.1 does not properly sanitise and escape user input before reflecting it back in the response of an unauthenticated AJAX action, leading to a Reflected Cross-Site Scripting vulnerability which can be exploited against unauthenticated users.	6.1	More Details
CVE-2026-65899	DOMPurify 3.0.0 before 3.4.9 does not reset the retained Trusted Types policy when clearConfig() is called, so a DOMPurify instance reused across trust boundaries stays bound to a previously supplied TRUSTED_TYPES_POLICY. A later caller that requests RETURN_TRUSTED_TYPE output receives a TrustedHTML object created by the old (potentially unsafe) policy rather than a clean default, which can lead to script execution at a Trusted Types sink. Passing TRUSTED_TYPES_POLICY: null on the later call also does not clear the retained policy.	6.1	More Details
CVE-2026-65900	DOMPurify versions >=3.0.0 and before 3.4.8, when configured with SAFE_FOR_TEMPLATES together with a DOM output mode (RETURN_DOM, RETURN_DOM_FRAGMENT, or IN_PLACE), fail to strip template expressions (e.g. ${evil}, {{evil}}, <%evil%>) inside <template> element content. The final normalization/scrub pass (_scrubTemplateExpressions) uses a NodeIterator and node.normalize() that do not descend into template.content, so expressions that only form after adjacent text nodes merge survive sanitization. This bypasses SAFE_FOR_TEMPLATES and can allow a downstream template engine to evaluate attacker-supplied expressions. The string output path is not affected.	6.1	More Details
CVE-2026-65901	DOMPurify through 3.4.6 contains a cross-site scripting vulnerability in IN_PLACE mode that trusts attacker-controlled nodeName on live non-form nodes. Attackers can supply hostile live DOM objects with real script children whose observable nodeName is clobbered to appear as allowed elements, causing scripts to execute when the sanitized tree is inserted into a live document.	6.1	More Details
CVE-2026-65902	DOMPurify before 3.4.7 (affected versions <= 3.4.5) passes direct references to the module-level DEFAULT_ALLOWED_TAGS and DEFAULT_ALLOWED_ATTR sets to the uponSanitizeElement and uponSanitizeAttribute hooks via data.allowedTags / data.allowedAttributes when sanitize is called without an explicit cfg.ALLOWED_TAGS / cfg.ALLOWED_ATTR array. A hook that mutates these fields permanently widens the default allow-lists for the lifetime of the DOMPurify instance, so all subsequent default-config sanitize calls inherit the widened defaults and attacker payloads using the poisoned tag/attribute name survive sanitization. removeAllHooks(), clearConfig(), and passing a fresh cfg do not recover the state; only constructing a new DOMPurify instance does.	6.1	More Details

CVE-2026-65903	DOMPurify before 3.4.0 contains a logic error in the ADD_TAGS function where short-circuit evaluation allows forbidden tags to bypass FORBID_TAGS restrictions. Attackers can craft input containing tags listed in FORBID_TAGS that are also added via ADD_TAGS function, causing them to be retained in sanitized output.	6.1	More Details
CVE-2026-65911	In DOMPurify through 3.3.3, function predicates supplied via ADD_ATTR or ADD_TAGS to DOMPurify.sanitize() persist in internal state (EXTRA_ELEMENT_HANDLING) across subsequent sanitize() calls on the same instance. If a later call on the same instance provides ADD_ATTR or ADD_TAGS as an array rather than a function, the previously set function handler is neither cleared nor overwritten, so it continues to approve attacker-controlled attributes or tags. This can allow dangerous event-handler attributes or forbidden tags (bypassing FORBID_TAGS) to survive sanitization, resulting in cross-site scripting. The vendor (Cure53) considers this an edge case outside DOMPurify's threat model; the referenced advisory lists 3.4.0 as the patched version.	6.1	More Details
CVE-2026-65912	DOMPurify before 3.3.2 contains a URI validation bypass vulnerability when ADD_ATTR is provided as a predicate function via EXTRA_ELEMENT_HANDLING.attributeCheck. Attackers can supply a predicate that accepts specific attribute and tag combinations to bypass URI-safe validation, allowing unsafe protocols like javascript: to survive sanitization and execute as DOM-based XSS when the link is activated.	6.1	More Details
CVE-2026-10082	The Advanced Ads WordPress plugin before 2.0.23 does not sanitize and escape a shortcode parameter before outputting it in the page, allowing users with the Contributor role and above to inject arbitrary web scripts that execute when the affected content is viewed, including by higher-privileged users.	6.1	More Details
CVE-2026-66010	DOMPurify before 3.4.12 fails to execute afterSanitizeElements hook for custom elements allowed via CUSTOM_ELEMENT_HANDLING.tagNameCheck, allowing attributes to bypass application security policies. Attackers can preserve sensitive attributes on custom elements that later re-inject them into innerHTML sinks, creating second-order XSS gadgets.	6.1	More Details
CVE-2026-65914	DOMPurify before 3.3.2 contains a mutation-XSS vulnerability when sanitized HTML is reinserted into special parsing contexts using innerHTML with wrappers like script, xmp, iframe, noembed, noframes, or noscript. Attackers can craft payloads with closing sequences that break out of the wrapper context during reparsing, reactivating dangerous markup with event handlers to execute JavaScript.	6.1	More Details
CVE-2026-15346	The VikBooking Hotel Booking Engine & PMS plugin for WordPress is vulnerable to Reflected Cross-Site Scripting via the 'category_id' parameter in all versions up to, and including, 1.8.13 due to insufficient input sanitization and output escaping. This makes it possible for unauthenticated attackers to inject arbitrary web scripts in pages that execute if they can successfully trick a user into performing an action such as clicking on a link. This is limited to browsers that support access keys as the injection is in a hidden element.	6.1	More Details
CVE-2026-8167	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in THEWP Digital Solutions News Theme V8 allows Reflected XSS. This issue affects News Theme V8: through 16.06.2026.	6.1	More Details
CVE-2026-53666	React Router is a router for React. In versions 6.4.0 through 7.17.0, if application code was written in a way that allows attacker-supplied input to overwrite certain aspects of errors caught by the SSR process, then it was possible for an attacker to trigger unexpected constructor execution on the client, which would in turn trigger an outbound network request. This is only possible with very specific (and unlikely) application-layer code. Note that this does not impact an application if it is using Declarative Mode. It only impacts Framework Mode and Data Mode applications that perform manual SSR/hydration. This issue has been fixed in version 7.18.0.	6.1	More Details
CVE-2026-51565	Cross-site scripting (XSS) vulnerability in Modules/Docs/DocsController.php in Milk admin <=0.9.8 allows remote attackers to inject arbitrary web script or HTML via the action parameter in a crafted request	6.1	More Details
CVE-2026-14515	IBM WebSphere Application Server 8.5, and 9.0 traditional could allow a remote attacker to conduct a cross-site scripting attack.	6.1	More Details
CVE-2026-17528	Versions of the package nice-select2 before 2.4.1 are vulnerable to Cross-site Scripting (XSS) via the <select> element. An attacker can supply a malicious payload that is rendered directly into the DOM without proper sanitization, causing arbitrary script execution in a victim’s browser when they view or interact with the affected page.	6.1	More Details
CVE-2026-9066	The WP Compress WordPress plugin before 7.10.04 does not validate the value of a query parameter that controls the asset CDN host before using it to build the URLs of JavaScript files emitted on the page, leading to Reflected XSS. When a visitor follows a crafted link, the WP Compress WordPress plugin before 7.10.04's loader injects script elements pointing to an attacker-controlled origin, which lets the attacker execute arbitrary JavaScript in the visitor's session on the target site.	6.1	More Details
CVE-2026-64828	Froiden TableTrack through 1.3.10 contains a stored cross-site scripting vulnerability that allows unauthenticated attackers to inject arbitrary HTML and JavaScript through the order notes field without sanitization. Attackers can craft malicious payloads in customer order placement that execute in the admin's browser session when viewing order details, enabling session token theft or unauthorized administrative actions.	6.1	More Details
CVE-2026-14171	An unauthenticated remote attacker can abuse the improper validation of the post-login redirect of the web-UI to trick users to a malicious website. This can result in a loss of confidentiality and availability.	6.1	More Details
CVE-	DOMPurify before 3.3.2 contains a prototype pollution vulnerability in USE_PROFILES mode that allows attackers to		

2026-65913	bypass attribute filtering by polluting Array.prototype properties. Attackers can set Array.prototype properties like onclick to true, causing DOMPurify to accept event handlers as allowlisted attributes and resulting in DOM-based XSS when sanitized markup is rendered.	6.1	More Details
CVE-2026-42494	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] The directory and Rock Ridge / SUSP walk in libfsimage's iso9660 driver derives several lengths directly from attacker-controlled on-disk fields without validating them: * The directory loop itself assumes a good record length. This is CVE-2026-42494. * The calculation of the System Use area may underflow. This is CVE-2026-42495. * The Rock Ridge extension loop assumes a good (inner) record length. This is CVE-2026-62423. * The Rock Ridge NM record processing assumes a good entry length. This is CVE-2026-62424. * The Rock Ridge CE record processing assumes a good size and offset. This is CVE-2026-62425.	6.1	More Details
CVE-2026-65756	Joomla Extension - regularlabs.com - XSS vector in Keyboard Shortcuts extension - Shortcut configuration accepted arbitrary inline JavaScript.	6.1	More Details
CVE-2026-8308	Improper neutralization of input during web page generation ('cross-site scripting') vulnerability in Polen Media Software and Information Services Website Template allows Reflected XSS. This issue affects Website Template: before v2.	6.1	More Details
CVE-2026-65882	Joomla Extension - joomdle.com - Reflected XSS vulnerability in Joomdle < 3.1.1 - The goto url parameter of the moodle wrapper endpoint allowed a reflected XSS vector.	6.1	More Details
CVE-2026-44276	Dell PowerProtect Data Manager, versions prior to 20.2.0.0, contain(s) an Exposure of Sensitive Information to an Unauthorized Actor vulnerability in the REST API. A high privileged attacker with local access could potentially exploit this vulnerability, leading to Information exposure.	6.0	More Details
CVE-2026-66475	Shop manager Cross Site Scripting (XSS) in Checkout Field Editor for WooCommerce – Checkout Manager <= 3.0.5 versions.	5.9	More Details
CVE-2026-66053	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift Python bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue. This replaces CVE-2026-41603	5.9	More Details
CVE-2026-65563	Author Cross Site Scripting (XSS) in Orbit Fox by ThemeIsle <= 3.0.7 versions.	5.9	More Details
CVE-2026-65557	Shop manager Cross Site Scripting (XSS) in Abandoned Cart Lite for WooCommerce <= 6.8.0 versions.	5.9	More Details
CVE-2026-50046	In NLnet Labs Unbound 1.15.0 up to and including 1.25.1, the TLS server name used for DNS-over-TLS (DoT) forwarded queries is tied to a struct's ('serviced_query') lifetime but also referenced by another struct ('waiting_tcp'). When the owning struct is jostled out of the mesh while the DoT TCP stream is still handshaking it frees the storage behind the referenced string and if the TLS stream then errors out, it dereferences the freed pointer. The dereference is read-only and the practical impact is a daemon crash resulting in denial of service. A malicious actor that knows a DoT forwarding/stub Unbound's configuration could exploit the vulnerability by quering records in the appropriate zone while keeping Unbound uder pressure so that the jostle logic kicks in. If answers for the vulnerable zone are slow, the likelihood of jostling such queries is higher, although the timing of the jostle needs to be precise. Requirements for a vulnerable Unbound is the existence of a stub/forward zone configured for DoT together with a configured '#authname' suffix on the server identification. The connectivity to the server needs to exhibit a transient failure at the correct time in order to kick off the vulnerable error path.	5.9	More Details
CVE-2026-56444	In NLnet Labs Unbound 1.20.0 up to and including 1.25.1, when Unbound is configured with 'serve-expired: yes' and 'serve-expired-client-timeout > discard-timeout > 0' (contrary to the suggested values), the discard-timeout branch during the serve expired logic drops an aged client reply without performing the correct accounting for the number of reply addresses for the query. Other identical branches outside of serve expired perform the correct decrement. Since the counter is never decremented in such scenario, it can reach the maximum limit and new clients for duplicate in-flight queries are silently dropped resulting in degradation of resolution service. A malicious actor can exploit the vulnerability by querying the resolver for a client-controlled slow-on-demand authoritative zone that can drive the counter past the threshold. Shipped defaults for 'serve-expired-client-timeout: 1800' and 'discard-timeout: 1900' make the branch unreachable.	5.9	More Details
CVE-2026-55990	In NLnet Labs Unbound 1.7.0 up to and including 1.25.1, when the 'dnscrypt:' clause lists more 'dnscrypt-provider-cert:' files than there are matching 'dnscrypt-secret-key:' files, Unbound fills only the matched prefix and leaves the tail slots at the '0xdb' fill that libsodium's allocator writes into every allocation. Unbound would then iterate over the number of cert files, not the actual slots, so it walks into a slot with garbage data filled with '0xdb' bytes. Any unauthenticated client that sends one UDP datagram of ≥ 68 bytes whose first 8 bytes are '0xdb' to 'dnscrypt-port' will use that garbage entry which leads to a garbage dereference killing the server. This is a silent faulty configuration that goes unnoticed until triggered with the right client query. Unbound needs to be compiled with DNSCrypt support ('--enable-dnscrypt').	5.9	More Details
	In NLnet Labs Unbound 1.10.0 up to and including 1.25.1, when 'serve-expired: yes' is set together with a 'response-ip: <net> redirect' /'response-ip-data: <net> CNAME <target>' rule (or the RPZ 'rpz-cname-override' equivalent), a		

CVE-2026-55717	remote client who controls any delegated domain can crash the daemon. The serve-expired-client-timeout callback runs a two-pass loop to chase the respip-generated CNAME alias; on the second pass it resets 'alias_rrset' but not 'partial_rep'. Later, this inconsistency leads to a NULL pointer dereference and an eventual crash. A malicious actor can exploit the vulnerability by controlling any zone that replies with an A/AAAA record that falls inside the configured response-ip/rpz subnet. By delaying the answer when the previous record has expired, the vulnerable path of 'serve-expired-client-timeout' is taken leading to denial of service via the server crash.	5.9	More Details
CVE-2026-52863	In NLnet Labs Unbound 1.25.0 up to and including 1.25.1, a fix that makes the 'respip' and 'dns64' modules work together, creates a shallow copy of the view name in effect that could lead to memory corruption if the owner of the original view name is jostled out when Unbound is under pressure. Unbound needs to be configured with one of 'respip'/'rpz' modules, together with a module that can attach subqueries (respip CNAME redirection, dns64, subnetcache) and a configured 'access-control-view' while Unbound is under pressure so that joslte logic kicks in and starts dropping slow queries. The subquery is getting a shallow copy of the view name and if the super query which owns the view name is jostled out, memory corruption can occur. Likelihood of a crash is low, since it relies heavily on the underlying memory allocator and the memory layout. Debug memory builds (e.g., ASAN) that catch the free terminate the server.	5.9	More Details
CVE-2026-16107	IBM TS4500 CLI tool Versions: 0.1.31 through 1.12.0.0 does not validate or improperly validates TLS certificate validation, which could allow an attacker to obtain sensitive information using man in the middle techniques.	5.9	More Details
CVE-2025-68081	Administrator Cross Site Scripting (XSS) in WP-Polls <= 2.77.3 versions.	5.9	More Details
CVE-2026-24628	Administrator Cross Site Scripting (XSS) in Photo Gallery by Supsystic <= 1.16.3 versions.	5.9	More Details
CVE-2026-44621	With NLnet Labs Unbound up to and including version 1.25.1, applications using libunbound and configured with 'unwanted-reply-threshold', could eventually be abruptly terminated if the threshold is reached and libunbound needs to call 'libworker_alloc_cleanup' since the function is absent from the function call allow list. When an application using libunbound sets 'unwanted-reply-threshold' to any non-zero value and the iterator queries an authoritative that replies with enough wrong-transaction-ID UDP datagrams to cross the threshold, the 'libworker_alloc_cleanup' will eventually be called. Since the function is absent from the function call allow list, this leads to a fatal exit of libunbound and eventual termination of the embedding application.Unbound itself is not affected since its relevant function 'worker_alloc_cleanup' is registed in the allow list and proceeds to perform the documented cache flush.	5.9	More Details
CVE-2026-55991	In NLnet Labs Unbound 1.22.0 up to and including 1.25.1, a remote unauthenticated client can trigger a libngtcp2 assertion (if compiled with assertions on) and terminate the entire Unbound process using a single DNS-over-QUIC (DoQ) connection and one normal DNS query. This is caused by an erroneous error value passed to libngtcp2. When 'ngtcp2_conn_writev_stream()' returns 'NGTCP2_ERR_STREAM_DATA_BLOCKED', Unbound continues to call 'ngtcp2_ccerr_set_application_error()' with a '-1' error value. The 'int' literal '-1' is implicitly converted to the function's 'uint64_t error_code' parameter as '0xFFFFFFFFFFFFFFFF'. The follow-on 'ngtcp2_conn_write_connection_close()' serialises that value as a QUIC variable-length integer; because '2^64-1' exceeds the 62-bit varint ceiling, 'ngtcp2_put_uvarintlen()' fails 'assert(n < 4611686018427387904ULL)' and the whole resolver process aborts. A remote, unauthenticated DoQ client can trigger this deterministically with a single QUIC connection by advertising 'initial_max_stream_data_bidi_local = 1' in its transport parameters and sending one DoQ query without ever reading the stream.	5.9	More Details
CVE-2026-14586	In NLnet Labs Unbound 1.22.0 up to and including 1.25.1, in DNS-over-QUIC environments, with high concurrency and under pressure, an assertion in libngtcp2 about monotonic timestamps could trigger and result in server termination and thus denial of service. When interfacing with libngtcp2, for DNS-over-QUIC support in Unbound, it is expected to use monotonic time. Unbound was using realtime instead, and in DoQ environments with high concurrency and under pressure, an assert in libngtcp2 for the quic timestamp would trigger and terminate the server.This vulnerability needs Unbound to be compiled with DoQ support ('--with-libngtcp2') and the 'quic-port' to be configured for the listening interfaces.	5.9	More Details
CVE-2026-66754	Rouille 0.1.6 through 3.6.2 contains a reachable assertion vulnerability in the Request::remove_prefix function that allows remote unauthenticated attackers to crash the server by sending a crafted percent-encoded URL. Attackers can send a request whose decoded path matches a configured prefix while the raw percent-encoded path does not, causing the assert! to fail and triggering either a 500 error or full process termination depending on the panic configuration.	5.9	More Details
CVE-2026-13188	In Progress® Telerik® UI for AJAX prior to v2026.2.708, DialogHandler request parameters may be tampered with, potentially altering dialog server-side behavior and enabling chained exploitation.	5.9	More Details
CVE-2026-65483	Author Cross Site Scripting (XSS) in HashThemes Demo Importer <= 1.4.2 versions.	5.9	More Details
CVE-2026-65534	Author Cross Site Scripting (XSS) in Custom links in Elementor Image Carousel <= 1.1.1 versions.	5.9	More Details
CVE-			

2026-65538	Author Cross Site Scripting (XSS) in Machete <= 5.2 versions.	5.9	More Details
CVE-2026-65550	Shop Manager Cross Site Scripting (XSS) in Tabs <= 2.5 versions.	5.9	More Details
CVE-2026-9680	Improper exposure of the MCP server in alibabacloud-rds-openapi-mcp-server allows remote attackers to invoke exposed MCP tools via network access to an MCP endpoint listening on all network interfaces by default.	5.8	More Details
CVE-2026-59921	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, HttpPostRequestEncoder constructs multipart HTTP request bodies by directly concatenating user-supplied filenames and field names into Content-Disposition MIME headers without validating or sanitizing CRLF characters (\r\n). Since MIME headers are delimited by CRLF, an attacker who controls the filename can inject arbitrary MIME headers into the multipart body part. The root cause is that neither the encoder nor the FileUpload implementations' setFilename() methods, which only check for null, neutralize CRLF characters before the filename is embedded into the header. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.	5.7	More Details
CVE-2026-15003	A flaw was found in the GNU Binutils (Binary Utilities) linker. This vulnerability, a heap-buffer-overflow read (CWE-125), occurs when the linker processes a specially crafted 32-bit XCOFF (Extended Common Object File Format) object file. An attacker could exploit this by providing a malicious file, leading to an out-of-bounds read of memory. This can result in information disclosure, potentially revealing sensitive heap data, and a Denial of Service (DoS) due to the linker crashing.	5.6	More Details
CVE-2026-63317	Arbitrary Class Instantiation via XML Feature Generator Descriptor and Format Name in Apache OpenNLP Versions Affected: - before 2.5.10 - before 3.0.0-M5 Description: Three code paths in Apache OpenNLP load a class by its fully-qualified name via Class.forName() and invoke its no-arg constructor without any prior validation of the class name or its type. The affected paths are: (1) GeneratorFactory, which reads the class attribute of generator elements in an XML feature generator descriptor; such descriptors are embedded as artifacts in model archives (e.g. TokenNameFinder and POSTagger models) and are parsed during model loading, so an attacker who can supply a crafted model archive controls the class name directly. (2) StreamFactoryRegistry.getFactory(Class, String), which falls back to interpreting an unregistered format name as the fully-qualified class name of an ObjectStreamFactory; this is exploitable in applications that pass untrusted format names (e.g. exposing the -format parameter of the command-line tooling to external input). (3) StringInterners, which instantiates the interner implementation named by the opennlp.interner.class system property; this value is normally deployer-controlled, so it is hardened as defense in depth rather than being independently attacker-reachable. Exploitation requires a class with attacker-useful side effects in its static initializer or no-arg constructor (JNDI lookup, outbound network I/O, filesystem access) to be present on the classpath, so this is not drop-in remote code execution. T Mitigation: Upgrade to a fixed release. The fix routes all three paths through ExtensionLoader.instantiateExtension(...), which consults a package-prefix allowlist before Class.forName() is invoked, so a disallowed class is never loaded, initialized, or constructed. Classes under the opennlp. prefix remain permitted by default. Deployments that load models referencing feature generator factories, object stream factories, or string interners outside opennlp.* must opt those packages in, either programmatically via ExtensionLoader.registerAllowedPackage(String) before the first model load, or by setting the OPENNLP_EXT_ALLOWED_PACKAGES system property to a comma-separated list of allowed package prefixes. Users who cannot upgrade immediately should ensure all model files and format names are sourced from trusted origins and should audit their classpath for classes with side-effecting static initializers or constructors.	5.6	More Details
CVE-2026-16743	A flaw was found in accountsservice. The systemd-homed code path for SetIconFile opens a user-supplied filename as root without the validation and privilege drop performed by the classic handler. A local attacker with a systemd-homed-managed account can read arbitrary files accessible to the accounts-daemon process.	5.5	More Details
CVE-2026-28900	A file quarantine bypass was addressed with additional checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. A maliciously crafted ZIP archive may bypass Gatekeeper checks.	5.5	More Details
CVE-2026-7775	IBM Sterling B2B Integrator 6.2.0.0 through 6.2.0.6, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.6, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 is vulnerable to stored cross-site scripting. This vulnerability allows a privileged user to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session.	5.5	More Details
CVE-2026-17534	Kimi Code (@moonshot-ai/kimi-code) before 0.27.0 implements FetchURL SSRF hardening as a static hostname and IP-literal denylist in assertSafeFetchTarget, without resolving DNS or re-validating hosts after HTTP redirects. An attacker who can influence a FetchURL call (for example via prompt injection) can supply a crafted public hostname that resolves to loopback or another internal address, or a public URL that redirects to such a target, and thereby reach internal network services that the denylist was intended to block. FetchURL is included in the default auto-approve tool set, so the call does not require interactive user confirmation in manual mode.	5.5	More Details
CVE-2026-66757	A flaw was found in the file-sgi plugin in GIMP. When processing an RLE-compressed SGI image, the plugin allocates memory for a row table. The image header dimensions (ysize and zsize) are read as 16-bit unsigned integers. If a crafted file sets both dimensions to their maximum value (65535), the multiplication ysize * zsize overflows the standard 32-bit int boundary before being passed to calloc. This integer overflow issue results in undefined behavior, aborting the plugin and causing a denial of service.	5.5	More Details
CVE-	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.2,		

2026-47768	newly-minted operator API key exposed in redirect URL (Referer, history, proxy logs). This issue has been patched in version 0.3.2.	5.5	More Details
CVE-2026-28932	A logic issue existed resulting in memory corruption. This was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause a denial of service.	5.5	More Details
CVE-2026-54422	In OpenStack Ironic Python Agent through 11.5.0, a malicious bootc container, when deployed using ironic-python-agent, may be able to extract the credentials used to download it.	5.5	More Details
CVE-2026-38763	An issue in Unistal Systems Pvt. Ltd.Protegent 360 v2.0.0.4 allows a local attacker to cause a denial of service via the function sub_13828	5.5	More Details
CVE-2026-16910	A flaw was found in Red Hat Quay's notification webhook feature. The Slack and generic webhook notification handlers accept user-supplied URLs without SSRF validation, allowing a repository administrator to make the Quay worker issue POST requests to internal network addresses or cloud infrastructure endpoints that should not be reachable from the application.	5.5	More Details
CVE-2026-42495	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] The directory and Rock Ridge / SUSP walk in libfsimage's iso9660 driver derives several lengths directly from attacker-controlled on-disk fields without validating them: * The directory loop itself assumes a good record length. This is CVE-2026-42494. * The calculation of the System Use area may underflow. This is CVE-2026-42495. * The Rock Ridge extension loop assumes a good (inner) record length. This is CVE-2026-62423. * The Rock Ridge NM record processing assumes a good entry length. This is CVE-2026-62424. * The Rock Ridge CE record processing assumes a good size and offset. This is CVE-2026-62425.	5.5	More Details
CVE-2026-43665	This issue was addressed with additional entitlement checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. A local attacker may be able to determine the legacy VNC password configured for Screen Sharing.	5.5	More Details
CVE-2026-62423	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] The directory and Rock Ridge / SUSP walk in libfsimage's iso9660 driver derives several lengths directly from attacker-controlled on-disk fields without validating them: * The directory loop itself assumes a good record length. This is CVE-2026-42494. * The calculation of the System Use area may underflow. This is CVE-2026-42495. * The Rock Ridge extension loop assumes a good (inner) record length. This is CVE-2026-62423. * The Rock Ridge NM record processing assumes a good entry length. This is CVE-2026-62424. * The Rock Ridge CE record processing assumes a good size and offset. This is CVE-2026-62425.	5.5	More Details
CVE-2026-28849	The issue was addressed with improved checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. A maliciously crafted ZIP archive may bypass Gatekeeper checks.	5.5	More Details
CVE-2026-20672	An information disclosure issue was addressed with improved privacy controls. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-62424	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] The directory and Rock Ridge / SUSP walk in libfsimage's iso9660 driver derives several lengths directly from attacker-controlled on-disk fields without validating them: * The directory loop itself assumes a good record length. This is CVE-2026-42494. * The calculation of the System Use area may underflow. This is CVE-2026-42495. * The Rock Ridge extension loop assumes a good (inner) record length. This is CVE-2026-62423. * The Rock Ridge NM record processing assumes a good entry length. This is CVE-2026-62424. * The Rock Ridge CE record processing assumes a good size and offset. This is CVE-2026-62425.	5.5	More Details
CVE-2026-62425	[This CNA information record relates to multiple CVEs; the text explains which aspects/vulnerabilities correspond to which CVE.] The directory and Rock Ridge / SUSP walk in libfsimage's iso9660 driver derives several lengths directly from attacker-controlled on-disk fields without validating them: * The directory loop itself assumes a good record length. This is CVE-2026-42494. * The calculation of the System Use area may underflow. This is CVE-2026-42495. * The Rock Ridge extension loop assumes a good (inner) record length. This is CVE-2026-62423. * The Rock Ridge NM record processing assumes a good entry length. This is CVE-2026-62424. * The Rock Ridge CE record processing assumes a good size and offset. This is CVE-2026-62425.	5.5	More Details
CVE-2026-7521	Mattermost versions 11.8.x <= 11.8.0, 11.7.x <= 11.7.3, 11.6.x <= 11.6.5, 10.11.x <= 10.11.20 fail to verify file deletion path which allows an admin with SAML system-console write permissions to delete arbitrary files outside the config directory from the server via the remove file endpoint.. Mattermost Advisory ID: MMSA-2026-00666	5.5	More Details
CVE-2026-16730	A flaw was found in dbus-broker. When the process file-descriptor limit is reached, EMFILE/ENFILE errors during peer setup (notably SO_PEERPIDFD) are handled as fatal failures, causing the broker to exit. A local attacker who can open many connections to the user session bus can trigger this and deny service to the desktop session. Flatpak applications can reach the host session bus through the dbus proxy.	5.5	More Details
CVE-2026-17048	A flaw was found in the Keycloak Admin REST API, which is used to manage security realms and clients. The issue occurs when the system processes requests for rotated client secrets that are stored in a secure vault. Due to improper boundary enforcement, a delegated administrator with view-only permissions can retrieve the actual	5.5	More Details

	resolved secret instead of the vault placeholder, leading to the exposure of sensitive credentials.		
CVE-2026-64776	The issue was addressed with improved bounds checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to disclose kernel memory.	5.5	More Details
CVE-2026-43760	An access issue was addressed with improved access restrictions. This issue is fixed in macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access user-sensitive data.	5.5	More Details
CVE-2026-64755	An authorization issue was addressed with improved state management. This issue is fixed in iOS 26.6 and iPadOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43800	An information disclosure issue was addressed by removing the vulnerable code. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-64693	A type confusion issue was addressed with improved checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted image may lead to a denial-of-service.	5.5	More Details
CVE-2026-43819	An access issue was addressed with additional sandbox restrictions. This issue is fixed in macOS Tahoe 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43754	This issue was addressed with improved redaction of sensitive information. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to leak sensitive kernel state.	5.5	More Details
CVE-2026-43817	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	5.5	More Details
CVE-2026-43816	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	5.5	More Details
CVE-2026-43758	An authorization issue was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43806	A denial of service issue was addressed by removing the vulnerable code. This issue is fixed in macOS Tahoe 26.6. A local attacker may be able to cause a denial of service.	5.5	More Details
CVE-2026-43801	This issue was addressed with improved checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43759	An authorization issue was addressed with improved state management. This issue is fixed in macOS Tahoe 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-64707	A permissions issue was addressed with improved validation. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, visionOS 26.6. An app may be able to delete files for which it does not have permission.	5.5	More Details
CVE-2026-43797	This issue was addressed with improved checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6. An app may be able to access information about a user's contacts.	5.5	More Details
CVE-2026-43796	This issue was addressed with improved data protection. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43782	This issue was addressed with improved checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43775	An authorization issue was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Tahoe 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43774	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access sensitive user data.	5.5	More Details

CVE-2026-43768	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	5.5	More Details
CVE-2026-16492	A weakness has been identified in umijs umi up to 4.6.63. The affected element is the function git.getFileCreateInfo of the file packages/utils/src/getFileGitIno.ts of the component GIT File Helper. This manipulation causes os command injection. The exploit has been made available to the public and could be used for attacks. Upgrading to version 4.6.64 is sufficient to fix this issue. Patch name: b6da12c17b024a43badb1fa565720c38cf42e647. Upgrading the affected component is advised.	5.5	More Details
CVE-2026-43765	This issue was addressed with improved handling of symlinks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to modify protected parts of the file system.	5.5	More Details
CVE-2026-64699	A memory initialization issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to disclose kernel memory.	5.5	More Details
CVE-2026-43756	A logic issue was addressed with improved validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access user-sensitive data.	5.5	More Details
CVE-2026-43763	A permissions issue was addressed by removing the vulnerable code. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to read files outside of its sandbox.	5.5	More Details
CVE-2026-64718	A use-after-free issue was addressed with improved memory management. This issue is fixed in Safari 26.6, iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing maliciously crafted web content may lead to an unexpected Safari crash.	5.5	More Details
CVE-2026-64734	The issue was addressed with improved checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted contact may leak sensitive data.	5.5	More Details
CVE-2026-64744	An information leakage was addressed with additional validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to disclose kernel memory.	5.5	More Details
CVE-2026-64724	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An attacker on the local network may be able to cause a denial-of-service.	5.5	More Details
CVE-2026-64723	A logic issue was addressed with improved checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-43744	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing an audio stream in a maliciously crafted media file may terminate the process.	5.5	More Details
CVE-2026-64754	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. Processing a maliciously crafted file may lead to a denial-of-service.	5.5	More Details
CVE-2026-43714	The issue was addressed with improved input sanitization. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, visionOS 26.6, watchOS 26.6. A malicious app may be able to access protected user data.	5.5	More Details
CVE-2026-64722	A buffer overflow issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Tahoe 26.6. Processing a 3D model may result in disclosure of process memory.	5.5	More Details
CVE-2026-64721	This issue was addressed through improved state management. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to access sensitive user data.	5.5	More Details
CVE-2026-64711	This issue was addressed with additional entitlement checks. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to leak sensitive user information.	5.5	More Details
CVE-2026-64710	A privacy issue was addressed by removing sensitive data. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to leak sensitive user information.	5.5	More Details
CVE-2026-64709	The issue was addressed with improved memory handling. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to disclose kernel memory.	5.5	More Details

CVE-2026-43738	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8. Processing a maliciously crafted asset catalog may result in disclosure of process memory.	5.5	More Details
CVE-2026-64708	A file quarantine bypass was addressed with additional checks. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may bypass Gatekeeper checks.	5.5	More Details
CVE-2026-43739	An out-of-bounds write issue was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Tahoe 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to cause unexpected system termination.	5.5	More Details
CVE-2026-64741	A permissions issue was addressed with additional restrictions. This issue is fixed in iOS 26.6 and iPadOS 26.6, tvOS 26.6, visionOS 26.6, watchOS 26.6. An app may be able to read a persistent device identifier.	5.5	More Details
CVE-2026-65593	n8n versions before 1.123.64, 2.29.8, and 2.30.1 contain a server-side request forgery vulnerability in the dynamic-node-parameters endpoints that lack authorization scopes. Authenticated attackers can supply absolute URLs in routing configuration to override baseURL restrictions and make the n8n server issue HTTP requests to arbitrary internal targets when SSRF protection is disabled.	5.4	More Details
CVE-2026-48530	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the Classification Rules configuration that allows authenticated attackers to inject arbitrary web script or HTML via the rule name and email criteria parameters to /Archiver/CategorizationPolicyWizard.aspx. The injected payload is stored by CategorizationPolicyWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the Classification Rules page.	5.4	More Details
CVE-2026-65478	Subscriber Broken Access Control in ListingPro <= 2.9.10 versions.	5.4	More Details
CVE-2026-65597	n8n before 1.123.64, 2.x before 2.29.8, and before 2.30.1 contains a DOM-based cross-site scripting vulnerability in the HTML preview, which renders execution output into an iframe srcdoc without the sandbox attribute. A sanitizer bypass allows injected script to execute same-origin as the editor. When a victim opens the preview, the script can call authenticated APIs using the victim's session. An account with global:member privileges can exploit the issue.	5.4	More Details
CVE-2026-48531	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the Retention Policy configuration that allows authenticated attackers to inject arbitrary web script or HTML via the policy name parameter to /Archiver/RetentionPolicyWizard.aspx. The injected payload is stored by RetentionPolicyWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the Retention and Spam Policies page.	5.4	More Details
CVE-2026-48052	Papra is a minimalistic document management and archiving platform. Prior to version 26.5.0, an authenticated user who is a member of any organization can delete or rename tags belonging to a different organization, given the target tag's ID. The route handler verifies the caller's membership of the ":organizationId" in the URL, but the repository write filters on tag.id alone, so the URL-level org scope never reaches the database. This issue has been patched in version 26.5.0.	5.4	More Details
CVE-2026-65512	Unauthenticated Cross Site Request Forgery (CSRF) in WP Activity Log <= 5.6.4 versions.	5.4	More Details
CVE-2026-48532	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the File History Retention Policy configuration that allows authenticated attackers to inject arbitrary web script or HTML via the policy name parameter to /Archiver/FAARetentionPolicyWizard.aspx. The injected payload is stored by RetentionPolicyWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the File History Retention Policies page.	5.4	More Details
CVE-2026-65479	Subscriber Broken Access Control in Reviewer <= 3.14.2 versions.	5.4	More Details
CVE-2025-50325	BandiZip v.7.37 is affected by a Authentication Bypass Vulnerability. This vulnerability allows remote attackers to bypass the Mark-of-the-Web protection mechanism on affected installations of BandiZip	5.4	More Details
CVE-2026-65464	Unauthenticated Cross Site Request Forgery (CSRF) in GiveWP <= 4.16.3 versions.	5.4	More Details
CVE-2026-65463	Subscriber Insecure Direct Object References (IDOR) in Masteriyo - LMS <= 2.3.1 versions.	5.4	More Details
CVE-2026-57978	Origin validation error in Microsoft Edge (Chromium-based) allows an unauthorized attacker to perform spoofing over a network.	5.4	More Details

CVE-2026-27391	Subscriber Broken Access Control in uListing <= 2.2.0 versions.	5.4	More Details
CVE-2026-25427	Subscriber Broken Access Control in eRoom <= 1.7.1 versions.	5.4	More Details
CVE-2026-64871	Joomla Extension - regularlabs.com - Inconsistent CSRF token checks / privilege checks in Cache Cleaner extension - Administrator URL purges did not consistently require a valid token and cache-management permission.	5.4	More Details
CVE-2026-57511	SuperPlane before 0.30.0 contains an SMTP header injection vulnerability that allows unauthenticated attackers to inject arbitrary SMTP headers by including CRLF sequences in the event payload title field delivered via webhook. Attackers can manipulate the unsanitized title field passed to the SMTP DATA command to add Bcc recipients for content exfiltration, forge the From address to bypass SPF and DKIM checks, or inject Content-Type and MIME boundary headers to corrupt message bodies for phishing.	5.4	More Details
CVE-2026-66442	Subscriber Broken Access Control in YayPricing <= 3.5.6 versions.	5.4	More Details
CVE-2026-65592	n8n before 1.123.64, 2.29.8, and 2.30.1 contains a stored DOM cross-site scripting vulnerability in the Resource Locator component, which passes the workflow-persisted cachedResultUrl parameter to window.open() without scheme validation. An attacker with workflow creation/editing privileges can craft a workflow with a malicious (e.g., javascript:) scheme in cachedResultUrl; when a victim opens the crafted workflow and interacts with external links, the payload executes in the victim's browser.	5.4	More Details
CVE-2026-64795	Joomla Extension - regularlabs.com - XSS vectors in tag-provided inputs in various Regular Labs extensions - Tag-provided custom HTML, module content/title overrides and decoded modal or tooltip values could execute unsafe markup. A content author could inject JavaScript that ran in visitors’ browsers.	5.4	More Details
CVE-2026-61981	Unauthenticated Cross Site Request Forgery (CSRF) in Simple Link Directory Pro <= 15.0.8 versions.	5.4	More Details
CVE-2026-48534	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the IMAP Server configuration that allows authenticated attackers to inject arbitrary web script or HTML via the server URL parameter to /Archiver/ImapServerWizard.aspx. The injected payload is stored by ImapServerWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the IMAP Server configuration page.	5.4	More Details
CVE-2026-48535	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the Call Home proxy server configuration that allows authenticated attackers to inject arbitrary web script or HTML via the proxy server address parameter to /Archiver/CallHomeSettingsWizard.aspx. The injected payload is stored by CallHomeSettingsWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the General Settings Additional Settings page.	5.4	More Details
CVE-2026-66029	Ekushey Project Manager CRM through version 5.0 contains a stored cross-site scripting vulnerability that allows authenticated client users to inject arbitrary HTML and JavaScript by entering malicious payloads into the client Name field on the Edit Profile page without sanitization. Attackers can craft and store malicious scripts that execute in the browser sessions of Staff or Administrator users who view the Manage Clients or Manage Client Projects pages where client names are rendered unsanitized.	5.4	More Details
CVE-2026-66338	A flaw was found in libsoup. The chunked transfer encoding parser uses a permissive parsing function for chunk sizes that silently accepts inputs violating RFC 9112, including leading whitespace, plus sign prefixes, and trailing invalid characters. When libsoup operates behind a strict frontend proxy, this parsing differential can be exploited to smuggle HTTP requests.	5.4	More Details
CVE-2026-57531	Milkdown before 7.21.3 contains a DOM cross-site scripting vulnerability in the @milkdown/plugin-emoji package that allows unauthenticated attackers to execute arbitrary JavaScript in the host application's origin by causing a victim to paste attacker-controlled content. The parseDOM.getAttrs handler stores raw innerHTML of pasted span elements with data-type="emoji" without sanitization, and the toMarkdown runner subsequently assigns this unsanitized value directly to a live DOM element's innerHTML, bypassing the DOMPurify sanitization used in the toDOM path, causing payload execution on every markdown serialization cycle.	5.4	More Details
CVE-2026-57530	Milkdown before 7.21.3 contains a stored cross-site scripting vulnerability in the @milkdown/preset-commonmark and @milkdown/components packages that allows attackers with document write access to execute arbitrary JavaScript in the browser context of any user who opens the document or clicks a rendered link. The parseMarkdown runner stores raw URL values from the remark AST as href mark attributes without URL scheme validation, and the ineffective DOMPurify.sanitize call in edit-view.ts treats the bare URL string as a text node and returns it unchanged, allowing javascript: payloads to pass through the link-tooltip preview component and read-only mode anchor elements unmodified.	5.4	More Details
CVE-2026-65558	Unauthenticated Server Side Request Forgery (SSRF) in AffiliateX <= 2.3.5 versions.	5.4	More Details

CVE-2026-12689	The ProfileGrid WordPress plugin before 5.9.9.7 does not perform any authorization or ownership check on some of its private-message thread actions, allowing authenticated users with Subscriber-level access and above to soft-delete, tamper with the metadata of, and mark as read other users' private message threads.	5.4	More Details
CVE-2026-48536	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the General Settings SMTP configuration that allows authenticated attackers to inject arbitrary web script or HTML via the SMTP server address parameter to /Archiver/GeneralSettingsWizard.aspx. The injected payload is stored by GeneralSettingsWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the General Settings page.	5.4	More Details
CVE-2026-66746	Rouille 0.4.0 through 3.6.2 contains an HTTP response splitting vulnerability that allows remote attackers to inject arbitrary response headers by embedding carriage return (0x0D) or line feed (0x0A) bytes into attacker-controlled input. Attackers can exploit percent-decoded query parameters reflected into response headers or inject bare LF characters into Cookie header values that are interpolated directly into Set-Cookie response headers, enabling cache poisoning, session fixation, and security header override attacks such as bypassing CSP or CORS policies.	5.4	More Details
CVE-2026-66751	Let's Chat 0.3.0 through 0.4.8 contains an improper authorization vulnerability that allows any authenticated user to archive any room on the server by sending a DELETE request to the rooms handler without ownership verification. Attackers can enumerate room IDs via the rooms listing endpoint and permanently archive private or password-protected rooms they cannot access, with no application-level recovery path requiring direct database intervention to restore.	5.4	More Details
CVE-2026-66031	Ekushey Project Manager CRM through version 5.0 contains a stored cross-site scripting vulnerability that allows authenticated client users to inject arbitrary HTML and JavaScript by entering malicious payloads into the Reply Ticket field. Attackers can craft and store malicious scripts that execute in the browser sessions of Staff or Administrator users who view the Support Ticket detail page.	5.4	More Details
CVE-2026-66752	tiny-http through 0.12.0 contains an HTTP request smuggling vulnerability that allows remote attackers to desynchronize request framing by sending a Transfer-Encoding header with any value, including non-chunked codings, which causes the library to unconditionally apply chunk-decoding and discard Content-Length. Attackers can exploit the discrepancy between tiny_http's improper Transfer-Encoding parsing and a correctly-implemented front-end proxy to produce two distinct interpretations of a single byte stream, enabling request smuggling, and can additionally send non-chunked bodies with non-chunked Transfer-Encoding values to cause failed body reads that tie up connections and consume worker threads without signaling errors to clients.	5.4	More Details
CVE-2026-67181	Rouille 0.3.3 through 3.6.2 contains an HTTP request smuggling vulnerability that allows remote attackers to desynchronize HTTP message boundaries by exploiting improper header forwarding in the proxy implementation. The proxy in src/proxy.rs forwards the client's Transfer-Encoding header to upstream backends unchanged while transmitting a body already de-chunked by tiny_http, enabling CL.TE desynchronization attacks where attackers control where the backend believes the request body ends.	5.4	More Details
CVE-2026-66030	Ekushey Project Manager CRM through version 5.0 ccontains a stored cross-site scripting vulnerability that allows authenticated client users to inject arbitrary HTML and JavaScript by entering malicious payloads into the Ticket Title field on the Create New Ticket page. Attackers can craft and store malicious scripts that execute in the browser sessions of Staff or Administrator users who view the Client Support page where ticket titles are rendered unsanitized.	5.4	More Details
CVE-2026-62828	Improper input validation in Microsoft Edge for Android allows an unauthorized attacker to perform tampering over a network.	5.4	More Details
CVE-2026-48539	GFI Archiver before 15.13 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 /Archiver/MailInsights.aspx. The injected payload is stored by ReportScheduling.btnSaveReport_Click() 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-2026-48538	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the default import settings configuration that allows authenticated attackers to inject arbitrary web script or HTML via the configured folders parameter to /Archiver/ImportSettingsWizard.ashx. The injected payload is stored by ImportSettingsWizard.SaveAllConfigSettings() without output encoding and is executed in the browsers of users who subsequently view the Archive Assistant default import settings.	5.4	More Details
CVE-2026-48537	GFI Archiver before 15.13 contains a stored cross-site scripting vulnerability in the File Archive Assistant configuration that allows authenticated attackers to inject arbitrary web script or HTML via the excluded extensions parameter to /Archiver/FileArchiveAssistantWizard.aspx. The injected payload is stored by FileArchiveAssistantWizard.btnSave_Click() without output encoding and is executed in the browsers of users who subsequently view the File Archive Assistant settings page.	5.4	More Details
CVE-2026-65696	Overseerr through 1.35.0 contains an authorization bypass through user-controlled key vulnerability in the push subscription API that allows authenticated users to list, read, and delete any other user's push subscriptions by supplying an arbitrary userId in the path parameters. Attackers can exploit the missing ownership check in the affected handlers to access target user records without the filteredFields filter, leaking sensitive data including email addresses and plexId values.	5.4	More Details
CVE-2026-	A vulnerability was detected in oclif up to 4.23.16. Affected by this vulnerability is the function child_process.exec of the component JIT Plugin Entry Handler. Performing a manipulation of the argument jitPlugins results in os command	5.3	More

16628	injection. The attack is only possible with local access. The exploit is now public and may be used. The patch is named 939b045725e065baebc4587b8bccfd56731eed3d. To fix this issue, it is recommended to deploy a patch.		Details
CVE-2026-3482	IBM Sterling B2B Integrator and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 could allow an unauthenticated user to read sensitive information by bypassing authentication through a specially crafted HTTP request.	5.3	More Details
CVE-2026-13057	An issue in the server's Atlas Search integration allows an authenticated user to bypass per-user access controls. In sharded topologies, the $search and $searchMeta aggregation stages use internal routing that is normally populated only by the trusted router during sharded search planning. Due to insufficient input validation, an authenticated client can supply these fields directly.	5.3	More Details
CVE-2026-13070	A MongoDB server initiating an outbound TLS connection may terminate abnormally when processing a malformed OCSP response from a remote peer during the TLS handshake. OCSP stapling validation is enabled by default for outgoing TLS connections. Affected scenarios require the remote peer to hold a certificate issued by the cluster's trusted certificate authority, or for the connection to traverse an untrusted network path.	5.3	More Details
CVE-2026-13074	An unauthenticated remote client can cause excessive CPU consumption on a MongoDB server by sending a specific combination of parameters to the awaitable hello command in exhaust mode. The server's handling of this combination results in a response loop that bypasses normal throttling, allowing a small number of connections to degrade server availability.	5.3	More Details
CVE-2026-65564	Unauthenticated Sensitive Data Exposure in MapPress Maps for WordPress <= 2.97.6 versions.	5.3	More Details
CVE-2026-65567	Unauthenticated Broken Access Control in Event Tickets <= 5.29.0.1 versions.	5.3	More Details
CVE-2026-62434	A guest started with Populated on Demand enabled (PoD) can attempt to reclaim pages which aren't regular guest RAM. This can cause corruption of memory management state in Xen.	5.3	More Details
CVE-2026-7484	External control of Assumed-Immutable web parameter vulnerability in ABIS Technology Ltd. Co. AVESİS allows Accessing Functionality Not Properly Constrained by ACLs. This issue affects AVESİS: before 202606251646.	5.3	More Details
CVE-2026-65525	Unauthenticated Broken Access Control in Civi Framework <= 2.2.0 versions.	5.3	More Details
CVE-2026-16629	A vulnerability was identified in danger danger-js up to 13.0.7. Impacted is the function danger.git.diffForFile of the file source/platforms/git/localGetFileAtSHA.ts of the component CLI. Such manipulation of the argument File leads to os command injection. The attack needs to be performed locally. Upgrading to version 13.0.8 is recommended to address this issue. The name of the patch is 087a7290264cc6fb7154ea8c2552a7b2cb8b33a3. It is advisable to upgrade the affected component.	5.3	More Details
CVE-2026-66438	Unauthenticated Sensitive Data Exposure in Exclusive Addons Elementor <= 2.8.0 versions.	5.3	More Details
CVE-2026-46452	Improper Input Validation vulnerability in Apache NimBLE in Mesh Proxy SAR reassembly could result in passing broken data toward application resulting in memory pressure and unstable parsing behavior. This issue affects Apache NimBLE: through 1.9.0. Users are recommended to upgrade to version 1.10.0, which fixes the issue.	5.3	More Details
CVE-2026-16630	A security vulnerability has been detected in syncfusion ej2-javascript-ui-controls up to 33.2.3. This affects the function child_process.exec of the file package.json. The manipulation leads to os command injection. An attack has to be approached locally. The exploit has been disclosed publicly and may be used.	5.3	More Details
CVE-2026-12654	The Payment Plugins for Stripe WooCommerce plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 4.0.7. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for unauthenticated attackers to mark arbitrary pending asynchronous WooCommerce orders as paid by forging a charge.pending event with attacker-controlled metadata.order_id, metadata.gateway_id, and a charge object carrying status=succeeded and captured=true, triggering payment_complete() and downstream fulfillment flows with an attacker-supplied transaction ID. Exploitation requires the merchant to have left the webhook_secret_test or webhook_secret_live option blank, which is the plugin's default state until a Stripe-issued whsec_ value is manually configured; once a non-empty secret is set, the signature verification cannot be bypassed.	5.3	More Details
CVE-2026-16631	A vulnerability was detected in publint up to 0.1.4. This impacts the function child_process.exec of the file src/node/pack.js of the component package-manager Command Handler. The manipulation results in os command injection. Attacking locally is a requirement. The exploit is now public and may be used. The patch is identified as adf2d9a09945fc98c85a2520a89f441d78b2dbd8. It is advisable to implement a patch to correct this issue. The project maintainer explains: "I think it's very rare for someone to use this package with untrusted input".	5.3	More Details
CVE-	A security flaw has been discovered in boazsegev facil.io up to 0.7.58. This affects the function http_sendfile2 of the		

2026-16653	file lib/facil/http/http.c of the component Public Folder Handler. Performing a manipulation results in path traversal. Remote exploitation of the attack is possible. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-21723	The alertmanager templates test endpoint (/api/alertmanager/grafana/config/api/v1/templates/test) can execute templates with no memory limits. Mass-executing templates in a short period causes OOM and crashes the Grafana service. The endpoint requires very low privileges and is exploitable with anonymous access enabled.	5.3	More Details
CVE-2026-17514	A vulnerability was determined in ZJONSSON node-unzipper up to 0.12.3. Affected by this vulnerability is the function Extract of the file lib/extract.js. This manipulation causes path traversal. The attack requires local access. The exploit has been publicly disclosed and may be utilized. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-66004	BlenderMCP before commit 30a3308 contains a path traversal vulnerability in the download_polyhaven_asset method that allows attackers to write arbitrary files by injecting traversal sequences in API response include keys. Attackers performing MITM attacks or prompt injection can supply malicious paths like '../../.bashrc' to overwrite sensitive files and achieve persistent code execution.	5.3	More Details
CVE-2026-65012	InvokeAI before 6.13.7 contains an unauthenticated directory enumeration vulnerability in the GET /api/v2/models/scan_folder endpoint that accepts attacker-controlled scan_path parameters. Unauthenticated attackers can recursively enumerate arbitrary server filesystem directories and use HTTP response codes to determine file existence and readability, bypassing multi-user mode access controls.	5.3	More Details
CVE-2026-14865	In Progress® Telerik® UI for AJAX prior to v2026.2.708, the internal LayoutBuilder control processes client-state XML without disabling DTD processing, allowing unauthenticated denial of service via recursive XML entity expansion.	5.3	More Details
CVE-2026-16489	A vulnerability was identified in jsforce up to 3.10.16. This issue affects the function _execCommand in the library lib/registry/sfdx.js of the component SFDX Connection Registry. The manipulation leads to os command injection. The attack can only be performed from a local environment. 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-2026-14322	The Timetics WordPress plugin before 1.0.57 does not enforce a pending or unpaid status for new bookings created through a payment method other than its recognised gateways, allowing unauthenticated users to create fully-approved bookings for priced appointments without making any payment.	5.3	More Details
CVE-2026-61392	There is a information disclosure vulnerability in some Hikvision cameras, allowing unauthenticated attackers to obtain partial information from the device’s memory.	5.3	More Details
CVE-2026-65014	n8n before 2.28.0 (and before 2.27.4 on the 2.27.x branch) registers the DELETE /${restEndpoint}/test-webhook/:id endpoint before authentication middleware is applied, allowing any unauthenticated network caller who knows a workflow ID to cancel that workflow's active test webhook registration. The impact is limited to disrupting in-progress test sessions; production webhooks, persistent workflow state, and stored data are not affected.	5.3	More Details
CVE-2026-12124	The PDFDraft – Drag & Drop PDF Builder, PDF Viewer, Embed & Download PDF, Certificate & Invoice Designer plugin for WordPress is vulnerable to unauthorized access of data due to a missing capability check on the serveTemplatePdfAjax() function and the serveTemplatePdf() REST route (which is registered with `permission_callback => '__return_true'`) in versions up to, and including, 1.1.0. This makes it possible for unauthenticated attackers to download stored template PDFs — which may contain customer PII, invoice, order, and certificate data — by requesting the publicly registered admin-ajax action `pdfdraft_embed_pdf` or the REST endpoint `/wp-json/pdfdraft/v1/embed-pdf/templates/{slug}/pdf` with a known or guessable design slug, bypassing the plugin's own .	5.3	More Details
CVE-2026-17500	A vulnerability was detected in ggml-org llama.cpp d006858/e15efe0. This affects the function _visit_pattern of the file common/json-schema-to-grammar.cpp. The manipulation results in null pointer dereference. The attack can be launched remotely. The pull request to fix this issue awaits acceptance.	5.3	More Details
CVE-2026-17501	A flaw has been found in ggml-org llama.cpp e15efe0. This vulnerability affects the function transform of the file common/json-schema-to-grammar.cpp of the component JSON-Schema-to-GBNF Conversion. This manipulation causes allocation of resources. The attack may be initiated remotely. The pull request to fix this issue awaits acceptance.	5.3	More Details
CVE-2026-15012	The Demi – One Click Demo Import, WP Backup & Site Migration plugin for WordPress is vulnerable to Arbitrary Directory Copy in all versions up to, and including, 0.0.8 via the handle_restore_step function. This is due to missing HTTP access controls on the wp-content/uploads/demi-backup-state/ directory, which exposes the cryptographic restore key used to both authenticate the unauthenticated AJAX handler and forge signed restore-state envelopes. This makes it possible for unauthenticated attackers to copy arbitrary files to attacker-controlled destinations on the server. An active restore operation must have been initiated, which writes the .restore_key and .restore_step_token files to the public upload directory, before the exposed secrets can be harvested and chained to achieve unauthenticated arbitrary file copy.	5.3	More Details
CVE-2026-16560	A heap-buffer-overflow flaw was found in Directory Server (389-ds-base). When a DN contains a legacy-quoted value, the server won't close the heap allocation allowing another call to refer to the same memory pointer causing a denial of service or an arbitrary memory write operation.	5.3	More Details
	lakeFS through 1.83.0, fixed in commit 71a45ee, contains an authentication bypass vulnerability in the		

CVE-2026-66006	/setup_comm_prefs endpoint that allows unauthenticated attackers to overwrite operator metadata including email, name, and company after setup completion. Attackers can POST to this endpoint to modify security update preferences, disable security communications, and trigger falsified telemetry events using the legitimate installation ID.	5.3	More Details
CVE-2026-50045	In NLnet Labs Unbound 1.22.0 up to and including 1.25.1, a single client query for a deeply nested name under a DNSSEC-signed parent can cause Unbound to send more upstream packets per client query than the configured 'max-global-quota'. This effectively bypasses a security configuration that limits upstream amplification traffic.	5.3	More Details
CVE-2026-13110	The Storegrowth Sales Booster plugin for WordPress is vulnerable to Missing Authorization in versions up to and including 2.1.0. This is due to a missing capability check on the bogo_category_msg_create() AJAX handler, which is registered for both authenticated (wp_ajax_) and unauthenticated (wp_ajax_nopriv_) users and only validates a nonce ('ajd_protected') that is emitted publicly via wp_localize_script() on every frontend page through front_scripts() . This makes it possible for unauthenticated attackers to modify the plugin's BOGO category-message configuration stored in the spsg_bogo_general_settings option by reading the nonce from any public page and POSTing attacker-controlled data to admin-ajax.	5.3	More Details
CVE-2026-49447	Cosmos provides users the ability self-host a home server by acting as a secure gateway to your application, as well as a server manager. In 0.22.18, `GET /cosmos/api/constellation/public-devices` discloses Constellation device metadata to a requester that supplies any non-empty Authorization header. The handler strips the string Bearer from the header but never validates the resulting token and never uses it in the database query. This vulnerability is fixed in 0.22.19.	5.3	More Details
CVE-2026-50251	In NLnet Labs Unbound up to and including version 1.25.1, when 'unwanted-reply-threshold' is enabled (set to any value greater than zero), glue records of 0.0.0.0/::0 can short-circuit Unbound, on systems that can direct such traffic, by issuing DNS queries and receiving seemingly unwanted replies since the remote IP does not match the original source IP of 0.0.0.0/::0. This behavior keeps on looping for the glue records and pushing the counter to the configured 'unwanted-reply-threshold' that triggers a defensive cache clear. A malicious actor who controls a delegation that returns in-bailiwick glue of 0.0.0.0/::0 can drive the counter to the limit of 'unwanted-reply-threshold' to the threshold and trigger a cache clean of the message and rrset caches; at will, indefinitely, without sending a single spoofed packet. The iterator uses the 0.0.0.0/::0 glue, and a system that can route this (e.g., Linux kernel routes the datagram over loopback), Unbound's own listener answers from 127.0.0.1. Because of the mismatch of 0.0.0.0 and 127.0.0.1, in this example, Unbound accounts the reply as an unwanted (probably spoofed) answer. The counter resets to zero on every cache flush, so the attack loops forever.	5.3	More Details
CVE-2026-15411	The StoreGrowth: Smart Sales Booster for WooCommerce | BOGO, Upsells, Direct Checkout, Quick View, Side Cart plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.1.0. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for unauthenticated attackers to overwrite the spsg_popup_products option with arbitrary attacker-controlled data. The 'ajd_protected' nonce used as the sole gate is exposed to unauthenticated visitors on every frontend page through the BoGo module's wp_localize_script call, rendering it ineffective as an authorization barrier.	5.3	More Details
CVE-2026-16581	In igloohome Smart Lock Mobile App versions 3.2.3 and prior, an Inclusion of Sensitive Information in Source Code vulnerability could allow an unauthorized actor to access functions or backend services that were not sufficiently protected by authentication controls.	5.3	More Details
CVE-2026-16773	The WPBot – AI ChatBot for Live Support, Lead Generation, AI Services plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 8.5.9 via the wpbot_send_email_transcript_free. This makes it possible for unauthenticated attackers to exfiltrate full chat transcripts and associated user PII — including names, email addresses, and phone numbers — stored in the wpbot_user and wpbot_conversation tables to an attacker-controlled email address.	5.3	More Details
CVE-2026-16774	The Chatbot plugin for WordPress is vulnerable to Missing Authorization in versions up to, and including, 8.5.9 via the wpcs_send_email() AJAX handler. This is due to the wpcs_send_email() function being registered on both wp_ajax_wpcs_send_email and wp_ajax_nopriv_wpcs_send_email with no nonce verification, capability check, or rate limiting, while forwarding attacker-controlled recipient, subject, and body directly to wp_mail(). This makes it possible for unauthenticated attackers to send arbitrary emails to any recipient from the site's domain, enabling spam, phishing, and abuse that can lead to the site's IP/domain being blacklisted.	5.3	More Details
CVE-2026-65521	Unauthenticated Sensitive Data Exposure in WP Social Ninja <= 4.3.0 versions.	5.3	More Details
CVE-2026-7120	@fastify/static evaluates the allowedPath callback before normalizing dot segments and duplicate path separators in the pathname used for file resolution. Versions up to and including 10.1.1 are affected. An unauthenticated attacker can bypass allowedPath restrictions by requesting equivalent non-canonical pathnames, causing files that were intended to be denied to be served anyway. The bypass does not allow access outside the configured static root by itself, it defeats path-based filtering only. The issue is patched in @fastify/static 10.1.2.	5.3	More Details
CVE-2026-13464	The Kirki – Freeform Page Builder, Website Builder & Customizer plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 6.0.14 via the 'context' parameter due to missing validation on a user controlled key. This makes it possible for unauthenticated attackers to read the full title, content, and excerpt of any WordPress post — including drafts, pending, privately published, password-protected, and trashed posts — regardless of author, by supplying an arbitrary post ID via the context parameter alongside an attacker-controlled block template.	5.3	More Details

CVE-2026-65487	Unauthenticated Broken Access Control in Photography <= 7.7.6 versions.	5.3	More Details
CVE-2026-65472	Unauthenticated Broken Access Control in Kit (formerly ConvertKit) <= 3.3.5 versions.	5.3	More Details
CVE-2026-65474	Unauthenticated Sensitive Data Exposure in Ninja Tables <= 5.2.10 versions.	5.3	More Details
CVE-2026-65476	Unauthenticated Broken Access Control in Civi <= 2.2.4 versions.	5.3	More Details
CVE-2026-65698	Void through 1.3.4 contains a path traversal vulnerability in the AI agent file-reading tools that allows network-adjacent attackers to read arbitrary host files outside the open workspace by injecting instructions into content the agent processes. Attackers can supply absolute paths or file:// URIs through the read_file, ls_dir, get_dir_tree, and search_* tools, which lack workspace confinement and bypass the approval gate, enabling silent exfiltration of sensitive files such as SSH private keys or cloud credentials via subsequent tool calls.	5.3	More Details
CVE-2026-12353	An unauthenticated attacker could trigger an Out of Memory condition to crash the Java process for RHCS by repeatedly sending HTTP requests to the TLS endpoint. Depending on how the RHCS server is configured, a manual intervention to restart it may prove necessary.	5.3	More Details
CVE-2026-65529	Unauthenticated Broken Access Control in Graphina <= 3.1.12 versions.	5.3	More Details
CVE-2026-65485	Unauthenticated Broken Access Control in Content Control <= 2.6.5 versions.	5.3	More Details
CVE-2026-65486	Unauthenticated Broken Access Control in Event post <= 6.0.1 versions.	5.3	More Details
CVE-2026-65489	Unauthenticated Broken Access Control in LA-Studio Element Kit for Elementor <= 1.6.2 versions.	5.3	More Details
CVE-2026-65468	Unauthenticated Broken Access Control in JetBooking <= 4.1.2 versions.	5.3	More Details
CVE-2026-65490	Unauthenticated Sensitive Data Exposure in Create by Mediavine <= 2.5.3 versions.	5.3	More Details
CVE-2026-65498	Unauthenticated Sensitive Data Exposure in Complianz <= 7.5.0 versions.	5.3	More Details
CVE-2026-65501	Unauthenticated Insecure Direct Object References (IDOR) in Shiptastic for WooCommerce <= 5.1.0 versions.	5.3	More Details
CVE-2026-14820	The Quiz and Survey Master (QSM) WordPress plugin before 11.1.3 does not implement rate limiting or standard failed-login auditing on its front-end credential-check functionality and returns distinct responses for valid and invalid accounts, allowing unauthenticated attackers to enumerate valid usernames and to brute-force passwords while bypassing brute-force protection Quiz and Survey Master (QSM) WordPress plugin before 11.1.3.	5.3	More Details
CVE-2026-65505	Unauthenticated Sensitive Data Exposure in Ultimate Store Kit Elementor Addons <= 3.0.5 versions.	5.3	More Details
CVE-2026-65506	Unauthenticated Broken Access Control in MP3 Audio Player for Music, Radio & Podcast by Sonaar <= 5.12 versions.	5.3	More Details
CVE-2026-47769	APIFold reads an OpenAPI 3.x or Swagger 2.x specification and generates a live, production-ready MCP server endpoint. Prior to commit 7f19b52280f414f57af2b79a95333d1c8fbeece5, the `/webhooks/:serverSlug/:eventName` endpoint accepts arbitrary unauthenticated JSON and stores it in Redis and the `webhook_events` PostgreSQL table without any signature check or authentication requirement. The root cause is that `createWebhookRouter` is called at `server.ts:188` without a `validators` map, so `receivers.ts:80`'s optional-chaining guard evaluates to `undefined` and the signature-validation block (`receiver.ts:81–95`) is unconditionally skipped. Any unauthenticated network	5.3	More Details

	client that knows a valid server slug can inject arbitrary payloads, which are subsequently served as trusted resource state to legitimate MCP clients. Commit 7f19b52280f414f57af2b79a95333d1c8fbeece5 patches the issue.		
CVE-2026-17433	A vulnerability was detected in nanocoai NanoClaw up to 2.0.64. This impacts the function createChatSdkBridge.setup of the file src/channels/chat-sdk-bridge.ts of the component MCP Server Approval. Performing a manipulation results in improper authorization. The attack needs to be approached locally. The exploit is now public and may be used. The project was informed of the problem early through an issue report but has not responded yet.	5.3	More Details
CVE-2026-66477	Unauthenticated Broken Access Control in Gillion <= 4.13 versions.	5.3	More Details
CVE-2026-65469	Unauthenticated Broken Access Control in AWP Classifieds <= 4.4.7 versions.	5.3	More Details
CVE-2026-13390	The Events Calendar WordPress plugin before 6.16.5.1 does not perform an authorization check on one of its Event Aggregator import REST API routes and skips an integrity check for a particular status value, allowing unauthenticated attackers to mark existing import records as failed and to store arbitrary content in a hidden comment record.	5.3	More Details
CVE-2026-27422	Unauthenticated Broken Access Control in YT Player <= 2.0.9 versions.	5.3	More Details
CVE-2026-15827	The GutenKit Blocks plugin for WordPress is vulnerable to unauthorized access of data due to a missing capability check on the /wp-json/gutenkit/v1/mailchimp/get/lists and /wp-json/gutenkit/v1/mailchimp/get/interests REST API endpoints in versions up to, and including, 2.4.12. Both endpoints are registered with permission_callback => '__return_true', and their callbacks read the site's stored Mailchimp API key from the gutenkit_settings_list option and proxy Mailchimp audience/list, merge-field, interest-category, interest-name, and subscriber-count metadata back to the caller with no login, nonce, or capability check. This makes it possible for unauthenticated attackers to retrieve private Mailchimp audience configuration information from any site that has configured the GutenKit Mailchimp integration.	5.3	More Details
CVE-2026-11354	The Participants Database plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 2.7.8.3 via the 'id' parameter. This makes it possible for unauthenticated attackers to overwrite arbitrary participant records by numeric ID and redirect the private_id-bearing record-access link to an attacker-controlled email address, granting full read and edit access to the victim's stored personally identifiable information including names, email addresses, phone numbers, and any other fields collected in the participant database. An attacker can harvest a valid nonce with a plain unauthenticated GET request to any page rendering the public signup or record form, then POST action=update with an arbitrary id value to overwrite any record; chaining a subsequent action=retrieve then delivers the private-access link to the attacker-controlled mailbox.	5.3	More Details
CVE-2026-44955	Pronetiqs IntraVUE versions 3.2.1a14 and prior have an exposure of sensitive system information to an unauthorized control sphere vulnerability which could allow for asset discovery by unauthenticated users.	5.3	More Details
CVE-2026-16763	A vulnerability was identified in localstack serverless-localstack up to 1.4.0. The affected element is an unknown function of the file src/index.js of the component Configuration Handler. The manipulation of the argument custom.localstack.docker.compose_file leads to os command injection. An attack has to be approached 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-2026-25466	Unauthenticated Broken Access Control in WP Go Maps <= 10.1.04 versions.	5.3	More Details
CVE-2026-27355	Unauthenticated Broken Access Control in Ditty <= 3.1.66 versions.	5.3	More Details
CVE-2026-16735	A security vulnerability has been detected in release-it conventional-changelog up to 11.0.1. This affects the function writeChangelog of the file index.js of the component Changelog File Handler. Such manipulation of the argument infile leads to os command injection. The attack must be carried out locally. 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.	5.3	More Details
CVE-2026-27399	Unauthenticated Broken Access Control in MarketKing <= 2.1.40 versions.	5.3	More Details
CVE-2026-16768	A flaw was found in gdk-pixbuf. When parsing a specially crafted ICO file with pixel values that exceed the defined palette range, an out-of-bounds read can occur due to improper bounds checking against the actual palette size. This vulnerability causes heap bytes to be interpreted as valid palette indices and rendered as RGB pixel values in the output image, allowing an attacker to extract heap content via the generated output, such as a thumbnail.	5.3	More Details
CVE-2026-	Unauthenticated Broken Access Control in WP Fast Total Search <= 1.81.282 versions.	5.3	More Details

27418			
CVE-2026-57716	Unauthenticated Arbitrary File Deletion in Broadcast Live Video <= 7.2.4 versions.	5.3	More Details
CVE-2026-61972	Unauthenticated Broken Access Control in ShopLentor Pro <= 2.8.5 versions.	5.3	More Details
CVE-2026-64785	SwiftNIO HTTP/2 was missing validation on inbound HEADERS frames that let CR, LF, NUL, SP and other control characters reach an HTTP/1.1 backend through NIOHTTP2's HTTP/2-to-HTTP/1 codec, enabling HTTP request smuggling or response splitting. This vulnerability is addressed in swift-nio-http2 version 1.45.0.	5.3	More Details
CVE-2026-65452	Unauthenticated Broken Access Control in Ebook Store <= 6.19 versions.	5.3	More Details
CVE-2026-65453	Unauthenticated Broken Access Control in Ebook Store <= 6.19 versions.	5.3	More Details
CVE-2026-16733	A weakness has been identified in bahmutov find-cypress-specs up to 1.54.12. The impacted element is the function shell.exec of the file src/index.js of the component Branch Handler. This manipulation of the argument --branch causes os command injection. The attack is restricted to local execution. 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.	5.3	More Details
CVE-2026-66064	goshs is a feature-rich single-binary file server for red teamers and developers. Prior to 2.1.5, the httpserver/handler.go sendFile handler opened files using a cleaned path but derived the authorization filename from raw req.URL.Path, so a trailing slash could bypass .goshs ACL-file protection and block-list checks. This issue is fixed in version 2.1.5.	5.3	More Details
CVE-2026-11804	Improper handling of insufficient permissions or privileges vulnerability in Tridium Niagara Framework on Windows, Linux, QNX, Tridium Niagara Enterprise Security on Windows, Linux, QNX allows Privilege Abuse. This issue affects Niagara Framework: before 4.14.6, before 4.15.5; Niagara Enterprise Security: before 4.14.6, before 4.15.5.	5.2	More Details
CVE-2026-11598	The Shortcodify plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'name' Shortcode Attribute in all versions up to, and including, 1.4.3 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with contributor-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page.	5.0	More Details
CVE-2026-65568	Contributor Broken Access Control in Visual Composer Website Builder <= 45.15.0 versions.	5.0	More Details
CVE-2026-17531	A weakness has been identified in unitedbyai droidclaw up to 0.5.3. Affected by this issue is some unknown functionality of the file server/src/routes/goals.ts of the component Unsigned Scheduled Callback. This manipulation causes authorization bypass. Remote exploitation of the attack is possible. The attack is considered to have high complexity. The exploitation is known to be difficult. 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.	5.0	More Details
CVE-2026-43767	The issue was addressed with improved memory handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to cause unexpected system termination.	5.0	More Details
CVE-2026-16488	A vulnerability was determined in QUSETIONS MiniCode-Python 0.1.0. This vulnerability affects the function subprocess.Popen of the file minicode/config.py of the component Project File Handler. Executing a manipulation can lead to os command injection. The attack may be launched remotely. A high complexity level is associated with this attack. It is stated that the exploitability is difficult. The exploit has been publicly disclosed and may be utilized. Upgrading to version 0.1.0-rc1 is able to resolve this issue. This patch is called 9d868dc2550f426c6ddf8ee98f30ffe450ca5e32. It is suggested to upgrade the affected component.	5.0	More Details
CVE-2026-16799	Improper access control in the automation tests and workflows features in Devolutions PowerShell Universal 2026.2.2 and earlier allows an authenticated user with only the Reader role to execute automation tests and modify workflow properties via missing server-side authorization checks.	5.0	More Details
CVE-2026-17432	A vulnerability was detected in NousResearch hermes-agent 2026.6.5. Affected by this vulnerability is an unknown functionality of the file hermes-agent/plugins/platforms/simplex/adapter.py of the component SimpleX Gateway Authorization. The manipulation of the argument contactId results in improper access controls. The attack may be launched remotely. A high complexity level is associated with this attack. The exploitation appears to be difficult. The exploit is now public and may be used. The patch is identified as 490c486ff65b766d9de0fe0e6f26e1778aaa8fb3. Applying a patch is advised to resolve this issue.	5.0	More Details
CVE-2026-	The SpeedyCache plugin for WordPress is vulnerable to Arbitrary File Read via Path Traversal in all versions up to, and including, 1.3.8. This is due to a mismatch between CSS URL validation (which allows query strings like `.css?...`) and path resolution (which strips query strings), combined with no validation that the resolved file is actually a CSS file.	4.9	More

5114	This makes it possible for authenticated attackers, with Administrator-level access and above, to read arbitrary files from the server (including `wp-config.php` and `/etc/passwd`) by injecting crafted `<link>` tags into page content, with the file contents written to publicly accessible cache files.		Details
CVE-2026-16811	The ShopLentor – All-in-One WooCommerce Growth & Store Enhancement Plugin plugin for WordPress is vulnerable to time-based SQL Injection via the 'orderby' parameter in all versions up to, and including, 3.4.5 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-51564	An issue in the redirect parameter in Milk admin <=0.9.8 allows remote attackers to redirect users to arbitrary external URLs via a crafted request.	4.9	More Details
CVE-2026-15663	The Ninja Forms – The Contact Form Builder That Grows With You plugin for WordPress is vulnerable to generic SQL Injection via Import File 'settings' Key in all versions up to, and including, 3.14.9 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. The vulnerable keys originate from the 'settings' object in an attacker-controlled import file processed via file_get_contents() or base64-decoded/JSON-decoded blobs, bypassing wp_magic_quotes protections entirely; two distinct sinks are affected — _save_setting() in Model.php and insert_form_meta() in ImportForm.php — as only the value side is escaped while the key side receives no sanitization or parameterization at any point in the call chain.	4.9	More Details
CVE-2026-15671	The SMS Alert – SMS & OTP for WooCommerce, Order Notifications & Abandoned Cart Recovery plugin for WordPress is vulnerable to generic SQL Injection via the 'id' parameter in all versions up to, and including, 3.9.7 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-15670	The SMS Alert – SMS & OTP for WooCommerce, Order Notifications & Abandoned Cart Recovery plugin for WordPress is vulnerable to time-based SQL Injection via the 'orderby' parameter in all versions up to, and including, 3.9.7 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-66437	Contributor Server Side Request Forgery (SSRF) in Feedzy <= 5.2.4 versions.	4.9	More Details
CVE-2026-66476	Administrator Arbitrary File Deletion in Easy Digital Downloads <= 3.6.9 versions.	4.9	More Details
CVE-2026-15444	The Tutor LMS – eLearning and online course solution plugin for WordPress is vulnerable to generic SQL Injection via the 'coupon_code' parameter in all versions up to, and including, 4.0.1 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database.	4.9	More Details
CVE-2026-65467	Contributor Server Side Request Forgery (SSRF) in JetEngine <= 3.8.11 versions.	4.9	More Details
CVE-2026-1918	IBM Sterling B2B Integrator 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 stores potentially sensitive information in log files that could be read by a privileged user.	4.9	More Details
CVE-2026-65466	Custom role Server Side Request Forgery (SSRF) in JetBooking <= 4.1.2 versions.	4.9	More Details
CVE-2026-56416	In NLnet Labs Unbound up to and including version 1.25.1, when the validator builds the canonical RDATA form for an RRSIG-covered PX/RP/MINFO/SOA RRset, it computes the address of the second embedded domain name as 'datstart + dname_valid(datstart, ...)' and passes it straight to 'query_dname_tolower()' without checking that a second name is actually present in the RDATA. The wire-format parser accepts multi-dname RRs whose RDATA ends after the first name, so an attacker who runs a DNSSEC-signed authoritative server can deliver a record with an absent second domain name (e.g. SOA record) and cause 'query_dname_tolower()' to walk label-by-label through stale bytes in the per-worker 'env->scratch_buffer', past the end of that heap allocation if 'msg-buffer-size' has been lowered from the default. This leads to heap buffer overflow and on a release build the outcome relies heavily on the contents of the buffer tail and the adjacent heap chunk.	4.8	More Details
CVE-2026-9577	The Post Status Notifier Lite WordPress plugin before 1.13.0 does not properly escape the `mod` URL parameter before reflecting it into the admin settings page (`admin.php?page=post-status-notifier-lite`), leading to a Reflected Cross-Site Scripting vulnerability that fires in the administrator's session when they are tricked into following a crafted URL.	4.8	More Details

CVE-2026-63281	Joomla Extension - regularlabs.com - XSS vulnerability in Regular Labs conditions manager - Stored condition values could also execute HTML/JavaScript in administrator summaries.	4.8	More Details
CVE-2026-65531	Unauthenticated Broken Access Control in Qubely <= 1.8.14 versions.	4.8	More Details
CVE-2026-66139	OpenStack Zaqar through 22.0.0 allows authentication bypass via an EXTRA-SPEC header when a UUID is known.	4.8	More Details
CVE-2026-14203	The Smart Manager WordPress plugin before 8.92.0 does not properly encode a post field before rendering it into an HTML attribute in its management grid, allowing users with the Contributor role or above to inject JavaScript that executes in the browser session of an administrator who views the grid.	4.8	More Details
CVE-2026-43770	A race condition was addressed with additional validation. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6, tvOS 26.6. An app may be able to access sensitive user data.	4.7	More Details
CVE-2026-14236	The Contact Form 7 WordPress plugin before 2.5 does not validate the host of a user-supplied return URL before using it as the success and cancel redirect targets of a Stripe checkout, allowing an unauthenticated attacker to redirect a victim, via a crafted link, to an arbitrary external site after the checkout flow.	4.7	More Details
CVE-2026-65904	DOMPurify through 3.3.3 fails to sanitize DOM elements passed via IN_PLACE mode when the element originates from a different window/realm (e.g., an iframe's contentDocument). A cross-realm instanceof check in the private _isNode() function returns false for foreign-realm nodes, causing DOMPurify to stringify the element (yielding '[object HTMLDivElement]'), silently reset IN_PLACE to false, and return the unsanitized element unchanged with any XSS payloads intact.	4.7	More Details
CVE-2026-43781	A race condition was addressed with improved state handling. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An app may be able to access sensitive user data.	4.7	More Details
CVE-2026-43811	A race condition was addressed with improved checks. This issue is fixed in iOS 26.6 and iPadOS 26.6. An app may be able to modify protected parts of the file system.	4.7	More Details
CVE-2026-43753	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in iOS 26.6 and iPadOS 26.6, macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An attacker with physical access to a locked device may be able to view sensitive user information.	4.6	More Details
CVE-2026-43766	An authorization issue was addressed with improved state management. This issue is fixed in macOS Sequoia 15.7.8, macOS Sonoma 14.8.8, macOS Tahoe 26.6. An attacker with physical access to a locked device may be able to view sensitive user information.	4.6	More Details
CVE-2026-64732	This issue was addressed through improved state management. This issue is fixed in iOS 26.6 and iPadOS 26.6. An attacker with physical access may be able to access sensitive user data during iPhone Mirroring.	4.6	More Details
CVE-2026-7007	The Zephyr ext2 file system validates the on-disk superblock in ext2_verify_disk_superblock() (subsys/fs/ext2/ext2_impl.c) before completing a mount. The validator checked the magic number, block size, revision and feature flags, but did not verify that the on-disk fields s_blocks_per_group and s_inodes_per_group are non-zero. Both fields are read directly from the image and are later used as divisors during mount-time initialization. During mount, get_ngroups() divides and modulos s_blocks_count by s_blocks_per_group (reached via ext2_fetch_block_group() from ext2_init_fs()), and get_itable_entry() divides (ino - 1) by s_inodes_per_group when fetching the root inode (both in subsys/fs/ext2/ext2_diskops.c). A superblock with either field set to zero therefore causes an integer division by zero during the mount sequence. An attacker who can present a crafted ext2 image to a device that mounts ext2 — removable media such as an SD card or a USB mass-storage device — can trigger this. On ARMv7-M / ARMv8-M-mainline Cortex-M targets, divide-by-zero trapping is enabled (SCB_CCR_DIV_0_TRP), so the division raises a UsageFault that Zephyr treats as a fatal error, producing a denial of service. The impact is limited to availability; the malformed value is consumed only as a divisor. The fix rejects a zero s_blocks_per_group or s_inodes_per_group in the superblock validator, returning -EINVAL so the mount fails before any block-group or inode I/O occurs.	4.6	More Details
CVE-2026-8058	IBM OPENBMC FW1110.00 through FW1110.20, and FW1060.00 through FW1060.71 allows a user to supply a password with a resource dump request stores that password into the BMC audit log where an admin user can see it.	4.5	More Details
CVE-2026-15673	The SMS Alert – SMS & OTP for WooCommerce, Order Notifications & Abandoned Cart Recovery plugin for WordPress is vulnerable to generic SQL Injection via 'checkout_payment_plans' and 'order_status' Settings in all versions up to, and including, 3.9.7 due to insufficient escaping on the user supplied parameter and lack of sufficient preparation on the existing SQL query. This makes it possible for authenticated attackers, with administrator-level access and above, to append additional SQL queries into already existing queries that can be used to extract sensitive information from the database. This is a second-order SQL injection: the malicious payload is stored in the 'checkout_payment_plans' and 'order_status' settings via update_option() and executed later when the	4.4	More Details

	cod_to_prepaid_cart_notification_sendsms_hook WP-Cron event fires SA_CodTOPrepaid::sendSms().		
CVE-2026-65496	Author Server Side Request Forgery (SSRF) in Complianz <= 7.5.0 versions.	4.4	More Details
CVE-2026-15786	The WP Encryption – One Click Free SSL Certificate & SSL / HTTPS Redirect, Security & SSL Scan plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 7.8.6.6 via the 'imploded' parameter parameter. This makes it possible for authenticated attackers, with administrator-level access and above, to read the contents of arbitrary files on the server, which can contain sensitive information. Although file write content is passed through esc_html(), which encodes angle brackets and prevents direct PHP execution, plaintext configuration files such as .htaccess are fully writable and exploitable for denial-of-service or redirect attacks. This is only exploitable when the premium version of the software is enabled and active.	4.4	More Details
CVE-2026-15647	The Brands for WooCommerce plugin for WordPress is vulnerable to Stored Cross-Site Scripting via 'br_brand_tooltip' Term Meta Field in all versions up to, and including, 3.8.8 due to insufficient input sanitization and output escaping. This makes it possible for authenticated attackers, with custom-level access and above, to inject arbitrary web scripts in pages that will execute whenever a user accesses an injected page. Because the payload is stored in term meta rather than post content, the WordPress unfiltered_html capability exception does not apply, meaning Shop Manager-level users — who normally lack unfiltered_html — can fully exploit this vulnerability.	4.4	More Details
CVE-2026-24639	Author Server Side Request Forgery (SSRF) in Photo Block <= 1.7.1 versions.	4.4	More Details
CVE-2026-61973	Subscriber Broken Access Control in ShopLentor Pro <= 2.8.5 versions.	4.3	More Details
CVE-2026-17459	A vulnerability was determined in perwendel spark up to 2.9.4. This vulnerability affects the function staticFiles.externalLocation of the file src/main/java/spark/resource/ExternalResourceHandler.jav of the component SparkJava. Executing a manipulation can lead to symlink following. It is possible to launch the attack remotely. The exploit has been publicly disclosed and may be utilized. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2026-65920	Diffusers through 0.39.0, fixed in commit cee298c, contains a path traversal vulnerability in the _get_checkpoint_shard_files function that allows attackers to read arbitrary files by supplying malicious weight_map values in model index JSON. Attackers can use ../ sequences or absolute paths in weight_map entries to escape the model directory and read safetensors files outside the intended location during model loading.	4.3	More Details
CVE-2026-65457	Subscriber Broken Access Control in ЮKassa для WooCommerce <= 2.16.1 versions.	4.3	More Details
CVE-2026-65456	Contributor Insecure Direct Object References (IDOR) in Product Slider for WooCommerce <= 1.13.62 versions.	4.3	More Details
CVE-2026-3158	IBM Sterling B2B Integrator 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 is vulnerable to an information disclosure due to sensitive information being included in the source code comments of a dashboard component.	4.3	More Details
CVE-2026-3157	IBM Sterling B2B Integrator 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.0.0 through 6.2.0.5_2, 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 is vulnerable to an information disclosure due to sensitive information being included in the source code comments of a mailbox component.	4.3	More Details
CVE-2026-64810	In JetBrains IntelliJ IDEA before 2026.2 hTML injection was possible in an IDE notification, allowing silent user activity tracking	4.3	More Details
CVE-2026-17457	A vulnerability has been found in mf-yang openclaw-cn up to 0.2.1. Affected by this issue is the function assertBrowserNavigationAllowed of the file src/browser/navigation-guard.ts of the component Scheme Handler. Such manipulation of the argument url leads to information disclosure. The attack may be performed from remote. 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-2026-24537	Unauthenticated Cross Site Request Forgery (CSRF) in WP Accessibility Helper (WAH) <= 0.6.6 versions.	4.3	More Details
CVE-2026-8287	Allocation of resources without limits or throttling vulnerability in BizimHesap Information Systems Industry and Trade Inc. Online Pre-Accounting Software allows Excessive Allocation. This issue affects Online Pre-Accounting Software: through 17072026.	4.3	More Details
CVE-			

2026-25424	Contributor Broken Access Control in Mediavine Control Panel <= 2.10.10 versions.	4.3	More Details
CVE-2026-65537	Subscriber Broken Access Control in Cyr to Lat reloaded – transliteration of links and file names <= 1.3.3 versions.	4.3	More Details
CVE-2026-27423	Subscriber Broken Access Control in Participants Database <= 2.7.8.4 versions.	4.3	More Details
CVE-2026-65491	Subscriber Broken Access Control in Query Wrangler <= 1.5.57 versions.	4.3	More Details
CVE-2026-65535	Contributor Sensitive Data Exposure in TinyMCE Templates <= 4.8.1 versions.	4.3	More Details
CVE-2026-27392	Contributor Broken Access Control in uListing <= 2.2.0 versions.	4.3	More Details
CVE-2026-7362	IBM Sterling B2B Integrator 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 and IBM Sterling File Gateway 6.2.1.0 through 6.2.1.1_2, and 6.2.2.0 through 6.2.2.0_1 could allow an authenticated user to obtain sensitive information that should only be available to a privileged user.	4.3	More Details
CVE-2026-65524	Contributor Broken Access Control in Avada Custom Branding <= 1.2 versions.	4.3	More Details
CVE-2026-65530	Subscriber Broken Access Control in TemplateSpare <= 4.2.2 versions.	4.3	More Details
CVE-2026-10600	Mattermost versions 11.8.x <= 11.8.0, 11.7.x <= 11.7.3, 11.6.x <= 11.6.5, 10.11.x <= 10.11.20 fail to bound the time and resource consumption of server-side document content extraction which allows an authenticated user with file-upload permission to degrade file uploads for all users on the server via repeatedly uploading small documents that are cheap to upload but expensive to extract, saturating the shared extraction worker pool.. Mattermost Advisory ID: MMSA-2026-00694	4.3	More Details
CVE-2026-65458	Contributor Sensitive Data Exposure in Polylang <= 3.8.5 versions.	4.3	More Details
CVE-2026-59728	Astro is a web framework for content-driven websites. In versions 1.0.0 through 4.0.18, the source.title and enclosure.type item fields in packages/astro-rss/src/index.ts are interpolated directly into XML template strings without XML-character escaping before being parsed by fast-xml-parser. Both fields are validated only as z.string(), placing no restriction on XML special characters. An attacker who controls these values can inject arbitrary XML into the generated RSS feed: a value containing " can break out of an attribute (as with enclosure.type), and a value containing </source> can close an element early and inject additional nodes (as with source.title). This corrupts feed structure, injects false metadata (for example, a fake <link> pointing to a malicious URL), and can cause feed readers to misparse or display attacker-controlled content. In SSR mode (output: 'server'), the poisoned feed is served on every request to all subscribers. This issue has been fixed in version 4.0.19.	4.3	More Details
CVE-2026-65650	Elgg before 7.0.0 does not check image dimensions to prevent denial of service via a large avatar upload.	4.3	More Details
CVE-2026-66474	Unauthenticated Cross Site Request Forgery (CSRF) in Insert Headers and Footers Code – HT Script <= 1.1.8 versions.	4.3	More Details
CVE-2026-16473	A flaw was found in the sbc library (BlueZ SBC codec). An off-by-one error in the SBC frame decoder allows a crafted audio payload to trigger a one-byte heap out-of-bounds read. This could allow an adjacent attacker streaming Bluetooth audio to read a single byte of adjacent heap memory.	4.3	More Details
CVE-2026-66428	Unauthenticated Cross Site Request Forgery (CSRF) in WP Google Review Slider <= 18.4 versions.	4.3	More Details
CVE-2026-16797	The ShopLentor – All-in-One WooCommerce Growth & Store Enhancement Plugin plugin for WordPress is vulnerable to Insecure Direct Object Reference in all versions up to, and including, 3.4.5 via the 'optionSection' parameter due to missing validation on a user controlled key. This makes it possible for authenticated attackers, with contributor-level access and above, to read arbitrary wp_options rows — including internal plugin news feed data, WooCommerce block pattern transients, and third-party configuration records — whose values are stored as arrays-of-arrays containing 'title' keys, enabling cross-plugin data leakage.	4.3	More Details

CVE-2026-15420	The Nexter Blocks – Gutenberg Blocks, Page Builder & AI Website Builder plugin for WordPress is vulnerable to Directory Traversal in all versions up to, and including, 5.0.0 via the 'plus_name' parameter. This makes it possible for authenticated attackers, with subscriber-level access and above, to delete arbitrary JS/CSS files on the server, which can lead to denial of service or destruction of critical plugin and theme assets.	4.3	More Details
CVE-2026-16587	The Advanced Form Integration — Connect Forms to 200+ Apps plugin for WordPress is vulnerable to authorization bypass in all versions up to, and including, 2.6.0. This is due to the plugin not properly verifying that a user is authorized to perform an action. This makes it possible for authenticated attackers, with subscriber-level access and above, to overwrite the site's stored MailUp OAuth tokens in the adfoin_mailup_keys option with attacker-controlled tokens, hijacking future form-submission data to a MailUp account they control or nulling the tokens to break the integration entirely. This is exploitable by any authenticated user who can reach /wp-admin/profile.php, as admin_init fires for all logged-in users visiting any wp-admin page.	4.3	More Details
CVE-2026-15136	The Cookie Banner for GDPR / CCPA – WPLP Cookie Consent plugin for WordPress is vulnerable to Cross-Site Request Forgery in all versions up to, and including, 4.3.7. This is due to missing or incorrect nonce validation on the process_bulk_action function. This makes it possible for unauthenticated attackers to permanently delete or forcibly resolve arbitrary GDPR data request records stored in the wpl_data_req table via a forged request granted they can trick a site administrator into performing an action such as clicking on a link.	4.3	More Details
CVE-2026-65011	Graylog2 Server before commit 46a2eeb contains a missing per-entity permission check in the POST /events/definitions/{definitionId}/duplicate endpoint that allows authenticated users to clone any event definition. Attackers with the low-privilege eventdefinitions:create capability can read private event definitions including detection queries, aggregation thresholds, grouping fields, schedules, and notification bindings by duplicating them.	4.3	More Details
CVE-2026-18038	A flaw has been found in nextlevelbuilder GoClaw up to 3.13.2. Affected by this vulnerability is the function ExecTool.Execute of the file goclaw/internal/http/tools_invoke.go of the component jq Handler. Executing a manipulation can lead to information disclosure. The attack can be launched remotely. The exploit has been published and may be used. This patch is called 1230. It is advisable to implement a patch to correct this issue.	4.3	More Details
CVE-2026-66750	Let's Chat 0.3.0 through 0.4.8 contains a broken access control vulnerability that allows authenticated attackers to download file attachments from private and password-protected rooms they are not a member of by exploiting missing room membership checks in the file retrieval route. Attackers can enumerate adjacent MongoDB ObjectIds derived from a known file ID to recover files uploaded by other users, as the GET /files/:id/:name route in app/controllers/files.js only enforces login authentication without consulting room membership or the Room.canJoin check.	4.3	More Details
CVE-2026-55985	The web management interface in Tycon Systems TPDIN-Monitor-WEB2 stores and displays system credentials in cleartext on a certain configuration page accessible to authenticated users. Any party with access to the administrative dashboard can immediately read these credentials, which may be used to compromise other systems on the local network.	4.3	More Details
CVE-2026-48012	Shopware is an open commerce platform. Versions 6.7.3.0 through 6.7.10.0 have an open redirect in Shopware's public SSO entry point at `GET /api/oauth/sso/auth`. When the endpoint is reached without the expected SSO session state, the application falls back to the request's `Referer` header and uses that value as the redirect destination. In the validated behavior, the server does not restrict that fallback target to same-origin URLs, does not require a relative path, and does not reject dangerous schemes such as `javascript:`. As a result, an unauthenticated request can turn this endpoint into a reusable redirect primitive whose destination is fully controlled by attacker-supplied request metadata. The security problem is not limited to a harmless navigation mismatch. The endpoint sits under `/api/oauth/`, which gives the redirect a trustworthy application-controlled origin and makes it suitable for phishing chains, branded redirect abuse, and cases where client software automatically follows redirects issued by a trusted host. The attached evidence also shows that the response is not only an HTTP `302` with a user-controlled `Location` header. The HTML body contains a matching meta refresh tag and redirect link built from the same attacker-controlled value. In the validated proof, the endpoint redirects to `https://attacker.example/poc` when that URL is supplied through `Referer`, and it also reflects `javascript:alert(1)` into `Location` and the HTML redirect body without any scheme filtering. This report therefore stays conservative and claims an open redirect with arbitrary redirect targets, while noting that the lack of scheme restrictions makes the behavior materially worse than a same-scheme external redirect. Version 6.7.10.1 fixes the issue.	4.3	More Details
CVE-2026-13061	An authenticated user may be able to view session metadata belonging to other users on the system through the $listSessions aggregation stage. This information is normally restricted to users with cluster-level administrative privileges, and includes active session identifiers, associated usernames, and activity timestamps.	4.3	More Details
CVE-2026-13063	An authenticated user with standard read/write privileges can cause the mongod process to terminate due to an out-of-memory condition by sending a crafted aggregation command. MongoDB's libmongocrypt library insufficiently validates payload-supplied values, which can result in an excessively large memory allocation.	4.3	More Details
CVE-2026-65460	Unauthenticated Cross Site Request Forgery (CSRF) in Zarinpal Gateway <= 5.1.0 versions.	4.3	More Details
CVE-2026-13073	An authenticated user with read-only privileges can cause the mongod process to terminate abnormally by issuing a crafted aggregation command, resulting in denial of service for all connected clients until the process is restarted. The issue stems from an internal engine selection inconsistency triggered by a specific combination of aggregation options.	4.3	More Details

CVE-2026-63262	Missing Authorization (CWE-862) in Kibana can lead to unauthorized cross-space information disclosure via user-supplied input that circumvents space-level access control.	4.3	More Details
CVE-2026-17569	Improper access control in the NetBox synchronizer in Devolutions Server allows an authenticated user with view-only permission on an entry to obtain a stored API token via the partial connection endpoint. This issue affects : * Devolutions Server 2026.2.4.0 through 2026.2.12.0 * Devolutions Server 2026.1.23.0 and earlier	4.3	More Details
CVE-2026-17570	Improper access control in the PAM password history endpoints in Devolutions Server allows an authenticated low-privileged user to disclose plaintext credential secrets via crafted API requests. This issue affects : * Devolutions Server 2026.2.4.0 through 2026.2.12.0 * Devolutions Server 2026.1.23.0 and earlier	4.3	More Details
CVE-2026-58246	SAP NetWeaver Application Server for ABAP and ABAP Platform writes sensitive session identifier information into a diagnostic trace when the trace is activated by a privileged user. An attacker with access to the resulting trace data could obtain identifiers that allow impersonation of legitimate users during their validity period. This leads to high impact on confidentiality. Integrity and availability are not impacted.	4.3	More Details
CVE-2026-14926	The FluentCart A New Era of eCommerce WordPress plugin before 1.4.0 does not verify that a subscription belongs to the requesting customer in several of its payment-method endpoints, allowing any authenticated customer to act on another customer's subscription (changing its payment method, or cancelling and re-binding it) when they know the target subscription identifier.	4.2	More Details
CVE-2026-4932	IBM PowerVM Hypervisor FW1110.00 through FW1110.20, and FW1060.00 through FW1060.71 could allow an attacker with physical access to the Transparent Memory Encryption (TME) hardware to decrypt encrypted memory due to insufficient cryptographic entropy.	4.2	More Details
CVE-2026-13068	An authenticated user holding cursor termination privileges on one database may incorrectly be permitted to terminate active cursors on a separate database, disrupting ongoing query operations for other users. The behavior stems from an authorization check that does not correctly scope privileges to the appropriate namespace.	4.2	More Details
CVE-2026-65699	AgentGPT through 1.0.0 contains an authorization bypass through user-controlled key vulnerability that allows authenticated users to attach tasks to another user's agent run by supplying a target run_id in the request body without ownership verification. The AgentCRUD.create_task and validate_task_count functions look up the target AgentRun using the client-supplied run_id without confirming the run belongs to the requesting user, enabling an attacker who obtains a valid run_id to corrupt task history, exhaust the per-run loop budget, and drive LLM costs against the victim's run.	4.2	More Details
CVE-2026-4912	The Media Cleaner: Clean your WordPress! plugin for WordPress is vulnerable to Server-Side Request Forgery in all versions up to, and including, 7.0.3. This is due to the `get_urls_from_html()` function using `DOMDocument::loadHTMLFile()` to fetch iframe source URLs with an insufficient hostname validation check that relies on a substring match against the site's server name. This makes it possible for authenticated attackers, with Administrator-level access and above, to make web requests to arbitrary locations originating from the web application, which can be used to query and interact with internal services.	4.1	More Details
CVE-2026-48013	Shopware is an open commerce platform. Prior to 6.6.10.18 and 6.7.10.1, the `/api/_action/media/external-link` endpoint allows authenticated admin users to make server-side HTTP HEAD requests to arbitrary internal IP addresses. While the parallel `uploadFromURL` flow validates target IPs against private/reserved ranges via `FileUrlValidator`, the `linkURL` flow only performs a URL format check (regex for `http://` or `https://` prefix), allowing SSRF to internal network services and cloud metadata endpoints. This issue is fixed in versions 6.6.10.18 and 6.7.10.1.	4.1	More Details
CVE-2026-12690	The ProfileGrid WordPress plugin before 5.9.9.7 does not perform a capability check on its license management actions, relying only on a nonce that is exposed to any logged-in user, allowing authenticated users with Subscriber-level access and above to overwrite the site's premium license settings.	3.8	More Details
CVE-2026-14189	The WPBot WordPress plugin before 8.5.2 does not validate administrator-configured field identifiers before using them in a SQL query, allowing users with administrator access to perform SQL injection that executes when a visitor triggers a search.	3.8	More Details
CVE-2026-41637	In NLnet Labs Unbound 1.22.0 up to and including 1.25.1, client terminated DNS-over-QUIC (DoQ) queries are not accounted properly by Unbound resulting in low-cost inflation of the waiting number of replies for already in-flight resolution queries. This results in degradation of resolution service for new clients for already in-flight queries. A malicious actor can exploit the vulnerability by issuing DoQ queries for query names that need resolution and proceeding on immediately terminating the query by one of STOP_SENDING/RESET_STREAM/CONNECTION_CLOSE QUIC frames. Those terminated DoQ queries are not properly counted for and keep inflating the number of waiting replies for in-flight queries. When the maximum is reached, it results in silent query drops for new clients needing resolution for already in-flight queries. This vulnerability needs Unbound to be compiled with DoQ support ('--with-libngtcp2') and the 'quic-port' to be configured for the listening interfaces. Additionally, a malicious actor needs access to multiple source IPs to bypass the by-default configured 'wait-limit' option.	3.7	More Details
CVE-2026-42955	In NLnet Labs Unbound 1.16.2 up to and including 1.25.1, a similar vulnerability as with CVE-2026-40622 in the 'ghost domain names' family of attacks was found in Unbound that could extend the ghost domain window by up to one cached TTL configured value for A/AAAA glue records. Similar to other 'ghost domain names' attacks, an adversary needs to control a (ghost) zone and be able to query a vulnerable Unbound. A single client A/AAAA query can cause Unbound to overwrite the cached expired parent-side glue rrset and essentially extend the ghost domain window by up to one cached TTL configured value ('cache-max-ttl'). In configurations where 'harden-referral-path: yes' is used	3.7	More Details

	(non-default configuration), no client query is required since Unbound implicitly performs that query. This is a variant of CVE-2026-40622 which only addressed the NS query.		
CVE-2026-44687	In NLnet Labs Unbound 1.13.2 up to and including 1.25.1, stub or forward zones where the name is below an intermediate labed below a DNSSEC signed zone could be shadowed by the intermediate label's secure NXDOMAIN answer from the parent. This is caused by an off-by-one error in 'harden-below-nxdomain' logic; enabled by default. It effectively bypasses the configuration and the configured stub/forward zone is never contacted. 'harden-below-nxdomain' does an upward DNS cache walk together with a delegation point guard that does not allow NXDOMAIN synthesis above stub/forward zones. The guard tests the domain name but before stripping a label. This results in an iteration where the domain name equals the configured stub/forward zone apex that passes the guard, strips one more label, and probes the cache at the apex's immediate public parent. If that parent has a cached DNSSEC-secure NXDOMAIN, which it will for any private namespace nested two or more labels under a signed public name, the walk returns it and the configured stub/forward upstream is never contacted. This can only be triggered by the query for the intermediate label (between the stub/forward apex and the DNSSEC parent zone).	3.7	More Details
CVE-2026-55403	datamodel-code-generator generates Python data models from schema definitions. Prior to 0.63.0, src/datamodel_code_generator/http.py get_body reuses Authorization, Cookie, and Proxy-Authorization headers when following cross-origin redirects while fetching remote schemas, allowing credentials scoped to one schema host to be leaked to another redirect target. This issue is fixed in version 0.63.0.	3.7	More Details
CVE-2026-46582	In NLnet Labs Unbound 1.6.0 up to and including 1.25.1, a replay of a wildcard rrset as another piece of data, could be briefly considered DNSSEC secure based only on the RRSIG validation and stored into cache, before later validation treats it as bogus based on NSEC validation. When the resolving thread puts secure on the rrset, and another thread that is on the serve expired path then picks up the updated rrset contents with the secure status for a reply, it can be used to change a specific record, next to a wildcard that could be covered by the wildcard, into the wildcard. A malicious actor can exploit the possible poisonous effect by having any DNSSEC-singed domain (irrelevant to the victim domain) and a CNAME wrapper record that points to a record next to a wildcard (that could be covered by the wildcard). Then quering Unbound for the wildcard sibling record would seed the secure message. A later (after expiry) query for the CNAME wrapper would need to resolve the target sibling record. If the wildcard replay is injected into the response, the wildcard rrset will update the expired sibling record with a secure status before completing proper wildcard validation with NSEC records and eventually treating the CNAME wrapper answer as bogus. The updated poisoned rrset is now secure and points to the wildcard. This vulnerability is explicit for the serve expired path and needs injection of the signed wildcard rrset without the NSEC accompanying rrset.	3.7	More Details
CVE-2026-66753	tiny-http through 0.12.0 contains an HTTP header injection vulnerability that allows attackers to inject carriage return (0x0D) and line feed (0x0A) bytes into HTTP header values on both request and response sides due to insufficient validation in header parsing and serialization. Attackers can exploit this injection primitive to perform response splitting, cache poisoning, session fixation via Set-Cookie injection, security header override, and request smuggling against line-feed-tolerant backends.	3.7	More Details
CVE-2026-50243	In NLnet Labs Unbound 1.6.2 up to and including 1.25.1, when Unbound is configured with the 'respip' module in front of the validator together with a 'response-ip' redirect rule or an RPZ file with an RPZ-IP trigger, the rewriting handler does not check the security status of the upstream answer and can instead rewrite a BOGUS A/AAAA answer to point to an operator's configured IP. If the validator finds an expired or otherwise invalid RRSIG on an answer whose A record falls within a 'response-ip'/RPZ configuration, the answer is still rewritten and given a hard coded security level of INSECURE. This results in the client receiving an INSECURE NOERROR reply rewritten by the operator's configured IP. A malicious actor can exploit the possible poisonous effect by spoofing a BOGUS A/AAAA answer that falls inside the operator's configured subnet rewrites. Such DNSSEC protected answers are then insecurely redirected to the operator's configured target.	3.7	More Details
CVE-2026-52686	The issue is a DNSSEC validation bypass where wildcard expansion proofs (NSEC/NSEC3 records) are accepted without signature validation when the wildcard answer is a CNAME or DNAME record.	3.7	More Details
CVE-2026-52684	If the auth responds very slowly and the records expire in between, the capping of TTLs is not enforced for lack of data. This does not happen on regular resolve as then then the child records are used immediately if not expired and thus valid, or the records are expired, and in that case not used. So this case can only happen if almost expired records are used to refresh the authoritative NS records.	3.7	More Details
CVE-2026-54478	In NLnet Labs Unbound 1.18.0 up to and including 1.25.1, when Unbound listens on a 'proxy-protocol-port' interface with 'answer-cookie: yes', the RFC 9018 server-cookie SipHash is computed over the proxy's wire address instead of the PROXYv2-declared client. One server cookie obtained through a given proxy node therefore validates for every PROXYv2-declared source behind that node. On a UDP+proxy-protocol front, an off-path attacker can harvest one cookie with a single legitimate query, then replay it under any spoofed source and pass DNS Cookie checks that were deployed to defeat this in the first place.	3.7	More Details
CVE-2026-14819	The Event Tickets and Registration WordPress plugin before 5.28.4 does not properly escape event titles before outputting them in a ticket history log, allowing users with the Editor role and above to perform Stored Cross-Site Scripting attacks that execute against higher-privileged users on multisite installations.	3.5	More Details
CVE-2026-64800	In JetBrains GoLand before 2026.2 sensitive configuration values written to log files by default	3.5	More Details
	Papra is a minimalistic document management and archiving platform. Prior to version 26.5.0, Papra's webhook delivery system contains an SSRF protection bypass that allows any authenticated organisation member to cause the		

CVE-2026-48051	server to make HTTP requests to internal addresses — loopback, link-local, and RFC-1918 ranges. The SSRF protection validates the registered webhook URL but ignores redirect destinations. The HTTP client (ofetch) follows 3xx responses automatically, and the redirect target is never checked against the blocklist. An attacker registers a webhook pointing to an attacker-controlled server, which redirects incoming POSTs to any internal address. Exploitation was confirmed by live test against the official Docker image. The fix is a single-line change to the webhook HTTP client. This issue has been patched in version 26.5.0.	3.5	More Details
CVE-2026-56537	HCL Connections is vulnerable to information disclosure which could allow a user to obtain sensitive information they are not entitled to, caused by improper handling of request data.they are not entitled to, caused by improper handling of request data.	3.5	More Details
CVE-2026-56538	An endpoint in HCL Connections is vulnerable to information disclosure. In certain scenarios this might lead to disclosing sensitive information to unauthorized users.	3.5	More Details
CVE-2026-17512	A vulnerability has been found in ggml-org whisper.cpp 1.8.4-58. This impacts the function log_mel_spectrogram of the file src/whisper.cpp. The manipulation leads to out-of-bounds read. The attack needs to be performed locally. The pull request to fix this issue awaits acceptance.	3.3	More Details
CVE-2026-66011	ImageMagick before 7.1.2-27 contains a memory leak vulnerability in the magick command-line interface when invalid options are provided. Attackers can trigger memory exhaustion by repeatedly supplying malformed command-line arguments to consume system resources.	3.3	More Details
CVE-2026-17072	A flaw was found in GStreamer's gst-plugins-good. A heap-based out-of-bounds read of 4 bytes can occur when parsing FLAC audio stream headers embedded in a Matroska or WebM container file. The vulnerability is triggered by a boundary check that does not account for the full size of the data being copied, allowing a small read past the end of the allocated buffer. An attacker could exploit this by crafting a malicious Matroska or WebM file and tricking a user into opening it, potentially leaking a small amount of adjacent heap memory.	3.3	More Details
CVE-2026-55977	Successful exploitation of this vulnerability could allow an attacker with local network access to bypass the application's rate-limiting mechanism, enabling brute-forcing of the screen-sharing code and potentially displaying harmful content on the affected screen.	3.3	More Details
CVE-2026-44187	A flaw was found in the Ansible Lightspeed extension for Visual Studio Code. This vulnerability allows an attacker with local access to the workstation, or malware running with the user's privileges, to read the Google Gemini API key. The extension insecurely stores the API key in plain text within the user's configuration file and writes it to output log files. This information disclosure can lead to the attacker obtaining the API credential and potentially consuming the user's API quota.	3.3	More Details
CVE-2026-17513	A vulnerability was found in ggml-org whisper.cpp 95ea8f9b. Affected is the function ggml_ftype_to_ggml_type of the file ggml/src/ggml.c. The manipulation of the argument ftype results in reachable assertion. The attack requires a local approach. The project was informed of the problem early through an issue report but has not responded yet.	3.3	More Details
CVE-2026-55708	In NLnet Labs Unbound 1.6.0 up to and including 1.25.1, the 'view_local_data' and 'view_local_datas' commands of 'unbound-control' create a bare local zones tree for an already configured named view when the view is configured with no local data to begin with. However, the creation through the control interface omits adding the default-protected zones (e.g., RFC 1918 reverse, AS112 zones, .onion, .localhost). Once the local zone tree exists without the defaults, every query for a default-protected name from a client mapped to that view escapes to the public DNS via the iterator instead of being answered locally, bypassing local policy expectations.	3.1	More Details
CVE-2026-17039	A flaw was found in pki-core. The certificate authority (CA) renewal request path does not perform the realm-based authorization check that the enrollment path performs, allowing an authenticated user entitled to one realm to cause a certificate belonging to a different realm to be renewed without that realm's authorization.	3.1	More Details
CVE-2026-14821	The Quiz and Survey Master (QSM) WordPress plugin before 11.1.5 does not perform a capability check before deleting output templates, allowing users with contributor-level access and above to delete arbitrary templates.	2.7	More Details
CVE-2026-64745	This issue was addressed with additional restrictions on the lock screen. This issue is fixed in macOS Sequoia 15.7.8, macOS Tahoe 26.6. A person with physical access to a locked device may be able to access contacts and photos.	2.4	More Details
CVE-2026-10683	In the Synopsys DesignWare I2C driver (drivers/i2c/i2c_dw.c) operating in target/slave mode, the rx_full interrupt handler gates the write_requested() callback on dw->state != CMD_SEND, and dw->state is only reset to READY on a STOP interrupt. The START_DET interrupt, whose handler in i2c_dw_slave_read_clear_intr_bits() would reset the state on every (re)START, was never added to the enabled interrupt mask in i2c_dw_slave_register(), so that recovery path was dead code. As a result, if the STOP interrupt is lost (bus glitch/reset, or a concurrent master driving STOP) or the bus master issues a legal WRITE-repeated-START-WRITE sequence with the same direction, the driver remains in CMD_SEND permanently and never invokes write_requested() again for the life of the target. An I2C master on the same physical bus can deliberately trigger this, causing the I2C target function to malfunction for all subsequent write transactions and desynchronizing consumer framing state (e.g. MCTP-over-I2C), a recoverable-by-reset denial of service of the target peripheral. The fix unmasks START_DET so the state is reset at every bus (re)START. Impact is availability-only over a local board-level bus; no memory corruption results in the in-tree consumer, whose per-byte buffer write is independently bounds-checked.	2.4	More Details
CVE-	The Activity zte.com.cn.filer/zte.com.cn.filer.FilePreViewActivity within ZTE File Manager is designed to preview compressed files. Third-party applications can launch this Activity and supply arbitrary file paths (e.g.,		

2026-40000	content://zte.com.cn.filer.fileprovider/root_path), enabling file access with the privilege level of ZTE File Manager. This allows unrooted devices to read files under certain system directories such as /data/data and /data/local/tmp. If access restrictions do not block untrusted applications, additional directories may also be accessible.	1.8	More Details
CVE-2026-21047	Out-of-bounds write in ImsService prior to SMR Jul-2026 Release 1 allows remote attackers to potentially execute arbitrary code.	N/A	More Details
CVE-2026-64237	In the Linux kernel, the following vulnerability has been resolved: Input: elan_i2c - validate firmware size before use Ensure that the firmware file is large enough to contain the expected number of pages and the signature (which resides at the end of the firmware blob) before accessing them to prevent potential out-of-bounds reads.	N/A	More Details
CVE-2026-64236	In the Linux kernel, the following vulnerability has been resolved: i2c: davinci: fix division by zero on missing clock-frequency When the 'clock-frequency' property is missing from the device tree, the driver falls back to DAVINCI_I2C_DEFAULT_BUS_FREQ. However, this macro was defined in kHz (100), whereas the device tree property is expected in Hz. The probe function divided the fallback value by 1000, causing integer truncation that resulted in dev->bus_freq = 0. This triggered a deterministic division-by-zero kernel panic when calculating clock dividers later in the probe sequence. Fix this by redefining DAVINCI_I2C_DEFAULT_BUS_FREQ in Hz (100000) to match the expected device tree property unit, allowing the existing division logic to work correctly for both cases.	N/A	More Details
CVE-2026-64234	In the Linux kernel, the following vulnerability has been resolved: tty: serial: pch_uart: add check for dma_alloc_coherent() Add a check for dma_alloc_coherent() failure to prevent a potential NULL pointer dereference in dma_handle_rx(). Properly release DMA channels and the PCI device reference using a goto ladder if the allocation fails.	N/A	More Details
CVE-2026-65880	Joomla Extension - balbooa.com - Unauthenticated remote code execution in Balbooa Forms < 2.4.3 - An insecure form processing logic allowed code execution for forms that include the signature field type.	N/A	More Details
CVE-2026-63303	A Path Traversal vulnerability exists in Quick.CMS through the URI path component of HTTP requests, where the server fails to normalize dot-dot-slash (../) sequences before resolving and serving the requested file. An authenticated attacker with admin privileges can use this vulnerability to read contents of files located in the sibling directory of the webroot via a crafted HTTP request containing ../ sequences in the URI. The vendor assessed the likelihood of exploitation as very low and determined that a fix is not necessary.	N/A	More Details
CVE-2026-63302	Quick.CMS is vulnerable to Local File Inclusion (LFI) in the admin.php endpoint via the p parameter. An authenticated attacker with admin privileges can include arbitrary files located within the application's directory structure via a crafted HTTP request. Successful exploitation allows disclosure of the server's directory structure and absolute file paths (path disclosure). The vendor assessed the likelihood of exploitation as very low and determined that a fix is not necessary.	N/A	More Details
CVE-2026-63301	In Quick.CMS, the administrative user interface restricts deletion of the primary language by omitting the corresponding option from the interface; however, the underlying language-deletion API endpoint does not enforce an equivalent server-side authorization check. As a result, an authenticated administrator can bypass the UI-level restriction and delete the primary language by sending a direct HTTP request to the API endpoint. Successful deletion of the primary language results in a Denial of Service (DoS) of application. Critically, when combined with a separate Cross-Site Request Forgery (CSRF) vulnerability (CVE-2026-1468) an unauthenticated remote attacker can craft a malicious link, which if visited by an authenticated administrator, will trigger the DoS condition without direct access to the application The vendor assessed the likelihood of exploitation as very low and determined that a fix is not necessary.	N/A	More Details
CVE-2026-65600	Traefik versions <= v2.11.51, >= v3.6.0 <= v3.6.22, and >= v3.7.0 <= v3.7.6 contain an authentication bypass via path traversal in the ReplacePathRegex middleware. When ReplacePathRegex is configured with a regex that captures user-controlled path segments without a mandatory path separator (e.g. regex "^/api(.*)", replacement "/$1"), the middleware forwards the replaced path to the backend without validating that it matches its normalized form. An unauthenticated remote attacker can send a crafted request (e.g. GET /api../admin) that produces an un-normalized path such as /../admin, which a backend that normalizes paths resolves to a protected route, bypassing authentication middleware. Fixed in v2.11.52, v3.6.23, and v3.7.7.	N/A	More Details
CVE-2026-65601	Traefik versions 3.7.0 through 3.7.6 contain a namespace confusion vulnerability in the Kubernetes Gateway API provider. When resolving HTTPRoute.spec.rules[].backendRefs[].filters[].extensionRef, Traefik used the backend Service namespace instead of the HTTPRoute namespace. A low-privileged route author holding a ReferenceGrant for a cross-namespace Service could therefore bind a Traefik Middleware from the backend namespace without a separate grant for that middleware, potentially injecting trusted reverse-proxy identity headers into downstream requests. The issue is fixed in version 3.7.7.	N/A	More Details
CVE-2026-64233	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: uvc: hold opts->lock across XU walks in uvc_function_bind uvc_function_bind() walks &opts->extension_units twice without holding opts->lock: - directly, for the iExtension string-descriptor fixup loop; - indirectly, four times via uvc_copy_descriptors() (once per speed), where the helper iterates uvc->desc.extension_units (which aliases &opts->extension_units) to size and emit XU descriptors. The configfs side (uvcg_extension_make / uvcg_extension_drop, in drivers/usb/gadget/function/uvc_configfs.c) takes opts->lock around its list_add_tail / list_del operations. A privileged userspace process that holds the configfs subtree open and writes the gadget UDC name to bind the function while concurrently rmdir()'ing an extensions subdir can race uvcg_extension_drop() against the bind-time list walks and dereference a freed struct uvcg_extension. Hold opts->lock from the start of the XU string-descriptor fixup through the last uvc_copy_descriptors() call, releasing on the descriptor-error path via a new error_unlock label that drops the	N/A	More Details

	lock before falling through to the existing error label. This matches the locking discipline of the configfs callbacks and removes the only remaining unsynchronised reader of the XU list during bind. Reachability: only privileged processes that can mount configfs and write to gadget UDC files can trigger the race, so this is a correctness fix rather than a security boundary.		
CVE-2026-64231	In the Linux kernel, the following vulnerability has been resolved: drm/msm/dsi: don't dump registers past the mapped region On DSI 6G platforms the IO address space is internally adjusted by io_offset. Later this adjusted address might be used for memory dumping. However the size that is used for memory dumping isn't adjusted to account for the io_offset, leading to the potential access to the unmapped region. Lower ctrl_size by the io_offset value to prevent access past the mapped area. msm_disp_snapshot_add_block+0x1d4/0x3c8 [msm] (P) msm_dsi_host_snapshot+0x4c/0x78 [msm] msm_dsi_snapshot+0x28/0x50 [msm] msm_disp_snapshot_capture_state+0x74/0x140 [msm] msm_disp_snapshot_state_sync+0x60/0x90 [msm] _msm_disp_snapshot_work+0x30/0x90 [msm] kthread_worker_fn+0xdc/0x460 kthread+0x120/0x140 Patchwork: https://patchwork.freedesktop.org/patch/721747/	N/A	More Details
CVE-2026-64238	In the Linux kernel, the following vulnerability has been resolved: gpio: shared: fix deadlock on shared proxy's parent removal Commit 710abda58055 ("gpio: shared: call gpio_chip::of_xlate() if set") used the mutex embedded in struct gpio_shared_entry to protect the offset field which now can be modified after assignment. The critical section however is too wide and introduced a potential deadlock on the removal of the shared GPIO proxy's parent. Make the critical section shorter - only protect the offset when it's being read. While at it: mention the fact that the entry lock is now also used to protect against concurrent access to the offset field in the structure's documentation.	N/A	More Details
CVE-2026-47725	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.3, every /ui/* POST / PUT / PATCH / DELETE route processes the request as soon as the session cookie validates. SameSite=Lax on the session cookie prevents most cross-site form submits but does not protect: top-level form-submit navigations from third-party pages (some browsers still send Lax cookies on top-level POSTs); same-registrable-domain attackers (sibling-subdomain XSS, subdomain takeover); the GET /ui/logout route, which a third-party <img src=".../ui/logout"> can force-trigger. This issue has been patched in version 0.3.3.	N/A	More Details
CVE-2026-4648	Use of an insecure cryptographic algorithm in the cashless payment system using NFC wristbands from CasfID Servicios Tecnológicos S.L.U. (version used at Resurrection Fest 2025), which employs cards based on MIFARE Classic technology (FM11RF08S). The cryptographic weakness of the authentication algorithm allows an attacker to retrieve access keys using techniques known as Backdoored Nested Attack, read the wristband’s entire contents, and clone its credentials onto a compatible rewritable card. Exploitation of this vulnerability could enable the impersonation of other attendees, the fraudulent use of the balance associated with their wristbands, and financial losses for both the affected users and the event organizers.	N/A	More Details
CVE-2026-64215	In the Linux kernel, the following vulnerability has been resolved: drm/msm/a6xx: Check kzalloc return in a8xx_hfi_send_perf_table Check the return value of kzalloc() to prevent a NULL pointer dereference on allocation failure. Patchwork: https://patchwork.freedesktop.org/patch/721342/	N/A	More Details
CVE-2026-55732	Out-of-bounds Read (CWE-125) in BACnet packet parsing (`bacdt_datetime_to_tod`) in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.18 on LINX-A64 allows an unauthenticated remote attacker to crash `linx_a64.exe` and ultimately reboot the device via a malformed BACnet TimeSynchronization or UTC-TimeSynchronization packet with an invalid month value. The same vulnerability affects multiple other Loytec products.	N/A	More Details
CVE-2026-16270	Open Mercato does not validate regex rules. An attacker with privileges to create the regex rule can add an unsafe regex to a field. When someone provide the proper string it can result in a DoS attack. This issue was fixed in version 0.6.4.	N/A	More Details
CVE-2026-64209	In the Linux kernel, the following vulnerability has been resolved: phy: qcom: qmp-usbc: Fix out-of-bounds array access in dp swing config swing_tbl and pre_emphasis_tbl are 4x4 arrays (valid indices 0-3), but the boundary check uses "> 4" instead of ">= 4", allowing index 4 to cause an out-of-bounds access.	N/A	More Details
CVE-2026-64211	In the Linux kernel, the following vulnerability has been resolved: srcu: Don't queue workqueue handlers to never-online CPUs While an srcu_struct structure is in the midst of switching from CPU-0 to all-CPUs state, it can attempt to invoke callbacks for CPUs that have never been online. Worse yet, it can attempt in invoke callbacks for CPUs that never will be online, even including imaginary CPUs not in cpu_possible_mask. This can cause hangs on s390, which is not set up to deal with workqueue handlers being scheduled on such CPUs. This commit therefore causes Tree SRCU to refrain from queueing workqueue handlers on CPUs that have not yet (and might never) come online. Because callbacks are not invoked on CPUs that have not been online, it is an error to invoke call_srcu(), synchronize_srcu(), or synchronize_srcu_expedited() on a CPU that is not yet fully online. However, it turns out to be less code to redirect the callbacks from too-early invocations of call_srcu() than to warn about such invocations. This commit therefore also redirects callbacks queued on not-yet-fully-online CPUs to the boot CPU.	N/A	More Details
CVE-2026-64212	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mld: don't dereference a pointer before NULL checking it In iwl_mld_remove_link, the link->fw_id is saved at the beginning of the function so we have it after we freed the link. But the link pointer can be NULL, and is not checked when the fw_id is stored. Fix it by simply freeing the link at the end of the function. fFixes: 0e66a39f4f0e ("wifi: iwlwifi: fix potential use after free in iwl_mld_remove_link()")	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: hwmon: (lm90) Add lock protection to lm90_alert Sashiko reports: lm90_alert() executes in the smbus alert context and calls lm90_update_confreg() to disable the hardware alert line, without acquiring hwmon_lock. Concurrently, sysfs write operations (such as lm90_write_convrate) hold the hwmon_lock, temporarily modify data->config, and then restore it. If an alert interrupt	N/A	More Details

64213	occurs concurrently with a sysfs write, the sysfs path will overwrite the alert handler's modifications to data->config and the hardware register. This unintentionally re-enables the hardware alert line while the alarm is still active, causing an interrupt storm. Add the missing lock to lm90_alert() to solve the problem.		
CVE-2026-64214	In the Linux kernel, the following vulnerability has been resolved: powerpc/time: Remove redundant preempt_disable|enable() calls from arch_irq_work_raise() A kernel panic is observed when handling machine check exceptions from real mode. BUG: Unable to handle kernel data access on read at 0xc00000006be21300 Oops: Kernel access of bad area, sig: 11 [#1] MSR: 8000000000001003 <SF,ME,RI,LE> CR: 88222248 XER: 00000005 CFAR: c00000000003ffc4 DAR: c00000006be21300 DSISR: 40000000 IRQMASK: 0 NIP [c000000000029e40] arch_irq_work_raise+0x10/0x70 LR [c00000000003ffc8] machine_check_queue_event+0xa8/0x150 Call Trace: [c0000000179d3c70] [c00000000003ff64] machine_check_queue_event+0x44/0x150 [c0000000179d3d30] [c0000000000084e0] machine_check_early_common+0x1f0/0x2c0 The crash occurs because arch_irq_work_raise() calls preempt_disable() from machine check exception (MCE) handlers running in real mode. In this context, accessing the preempt_count can fault, leading to the panic. The preempt_disable()/preempt_enable() pair in arch_irq_work_raise() was originally added by commit 0fe1ac48bef0 ("powerpc/perf_event: Fix oops due to perf_event_do_pending call") to avoid races while raising irq work from exception context. Later, commit 471ba0e686cb ("irq_work: Do not raise an IPI when queueing work on the local CPU") added preemption protection in irq_work_queue() path, while commit 20b876918c06 ("irq_work: Use per cpu atomics instead of regular atomics") added equivalent protection in irq_work_queue_on() before reaching arch_irq_work_raise(): irq_work_queue() / irq_work_queue_on() -> preempt_disable() -> __irq_work_queue_local() -> irq_work_raise() -> arch_irq_work_raise() As a result, callers other than mce_irq_work_raise() already execute with preemption disabled, making the additional preempt_disable()/preempt_enable() pair in arch_irq_work_raise() redundant. The arch_irq_work_raise() function executes in NMI context when called from MCE handler. Hence we will not be preempted or scheduled out since we are in NMI context with MSR[EE]=0. Therefore, it is safe to remove the preempt_disable()/preempt_enable() calls from here. Remove it to avoid accessing preempt_count from real mode context. [Maddy: Fixed the commit title]	N/A	More Details
CVE-2026-64220	In the Linux kernel, the following vulnerability has been resolved: device property: set fwnode->secondary to NULL in fwnode_init() If a firmware node is allocated on the stack (for instance: temporary software node whose life-time we control) or on the heap - but using a non-zeroing allocation function - and initialized using fwnode_init(), its secondary pointer will contain uninitalized memory which likely will be neither NULL nor IS_ERR() and so may end up being dereferenced (for example: in dev_to_swnode()). Set fwnode->secondary to NULL on initialization.	N/A	More Details
CVE-2026-64230	In the Linux kernel, the following vulnerability has been resolved: regulator: tps65219: fix irq_data.rdev not being assigned Commit 64a6b577490c ("regulator: tps65219: Remove debugging helper function") removed the tps65219_get_rdev_by_name() helper along with the irq_data.rdev assignment that depended on it. This left irq_data.rdev uninitialized for all IRQs, causing undefined behavior when regulator_notifier_call_chain() is called from the IRQ handler: Internal error: Oops: 0000000096000004 pc : regulator_notifier_call_chain lr : tps65219_regulator_irq_handler Call trace: regulator_notifier_call_chain tps65219_regulator_irq_handler handle_nested_irq regmap_irq_thread irq_thread_fn irq_thread kthread ret_from_fork Instead of restoring a dedicated lookup array, restructure the probe function to combine regulator registration with IRQ registration in the same loop. This way the rdev returned by devm_regulator_register() is naturally available for assigning to irq_data.rdev without any auxiliary data structure. Non-regulator IRQs (SENSOR, TIMEOUT) that don't correspond to any registered regulator are registered with rdev=NULL, and the IRQ handler is protected with a NULL check to avoid crashing.	N/A	More Details
CVE-2026-64224	In the Linux kernel, the following vulnerability has been resolved: octeontx2-pf: fix double free in rvu_rep_rsrc_init() rvu_rep_rsrc_init() allocates queue memory before calling otx2_init_hw_resources(). When hardware resource setup fails, otx2_init_hw_resources() already unwinds the partially initialized SQ, CQ, and aura state before returning an error. The representor error path then calls otx2_free_hw_resources() again and can free the same resources a second time. Fix this by splitting the cleanup labels so that a failure from otx2_init_hw_resources() only releases queue memory. Keep the otx2_free_hw_resources() call for failures that happen after hardware resource initialization completed successfully. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1-rc3. Runtime validation was not performed because reproducing this path requires OcteonTX2 representor hardware.	N/A	More Details
CVE-2026-64225	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: CGX: add bounds check to cgx_speed_mbps index cgx_speed_mbps has 13 elements but RESP_LINKSTAT_SPEED can yield values 0-15. If it returns a value >= 13, this causes an out-of-bounds array access. Add a bounds check and default to speed 0 if the index is out of range.	N/A	More Details
CVE-2026-64227	In the Linux kernel, the following vulnerability has been resolved: ACPI: driver: Check ACPI_COMPANION() against NULL during probe Since every platform driver can be forced to match a device that doesn't match its list of device IDs because of device_match_driver_override(), platform drivers that rely on the existence of a device's ACPI companion object should verify its presence. Accordingly, add requisite ACPI_COMPANION() or ACPI_HANDLE() checks against NULL to 13 platform drivers handling core ACPI devices. Also change the value returned by the ACPI thermal zone driver when the device's ACPI companion is not present to -ENODEV for consistency with the other drivers.	N/A	More Details
CVE-2026-41874	Quick.Cart stores hard-coded, plaintext admin credentials in a configuration file. This flaw allows attackers with access to the server file system to retrieve authentication details, potentially leading to privilege escalation. The vendor assessed the likelihood of exploitation as very low and determined that a fix is not necessary. Only version 6.7 was tested but all versions should be considered as vulnerable.	N/A	More Details
CVE-2026-18029	Our payment integration with GiroCheckout did not properly validate payment status responses. An attacker could use a successful payment status response from one payment and supply it to the system for a different payment, gaining access to multiple valid tickets with only one payment.	N/A	More Details

CVE-2026-65602	Traefik 3.6.0 through 3.6.22 and 3.7.0 through 3.7.6 fail to enforce the crossProviderNamespaces allowlist for IngressRouteTCP service serversTransport references (the allowlist was only enforced for HTTP serversTransport references). A low-privileged Kubernetes user in a namespace not listed in crossProviderNamespaces can set serversTransport: foo@file on an IngressRouteTCP service, causing Traefik to accept the forbidden cross-provider reference and use a file-provider TCPServersTransport — including privileged backend mTLS client certificates, SPIFFE identity, or PROXY-protocol settings. This is fixed in 3.6.23 and 3.7.7.	N/A	More Details
CVE-2026-64229	In the Linux kernel, the following vulnerability has been resolved: x86/mm: Disable broadcast TLB flush when PCID is disabled Booting with "nopcid" clears X86_FEATURE_PCID and keeps CR4.PCIDE from being set to one. On AMD CPUs that support INVLPGB, broadcast TLB flushing remains enabled. There are two checks that decide whether the global ASID code runs, mm_global_asid() and consider_global_asid(), that key off of the X86_FEATURE_INVLPGB feature. Once an mm becomes active on more than three CPUs, consider_global_asid() assigns it a global ASID, after which flush_tlb_mm_range() takes the broadcast_tlb_flush() path using a non-zero PCID. Issuing an INVLPGB with a non-zero PCID while CR4.PCIDE is not set results in a #GP: Oops: general protection fault, kernel NULL pointer dereference 0x1: 0000 [#1] SMP NOPTI CPU: 158 UID: 0 PID: 3119 Comm: snap Not tainted 7.1.0-rc3 #1 PREEMPT(full) Hardware name: ... RIP: 0010:broadcast_tlb_flush Code: ... 89 da 48 83 c8 07 <0f> 01 fe eb 08 cc cc cc ... Call Trace: <TASK> flush_tlb_mm_range ptep_clear_flush wp_page_copy ? _raw_spin_unlock __handle_mm_fault handle_mm_fault do_user_addr_fault exc_page_fault asm_exc_page_fault All processors that support broadcast TLB invalidation also have PCID support, so it is only the "nopcid" scenario that is of concern. In this situation just disable the broadcast TLB support using the CPUID dependency support by making X86_FEATURE_INVLPGB dependent on X86_FEATURE_PCID. [bp: Massage commit message.]	N/A	More Details
CVE-2026-64228	In the Linux kernel, the following vulnerability has been resolved: net: ethtool: phy: avoid NULL deref when PHY driver is unbound phydev->drv can become NULL while the phy_device is still attached to its net_device, namely after the PHY driver is unbound via sysfs: echo <mdio_id> > /sys/bus/mdio_bus/drivers/<phy_drv>/unbind phy_remove() clears phydev->drv but doesn't call phy_detach(), so the phy_device stays in the link topology xarray and ethnl_req_get_phydev() still hands it back. ETHTOOL_MSG_PHY_GET then oopses on: rep_data->drvname = kstrdup(phydev->drv->name, GFP_KERNEL); drvname is already treated as optional by phy_reply_size(), phy_fill_reply() and phy_cleanup_data(), so just skip the allocation when there is no driver bound.	N/A	More Details
CVE-2026-64239	In the Linux kernel, the following vulnerability has been resolved: mm/damon/sysfs-schemes: delete tried region in regions_rmdirs() DAMON sysfs maintains the DAMOS tried region directory objects via a linked list. When the user requests refresh of the directories, DAMON sysfs removes all the region directories first, and then generate updated regions directory on the empty space. The removal function (damon_sysfs_scheme_regions_rm_dirs()) only puts the kobj objects. Deletion of the container region object from the linked list is done inside the kobj release callback function. If somehow the callback invocation is delayed, the list will contain regions list that gonna be freed. If the updated region directories creation is started in this situation, the list can be corrupted and use-after-free can happen. Because the kobj objects are managed by only DAMON sysfs, the issue cannot happen in normal situation. But, such delays can be made on kernels that built with CONFIG_DEBUG_KOBJECT_RELEASE. On the kernel, the issue can indeed be reproduced like below. # damo start --damos_action stat # cd /sys/kernel/mm/damon/admin/kdamonds/0/ # for i in {1..10}; do echo update_schemes_tried_regions > state; done # dmesg | grep underflow [89.296152] refcount_t: underflow; use-after-free. Fix the issue by removing the region object from the list when decrementing the reference count. Also update damos_sysfs_populate_region_dir() to add the region object to the list only after the kobject_init_and_add() is success, so that fail of kobject_init_and_add() is not leaving the deallocated object on the list. The issue was discovered [1] by Sashiko.	N/A	More Details
CVE-2026-18028	The "quick setup" view presented to users after they first create an event allows to set up the most critical parts of an event in just a few clicks. This view did not properly check that the user has permission to change configuration for the given event. An attacker could use a well-timed request to create products, quotas, set bank transfer configuration, or connect a stripe account to an event they do not have access to.	N/A	More Details
CVE-2026-63264	Joomla Extension - joomshopping.com - Reflective XSS in JoomShopping < 5.9.3 - The Joomla extension JoomShopping is vulnerable to an reflected XSS vulnerability in the product frontend controller.	N/A	More Details
CVE-2026-65623	Inefficient Algorithmic Complexity vulnerability in mtrudel bandit allows unauthenticated remote denial of service via CPU exhaustion during WebSocket fragment reassembly. The size guard 'Elixir.Bandit.WebSocket.Connection':oversize_message?/2 called from handle_frame/3 in lib/bandit/websocket/connection.ex appends each non-final continuation frame to a left-nested iolist and then re-measures the entire accumulated buffer with IO.iodata_length/1 on every frame. Because the buffer grows by one element per frame and is fully re-traversed each time, reassembly work is quadratic (O(n^2)) in the number of continuation frames. The max_fragmented_message_size limit (default 8 MB) bounds total bytes but not frame count, and each frame can carry as little as one payload byte, so an attacker can send millions of tiny continuation frames using modest bandwidth to pin a CPU core for minutes to hours. Many concurrent connections can starve the whole server of CPU, denying service to legitimate users. The WebSocket read timeout does not help, because it is an idle timeout evaluated between reads and cannot preempt the synchronous reassembly work spent inside a single callback. This issue affects bandit: from 1.11.0 before 1.12.1.	N/A	More Details
CVE-2026-48032	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, IAM-role policy checks can be bypassed when the role trusts multiple OIDC providers. This issue has been patched in version 1.4.0.	N/A	More Details
CVE-2026-48033	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, policy packs can be bypassed by a forged Pulumi-URN logical name. This issue has been patched in version 1.4.0.	N/A	More Details

CVE-2024-14041	In Bouncy Castle for Java from 1.73 to before 1.78, three ML-KEM (CRYSTALS-Kyber) routines divided secret-derived polynomial coefficients by the modulus q: Poly.toMsg, which decodes the decrypted message, and the ciphertext compression routines Poly.compressPoly and PolyVec.compressPolyVec. An attacker able to measure the timing of a large number of decapsulations performed with the same long-term private key can recover that key. These are the KyberSlash1 (Poly.toMsg) and KyberSlash2 (ciphertext compression) divisions. Compression performed during encapsulation operates on values that become the public ciphertext and is not affected.	N/A	More Details
CVE-2026-48034	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, there is a bypass via decoy sibling resources targeting a different bucket. This issue has been patched in version 1.4.0.	N/A	More Details
CVE-2026-16551	Denial-of-Service in Thinkst Applied Research OpenCanary (MongoDB module) allows Excessive Allocation. This issue affects OpenCanary 0.9.8 only.	N/A	More Details
CVE-2026-63048	Joomla Extension - joomlack.fr - Improper access control in Page Builder CK < 3.6.2 - The Joomla extension Page Builder CK is vulnerable to an authenticated arbitrary file upload, leading to RCE.	N/A	More Details
CVE-2026-64253	In the Linux kernel, the following vulnerability has been resolved: kernel/fork: clear PF_BLOCK_TS in copy_process() PF_BLOCK_TS is only set in blk_time_get_ns() when current->plug is non-NULL, and blk_finish_plug() clears it via __blk_flush_plug() before NULLing the plug pointer. copy_process() breaks the invariant by inheriting PF_BLOCK_TS from the parent while resetting the child's plug to NULL. Clear PF_BLOCK_TS alongside that assignment so callers can rely on "PF_BLOCK_TS set implies current->plug != NULL" and dereference current->plug unguarded.	N/A	More Details
CVE-2026-48035	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, consumers using AccountFoundation could ship an AWS account whose CloudTrail / Config audit logs were deletable by any S3-delete-capable principal — while believing the startup-hardened tier guaranteed tamper-resistance. Sandbox-tier deployments had no audit immutability at all (defects 1 and 3 compounded). This issue has been patched in version 1.4.0.	N/A	More Details
CVE-2026-54659	Pagy is agnostic pagination in plain Ruby. From 43.0.0 until 43.5.6, Pagy::I18n.locale= in gem/lib/pagy/modules/i18n/i18n.rb stored locale values verbatim and later used them as <locale>.yml path components, allowing untrusted params[:locale] values with absolute paths or ../ sequences to create a file existence and readability oracle for YAML files. This issue is fixed in version 43.5.6.	N/A	More Details
CVE-2026-48036	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, consumers running drift detection in CI / cron could see transient adapter failures silently cached as "all clear" — masking real attacks for up to six hours — or see ordinary provider-version churn falsely promoted to incident severity. Either way, the verdict source was unreliable for downstream incident workflows that gate on it. This issue has been patched in version 1.4.0.	N/A	More Details
CVE-2026-45820	fflate through 0.8.2 is vulnerable to denial of service via an infinite loop in unzipSync(). A crafted ZIP archive with a central directory entry declaring compressed_size=0xFFFFFFFF (ZIP64 sentinel) but missing the required ZIP64 extra field tag 0x0001 causes z64e() to loop indefinitely due to out-of-bounds reads returning undefined, which coerces to 0, keeping the loop condition permanently true.	N/A	More Details
CVE-2026-48037	Hulumi is an open-source toolkit that ships secure-by-default cloud and platform infrastructure components for Pulumi. Prior to version 1.4.0, AccountFoundation reuse paths silently downgrade GuardDuty / Security Hub posture. This issue has been patched in version 1.4.0.	N/A	More Details
CVE-2026-56844	A vulnerability in the Veeam Updater component of the Veeam Software Appliance that could allow a local user to elevate their privileges and gain root-level access to the underlying operating system.	N/A	More Details
CVE-2026-64254	In the Linux kernel, the following vulnerability has been resolved: NTB: epf: Avoid pci_iounmap() with offset when PEER_SPAD and CONFIG share BAR When BAR_PEER_SPAD and BAR_CONFIG share one PCI BAR, the module teardown path ends up calling pci_iounmap() on the same iomem with some offset, which is unnecessary and triggers a kernel warning like the following: Trying to vunmap() nonexistent vm area (0000000069a5ffe8) WARNING: mm/vmalloc.c:3470 at vunmap+0x58/0x68, CPU#5: modprobe/2937 [...] Call trace: vunmap+0x58/0x68 (P) iounmap+0x34/0x48 pci_iounmap+0x2c/0x40 ntb_epf_pci_remove+0x44/0x80 [ntb_hw_epf] pci_device_remove+0x48/0xf8 device_remove+0x50/0x88 device_release_driver_internal+0x1c8/0x228 driver_detach+0x50/0xb0 bus_remove_driver+0x74/0x100 driver_unregister+0x34/0x68 pci_unregister_driver+0x34/0xa0 ntb_epf_pci_driver_exit+0x14/0xfe0 [ntb_hw_epf] [...] Fix it by unmapping only when PEER_SPAD and CONFIG use difference bars.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: MIPS: DEC: Prevent initial console buffer from landing in XKPHYS In 64-bit configurations calling the initial console output handler from a kernel thread other than the initial one will result in a situation where the stack has been placed in the XKPHYS 64-bit memory segment and consequently so has been the buffer allocated there that is used as the argument corresponding to the `%s' output conversion specifier for the firmware's printf() entry point. This 64-bit address will then be truncated by 32-bit firmware, resulting in an attempt to access the wrong memory location, which in turn will cause all kinds of unpredictable behaviour, such as a kernel crash: Console: colour dummy device 160x64 Calibrating delay loop... 49.36 BogoMIPS (lpj=192512) pid_max: default: 32768 minimum: 301 CPU 0 Unable to handle kernel paging request at virtual address 000000000203bd00, epc == ffffffffbfc08364, ra == ffffffffbfc08800 Oops[#1]: CPU: 0 PID: 0 Comm: swapper Not tainted 5.18.0-rc2-00254-gfb649bda6f56-dirty #121 $ 0 : 0000000000000000 0000000000000001		

CVE-2026-64252	0000000000000023 ffffffff80684ba0 $ 4 : 000000000203bd00 ffffffffbfc0f3b4 ffffffffffffffff 0000000000000073 $ 8 : 0a303d7469000000 0000000000000000 0000000000000073 ffffffffbfc0f473 $12 : 0000000000000002 0000000000000000 ffffffff80684c1c 0000000000000000 $16 : 0000000000000000 ffffffff80596dc9 0000000000000000 ffffffffbfc09240 $20 : ffffffff80684c40 ffffffffbfc0f400 000000000000002d 000000000000002b $24 : ffffffffffffffbf 000000000203bd00 $28 : ffffffff805f0000 ffffffff80684b58 0000000000000030 ffffffffbfc08800 Hi : 0000000000000000 Lo : 0000000000000aa8 epc : ffffffffbfc08364 0xffffffffbfc08364 ra : ffffffffbfc08800 0xffffffffbfc08800 Status: 140120e2 KX SX UX KERNEL EXL Cause : 00000008 (ExcCode 02) BadVA : 000000000203bd00 PrId : 00000430 (R4000SC) Modules linked in: Process swapper (pid: 0, threadinfo=(____ptrval____), task=(____ptrval____), tls=0000000000000000) Stack : 0000000000000000 0000000000000000 0000000000000000 0000004d0000004d 80684cc0806a2a40 80596dc80000004d 8061000000000000 bfc0850c80684c38 0000000000000000 000000000203bd00 0000000000000000 0000000000000000 0000000000000000 00000000bfc0f3b4 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000002500000000 0000000000000000 0000000000000000 802c1a7400000000 0203bd0080596dc8 0203bd4d69000000 6c61632000000018 5f746567646e6172 6c616320625f6d6f 5f736e5f6d6f7266 206361323778302b 303d74696e726320 806a0a38806b0000 806a0a38806b0000 00000000806b0000 80683c58806b0000 ... Call Trace: Code: a082ffff 03e00008 00601021 <80820000> 00001821 10400005 24840001 80820000 24630001 ---[end trace 0000000000000000]--- Kernel panic - not syncing: Fatal exception in interrupt KN04 V2.1k (PC: 0xa0026768, SP: 0x806848e8) >> In this case the pointer in $4 was truncated from 0x980000000203bd00 to 0x000000000203bd00. This may happen when no final console driver has been enabled in the configuration and consequently the initial console continues being used late into bootstrap or with an upcoming change that will switch the zs driver to use a platform device, which in turn will make the console handover happen only after other kernel threads have already been started. Fix the issue by making the buffer static and initdata, and therefore placed in the CKSEG0 32-bit compatibility segment, observing that the console output handler is called with the console lock held, implying no need for this code to be reentrant. Add an assertion to verify the buffer actually has been placed in a compatibility segment.	N/A	More Details
CVE-2026-65624	Allocation of Resources Without Limits or Throttling vulnerability in ninenines cowboy allows an unauthenticated remote attacker to exhaust connection process memory over HTTP/1.1. The HTTP/1.1 handler in cowboy_http enforces the max_headers limit by counting the number of distinct header names in a map (maps:size(Headers)). When a request contains multiple header lines with the same name, the values are concatenated into a single ever-growing binary stored under that one map key (", " for regular headers, "; " for cookies), so the map size stays at one and the max_headers cap (default 100) is never reached. Because no accumulator bounds the total number of header lines or the total byte size of the header block (only per-line max_header_name_length and max_header_value_length apply), an unauthenticated client can send an arbitrary number of header lines with the same name and grow the connection process's binary memory to arbitrary size within the request window. The impact per connection is bounded by request_timeout (default 5 seconds, not reset by header data), and by max_heap_size when set (the offending connection process is killed once its heap grows past the limit). When max_heap_size is left at the default (unset), sustained abuse can drive the Erlang VM into out-of-memory conditions. This issue affects cowboy from 2.0.0-pre.4 before 2.18.0.	N/A	More Details
CVE-2026-64245	In the Linux kernel, the following vulnerability has been resolved: fbdev: modedb: fix a possible UAF in fb_find_mode() If mode_option is NULL, it is assigned from mode_option_buf: if (!mode_option) { fb_get_options(NULL, &mode_option_buf); mode_option = mode_option_buf; } Later, name is assigned from mode_option: const char *name = mode_option; However, mode_option_buf is freed before name is no longer used: kfree(mode_option_buf); while name is still accessed by: if ((name_matches(db[i], name, namelen) || Since name aliases mode_option_buf, this may result in a use-after-free. Fix this by extending the lifetime of mode_option_buf until the end of the function by using scope-based resource management for cleanup.	N/A	More Details
CVE-2026-59248	Allocation of resources without limits vulnerability in ninenines cowlib allows an unauthenticated remote HTTP/2 or HTTP/3 peer to exhaust memory on the vulnerable server (or client) and cause a denial of service. The HPACK and QPACK prefixed-integer decoder cow_hpack_common:dec_big_int/3 in src/cow_hpack_common.hrl (invoked from cow_hpack:decode/2 in src/cow_hpack.erl and from cow_qpack:decode_field_section/3 in src/cow_qpack.erl) reads continuation octets until it sees one whose high bit is clear, evaluating Int + (Value bsl M) at each step with the shift M growing by seven per octet. No limit is enforced on the number of continuation octets, on the resulting bit width, or on the value; the decoder consumes whatever encoded length the peer supplies. Because Erlang integers are immutable, each intermediate Value bsl M and each accumulator update allocates a fresh bignum whose digit width grows linearly with the number of octets processed so far. Summed across the whole decode, the transient bignum digit materialization is on the order of the square of the encoded length. A single maximal HPACK indexed representation carried inside one HTTP/2 HEADERS plus one CONTINUATION frame at Cowboy's default max_frame_size_received can force hundreds of megabytes of transient allocation and garbage-collection churn before the resulting header-table index is rejected as invalid. Repeated or concurrent connections multiply the pressure and can drive the Erlang VM to memory exhaustion. Cowlib is the HTTP parser used by Cowboy, RabbitMQ's management plugin, and other Erlang and Elixir HTTP/2 and HTTP/3 servers and clients, so any exposed endpoint that accepts HPACK or QPACK from an untrusted peer is reachable. This issue affects cowlib: from 2.0.0 before 2.19.0.	N/A	More Details
CVE-2026-55730	Reflected Cross-Site Scripting (CWE-79) in LWEB802 in Loytec LWEB-802 before 5.0.8 on all platforms allows an unauthenticated remote attacker to execute arbitrary JavaScript in a victim's browser and perform actions with the victim's privileges via a crafted link containing a malicious `project` or `mspParams` parameter.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: media: rc: igorplugusb: fix control request setup packet Commit eac69475b01f ("media: rc: igorplugusb: heed coherency rules") changed the control request storage from an embedded struct to an allocated pointer so it can obey DMA coherency rules. However, the driver still passes &ir->request to usb_fill_control_urb(). That points the URB setup packet at the pointer field itself rather than at the	N/A	More Details

64240	allocated struct usb_ctrlrequest. USB core then interprets pointer bytes as the setup packet. This can produce an invalid bRequestType and trigger the control direction warning reported by syzbot: usb 2-1: BOGUS control dir, pipe 80003580 doesn't match bRequestType 0 Pass ir->request itself as the setup packet.		
CVE-2026-64241	In the Linux kernel, the following vulnerability has been resolved: gpio: rockchip: teardown bugs and resource leaks Address several teardown issues and resource leaks in the driver's remove path and error handling: 1. Debounce clock reference leak: The debounce clock (bank->db_clk) is obtained using of_clk_get() which increments the clock's reference count, but clk_put() is never called. Register a devm action to cleanly release it on unbind. Note that of_clk_get(..., 1) remains necessary over devm_clk_get() because the DT binding does not define clock-names, precluding name-based lookup. 2. Unregistered chained IRQ handler: The chained IRQ handler is not disconnected in remove(). If a stray interrupt fires after the driver is removed, the kernel attempts to execute a stale handler, leading to a panic. Fix this by clearing the handler in remove(). 3. IRQ domain leak: The linear IRQ domain and its generic chips are allocated manually during probe but never removed. Remove the IRQ domain during driver teardown to free the associated generic chips and mappings. [Bartosz: don't emit an error message on devres allocation failure]	N/A	More Details
CVE-2026-64242	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: net2280: Fix double free in probe error path usb_initialize_gadget() installs gadget_release() as the release callback for the embedded gadget device. The struct net2280 instance is therefore released through gadget_release() when the gadget device's last reference is dropped. The probe error path calls net2280_remove(), which tears down the partially initialized device and drops the gadget reference with usb_put_gadget(). Calling kfree(dev) afterwards can free the same object again. Drop the explicit kfree() and let the gadget device release callback handle the final free. This issue was found by a static analysis tool I am developing.	N/A	More Details
CVE-2026-64244	In the Linux kernel, the following vulnerability has been resolved: drivers/base/memory: set mem->altmap after successful device registration If __add_memory_block() fails at xa_store() (under memory pressure for example), device_unregister() is called, which eventually triggers memory_block_release() with mem->altmap still set, causing a WARN_ON(mem->altmap). This was triggered by modifying virtio-mem driver. Fix this by delaying the assignment of mem->altmap until after __add_memory_block() has succeeded.	N/A	More Details
CVE-2026-61376	ELECOM wireless LAN routers and access points devices contain an OS Command Injection vulnerability in Restore Settings. If this vulnerability is exploited, an arbitrary OS command may be executed by an attacker who can log in to the product.	N/A	More Details
CVE-2026-64250	In the Linux kernel, the following vulnerability has been resolved: LoongArch: Report dying CPU to RCU in stop_this_cpu() This is a port of MIPS commit 9f3f3bdc6d9dac1 ("MIPS: smp: report dying CPU to RCU in stop_this_cpu()"). smp_send_stop() parks all secondary CPUs in stop_this_cpu(). And the function marks the CPU offline for the scheduler via set_cpu_online(false) but never informs RCU, so RCU keeps expecting a quiescent state from CPUs that are now spinning forever with interrupts disabled. As long as nothing waits for an RCU grace period after smp_send_stop() this is harmless, which is why it went unnoticed. However, since commit 91840be8f710370 ("irq_work: Fix use-after-free in irq_work_single() on PREEMPT_RT"), irq_work_sync() calls synchronize_rcu() on architectures without an irq_work self-IPI, i.e. where arch_irq_work_has_interrupt() returns false. Any irq_work_sync() issued in the reboot/shutdown/halt path after smp_send_stop() then blocks on a grace period that can never complete, hanging the reboot: WARNING: CPU: 0 PID: 15 at kernel/irq_work.c:144 irq_work_queue_on ... rcu: INFO: rcu_sched detected stalls on CPUs/tasks: rcu: Offline CPU 1 blocking current GP. rcu: Offline CPU 2 blocking current GP. rcu: Offline CPU 3 blocking current GP. This issue needs some hacks to reproduce, and it was not noticed on LoongArch because arch_irq_work_has_interrupt() usually returns true. Call rcutree_report_cpu_dead() once interrupts are disabled, mirroring the generic CPU-hotplug offline path, so RCU stops waiting on the parked CPUs and grace periods can still complete. LoongArch shuts down all CPUs here without going through the CPU-hotplug mechanism, so this report is not otherwise issued.	N/A	More Details
CVE-2026-59764	ELECOM wireless LAN routers and access points devices contain an OS Command Injection vulnerability in WebUI. If this vulnerability is exploited, an arbitrary OS command may be executed by an attacker who can log in to the product.	N/A	More Details
CVE-2026-44387	ELECOM wireless LAN routers and access points devices contain a reflected cross-site scripting vulnerability in WebUI. If this vulnerability is exploited, an arbitrary script may be executed on a logged-in user's web browser.	N/A	More Details
CVE-2026-64246	In the Linux kernel, the following vulnerability has been resolved: power: reset: linkstation-poweroff: fix use-after-free in the linkstation_poweroff_init() Move of_node_put(dn) after the of_match_node() call, which still needs the node pointer. The node reference is correctly released after use.	N/A	More Details
CVE-2026-64248	In the Linux kernel, the following vulnerability has been resolved: MIPS: smp: report dying CPU to RCU in stop_this_cpu() smp_send_stop() parks all secondary CPUs in stop_this_cpu(). The function marks the CPU offline for the scheduler via set_cpu_online(false) but never informs RCU, so RCU keeps expecting a quiescent state from CPUs that are now spinning forever with interrupts disabled. As long as nothing waits for an RCU grace period after smp_send_stop() this is harmless, which is why it went unnoticed. Since commit 91840be8f710 ("irq_work: Fix use-after-free in irq_work_single() on PREEMPT_RT") however, irq_work_sync() calls synchronize_rcu() on architectures without an irq_work self-IPI, i.e. where arch_irq_work_has_interrupt() returns false. That is the asm-generic default used by MIPS. Any irq_work_sync() issued in the reboot/shutdown path after smp_send_stop() then blocks on a grace period that can never complete, hanging the reboot: WARNING: CPU: 0 PID: 15 at kernel/irq_work.c:144 irq_work_queue_on ... rcu: INFO: rcu_sched detected stalls on CPUs/tasks: rcu: Offline CPU 1 blocking current GP. rcu: Offline CPU 2 blocking current GP. rcu: Offline CPU 3 blocking current GP. This issue was noticed on several Realtek MIPS switch SoCs (MIPS interAptiv) and came up during kernel bump downstream in OpenWrt from 6.18.33 to 6.18.34, after the backport of the patch to the 6.18 stable branch. The patch also has been backported all the way back to 6.1. Call rcutree_report_cpu_dead() once interrupts are disabled, mirroring the generic CPU-hotplug offline	N/A	More Details

	path, so RCU stops waiting on the parked CPUs and grace periods can still complete. MIPS shuts down all CPUs here without going through the CPU-hotplug mechanism, so this report is not otherwise issued. Reporting a dying CPU to RCU outside the regular hotplug offline path is not unprecedented: arm64 does the same in cpu_die_early(). There it is an exception for a CPU that was coming online and is aborting bringup, rather than the default shutdown action as on MIPS.		
CVE-2026-64249	In the Linux kernel, the following vulnerability has been resolved: fpga: region: fix use-after-free in child_regions_with_firmware() Move of_node_put(child_region) after the error print to avoid accessing freed memory when pr_err() references child_region. [Yilun: Fix the Fixes tag]	N/A	More Details
CVE-2026-14170	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-55731	Unchecked input for loop condition (CWE-606) in the SNMP agent in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows an unauthenticated remote attacker to cause persistent denial of service (CPU exhaustion) via a crafted SNMP GETNEXT request with a large OID component.	N/A	More Details
CVE-2026-8152	Unblu Spark contains an open redirect vulnerability that can be escalated to a DOM-based cross-site scripting (XSS) attack. When Unblu Spark is deployed with com.unblu.identifier.siteEmbeddedSetup=true, it runs in the same origin as the host application. Any JavaScript injected through this vulnerability therefore executes with full access to the host application's cookies, DOM, and same-origin APIs — an attacker can reach all resources of the host application, not just Unblu's. This expanded blast radius is the reason on-premises deployments using this configuration are rated CRITICAL.	N/A	More Details
CVE-2026-55729	Exposure of Sensitive Information (CWE-200) in LWEB802 browser `localStorage` in Loytec LWEB-802 before 5.0.8 on all platforms allows an unauthenticated remote attacker to leak stored management credentials via a crafted link.	N/A	More Details
CVE-2026-47669	DbGate is cross-platform database manager. In versions 7.1.8 and prior, the `unzipDirectory()` function in `packages/api/src/shell/unzipDirectory.js` (line 27) does not validate that extracted file paths stay within the output directory. A malicious ZIP with `../` entries writes files anywhere on the filesystem. In the default Docker deployment, DbGate runs as root and the `none` auth provider issues JWT tokens without credentials via `POST /auth/login`, so this is exploitable by any network-adjacent attacker. Version 7.1.9 fixes the issue.	N/A	More Details
CVE-2026-65762	Joomla Extension - phoca.cz - Reflected XSS vulnerability in Phoca Guestbook 5.0.0-6.1.0 - Improper validation of user inputs lead to a reflective XSS vulnerability.	N/A	More Details
CVE-2026-65763	Joomla Extension - phoca.cz - Reflected XSS vulnerability in Phoca Maps 5.0.0-6.0.4 - Improper validation of user inputs lead to a reflective XSS vulnerability.	N/A	More Details
CVE-2026-18085	An Improper Input Validation in the BlackBerry UEM Management Console of BlackBerry UEM 12.23.0 QF8 and earlier allows Arbitrary File Download and Potential Denial of Service.	N/A	More Details
CVE-2026-18084	Improper Neutralization of Input During Web Page Generation vulnerability in BlackBerry UEM Management Console of BlackBerry UEM allows Cross-Site Scripting (XSS). This issue affects UEM: 12.23.0 QF8 or earlier.	N/A	More Details
CVE-2026-55555	Dompdf is an HTML to PDF converter for PHP. Versions 3.15 and prior are vulnerable to a File Existence Oracle attack through the manipulation of the CSS @font-face directive. By providing malicious HTML that references local files via the file:// protocol repeatedly, an attacker can trigger PHP memory exhaustion. Because Dompdf behaves differently depending on whether a referenced local file exists (an existing file is processed repeatedly until it triggers an "Allowed memory size exhausted" crash, whereas a missing file fails fast or is ignored and never hits the memory limit), an attacker can use this observable discrepancy as an oracle to enumerate sensitive files on the server regardless of CHROOT restrictions. Exploitation requires the attacker to supply unrestricted or unsanitized HTML in a request that permits large data, plus a configuration where Dompdf's memory limit is low enough to be exhausted (with $_dompdf_show_warnings=true making the overflow easier to reach). This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-59676	A Time-of-check Time-of-use (TOCTOU) Race Condition vulnerability in seunshare of selinux policycoreutils allows a user calling seunshare that is running in the unconfined SELinux domain to delete arbitrary root-owned files, This issue affects policycoreutils through 3.10.	N/A	More Details
CVE-2026-15687	A security issue was discovered in the Kubernetes Java client library where a compromised pod may be able to create new files in arbitrary locations on the client machine executing copy operations via non-tar copyDirectoryFromPod when enableTarCompressing is false.	N/A	More Details
CVE-2026-64256	In the Linux kernel, the following vulnerability has been resolved: xfs: don't wrap around quota ids in dqiterate LOLLM noticed that q_id is an unsigned 32-bit variable. If it happens to be set to XFS_DQ_ID_MAX due to a filesystem that actually has a dquot for ID_MAX, then this addition will truncate to zero and the iteration starts over. Fix this by casting to u64.	N/A	More Details
CVE-	gopacket provides packet processing capabilities for Go. In version 1.6.0 and earlier, the sFlow ExtendedGatewayFlow decoder in layers/sflow.go reads an attacker-controlled 32-bit community count and AS path member count and sizes		

2026-54332	a slice allocation from those counts without bounding them against the bytes remaining in the datagram, so a 104-byte UDP datagram can drive an allocation of up to 16 GiB and cause an unauthenticated remote denial of service. This issue is fixed in version 1.6.1.	N/A	More Details
CVE-2026-63226	Printers and Multifunction Printers (MFPs) provided by Ricoh Company, Ltd. do not implement restrictions on SSH port forwarding, allowing to connect to arbitrary destinations. When SSH is enabled on an affected product, SSH port forwarding may be leveraged to connect to other node on the LAN.	N/A	More Details
CVE-2026-47722	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.2, `internal/configgen/generator.go:86,108,119` interpolates the operator-supplied `ListenHost` and `TunDevice` fields raw into a `text/template` that produces the agent's `config.yml`. `internal/web/advanced.go:20-35` accepts both with only `strings.TrimSpace` — no character or shape validation. Version 0.3.2 fixes the issue.	N/A	More Details
CVE-2026-13058	An authenticated user with basic write privileges can cause the mongod process to terminate abnormally by sending a crafted transaction command with an incomplete set of required fields. The issue stems from inconsistent validation across related transaction command parameters, resulting in a fatal internal invariant failure and denial of service.	N/A	More Details
CVE-2026-22049	ONTAP versions 9.16.1 and higher with WebAuthn multi-factor authentication (MFA) configured are susceptible to a vulnerability related to the Relying Party ID which when successfully exploited could allow an attacker with valid credentials to bypass MFA.	N/A	More Details
CVE-2026-7328	Missing authorization in Caliptra Core Runtime Firmware (INVOKE_DPE_MLDSA87, CM_AES_GCM_DECRYPT_DMA, EXTERNAL_MAILBOX_CMD commands) in subsystem mode allows a privileged local attacker to cause a denial of service via mailbox commands containing unverified AXI addresses. The security impact beyond availability is integration-specific. This issue affects Core Runtime Firmware: 2.1.0.	N/A	More Details
CVE-2026-21653	Victor SSRF vulnerability in Johnson Controls CCure 9000 and victor application server allows Server Side Request Forgery. This issue affects CCure 9000 and victor application server: from 2.9 through 3.0.	N/A	More Details
CVE-2026-21655	Deserialization of untrusted data vulnerability in Johnson Control victor on Windows allows capec-586. This issue affects victor: from 2.9 before 3.0.	N/A	More Details
CVE-2026-55554	Dompdf is an HTML to PDF converter for PHP. In versions 3.15 and prior, the validateLocalUri() method enforces chroot boundaries with a strpos() prefix check after normalizing paths with realpath() . Because normalization strips the trailing directory separator from $chrootPath , the check only verifies that $chrootPath is a string prefix of $realfile, so a chroot of /var/www also matches sibling directories like /var/www2 , /var/www-admin, or /var/www_backup. An attacker who controls part of the rendered HTML could exploit this to escape the chroot and read sensitive files outside the allowed directory. This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-54345	gopacket provides packet processing capabilities for Go. In version 1.6.0 and earlier, the Diameter AVP decoder computes an AVP data length by subtracting a fixed header size from an attacker-controlled AVP Length field, so a vendor-flagged AVP whose Length is smaller than the 12-byte header underflows the unsigned 32-bit value and drives an unbounded allocation of roughly 4 GiB, and two such messages in succession OOM-kill a collector, causing an unauthenticated remote denial of service. This issue is fixed in version 1.6.1.	N/A	More Details
CVE-2026-55728	Stack-based Buffer Overflow (CWE-121) in `/usr/bin/ltsudo` `cmd_ipaddr_conflict` in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows a `superadmin`-group attacker to trigger a SUID-root process abort or potentially elevate privileges via an overly long interface-name argument.	N/A	More Details
CVE-2026-15037	Improper output neutralization (XML injection) in QDom comment, CDATA, and processing-instruction serialization in Qt XML from 4.0.0 through 6.11 allows untrusted text serialized by an application into those nodes to inject arbitrary XML markup, because the node terminators are not escaped under the default InvalidDataPolicy (AcceptInvalidChars). Fixed in Qt 6.12.	N/A	More Details
CVE-2026-48058	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.2, internal/web/session.go and internal/web/oidc.go set HttpOnly and SameSite=Lax on every cookie but never Secure. A single plaintext request to the origin (operator on a LAN, mistyped URL, HTTP→HTTPS not strictly enforced, reverse proxy misconfiguration) discloses the session. This issue has been patched in version 0.3.2.	N/A	More Details
CVE-2026-50735	pglogical's apply worker does not sufficiently validate the length of certain fields in incoming replication protocol messages before copying them, resulting in an out-of-bounds read. A party acting as the publisher for a subscription, for example a non-PostgreSQL endpoint that speaks the pglogical replication protocol, can return crafted messages that cause the subscriber's apply worker to read beyond the bounds of an allocated buffer, disclosing adjacent process memory or crashing the worker. To exploit the issue an attacker must be able to direct a subscription at an endpoint they control. In default installations this requires privileges normally reserved for a superuser, so the issue is most relevant to managed deployments where the ability to create subscriptions has been delegated to non-superuser roles.	N/A	More Details
CVE-2026-50736	The pglogical queue mechanism, used to convey out-of-band commands such as replicated DDL from a publisher to a subscriber, executes message payloads on the subscriber at the privilege level of the apply worker, which is equivalent to a PostgreSQL superuser in default installations. A party acting as the publisher can send crafted queue messages that cause arbitrary SQL to be executed on the subscriber as superuser, escalating from a role permitted to use pglogical to full superuser and breaking the isolation between tenants in shared deployments. To exploit the issue an attacker must be able to direct a subscription at an endpoint they control. In default installations this requires privileges normally reserved for a superuser, so the issue is most relevant to managed deployments where the ability	N/A	More Details

	to create subscriptions has been delegated to non-superuser roles.		
CVE-2026-50737	When applying replicated changes for a row that is missing one or more columns, pglogical evaluates the affected table's default expressions on the subscriber. Because the apply worker runs at a privilege level equivalent to a PostgreSQL superuser in default installations, any function invoked by such a default expression also runs at that privilege. A party acting as the publisher can use this path to cause functions to be executed on the subscriber as superuser, escalating from a role permitted to use pglogical to full superuser. This is a second, independent path to the same superuser escalation tracked under CVE-2026-50736 (the pglogical queue issue). To exploit the issue an attacker must be able to direct a subscription at an endpoint they control. In default installations this requires privileges normally reserved for a superuser, so the issue is most relevant to managed deployments where the ability to create subscriptions has been delegated to non-superuser roles.	N/A	More Details
CVE-2026-50738	A use-after-free condition exists in pglogical's worker signaling code, where a worker structure can be dereferenced after the underlying slot has been freed or recycled during normal worker lifecycle events. The condition is reachable during normal replication operation, including by a low-privileged user able to influence worker start, stop, and restart timing through permitted pglogical operations. In the typical case the condition crashes replication workers, causing an availability impact. In the worst case a use-after-free in a PostgreSQL backend can be leveraged as a remote code execution primitive at the privilege of that backend.	N/A	More Details
CVE-2026-65758	Joomla Extension - tassos.gr - Sensitive data exposure in Convert Forms extension 2.5.0-5.2.2 - The front-end Submissions view did not enforce access control. An unauthenticated visitor could therefore list a form's submissions.	N/A	More Details
CVE-2026-59678	An Incorrect Authorization vulnerability in Linux-Gaming PortProtonQt allows any users to mount and unmount arbitrary file systems and modify the network configuration via NetworkManager. This issue affects PortProtonQt before 0d0f0950ebd948cdf82e8c3e1ebd2bcb9b8bafbe.	N/A	More Details
CVE-2026-48025	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.7, internal/pki/resolver.go:36-64 constructs a CAManager with the plaintext ed25519.PrivateKey after unwrapping via the master key; internal/pki/ca.go:13-16 stores it. Callers at internal/api/enroll.go:116, internal/api/updates.go:297, and internal/api/mobile_bundle.go:40 use the manager for one Sign() and drop the reference on function return — but the underlying slice contents are not wiped before release. The keystore package's contract (internal/keystore/keystore.go doc: "Callers MUST zeroise the returned plaintext DEK as soon as it is no longer needed") is not met by the CAManager consumer. Decrypted CA private keys persist in process heap until Go's GC scavenges the underlying slice — minutes to hours under load, indefinitely on idle servers. This issue has been patched in version 0.3.7.	N/A	More Details
CVE-2026-44210	Kata Containers is an open source project focusing on a standard implementation of lightweight Virtual Machines (VMs) that perform like containers. Versions prior to 3.31.0 ship with a default configuration that allows pod creators to inject arbitrary command-line arguments into the virtiofsd process through the `io.katacontainers.config.hypervisor.virtio_fs_extra_args` pod annotation. By injecting `-o source=/` along with `--no-announce-submounts` and `--sandbox=none`, an attacker can override the virtiofsd shared directory to serve the entire host root filesystem into the guest VM. Combined with the `kernel_params` annotation (also enabled by default) to activate the agent debug console, the attacker can mount the host filesystem from inside the VM and read or write any file on the host, including /etc/shadow. Version 3.31.0 patches the issue.	N/A	More Details
CVE-2026-59677	A Missing Authorization vulnerability in selinux policycoreutils seunshares allows a user that is running in unconfined context to kill e.g. root-owned processes running also in unconfined context This issue affects policycoreutils through 3.10.	N/A	More Details
CVE-2026-48533	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2026-65759	Joomla Extension - joomshaper.com - unauthenticated payment/order forgery in Easy Store extension 1.0.0-2.0.1 - Critical order and payment information, including states, are processed from client side input, enabling unauthenticated attackers to manipulate payment and order states of arbitrary orders.	N/A	More Details
CVE-2026-54620	sqlite3 provides Ruby bindings for the SQLite3 embedded database. From 2.1.0 to 2.9.4, the callbacks used for SQLite aggregate functions can be freed while still referenced during aggregation, resulting in a use-after-free. This issue is fixed in version 2.9.5.	N/A	More Details
CVE-2026-54619	sqlite3 provides Ruby bindings for the SQLite3 embedded database. In version 2.9.4 and earlier, redefining a SQLite function with a different arity frees the previously registered function handler while SQLite may still reference it, resulting in a use-after-free. This issue is fixed in version 2.9.5.	N/A	More Details
CVE-2026-65760	Joomla Extension - joomshaper.com - cross-customer order and personal information disclosure in Easy Store extension 1.0.0-2.0.1 - Improper access checks allow logged in users to retreive order and customer information of any order in the system.	N/A	More Details
CVE-2026-65761	Joomla Extension - joomshaper.com - Unauthenticated SQL injection in Easy Store extension 1.0.0-2.0.1 - Improper validation of order parameters lead to an unauthenticated SQL injection in easystore, allowing full DB read access including credentials and sessions.	N/A	More Details
CVE-2026-34496	Cwe-269 vulnerability in Johnson Controls victor Web on Windows allows capec-233. This issue affects victor Web: before 7.1.	N/A	More Details

CVE-2026-16552	Rejected reason: The reported issue is invalid, as it requires root privileges to reproduce, and it is out of scope of the threat model of the affected component.	N/A	More Details
CVE-2026-47723	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.1, none of the response paths in `internal/web/` or `internal/api/` set the standard browser-security headers. `grep` for `Content-Security-Policy`, `X-Frame-Options`, `Strict-Transport-Security`, `X-Content-Type-Options`, `Referrer-Policy` returns zero matches across the codebase. Version 0.3.1 fixes the issue.	N/A	More Details
CVE-2026-12496	Stored Cross-Site Scripting (CWE-79) in the OPC XML-DA server statistics in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows an unauthenticated remote attacker to execute arbitrary JavaScript in an administrator's browser (session hijacking, credential theft, device reconfiguration) via a crafted `User-Agent` header in a `POST /da` request.	N/A	More Details
CVE-2026-15810	A Cross-Site Scripting (XSS) vulnerability in Google Cloud Looker versions prior to 25.6.103, 25.12.65, 25.18.68, 26.0.66, 26.2.47, 26.4.36, 26.6.28, and 26.8.7 on Looker-hosted and Self-hosted allows an attacker to execute arbitrary JavaScript leading to administrative account takeover using a maliciously crafted URL. Looker-hosted and Self-hosted were found to be vulnerable. This issue has already been mitigated for Looker-hosted instances. No user action is required for these. Self-hosted instances must be upgraded to the patched versions: 25.6.103+, 25.12.65+, 25.18.68+, 26.0.66+, 26.2.47+, 26.4.36+, 26.6.28+, or 26.8.7+.	N/A	More Details
CVE-2026-67173	Pivotick did not validate the URL scheme of node imagePath values derived from graph data before assigning them to SVG image resources. An attacker able to supply crafted graph data could set an image path to a malicious URI. When a victim rendered the affected graph, the browser could resolve the attacker-controlled URI and initiate an unintended request or invoke scheme-specific handling in the victim’s context. Depending on the URI, browser behaviour, and installed protocol handlers, exploitation could disclose limited client or network metadata, facilitate rendering-based tracking, or attempt to access local or internal resources. Exploitation requires a victim to load or render graph data containing the malicious imagePath. The patch normalizes ASCII whitespace and control characters in URI schemes and restricts image paths to relative URLs or the http, https, data, and blob schemes.	N/A	More Details
CVE-2026-66922	Pivotick used plain JavaScript objects as lookup tables indexed by caller-controlled graph node identifiers in its tree-layout and cycle-detection components. Node identifiers matching properties inherited from Object.prototype, such as constructor, toString, or __proto__, were not handled as ordinary identifiers. These values could be interpreted as existing inherited properties, resolve to values of an unexpected type, or—in the case of __proto__ assignments—modify the prototype of an internal lookup object. An attacker who can supply graph data containing crafted node identifiers could consequently cause nodes or edges to be silently omitted, produce incorrect hierarchy levels, bypass or corrupt cycle-detection results, or trigger an exception that interrupts graph processing and rendering. This affects the integrity of graph visualisations and analytics and may cause a client-side denial-of-service condition. The affected code also failed to safely handle edges whose source node was absent from the supplied node set. Furthermore, calculating the maximum tree depth by spreading all level values into Math.max() could exceed the JavaScript function-argument limit when processing a sufficiently large graph, resulting in an exception and denial of service. The patch replaces identifier-keyed plain objects with Map instances, ignores invalid edges during tree construction, and calculates the maximum depth iteratively.	N/A	More Details
CVE-2026-6881	A SQL Injection in the Giving Reports functionality in Ellucian Advance Web and Legacy Advance allows an authenticated attacker to extract sensitive information from databases via a crafted SQL query in the class credit field. This issue affects Advance Web: all versions; Legacy Advance: all versions. Ellucian CRM Advance is not impacted.	N/A	More Details
CVE-2026-66921	Pivotick’s Markdown node-reference renderer failed to HTML-escape the attacker-controlled nodeName value before interpolating it into both the data-node-name attribute and the body of a generated element. Because the node-reference tokenizer rejected only square brackets, a crafted node name could still contain quotation marks, angle brackets, or other HTML metacharacters. An attacker could therefore terminate the quoted attribute or inject additional HTML elements and event-handler attributes. When malicious node-reference content is rendered by a consumer that does not apply DOMPurify or equivalent sanitization, arbitrary JavaScript may execute in the victim’s browser in the security context of the application. Successful exploitation requires a victim to open or render a crafted graph or note and could allow the attacker to access same-origin information, modify displayed content, or perform actions using the victim’s session. The patch resolves the issue by applying context-appropriate HTML escaping to node names before inserting them into either HTML text or quoted attribute values. The shared escaping function now encodes ampersands, angle brackets, and both types of quotation marks.	N/A	More Details
CVE-2026-66008	Parse Server versions >= 9.0.0 before 9.10.0-alpha.6 and >= 8.2.2 before 8.6.87 disclose Pointer and Relation target class names through GraphQL validation and input-coercion error messages when public schema introspection is disabled (graphQLPublicIntrospection: false, the default). Because these errors are produced before authentication, authorization, or any resolver runs, an unauthenticated client possessing only the public application ID can trigger errors on Pointer or Relation fields to reconstruct hidden schema class names, partially defeating the schema-hiding protection. Only schema metadata (class names) is exposed; no object data, credentials, or user records are disclosed.	N/A	More Details
CVE-2026-66009	Parse Server versions >= 9.0.0 before 9.10.0-alpha.5 and >= 8.2.2 before 8.6.86 return GraphQL validation error messages that name required custom input fields even when public introspection is disabled (graphQLPublicIntrospection: false, the default). A client holding only the public application id — with no user session, master key, or maintenance key — can trigger validation errors to learn the names of required (non-null) custom fields on classes it already references by name, partially defeating the schema-hiding intent of disabling public introspection. No stored data, credentials, optional field names, unreferenced class names, or Cloud Code function	N/A	More Details

	names are exposed.		
CVE-2026-66920	Pivotick contains an uncontrolled-recursion vulnerability when processing caller-supplied graph and node data. The affected graph algorithms recursively traversed graph edges, while the JSON viewer recursively processed each level of a node’s data structure. A specially crafted graph containing an excessively long path, deeply nested properties, or circular object references could therefore exhaust the JavaScript call stack when Pivotick calculates a layout or displays a node in the inspection modal. Successful exploitation may cause an uncaught exception, freeze the affected page, or crash the browser tab, resulting in a client-side denial of service. No confidentiality or integrity impact has been identified. The patch replaces the recursive graph traversals with iterative stack-based implementations and limits the reachability calculation to 1,000,000 edge traversals. It also limits JSON rendering to 64 levels and detects circular references before descending further into an object.	N/A	More Details
CVE-2026-6924	A bug in the entropy initialization for SiWx917 causes the DRBG to use a predictable seed. As such, all random numbers generated in the Matter code use the same stream of numbers. This vulnerability was discovered after the impacted repository was already deprecated.	N/A	More Details
CVE-2026-66919	Pivotick contains a cross-site scripting vulnerability in the inspect and edit node modals. Node labels and descriptions originating from graph data were interpolated directly into HTML used to construct the modal headers. An attacker able to supply or modify graph data could insert a malicious HTML or JavaScript payload into a node’s label or description. The payload would be parsed and executed in the application’s origin when a user opened the affected node’s inspect or edit modal. Successful exploitation could allow the attacker to access information available to the victim, modify application data, or perform actions using the victim’s active session. The vulnerability has been addressed by creating the modal elements without embedding graph data in HTML and assigning node labels and descriptions through textContent.	N/A	More Details
CVE-2026-66918	Pivotick fails to sanitize attacker-controlled SVG markup supplied through the per-node style.svgIcon property before inserting it into the document. When rendering a graph node, the vulnerable code assigns the SVG icon markup directly to the innerHTML property of a live SVG element. An attacker able to influence graph data can provide crafted markup containing executable event handlers, such as an <image> element with an onerror attribute. When a victim loads or renders the malicious graph, the payload may execute arbitrary JavaScript in the security context of the application embedding Pivotick. Successful exploitation could allow the attacker to access application data available to the victim, modify displayed content, or perform actions using the victim’s authenticated session. Exploitation requires an application using Pivotick to render graph data that is controlled or modified by an attacker.	N/A	More Details
CVE-2026-66913	Lookyloo did not enforce limits on the decompressed size of uploaded capture archives and compressed HAR files. An attacker could submit a specially crafted ZIP, gzip, or zlib-compressed capture containing data that expands to a very large size during processing. Because the application decompressed this content directly in memory without first limiting the output size, processing the malicious capture could exhaust available memory, terminate a web or worker process, or make the Lookyloo instance unavailable. The vulnerability affects both full Lookyloo capture archive imports and API submissions containing gzip-compressed HAR data. Repeated exploitation could cause a persistent denial-of-service condition until the affected processes or instance are restarted. The patch introduces: * A 1 GB cumulative uncompressed-size limit for imported capture archives. * Size-limited gzip and zlib decompression for compressed HAR files. * Explicit detection and handling of suspected zip bombs. * An HTTP 400 response when an oversized compressed HAR file is submitted through the API.	N/A	More Details
CVE-2026-12502	Improper Privilege Management (CWE-269) in `/usr/bin/ltsudo` in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows a `superadmin`-group attacker to reset the password of any LARM user (including the `larmapp` service account) via the `set-passwd` subcommand.	N/A	More Details
CVE-2026-12503	Improper Link Resolution (CWE-59) in `/usr/bin/larm_starter` in Loytec L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows an authenticated `larmapp` attacker to make `/etc/passwd` writable by the `larmapp` group (leading to root privilege escalation) via a symlink attack on `/etc/lighttpd/ssl/server.pem`.	N/A	More Details
CVE-2026-12504	Improper Authentication (CWE-287) in the PAM configuration in Loytec LIP-ME201C, L-INX, L-GATE, L-ROC, L-IOB, L-DALI, L-VIS and L-PAD through 8.4.16 on LINX-A64 allows a local attacker to authenticate as a uid=0 account without a password and obtain a root shell via an `/etc/passwd` entry with an empty password field.	N/A	More Details
CVE-2026-47726	nebula-mesh is a self-hosted control plane for Slack Nebula mesh virtual private network. Prior to version 0.3.2, internal/api/audit.go:12 — handleGetAuditLog does no admin check. The route is bearer-auth gated only; any operator API key returns the full audit log via store.ListAuditEntries (up to limit=1000). This includes cross-tenant actor names, host/CA/operator IDs, action timestamps, and masked-IP entries from rate-limit refusals — enough surface for a tenant to enumerate the server's activity, infer staffing patterns, or identify high-value targets. This issue has been patched in version 0.3.2.	N/A	More Details
CVE-2026-15243	Apereo CAS Client accepts any CA-trusted certificate for any hostname, provided the URL the client is calling matches the configured allowlist or regex. An attacker with a MITM position (DNS poisoning, rogue Wi-Fi, malicious proxy, etc.) can provide any CA-signed certificate for a hostname that matches the configured allowlist or regex. This can lead to intercepting the CAS exchange, capturing the Ticket-Granting Ticket (TGT), and subsequently obtaining Service Tickets on behalf of the victim. Because maintainers contact attempts were unsuccessful, vulnerabilities have only been confirmed in version 4.1.0 (Java Apereo CAS Client) and 3.6.4 (Jasig CAS Client) but may also affect other versions.	N/A	More Details
CVE-2026-51261	Missing mutex synchronization in AudioBuffer::freeSpace() in schreibfaul1 ESP32-audioI2S 3.4.5 creates a race condition between concurrent tasks. The function calculates available buffer space without protecting shared read/write pointers, returning an incorrectly large value. Trusting this value leads to heap out-of-bounds write, memory corruption, device crash, and arbitrary code execution.	N/A	More Details

CVE-2026-10610	Local privilege escalation potentially allowed an attacker to execute arbitrary code as a privileged user.	N/A	More Details
CVE-2026-7483	Local privilege escalation potentially allowed an attacker to write an arbitrary file with fully controlled content as a privileged user.	N/A	More Details
CVE-2026-11922	A vulnerability in zenml-io/zenml versions 0.57.0 through 0.94.2 allows an attacker to bypass rate-limiting on the `POST /api/v1/login` and self password-change endpoints by rotating the `X-Forwarded-For` header. The rate limiter keys requests by `request.client.host`, which is derived from the `X-Forwarded-For` header when Uvicorn is launched with `--proxy-headers --forwarded-allow-ips *`. This configuration allows clients to control the value of `request.client.host`, effectively bypassing rate-limiting protections. This vulnerability leaves the affected endpoints open to unthrottled credential guessing attacks.	N/A	More Details
CVE-2026-53910	diff3 tool from GNU diffutils is vulnerable to a heap-based buffer overflow due to multiple signed integer overflows in line-mapping calculations. Incorrect arithmetic in mapping line ranges can result in corrupted values being used for memory allocation and loop bounds. When processing crafted diff output, these overflows may cause the application to allocate insufficient memory and subsequently perform out-of-bounds writes during internal processing. An attacker who can control the output of the diff program used by diff3 (e.g. via --diff-program pointing to a malicious script) can trigger out-of-bounds writes, resulting in a crash and potentially remote code execution depending on the environment. This issue has been fixed in commit 9ff04d5b84743e331e80b589335a52c5480d1815 NOTE: The project maintainers claim that this is not a security issue. They state that the worst outcome this issue can cause is a crash of diff and that it cannot be used to escalate privileges.	N/A	More Details
CVE-2026-51270	schreibfaul1 ESP32-audioI2S 3.4.5 has a heap-based buffer overflow vulnerability in the htmlToUTF8() HTML entity decoding function. The function parses attacker-controlled malicious HTML entities and uses memmove and memcpy to rearrange string content without validating buffer remaining size and boundary limits. Crafted oversized HTML entity strings can trigger heap out-of-bounds write, allowing remote code execution, memory information leakage, service crash or privilege escalation.	N/A	More Details
CVE-2026-56722	Dompdf is an HTML to PDF converter for PHP. In versions 3.15 and prior, aAn attacker who controls the HTML input can bypass this restriction by embedding a target file path inside an SVG image delivered through a data: URI, because dompdf processes the SVG twice and the second pass does not enforce the same protections as the first. When rendering, dompdf hands the SVG to the separate php-svg-lib library with external references forced on, and that library has no knowledge of the chroot directory, blocks only the phar:// scheme, and ultimately reads the referenced file with no path or protocol validation. This lets an external, unauthenticated attacker read arbitrary image files from the server's file system in the default configuration. This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-51268	schreibfaul1 ESP32-audioI2S 3.4.5 has a heap-based buffer overflow vulnerability in the host parsing logic. The dismantle_host() function parses untrusted host and URL input, and subsequent code uses clone_from() to copy parsed host, request host, extension and query_string segments into fixed heap buffers without boundary checking.	N/A	More Details
CVE-2026-59941	Dompdf is an HTML to PDF converter for PHP. Versions 3.15 and prior accept a BMP image and generates a PDF-compatible PNG based only on its declared header dimensions and never bounds width × height before the image is converted through GD. A 58-byte BMP whose header declares e.g. 6000×6000 is accepted and later drives imagecreatetruecolor($width, $height) (and PHP's native BMP decoder) to allocate the full pixel canvas. A payload can fit in a single HTTP request: the BMP can be inlined as a ... URI inside attacker-controlled HTML, so no upload, no remote fetch, and no chroot-reachable file is required. I measured a 169-byte request driving a dompdf render to ~412 MB peak RSS and ~4.8 s of CPU/wall time, versus ~34 MB for an identically-sized benign request — roughly a 12× memory amplification per request, repeatable and unauthenticated. This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-12702	In affected versions of Octopus Deploy Insufficient checks on the project trigger actions allows an unauthorized user to trigger a deployment.	N/A	More Details
CVE-2026-59942	Dompdf is an HTML to PDF converter for PHP. Versions 3.15 and prior are vulnerable to a Denial of Service (DoS) attack via resource exhaustion. An attacker can crash the PHP process by providing a specially crafted HTML document containing a single image with massive dimensions (e.g., 30,000x30,000 pixels). While Dompdf implements internal checks to validate image dimensions, these can be bypassed by using a high-entropy image (such as random noise) encoded in Base64 and wrapped in specific CSS containers. The vulnerability exists because the dimension validation happens early, but the resource allocation for calculating the object's bounding box and internal buffers during the rendering phase does not strictly limit the cumulative CPU time or memory usage for a single object that has passed the initial check. An unauthenticated remote attacker can cause a complete Denial of Service on the web server by submitting a crafted HTML string. This affects any application that allows users to provide HTML content or URLs that are subsequently converted to PDF using Dompdf. This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-24727	An unrestricted upload of file with dangerous type vulnerability in the e-paper draft upload function of SUNNET Corporate Training Management System through v10.3 allows remote authenticated users with administrator privileges to execute arbitrary commands by uploading a crafted ZIP archive containing a server-executable file.	N/A	More Details
CVE-	GNU coreutils uniq is vulnerable to an out-of-bounds read due to incorrect handling of multibyte input when the -w (--check-chars) option is used. The find_field() function miscalculates the byte length of characters by repeatedly processing a fixed pointer instead of advancing through the input, resulting in an inflated length value. This incorrect		

2026-56391	length is later used in a memcmp operation, causing reads beyond the allocated buffer when processing crafted multibyte input. When running GNU coreutils uniq with attacker-provided arguments, this behavior leads to a crash and potential adjacent heap memory exposure. This issue has been fixed in the commit d64e35a8a4c0e4608321433e0d84d917e4e36371.	N/A	More Details
CVE-2026-56392	GNU coreutils unexpand is vulnerable to a heap-based buffer overflow due to an integer overflow during buffer allocation when processing large tab stop (-t) values. The multiplication used to calculate the allocation size can wrap around, resulting in an undersized buffer. When processing crafted input, subsequent writes exceed the allocated memory, leading to an out-of-bounds heap write. When running GNU coreutils unexpand with attacker-provided large tab stop (-t) arguments, this behavior leads to a crash and potentially achieve a heap write primitive depending on memory layout. This issue has been fixed in the commit b60a159fdc5bfcf9988d3a4cb6f53abe8ad5d35d	N/A	More Details
CVE-2026-59943	Dompdf is an HTML to PDF converter for PHP. In versions 3.15 and prior, if a malicious actor can supply unrestricted content for rendering by Dompdf they can utilize the SVG rendering functionality to leak filesystem information when rendering PDF files using image references within a data-URI encoded SVG document. Using an <image> element inside a data-URI embedded SVG, an attacker can attempt to embed other files via the href or xlink:href attributes. When processing a file that does not exist (e.g. file:///DOESNOTEXIST), dompdf behaves differently than it does when accessing a file or directory that actually exists on the filesystem. This issue has been fixed in version 3.16.	N/A	More Details
CVE-2026-6879	`Element.findall()` and fully-consumed `Element.iterfind()` exhibit `O(n^2)` time complexity when using XPath index predicates (e.g. `[1]`, `[last()]`, `[last()-N]`) on XML documents with many same-tag siblings. `Element.find()` is only affected when the first match is near the end of the sibling list, such as with `[last()]` or `[last()-N]`; `.//item[1]` short-circuits after the first match.	N/A	More Details
CVE-2026-67178	MISP installation scripts generated an Apache HTTP virtual-host configuration containing an incorrectly formatted HTTP-to-HTTPS redirect: Redirect permanent / https://misp.example Apache's Redirect directive appends any portion of the requested path that follows the matched prefix to the configured destination URL. Because the destination did not end with /, attacker-controlled path content was appended directly to the hostname rather than to its URL path. For example, a request resembling: http://misp.example/@attacker.example/ could result in a redirect resembling: https://misp.example@attacker.example/ Under standard URL parsing, misp.example is interpreted as user information and attacker.example as the destination host. An unauthenticated remote attacker could therefore construct a URL hosted under the legitimate MISP domain that redirects users to an attacker-controlled website. The vulnerability could be used for phishing, credential collection, or potentially disclosing sensitive query-string information preserved during the redirect. Exploitation requires a user to follow the crafted HTTP URL. The fix adds the missing trailing slash to the redirect destination, ensuring that appended request data remains part of the path on the configured MISP host. Existing installationsExisting MISP installations should review their Apache HTTP virtual-host configuration and ensure that the HTTPS redirect destination ends with a trailing slash: Redirect permanent / https://misp.example/ After updating the configuration, validate it with apachectl configtest and reload or restart Apache for the change to take effect	N/A	More Details
CVE-2026-67174	Pivotick contains a DOM-based cross-site scripting vulnerability in its generic UI element resolution and icon-rendering utilities. The tryResolveHTMLElement function treated any resolved string as HTML markup by assigning it to a <template> element through innerHTML. Strings derived from untrusted graph properties or custom rendering callbacks could therefore introduce arbitrary HTML or SVG elements into the live document. The vulnerable function was used by multiple UI components, including headers, property panels, extra panels, and tooltips. Additionally, createIcon inserted caller-supplied svgIcon markup into a template without sanitization. An application integrating Pivotick and deriving icon markup from untrusted data could therefore expose a second script-execution path. An unauthenticated attacker able to provide a crafted graph, property value, rendering result, or SVG icon could execute JavaScript in another user's browser when the affected content is displayed or interacted with. Successful exploitation could allow the attacker to access information available to the victim, manipulate graph data or application state, and perform actions with the victim's privileges. The patch changes string rendering to use textContent, requiring callers to explicitly return an Element when HTML rendering is intended. It also sanitizes SVG icon markup before inserting it into the DOM.	N/A	More Details
CVE-2026-47670	DbGate is cross-platform database manager. Versions 7.1.8 and prior are vulnerable to authenticated Remote Code Execution (RCE). Any user with valid DbGate credentials can execute arbitrary OS commands as root by exploiting an unsanitized `functionName` parameter in the `/runners/load-reader` endpoint. The `require = null` mitigation is trivially bypassed via dynamic `import()`. Version 7.1.9 contains a patch.	N/A	More Details
CVE-2026-64285	In the Linux kernel, the following vulnerability has been resolved: KVM: SEV: Pin source page for write when adding CPUID data for SNP guest When populating a guest_memfd instance with the initial CPUID data for an SNP guest, acquire a writable pin on the source page as KVM will write back the "correct" CPUID information if the userspace provided data is rejected by trusted firmware. Because KVM writes to the source page using a kernel mapping, pinning for read could result in KVM clobbering read-only memory. Note, well-behaved VMMs are unlikely to be affected, as CPUID information is almost always dynamically generated by userspace, i.e. it's unlikely for the CPUID information to be backed by a read-only mapping. [sean: rewrite shortlog and changelog, tag for stable@]	N/A	More Details
CVE-2026-64258	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: remove request-less entries from ent_w_req_queue to fix NULL deref If a copy into the userspace ring buffer fails, a request will be terminated and fuse_uring_req_end() will set ent->fuse_req to NULL but it will leave the entry on ent_w_req_queue in FRRS_FUSE_REQ state. This can lead to a NULL deref if the request expiration logic scans ent_w_req_queue in the window before the entry is moved off it. Fix this by taking the entry off ent_w_req_queue and changing its state from FRRS_FUSE_REQ to FRRS_INVALID before terminating the request.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: x86/bugs: Enable IBPB flush on BPF JIT allocation		

CVE-2026-64507	Enable hardening against JIT spraying when Spectre-v2 mitigations are in use. Specifically, issue an IBPB flush on BPF JIT memory reuse. Skip enabling the IBPB flush if the BPF dispatcher is already using a retpoline sequence. This hardening applies only when BPF-JIT is in use. Guard the enabling under CONFIG_BPF_JIT so that bugs.c still builds with CONFIG_BPF_JIT=n.	N/A	More Details
CVE-2026-64493	In the Linux kernel, the following vulnerability has been resolved: iio: pressure: mpl115: fix runtime PM leak on read error mpl115_read_raw() takes a runtime PM reference with pm_runtime_get_sync() before reading the processed pressure or raw temperature, but on the read error path it returns without calling pm_runtime_put_autosuspend(). Each failed read therefore leaks a runtime PM reference and prevents the device from autosuspending. Drop the reference before checking the return value so both the success and error paths are balanced.	N/A	More Details
CVE-2026-64494	In the Linux kernel, the following vulnerability has been resolved: iio: light: gp2ap002: fix runtime PM leak on read error gp2ap002_read_raw() calls pm_runtime_get_sync() before reading the lux value, but if gp2ap002_get_lux() fails, it returns directly. This skips the pm_runtime_put_autosuspend() call at the "out" label, permanently leaking a runtime PM reference and preventing the device from autosuspending. Replace the direct return with a "goto out" to ensure the reference is properly dropped on the error path.	N/A	More Details
CVE-2026-64495	In the Linux kernel, the following vulnerability has been resolved: iio: gyro: bmg160: bail out when bandwidth/filter is not in table bmg160_get_filter() walks bmg160_samp_freq_table[] looking for the entry matching the bw_bits value read from the chip: for (i = 0; i < ARRAY_SIZE(bmg160_samp_freq_table); ++i) { if (bmg160_samp_freq_table[i].bw_bits == bw_bits) break; } *val = bmg160_samp_freq_table[i].filter; If no entry matches, i ends up equal to the array size and the next line reads one slot past the end. bmg160_set_filter() has the same shape, driven by 'val' instead of bw_bits. smatch flags both: drivers/iio/gyro/bmg160_core.c:204 bmg160_get_filter() error: buffer overflow 'bmg160_samp_freq_table' 7 <= 7 drivers/iio/gyro/bmg160_core.c:222 bmg160_set_filter() error: buffer overflow 'bmg160_samp_freq_table' 7 <= 7 Return -EINVAL when no entry matches. The set_filter() path is reachable from userspace via the sysfs in_anglvel_filter_low_pass_3db_frequency interface, so userspace can trivially trigger the out-of-bounds read with a value that is not in bmg160_samp_freq_table[].filter.	N/A	More Details
CVE-2026-64497	In the Linux kernel, the following vulnerability has been resolved: iio: chemical: scd30: Cleanup initializations and fix sign-extension bug Include linux/bitfield.h for FIELD_GET(). Create new macros for bit manipulation in combination with manual bit manipulation being replaced with FIELD_GET(). The current variable declaration and initializations are barely readable and use comma separations across multiple lines. Refactor the initializations so that mantissa and exp have separate declarations and sign gets initialized later. In addition (and due to the nature of the cleanup), fix a sign-extension bug where, float32 would get bitwise anded with ~BIT(31) (which is 0xFFFFFFFF7FFFFFFF) which corrupted the exponent.	N/A	More Details
CVE-2026-64498	In the Linux kernel, the following vulnerability has been resolved: iio: buffer: hw-consumer: free scan_mask on buffer release The scan_mask lifetime changed in commit 9a2e1233d38c ("iio: buffer: hw-consumer: remove redundant scan_mask flexible array"). Before that change, the scan mask storage was embedded in struct hw_consumer_buffer, so iio_hw_buf_release() could free the whole allocation with a single kfree(hw_buf). That commit moved the scan mask to a separate bitmap_zalloc() allocation stored in buffer.scan_mask, but left iio_hw_buf_release() unchanged. Free the scan mask in iio_hw_buf_release() before freeing the buffer wrapper.	N/A	More Details
CVE-2026-12001	A hardcoded credential vulnerability exists in the firmware of multiple TP-Link routers (TL-WR845N v4, TL-WR850N v3, Archer C20 v6 & Archer MR200 v5). Authentication-related credential material is embedded within a password file in the firmware image and may be recovered through firmware analysis. Successful exploitation could result in unauthorized access to privileged functions on affected devices.	N/A	More Details
CVE-2026-64499	In the Linux kernel, the following vulnerability has been resolved: iio: adc: ti-ads1119: fix PM reference leak in buffer preenable ads1119_triggered_buffer_preenable() resumes the device with pm_runtime_resume_and_get() before starting a conversion. If i2c_smbus_write_byte() fails, the function returns the error directly and leaves the runtime PM usage counter elevated. The matching postdisable callback is not called when preenable fails, so the reference is leaked and the device may remain runtime-active indefinitely. Store the I2C transfer result in ret and drop the runtime PM reference on failure before returning the error.	N/A	More Details
CVE-2026-64500	In the Linux kernel, the following vulnerability has been resolved: iio: adc: lpc32xx: Initialize completion before requesting IRQ In the report from Jaeyoung Chung: "lpc32xx_adc_probe() in drivers/iio/adc/lpc32xx_adc.c registers its interrupt handler with devm_request_irq() before it initializes st->completion with init_completion(). If an interrupt arrives after devm_request_irq() and before init_completion(), the handler calls complete() on an uninitialized completion, causing a kernel panic. The probe path, in lpc32xx_adc_probe(): iodev = devm_iio_device_alloc(&pdev->dev, sizeof(*st)); /* st kzalloc-zeroed */ ... retval = devm_request_irq(&pdev->dev, irq, lpc32xx_adc_isr, 0, LPC32XXAD_NAME, st); /* register handler */ ... init_completion(&st->completion); /* initialize completion */ lpc32xx_adc_isr() calls complete(): complete(&st->completion); If the device raises an interrupt before init_completion() runs, complete() acquires the uninitialized wait.lock and walks the zeroed task_list in swake_up_locked(). The zeroed task_list makes list_empty() return false, so swake_up_locked() dereferences a NULL list entry, triggering a KASAN wild-memory-access." Fix the chance of a spurious IRQ causing an uninitialized pointer dereference by moving init_completion() above devm_request_irq().	N/A	More Details
CVE-2026-64503	In the Linux kernel, the following vulnerability has been resolved: iio: accel: kxsd9: fix runtime PM imbalance on write_raw() error kxsd9_write_raw() takes a runtime PM reference with pm_runtime_get_sync() but returns -EINVAL directly when a scale with a non-zero integer part is requested, skipping the matching pm_runtime_put_autosuspend(). This leaks a runtime PM usage-counter reference on every such write, after which the device can no longer autosuspend. Set the error code and fall through to the existing put instead of returning early.	N/A	More Details

CVE-2026-64504	In the Linux kernel, the following vulnerability has been resolved: iio: accel: bmc150: clamp the device-reported FIFO frame count __bmc150_accel_fifo_flush() copies the number of samples the device reports in its hardware FIFO into an on-stack buffer u16 buffer[BMC150_ACCEL_FIFO_LENGTH * 3]; which is sized for at most BMC150_ACCEL_FIFO_LENGTH (32) samples. The frame count is read from the FIFO_STATUS register and only masked to its 7 valid bits: count = val & 0x7F; so it can be 0..127. The only other limit applied to it is the optional caller-supplied sample budget: if (samples && count > samples) count = samples; which does not constrain count on the flush-all path (samples == 0), and leaves it well above 32 whenever samples is larger. count samples are then transferred into buffer[]: bmc150_accel_fifo_transfer(data, (u8 *)buffer, count); bmc150_accel_fifo_transfer() reads count * 6 bytes through regmap, so a malfunctioning, malicious or counterfeit accelerometer (or an attacker tampering with the I2C/SPI bus) that reports up to 127 frames writes up to 762 bytes into the 192-byte buffer: a stack out-of-bounds write of up to 570 bytes that clobbers the stack canary, saved registers and the return address. Clamp count to BMC150_ACCEL_FIFO_LENGTH, the number of samples buffer[] is sized for, before the transfer, mirroring the watermark clamp already done in bmc150_accel_set_watermark(). A well-formed flush reports at most BMC150_ACCEL_FIFO_LENGTH frames, so legitimate devices are unaffected.	N/A	More Details
CVE-2026-64505	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: function: rndis: add length check for header Add a length check for the rndis header in rndis_rm_hdr, to ensure that MessageType, MessageLength, DataOffset, and DataLength fields are present before they are accessed.	N/A	More Details
CVE-2026-64506	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw89: correct drop logic for malformed AMPDU frames The previous commit aims to fix issue caused by malformed AMPDU frames. But the drop logic fails to deal with the first AMPDU packet paired with certain range of sequence number, and leads to unexpected packet drop. It is more likely to encounter this failure when there are busy traffic during rekey process and could lead to disconnection from the AP. Fix this by adding a initial state judgement and only reset status during pairwise rekey.	N/A	More Details
CVE-2026-64508	In the Linux kernel, the following vulnerability has been resolved: bpf: Support for hardening against JIT spraying The BPF JIT allocator packs many small programs into larger executable allocations and reuses space within those allocations as programs are loaded and freed. When fresh code is written into space that a previous program occupied, an indirect jump into the new program can reuse a branch prediction left behind by the old one. Flush the indirect branch predictors before reusing JIT memory so that indirect jumps into a newly written program don't reuse predictions from an old program that occupied the same space. Introduce bpf_arch_pred_flush_enabled static key and bpf_arch_pred_flush static call for flushing the branch predictors on JIT memory reuse. Architectures that need a flush, can update it to a predictor flush function. By default, its a NOP and does not emit any CALL. Allocations larger than a pack are not covered by this flush. That is safe because cBPF programs (the unprivileged attack surface) are bounded well below a pack size. Issue a warning if this assumption is ever violated while the flush is active.	N/A	More Details
CVE-2026-64446	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix heap buffer overflow in rtw_cfg80211_set_wpa_ie() supplicant_ie is a 256-byte array in struct security_priv. The WPA and WPA2 IE copy paths use: memcpy(padapter->securitypriv.supplicant_ie, &pwpa[0], wpa_ielen + 2); where wpa_ielen is the raw IE length field (u8, 0-255). When a local user supplies a connect request via nl80211 with a crafted WPA IE of length 255, wpa_ielen + 2 equals 257, overflowing the 256-byte buffer by one byte into the adjacent last_mic_err_time field. rtw_parse_wpa_ie() does not prevent this: its length consistency check compares *(wpa_ie+1) against (u8) (wpa_ie_len-2), which is (u8)(255) == 255 when wpa_ie_len = 257, so the check passes silently. Add explicit bounds checks for both the WPA and WPA2 paths before the memcpy, rejecting any IE whose total size (wpa_ielen + 2) exceeds the supplicant_ie buffer.	N/A	More Details
CVE-2026-64509	In the Linux kernel, the following vulnerability has been resolved: rust: block: fix GenDisk cleanup paths GenDiskBuilder::build() still has fallible work after __blk_mq_alloc_disk(), but its error path only recovers the foreign queue data. That leaks the temporary gendisk and request_queue until later teardown. If the caller moved the last Arc<TagSet<T>> into build(), the leaked queue can retain blk-mq state after the tag set is dropped. Fix the pre-registration failure path by dropping the temporary gendisk reference with put_disk() before recovering queue_data, so disk_release() can tear down the owned queue. Also pair GenDisk::drop() with put_disk() after del_gendisk(). Once a Rust GenDisk has been added with device_add_disk(), del_gendisk() only unregisters it; the final gendisk reference still has to be dropped to complete the release path.	N/A	More Details
CVE-2026-64511	In the Linux kernel, the following vulnerability has been resolved: ACPI: NFIT: core: Fix possible NULL pointer dereference After commit 9b311b7313d6 ("ACPI: NFIT: Install Notify() handler before getting NFIT table"), acpi_nfit_probe() installs an ACPI notify handler for the NFIT device before checking the presence of the NFIT table. If that table is not there, 0 is returned without allocating the acpi_desc object and setting the driver data pointer of the NFIT device. If the platform firmware triggers an NFIT_NOTIFY_UC_MEMORY_ERROR notification on the NFIT device at that point, acpi_nfit_uc_error_notify() will dereference a NULL pointer. Prevent that from occurring by adding an acpi_desc check against NULL to acpi_nfit_uc_error_notify().	N/A	More Details
CVE-2026-64512	In the Linux kernel, the following vulnerability has been resolved: ACPI: CPPC: Suppress UBSAN warning caused by field misuse The definition of reg->access_width changes depending on the reg->space_id type. Type ACPI_ADR_SPACE_PLATFORM_COMM uses access_width to indicate the PCC region, which can result in a UBSAN if the value is greater than 4. For example: UBSAN: shift-out-of-bounds in drivers/acpi/cppc_acpi.c:1090:9 shift exponent 32 is too large for 32-bit type 'int' CPU: 61 UID: 0 PID: 1220 Comm: (udev-worker) Not tainted 7.0.10-201.fc44.aarch64 #1 PREEMPT(lazy) Hardware name: To be filled by O.E.M. Call trace: ...(trimming) ubsan_epilogue+0x10/0x48 __ubsan_handle_shift_out_of_bounds+0xdc/0x1e0 cpc_write+0x4d0/0x670 cppc_set_perf+0x18c/0x490 cppc_cpufreq_cpu_init+0x1c8/0x380 [cppc_cpufreq] ... (trimming) Lets fix this by validating the region type, as well as whether access_width has a value. Then since we are returning bit_width directly for ACPI_ADR_SPACE_PLATFORM_COMM, drop the code correcting the size.	N/A	More Details
	Next.js is a React framework for building full-stack web applications. In versions 14.1.1 through 15.5.20 and 16.0.0		

CVE-2026-64649	through 16.2.10, when a Server Action forwards or redirects a request, an attacker can cause the server to send that outbound request to a malicious host (Server-Side Request Forgery). This requires the attacker's request to control Host-associated headers. In some configurations, it's also possible to obtain internal values that weaken middleware/proxy authorization. Applications that use Server Actions are affected when the incoming host header is not fixed to a trusted value. This typically occurs on custom servers, or on deployments not behind a proxy that pins the host. Managed hosting pins the host upstream and is not affected; next start and standalone output do the same from version 14.2 onward. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-64648	Next.js is a React framework for building full-stack web applications. In versions 12.0.0 through 15.5.20 and 16.0.0 through 16.2.10, a server-side fetch with a request body may return a cached response body from a different request to the same URL but different body. Confidential data in the POST's response body would then leak to unauthorized requests. Though the request itself will not be deduped. This only applies to fetch calls with a request that has a different init than the one passed to fetch. A safe request would be: fetch(new Request(init), init). An unsafe request would be: fetch(new Request(init), aDifferentInit). This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-59729	Astro is a web framework for content-driven websites. Versions prior to 7.0.6 are vulnerable to XSS through unescaped spread attribute names in renderHTMLElement. The fix for CVE-2026-54298 (GHSA-jrpj-wcv7-9fh9) added an INVALID_ATTR_NAME_CHAR guard to addAttribute() so that spread-prop attribute names containing "' >/= or whitespace are dropped. A second attribute-rendering path, renderHTMLElement() in packages/astro/src/runtime/server/render/dom.ts, has its own inline attribute loop that does not go through addAttribute() and was not updated. It interpolates the attribute name unescaped and only escapes the value, so untrusted prop keys spread onto a native-HTMLElement-subclass component can still break out of the attribute context. This issue has been fixed in version 7.0.6.	N/A	More Details
CVE-2026-59727	Astro is a web framework for content-driven websites. In versions 3.10.0 through 7.0.3, when a transition:persist, transition:scope, or transition:persist-props directive is applied to a client-hydrated (client:*) component, Astro copied the directive value onto the rendered <astro-island> element without HTML-escaping it. If a developer reflects attacker-controlled input into one of these directives, an attacker can break out of the attribute and inject arbitrary HTML/JavaScript into the server-rendered output, resulting in reflected cross-site scripting (XSS). Exploitation requires the application developer to have written a non-idiomatic pattern — passing untrusted, request-derived input directly into a transition directive. Astro applications that do not route untrusted input into these directives are unaffected. This issue has been fixed in version 7.0.4.	N/A	More Details
CVE-2026-64513	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Unconditionally recompute CR8 intercept on PPR update The TPR_THRESHOLD field in the VMCS is used by VMX to induce VM exits when the guest's virtual TPR falls under the specified threshold, allowing KVM to inject previously masked interrupts. KVM handles these VM exits in handle_tpr_below_threshold(). Commit eb90f3417a0c ("KVM: vmx: speed up TPR below threshold vmexits") optimized this function by calling apic_update_ppr() instead of raising KVM_REQ_EVENT. apic_update_ppr() then raises KVM_REQ_EVENT if there is a pending, deliverable interrupt. However, if there are no new interrupts pending, apic_update_ppr() does not issue the request. Thus, kvm_lapic_update_cr8_intercept() and vmx_update_cr8_intercept() are not called before VM entry, which results in a high, stale TPR_THRESHOLD. This is problematic due to the following sentence in 28.2.1.1 "VM-Execution Control Fields" in the SDM: The following check is performed if the “use TPR shadow” VM-execution control is 1 and the “virtualize APIC accesses” and “virtual-interrupt delivery” VM-execution controls are both 0: the value of bits 3:0 of the TPR threshold VM-execution control field should not be greater than the value of bits 7:4 of VTPR. This error condition is typically not observed when KVM runs on a bare metal system because modern processors support APICv, which enables virtual-interrupt delivery, and which KVM uses when possible. This causes the processor to no longer generate TPR-below-threshold exits and to no longer check TPR_THRESHOLD on entry. However, when running on older platforms, or under nested virtualization on a hypervisor that does not support virtual-interrupt delivery and enforces this check (like Hyper-V) this can cause a VM entry failure with hardware error 0x7, as seen in [1]. Call kvm_lapic_update_cr8_intercept() if apic_update_ppr() does not find a deliverable interrupt (and thus does not raise KVM_REQ_EVENT). Remove calls to kvm_lapic_update_cr8_intercept() on paths that end up in apic_update_ppr(), as they now become redundant. This ensures that any path that updates the guest's PPR also figures out if KVM needs to wait for a TPR change (using TPR_THRESHOLD on VMX or CR8 intercepts on SVM).	N/A	More Details
CVE-2026-64514	In the Linux kernel, the following vulnerability has been resolved: userfaultfd: gate must_wait writability check on pte_present() userfaultfd_must_wait() and userfaultfd_huge_must_wait() read the PTE without taking the page table lock and then apply pte_write() / huge_pte_write() to it. Those accessors decode bits from the present encoding only; on a swap or migration entry they read the offset bits that happen to share the same position and return an undefined result. The intent of the check is "is this fault still WP-blocked?". A non-marker swap entry means the page is in transit -- the userfault context the original fault delivered against is no longer the same, and the swap-in or migration completion path will re-deliver a fresh fault if userspace still needs to handle it. Worst case under the current code the garbage write bit says "wait", and the thread stays asleep until a UFFDIO_WAKE that may never arrive. Gate the writability check on pte_present() so the lockless re-check only inspects present-PTE bits when the entry is actually present. The non-present, non-marker case returns "don't wait" and lets the fault path retry.	N/A	More Details
CVE-2026-64517	In the Linux kernel, the following vulnerability has been resolved: drm/xe/gsc: Fix double-free of managed BO in error path The error path in xe_gsc_init_post_hwconfig() explicitly frees a BO allocated with xe_managed_bo_create_pin_map() via xe_bo_unpin_map_no_vm(). Since the managed BO already has a devm cleanup action registered, this causes a double-free when devm unwinds during probe failure. Remove the explicit free and let devm handle it, consistent with all other xe_managed_bo_create_pin_map() callers. (cherry picked from commit 71d61e3e299a17139e47f980a4d6f425b2c59bf7)	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: tcp: Fix out-of-bounds access for twsk in tcp_ao_established_key(). lockdep_sock_is_held() was added in tcp_ao_established_key() by the cited commit. It can		

2026-64518	be called from tcp_v[46]_timewait_ack() with twsk. Since it does not have sk->sk_lock, the lockdep annotation results in out-of-bound access. $ pahole -C tcp_timewait_sock vmlinux | grep size /* size: 288, cachelines: 5, members: 8 */ $ pahole -C sock vmlinux | grep sk_lock socket_lock_t sk_lock; /* 440 192 */ Let's not use lockdep_sock_is_held() for TCP_TIME_WAIT.	N/A	More Details
CVE-2026-64519	In the Linux kernel, the following vulnerability has been resolved: NFSD: Fix infinite loop in layout state revocation find_one_sb_stid() skips stids whose sc_status is non-zero, but the SC_TYPE_LAYOUT case in nfsd4_revoke_states() never sets sc_status before calling nfsd4_close_layout(). The retry loop therefore finds the same layout stid on every iteration, hanging the revoker indefinitely.	N/A	More Details
CVE-2026-64492	In the Linux kernel, the following vulnerability has been resolved: iio: temperature: tmp006: use devm_iio_trigger_register tmp006_probe() allocates the DRDY trigger with devm_iio_trigger_alloc() but registers it with plain iio_trigger_register(). The driver has no .remove() callback, so on module unload the trigger stays in the global trigger list while its memory is freed by devm, leaving a dangling entry. Switch to devm_iio_trigger_register() so the registration is undone in the same devm scope as the allocation.	N/A	More Details
CVE-2026-64491	In the Linux kernel, the following vulnerability has been resolved: ALSA: usx2y: us144mkii: fix work UAF on disconnect tascam_disconnect() cancels capture_work and midi_in_work before usb_kill_anchored_urbs() kills the capture/MIDI-in URBs. Those URBs self-resubmit, and their completion handlers reschedule the work. A URB that completes in the small window between cancel_work_sync() and usb_kill_anchored_urbs() therefore re-arms the work after its only cancel. Nothing cancels it again before snd_card_free() frees the card-private tascam structure, so the work handler then runs on freed memory. Kill the anchored URBs before cancelling the work; once the work is cancelled no remaining URB can complete to re-arm it.	N/A	More Details
CVE-2026-64489	In the Linux kernel, the following vulnerability has been resolved: ALSA: ymfpci: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. snd_ymfpci_create_spdif_controls() does not check the return value before dereferencing kctl->id.device, which can lead to a NULL pointer dereference. Add NULL checks after snd_ctl_new1() calls and return -ENOMEM if any fails.	N/A	More Details
CVE-2026-64488	In the Linux kernel, the following vulnerability has been resolved: ALSA: aoa: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. In layout.c, the function does not check the return value before dereferencing ctl->id.name or passing to aoa_snd_ctl_add(), which can lead to a NULL pointer dereference. Add NULL checks after snd_ctl_new1() calls and return early if any fails.	N/A	More Details
CVE-2026-64453	In the Linux kernel, the following vulnerability has been resolved: usb: misc: usbio: fix disconnect UAF in client teardown usbio_disconnect() walks usbio->cli_list in reverse and uninitializes each auxiliary device. auxiliary_device_uninit() drops the device reference, and for an unbound child that can run usbio_auxdev_release() and free the containing struct usbio_client. list_for_each_entry_reverse() advances after the loop body by reading client->link.prev. If the current client is freed by auxiliary_device_uninit(), the iterator dereferences freed memory. Use list_for_each_entry_safe_reverse() so the previous client is cached before the body can drop the final reference. This preserves reverse teardown order while keeping the next iterator cursor independent of the current client's lifetime. Validation reproduced this kernel report: BUG: KASAN: slab-use-after-free in usbio_disconnect+0x12e/0x150 Call Trace: <TASK> dump_stack_lvl+0x66/0xa0 print_report+0xce/0x630 ? usbio_disconnect+0x12e/0x150 ? srso_alias_return_thunk+0x5/0xfbef5 ? __virt_addr_valid+0x188/0x320 ? usbio_disconnect+0x12e/0x150 kasan_report+0xe0/0x110 ? usbio_disconnect+0x12e/0x150 usbio_disconnect+0x12e/0x150 usb_unbind_interface+0xf3/0x400 really_probe+0x316/0x660 __driver_probe_device+0x106/0x240 driver_probe_device+0x4a/0x110 __device_attach_driver+0xf1/0x1a0 ? __pfx___device_attach_driver+0x10/0x10 bus_for_each_drv+0xf9/0x160 ? __pfx_bus_for_each_drv+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? trace_hardirqs_on+0x18/0x130 ? srso_alias_return_thunk+0x5/0xfbef5 ? _raw_spin_unlock_irqrestore+0x44/0x60 __device_attach+0x133/0x2a0 ? __pfx___device_attach+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? do_raw_spin_unlock+0x9a/0x100 ? srso_alias_return_thunk+0x5/0xfbef5 device_initial_probe+0x55/0x70 bus_probe_device+0x4a/0xd0 device_add+0x9b9/0xc10 ? __pfx_device_add+0x10/0x10 ? _raw_spin_unlock_irqrestore+0x44/0x60 ? srso_alias_return_thunk+0x5/0xfbef5 ? lockdep_hardirqs_on_prepare+0xea/0x1a0 ? srso_alias_return_thunk+0x5/0xfbef5 ? usb_enable_lpm+0x3c/0x260 usb_set_configuration+0xb64/0xf20 usb_generic_driver_probe+0x5f/0x90 usb_probe_device+0x71/0x1b0 really_probe+0x46b/0x660 __driver_probe_device+0x106/0x240 driver_probe_device+0x4a/0x110 __device_attach_driver+0xf1/0x1a0 ? __pfx___device_attach_driver+0x10/0x10 bus_for_each_drv+0xf9/0x160 ? __pfx_bus_for_each_drv+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? trace_hardirqs_on+0x18/0x130 ? srso_alias_return_thunk+0x5/0xfbef5 ? _raw_spin_unlock_irqrestore+0x44/0x60 __device_attach+0x133/0x2a0 ? __pfx___device_attach+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? do_raw_spin_unlock+0x9a/0x100 ? srso_alias_return_thunk+0x5/0xfbef5 device_initial_probe+0x55/0x70 bus_probe_device+0x4a/0xd0 device_add+0x9b9/0xc10 ? __pfx_device_add+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? add_device_randomness+0xb7/0xf0 usb_new_device+0x492/0x870 hub_event+0x1b10/0x29c0 ? __pfx_hub_event+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_acquire+0x187/0x300 ? process_one_work+0x475/0xb90 ? srso_alias_return_thunk+0x5/0xfbef5 ? lock_release+0xc8/0x290 ? srso_alias_return_thunk+0x5/0xfbef5 process_one_work+0x4d7/0xb90 ? __pfx_process_one_work+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? srso_alias_return_thunk+0x5/0xfbef5 ? __list_add_valid_or_report+0x37/0xf0 ? __pfx_hub_event+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 worker_thread+0x2d8/0x570 ? __pfx_worker_thread+0x10/0x10 kthread+0x1ad/0x1f0 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x3c9/0x540 ? __pfx_ret_from_fork+0x10/0x10 ? srso_alias_return_thunk+0x5/0xfbef5 ? __switch_to+0x2e9/0x730 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1a/0x30 </TASK>	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: run gadget disconnect from sleepable suspend context dwc3_gadget_suspend() takes dwc->lock with IRQs disabled and then calls dwc3_disconnect_gadget(). For async callbacks that helper only uses plain spin_unlock()/spin_lock(), so the gadget -		

CVE-2026-64454	>disconnect() callback still runs with IRQs disabled and any sleepable callback trips Lockdep. This issue was found by our static analysis tool and then manually reviewed against the current tree. The grounded PoC kept the dwc3_gadget_suspend() -> dwc3_disconnect_gadget() -> gadget_driver->disconnect() chain, and Lockdep reported: BUG: sleeping function called from invalid context gadget_disconnect+0x21/0x39 [vuln_msv] dwc3_gadget_suspend.constprop.0+0x2b/0x42 [vuln_msv] Keep the disconnect callback selection in one common helper, but add a sleepable suspend-side wrapper which snapshots the callback under dwc->lock and then runs it after spin_unlock_irqrestore(). The regular event path still uses the existing spin_unlock()/spin_lock() window.	N/A	More Details
CVE-2026-64455	In the Linux kernel, the following vulnerability has been resolved: USB: chaoskey: Fix slab-use-after-free in chaoskey_release() The chaoskey driver has a use-after-free bug in its release routine. If the user closes the device file after the USB device has been unplugged, a debugging log statement will try to access the usb_interface structure after it has been deallocated: BUG: KASAN: slab-use-after-free in dev_driver_string (drivers/base/core.c:2406) Read of size 8 at addr ffff888168e8a0b8 by task chaoskey_raw_re/10106 Hardware name: QEMU Ubuntu 24.04 PC v2 (i440FX + PIIX, arch_caps fix, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 Call Trace: <TASK> dump_stack_lvl (lib/dump_stack.c:94 lib/dump_stack.c:120) print_report (mm/kasan/report.c:378 mm/kasan/report.c:482) kasan_report (mm/kasan/report.c:595) dev_driver_string (drivers/base/core.c:2406) __dynamic_dev_dbg (lib/dynamic_debug.c:906) chaoskey_release (drivers/usb/misc/chaoskey.c:323) __fput (fs/file_table.c:510) fput_close_sync (fs/file_table.c:615) __x64_sys_close (fs/open.c:1507 fs/open.c:1492 fs/open.c:1492) do_syscall_64 (arch/x86/entry/syscall_64.c:63 arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121) The driver's last reference to the interface structure is dropped in the chaoskey_free() routine, so the code must not use the interface -- even in a debugging statement -- after that routine returns. (Exception: If we know that another reference is held by someone else, such as the device core while the disconnect routine runs, there's no problem. Thanks to Johan Hovold for pointing this out.) Since the bad access is part of an unimportant debugging statement, we can fix the problem simply by removing the whole statement.	N/A	More Details
CVE-2026-64457	In the Linux kernel, the following vulnerability has been resolved: virtio_pci: fix vq info pointer lookup via wrong index Unbinding a virtio balloon device: echo virtio0 > /sys/bus/virtio/drivers/virtio_balloon/unbind triggers a NULL pointer dereference. The dmesg says: BUG: kernel NULL pointer dereference, address: 0000000000000008 [...] RIP: 0010:__list_del_entry_valid_or_report+0x5/0xf0 Call Trace: <TASK> vp_del_vqs+0x121/0x230 remove_common+0x135/0x150 virtballoon_remove+0xee/0x100 virtio_dev_remove+0x3b/0x80 device_release_driver_internal+0x187/0x2c0 unbind_store+0xb9/0xe0 kernfs_fop_write_iter.llvm.11660790530567441834+0xf6/0x180 vfs_write+0x2a9/0x3b0 ksys_write+0x5c/0xd0 do_syscall_64+0x54/0x230 entry_SYSCALL_64_after_hwframe+0x29/0x31 [...] </TASK> The virtio_balloon device registers 5 queues (inflate, deflate, stats, free_page, reporting) but only the first two are unconditional. The stats, free_page and reporting queues are each conditional on their respective feature bits. When any of these features are absent, the corresponding vqs_info entry has name == NULL, creating holes in the array. The root cause is an indexing mismatch introduced when vq info storage was changed to be passed as an argument. vp_find_vqs_msix() and vp_find_vqs_intx() store the info pointer at vp_dev->vqs[i], where 'i' is the caller's sparse array index. However, the virtqueue itself gets vq->index assigned from queue_idx, a dense index that skips NULL entries. When holes exist, 'i' and queue_idx diverge. Later, vp_del_vqs() looks up info via vp_dev->vqs[vq->index] using the dense index into the sparsely-populated array, and hits NULL. Fix this by storing info at vp_dev->vqs[queue_idx] instead of vp_dev->vqs[i], so the store index matches the lookup index (vq->index). Apply the fix to both the MSIX and INTX paths.	N/A	More Details
CVE-2026-64458	In the Linux kernel, the following vulnerability has been resolved: mm/damon/ops-common: handle extreme intervals in damon_hot_score() Fix three issues in damon_hot_score() that comes from wrong handling of extreme (zero or too high) monitoring intervals user setup. When the user sets sampling interval zero, damon_max_nr_accesses(), which is called from damon_hot_score(), causes a divide-by-zero. Needless to say, it is a problem. When the user sets the aggregation interval zero, the function returns zero. It is wrong, since the real maximum nr_acceses in the setup should be one. Worse yet, it can cause another divide-by-zero from its caller, damon_hot_score(), since it uses damon_max_nr_accesses() return value as a denominator. When the user sets the aggregation interval very high, damon_hot_score() could return a value out of [0, DAMOS_MAX_SCORE] range. Since the return value is used as an index to the regions_score_histogram array, which is DAMOS_MAX_SCORE+1 size, it causes out of bounds array access. The issues can be relatively easily reproduced like below. The sysfs write permission is required, though. # ./damo start --damos_action lru_prio --damos_quota_space 100M \ --damos_quota_interval 1s # cd /sys/kernel/mm/damon/admin/kdamonds/0 # echo 0 > contexts/0/monitoring_attrs/intervals/sample_us # echo 0 > contexts/0/monitoring_attrs/intervals/aggr_us # echo commit > state # dmesg [...] [131.329762] Oops: divide error: 0000 [#1] SMP NOPTI [...] [131.336089] RIP: 0010:damon_hot_score+0x27/0xd0 [...] Fix the divide-by-zero intervals problems by explicitly handling the zero intervals in damon_max_nr_accesses(). Fix the out-of-bound array access by applying [0, DAMOS_MAX_SCORE] bounds before returning from damon_hot_score(). The issue was discovered [1] by Sashiko.	N/A	More Details
CVE-2026-64461	In the Linux kernel, the following vulnerability has been resolved: PCI: mediatek: Fix IRQ domain leak when port fails to enable When mtk_pcie_enable_port() fails, mtk_pcie_port_free() removes the port from pcie->ports and frees the port structure. However, the IRQ domains set up earlier by mtk_pcie_init_irq_domain() are never freed. Fix this by refactoring mtk_pcie_irq_teardown() into a per-port helper, mtk_pcie_irq_teardown_port(), and calling it from mtk_pcie_setup() when mtk_pcie_enable_port() fails. Since the IRQ teardown must only happen in the probe error path (during resume, child devices may have active MSI mappings and the NOIRQ context prohibits sleeping locks), mtk_pcie_enable_port() is changed to return an error code so callers can distinguish the two paths and act accordingly. This issue was reported by Sashiko while reviewing the EcoNet EN7528 SoC support series.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: PCI: altera: Fix resource leaks on probe failure The chained IRQ handler is set during probe, but is only removed during the driver remove(). If pci_host_probe() fails, the handler and INTx IRQ domain remain set even though the devm-managed host bridge storage containing struct altera_pcie will be released, leaving the handler with a stale data pointer. Interrupts are also enabled before	N/A	More

64462	pci_host_probe() is called. If probe fails after that point, the controller interrupt source should be disabled before the chained handler and INTx domain are removed. So set the chained handler only after the INTx domain has been created. Disable controller interrupts during IRQ teardown, and tear the IRQ setup down if pci_host_probe() fails. [mani: commit log]		Details
CVE-2026-64464	In the Linux kernel, the following vulnerability has been resolved: xhci: sideband: fix ring sg table pages leak xhci_ring_to_sgtable() allocates a temporary pages array and uses it to build the returned sg_table with sg_alloc_table_from_pages(). The error paths free the pages array, but the success path returns the sg_table without freeing it. This leaks the temporary array every time a sideband client gets an endpoint or event ring buffer. Free the pages array after sg_alloc_table_from_pages() succeeds. The returned sg_table has its own scatterlist entries and does not depend on the temporary array after construction.	N/A	More Details
CVE-2026-64465	In the Linux kernel, the following vulnerability has been resolved: usb: xhci: Fix sleep in atomic context in xhci_free_streams() When a USB device with active stream endpoints is disconnected, xhci_free_streams() is called from the hub_event workqueue to free the stream resources. It calls xhci_free_stream_info() while holding xhci->lock with irqs disabled. xhci_free_stream_info() invokes xhci_free_stream_ctx(), which calls dma_free_coherent() for large stream context arrays. dma_free_coherent() can sleep (e.g. via vunmap), triggering a BUG when called from atomic context. Call trace: dma_free_attrs+0x174/0x220 xhci_free_stream_info+0xd0/0x11c xhci_free_streams+0x278/0x37c usb_free_streams+0x98/0xc0 usb_unbind_interface+0x1b8/0x2f8 device_release_driver_internal+0x1d4/0x2cc device_release_driver+0x18/0x28 bus_remove_device+0x160/0x1a4 device_del+0x1ec/0x350 usb_disable_device+0x98/0x214 usb_disconnect+0xf0/0x35c hub_event+0xab4/0x19ec process_one_work+0x278/0x63c Fix this by saving the stream_info pointers and clearing the ep references under the lock, then calling xhci_free_stream_info() outside the lock where sleeping is allowed.	N/A	More Details
CVE-2026-64466	In the Linux kernel, the following vulnerability has been resolved: rust_binder: clear freeze listener on node removal Generally userspace is supposed to explicitly clear freeze listeners before they drop the refcount on the node ref to zero, but there's nothing forcing that. Currently, in this scenario the freeze listener remains in the freeze_listeners rbtree and in the remote node's freeze listener list, even though the ref for which the listener is registered is gone. This could potentially lead to a memory leak due to a refcount cycle. Thus, remove the freeze listener in this scenario.	N/A	More Details
CVE-2026-64470	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btusb: fix use-after-free on marvell probe failure Make sure to stop any TX URBs submitted during Marvell OOB wakeup configuration on later probe failures to avoid use-after-free in the completion callback. This issue was reported by Sashiko while reviewing a fix for a wakeup source leak in the btusb probe errors paths.	N/A	More Details
CVE-2026-64471	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btusb: fix use-after-free on registration failure Make sure to release the sibling interfaces in case controller registration fails to avoid use-after-free and double-free when they are eventually disconnected. This issue was reported by Sashiko while reviewing a fix for a wakeup source leak in the btusb probe errors paths.	N/A	More Details
CVE-2026-64472	In the Linux kernel, the following vulnerability has been resolved: vfio/mlx5: Fix racy bitfields and tighten struct layout Bitfield operations are not atomic, they use a read-modify-write pattern, therefore we should be careful not to pack bitfields that can be concurrently updated into the same storage unit. This split takes a binary approach: flags that are only modified pre/post open/close remain bitfields, flags modified from user action, including actions that reach across to another device (ex. reset) use dedicated storage units. Note mlx5_vhca_page_tracker.status is relocated to fill the alignment hole this split exposes. Bitfield justifications: migrate_cap: written only in mlx5vf_cmd_set_migratable() at probe chunk_mode: written only in mlx5vf_cmd_set_migratable() at probe mig_state_cap: written only in mlx5vf_cmd_set_migratable() at probe Dedicated storage units: mdev_detach: written in the VF attach/detach event notifier mlx5fv_vf_event() at runtime log_active: written in mlx5vf_start_page_tracker()/ mlx5vf_stop_page_tracker() during runtime dirty tracking deferred_reset: written in mlx5vf_state_mutex_unlock()/ mlx5vf_pci_aer_reset_done() during runtime reset handling is_err: set by tracker error handling and dirty-log polling at runtime object_changed: set by tracker event handling and cleared by dirty-log polling at runtime	N/A	More Details
CVE-2026-64473	In the Linux kernel, the following vulnerability has been resolved: vfio: Remove device debugfs before releasing devres VFIO device debugfs files created with debugfs_create_devm_seqfile() store a devres allocated debugfs_devm_entry as inode private data. vfio_unregister_group_dev() currently calls vfio_device_del() before vfio_device_debugfs_exit(), but device_del() releases devres. This can leave debugfs entries visible with stale inode private data while unregister waits for userspace references to drain. Remove the per-device debugfs tree before vfio_device_del(). The debugfs view is diagnostic only, so losing it at the start of unregister is preferable to preserving entries whose backing storage may already have been released. Complete the teardown by clearing the per-device debugfs root after removal. This matches the global debugfs root cleanup and prevents future users from mistaking a removed dentry for a live debugfs tree during the remainder of unregister.	N/A	More Details
CVE-2026-64474	In the Linux kernel, the following vulnerability has been resolved: vfio: prevent infinite loop in vfio_mig_get_next_state() on blocked arc vfio_mig_get_next_state() walks vfio_from_fsm_table[] one step at a time, looping to skip optional states the device does not support until *next_fsm is supported. A blocked transition is encoded as VFIO_DEVICE_STATE_ERROR, which the trailing return reports as -EINVAL. The skip loop does not account for the ERROR sentinel. state_flags_table[ERROR] is ~0U and vfio_from_fsm_table[ERROR][*] is ERROR, so once *next_fsm becomes ERROR the loop condition stays true and *next_fsm never changes. The blocked arcs STOP_COPY -> PRE_COPY and STOP_COPY -> PRE_COPY_P2P map to ERROR yet pass the support check on a precopy-capable device, causing the loop to spin forever while holding the driver state mutex. This can result in a soft lockup, and a panic with softlockup_panic set. Terminate the skip loop on the ERROR sentinel so a blocked transition falls through to the existing return and reports -EINVAL.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: vfio/pci: Latch disable_idle_d3 per device When		

CVE-2026-64476	disable_idle_d3 was introduced in vfio-pci, it directly manipulated the device power state with pci_set_power_state(). There were no refcounts to maintain or balanced operations, we could unconditionally bring the device to D0 and conditionally move it to D3hot. Therefore the module parameter was made writable. Later, in commit c61302aa48f7 ("vfio/pci: Move module parameters to vfio_pci.c"), as part of the vfio-pci-core split, the writable aspect of the module parameter was nullified. The parameter value could still be changed through sysfs, but the vfio-pci driver latched the values into vfio-pci-core globals at module init. Loading the vfio-pci module, or unloading and reloading, with non-default or different values could change the globals relative to existing devices bound to vfio-pci variant drivers. Runtime PM was introduced in commit 7ab5e10eda02 ("vfio/pci: Move the unused device into low power state with runtime PM"), which marks the point where power states became refcounted. PM get and put operations need to be balanced, but the same module operations noted above can change the global variables relative to those devices already bound to vfio-pci variant drivers. This introduces a window where PM operations can now become unbalanced. To resolve this with a narrow footprint for stable backports, the disable_idle_d3 flag is latched into the vfio_pci_core_device at the time of initialization, such that the device always operates with a consistent value. NB. vfio_pci_dev_set_try_reset() now unconditionally raises the runtime PM usage count around bus reset to account for disable_idle_d3 becoming a per-device rather than global flag. When this flag is set, the additional get/put pair is harmless and allows continued use of the shared vfio_pci_dev_set_pm_runtime_get() helper.	N/A	More Details
CVE-2026-64477	In the Linux kernel, the following vulnerability has been resolved: x86,fs/resctrl: Prevent out-of-bounds access while offlining CPU when SNC enabled The architecture updates the cpu_mask in a domain's header to track which online CPUs are associated with the domain. When this mask becomes empty the architecture initiates offline of the domain that includes calling on resctrl fs to offline the domain. If it is a monitoring domain in which LLC occupancy is tracked resctrl fs forces the limbo handler to clear all busy RMID state associated with the domain. The limbo handler always reads the current event value associated with a busy RMID irrespective of it being checked as part of regular "is it still busy" check or whether it will be forced released anyway. When reading an RMID on a system with SNC enabled the "logical RMID" is converted to the "physical RMID" and this conversion requires the NUMA node ID of the resctrl monitoring domain that is in turn determined by querying the NUMA node ID of any CPU belonging to the monitoring domain. When the monitoring domain is going offline its cpu_mask is empty causing the NUMA node ID query via cpu_to_node() to be done with "nr_cpu_ids" as argument resulting in an out-of-bounds access. Refactor the limbo handler to skip reading the RMID when the RMID will just be forced to no longer be dirty in the domain anyway. Add a safety check to the architecture's RMID reader to protect against this scenario.	N/A	More Details
CVE-2026-64478	In the Linux kernel, the following vulnerability has been resolved: ALSA: usb-audio: avoid kobject path lookup in DualSense match The DualSense jack-detection input handler verifies that a matching input device belongs to the same physical controller by building kobject path strings for both the input device and the USB audio device, then comparing the path prefix. This was observed when a weak physical connection caused the controller to rapidly disconnect and reconnect. During that repeated hotplug, snd_dualsense_ih_match() can run while the controller's USB device is being disconnected. kobject_get_path() walks ancestor kobjects and dereferences their names; if the USB device kobject name is no longer valid, this can fault in strlen(): RIP: 0010:strlen+0x10/0x30 Call Trace: kobject_get_path+0x34/0x150 snd_dualsense_ih_match+0x49/0xd0 [snd_usb_audio] input_register_device+0x566/0x6a0 ps_probe+0xb89/0x1590 [hid_playstation] The same ownership check can be done without building kobject path strings. The input device is parented below the HID device, USB interface and USB device, so walking the input device parent chain and comparing against the mixer USB device preserves the check without dereferencing kobject names during disconnect.	N/A	More Details
CVE-2026-64479	In the Linux kernel, the following vulnerability has been resolved: ALSA: seq: Fix uninitialised heap leak in snd_seq_event_dup() snd_seq_event_dup() copies an incoming event into a pool cell and, in the UMP-enabled build, clears the trailing cell->ump.raw.extra word that the memcpy() did not cover. The guard deciding whether to clear it compares the copied size against sizeof(cell->event): memcpy(&cell->ump, event, size); if (size < sizeof(cell->event)) cell->ump.raw.extra = 0; For a legacy (non-UMP) event, size == sizeof(struct snd_seq_event) == sizeof(cell->event), so the condition is false and the extra word keeps stale data. The cell pool is allocated with kvmalloc() (not zeroed) and cells are reused via a free list, so that word holds uninitialised heap or leftover event data. When such a cell is delivered to a UMP client (client->midi_version > 0) that set SNDRV_SEQ_FILTER_NO_CONVERT -- so the legacy event reaches it unconverted -- snd_seq_read() reads it out as the larger struct snd_seq_ump_event and copies the stale word to user space, a 4-byte kernel heap infoleak to an unprivileged /dev/snd/seq client. Compare against sizeof(cell->ump) instead, so the trailing word is zeroed for every event shorter than the UMP cell.	N/A	More Details
CVE-2026-64480	In the Linux kernel, the following vulnerability has been resolved: ALSA: ice1712: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. The ice1712 driver calls snd_ctl_new1() without checking the return value before dereferencing the pointer in multiple places (ice1712.c, ice1724.c, aureon.c), which can lead to NULL pointer dereferences. Add NULL checks after snd_ctl_new1() calls and return -ENOMEM if any fails.	N/A	More Details
CVE-2026-64482	In the Linux kernel, the following vulnerability has been resolved: ALSA: gus: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. snd_gf1_pcm_volume_control() does not check the return value before dereferencing kctl->id.index, which can lead to a NULL pointer dereference. Add a NULL check after snd_ctl_new1() and return -ENOMEM if it fails.	N/A	More Details
CVE-2026-64483	In the Linux kernel, the following vulnerability has been resolved: ALSA: firewire: isight: bound the sample count to the packet payload isight_packet() takes the frame count from the device iso packet and checks it only against the device claimed iso length. count = be32_to_cpu(payload->sample_count); if (likely(count <= (length - 16) / 4)) isight_samples(isight, payload->samples, count); length is the iso header data_length. It can be up to 0xffff. So the gate allows a count up to about 16379. isight_samples() then copies count frames out of payload->samples into the PCM DMA buffer. payload->samples holds only 2 * MAX_FRAMES_PER_PACKET values. The device multiplexes two samples per frame. A count past MAX_FRAMES_PER_PACKET reads past the payload. A count past the buffer size writes past runtime->dma_area. The smallest PCM buffer is larger than MAX_FRAMES_PER_PACKET. Bounding the	N/A	More Details

	count to MAX_FRAMES_PER_PACKET keeps both the read and the write in range. A malicious or faulty Apple iSight on the FireWire bus reaches this during a normal capture. Add the MAX_FRAMES_PER_PACKET bound to the gate.		
CVE-2026-64484	In the Linux kernel, the following vulnerability has been resolved: ALSA: es1938: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. snd_es1938_mixer() does not check the return value before dereferencing the pointer, which can lead to a NULL pointer dereference. Add a NULL check after snd_ctl_new1() and return -ENOMEM if it fails.	N/A	More Details
CVE-2026-64486	In the Linux kernel, the following vulnerability has been resolved: ALSA: cmipci: check snd_ctl_new1() return value snd_ctl_new1() can return NULL when memory allocation fails. snd_cmipci_spdif_controls() does not check the return value before dereferencing kctl->id.device, which can lead to a NULL pointer dereference. Add NULL checks after snd_ctl_new1() calls and return -ENOMEM if any fails.	N/A	More Details
CVE-2026-64487	In the Linux kernel, the following vulnerability has been resolved: ALSA: caiaq: fix out-of-bounds read in the Traktor Kontrol S4 input parser snd_usb_caiaq_tks4_dispatch() decodes the Traktor Kontrol S4 input stream in fixed 16-byte (TKS4_MSGBLOCK_SIZE) message blocks. On every iteration it advances buf and subtracts the block size while looping on "while (len)". len is urb->actual_length. That value is supplied by the device and is not guaranteed to be a multiple of 16. When a final short block leaves len between 1 and 15, the loop runs once more, reads up to buf[15], and then does "len -= TKS4_MSGBLOCK_SIZE". As len is unsigned this underflows to a huge value. The loop then keeps iterating and walking buf far past the end of the 512-byte ep4_in_buf, reading out of bounds until a bogus block id happens to be hit. Iterate only while a full message block is available. This stops the unsigned underflow and silently drops any trailing partial block, which carries no complete control value anyway. The sibling endpoint-4 parsers are not affected. The Traktor Kontrol X1 and Maschine arms in snd_usb_caiaq_ep4_reply_dispatch() floor urb->actual_length before dispatching.	N/A	More Details
CVE-2026-64521	In the Linux kernel, the following vulnerability has been resolved: pinctrl: meson: amlogic-a4: fix deadlock issue Accessing the pinconf-pins sysfs node may deadlock. pinconf_pins_show() holds pctldev->mutex, and the platform driver calls pinctrl_find_gpio_range_from_pin(), which tries to acquire the same mutex again, leading to a deadlock. Use pinctrl_find_gpio_range_from_pin_nolock() to fix this issue.	N/A	More Details
CVE-2026-64525	In the Linux kernel, the following vulnerability has been resolved: xfrm: move policy_bydst RCU sync from per-netns .exit to .pre_exit The struct pernet_operations docstring in include/net/net_namespace.h explicitly warns against blocking RCU primitives in .exit handlers: Exit methods using blocking RCU primitives, such as synchronize_rcu(), should be implemented via exit_batch. [...] Please, avoid synchronize_rcu() at all, where it's possible. Note that a combination of pre_exit() and exit() can be used, since a synchronize_rcu() is guaranteed between the calls. xfrm_policy_fini() violates this: it calls synchronize_rcu() before freeing the policy_bydst hash tables (so no RCU reader is mid- traversal at free time), but runs from xfrm_net_ops.exit -- once per namespace -- so a cleanup_net() of N namespaces pays N full RCU grace periods serially. Use the documented pre_exit/exit split. Move the policy flush (and the workqueue drains it depends on) into a new .pre_exit handler; xfrm_policy_fini() then runs in .exit and frees the hash tables after the synchronize_rcu_expedited() that cleanup_net() guarantees between the two phases. Providing O(1) RCU grace periods per batch instead of O(N). Observed on Linux 6.18 with a workload doing unshare(CLONE_NEWNET) at ~13/sec sustained: cleanup_net() and the netns_wq rescuer kthread both stuck in xfrm_policy_fini()'s synchronize_rcu(), >300k struct net accumulated in the cleanup queue, Percpu in /proc/meminfo climbed to 130+ GB on 256-CPU hosts, and memcg OOMs followed. setup_net and __put_net counts were balanced, ruling out a refcount leak.	N/A	More Details
CVE-2026-64526	In the Linux kernel, the following vulnerability has been resolved: ethtool: tsconfig: fix missing ethnl_ops_complete() tsconfig_prepare_data() calls ethnl_ops_begin(), we need to call ethnl_ops_complete() before returning the error.	N/A	More Details
CVE-2026-65765	Joomla Extension - phoca.cz - Path Traversal vulnerability in Phoca Commander 1.0.0-6.1.1 - Improper limitation of paths for save and download actions lead to path traversal vulnerabilities.	N/A	More Details
CVE-2026-59250	Classic buffer overflow in the Erlang/OTP megaco flex scanner C driver allows a remote unauthenticated attacker to corrupt the driver's memory (and potentially achieve remote code execution or a denial-of-service crash) by sending a single text-encoded H.248/Megaco message containing an oversized property parm name. When tokenizing a Local/Remote descriptor, mfs_load_property_groups extracts the attacker-controlled property name (bounded only by the message length) and, when no value follows, formats it into a fixed 512-byte error_msg field of the MfsErlDrvData struct using an unchecked sprintf call. Names longer than roughly 452 bytes overflow into the immediately following struct fields (text_buf, text_ptr, term_spec, term_spec_size, term_spec_index), overwriting live pointers and counters with attacker-chosen bytes. Subsequent scanner code writes and frees through the corrupted pointers, producing arbitrary write and arbitrary free primitives inside the BEAM VM process, which can be leveraged for remote code execution. On builds compiled with _FORTIFY_SOURCE the overflow is detected at runtime and terminates the process with SIGABRT, resulting in denial of service. The overflow occurs in the flex scanner before any grammar or Megaco-level authentication processing, so exploitation requires only network reachability to the megaco transport port on a node configured with {scanner, flex}. This vulnerability is associated with program files lib/megaco/src/flex/megaco_flex_scanner_drv.flex.src and program routines mfs_load_property_groups. This issue affects OTP from OTP 17.0 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to megaco from 3.17.1 before 4.9.1, 4.8.3.1 and 4.7.2.2. Versions prior to OTP 17.0 are also affected but are not listed because the OTP version scheme is only defined from OTP 17.0 onwards.	N/A	More Details
	The Erlang/OTP ssl application does not detect cycles when reconstructing an incomplete peer certificate chain during a TLS or DTLS handshake. In ssl_certificate:handle_incomplete_chain/5, the received chain is passed to ssl_certificate:build_certificate_chain/5, which walks issuer relationships via ssl_certificate:do_certificate_chain/7 with no cycle detection and no depth limit. When the peer supplies two mutually cross-signed certificates in unordered		

CVE-2026-58227	form (A issues B, B issues A), the issuer lookup alternates between the two certificates and the pair of functions recurses indefinitely, growing the call stack and chain accumulator without bound. An unauthenticated remote attacker can send a crafted certificate chain in a TLS or DTLS Certificate handshake message to exhaust available memory and crash the BEAM node. Only a TCP connection and a partial handshake are required; no authentication or completed handshake is needed, and both TLS/DTLS servers and clients are affected when processing peer certificate messages. This issue affects OTP from OTP 23.2 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to ssl from 10.2 before 11.7.4, 11.6.0.4 and 11.2.12.11.	N/A	More Details
CVE-2026-55953	The Erlang/OTP ssl TLS 1.2 (and earlier) and DTLS client does not verify that the cipher suite selected by the server in ServerHello was among the suites offered by the client in ClientHello. The client-side tls_handshake:hello/5 handler validates the negotiated protocol version and the downgrade sentinel but hands the server-chosen suite directly to ssl_handshake:handle_server_hello_extensions/9, which installs it without a membership check. The TLS 1.3 client path performs this check (per RFC 8446), so it is not affected. An on-path attacker between the client and the intended server can respond with a ServerHello selecting an anonymous key exchange suite such as TLS_DH_anon_* or TLS_ECDH_anon_* that the client never offered. Anonymous suites do not require the server to present a certificate, so the entire verify_peer and cacerts configuration is bypassed: the attacker completes the handshake with its own ephemeral parameters, no certificate is validated, no hostname is checked, and ssl:connect returns {ok, Socket}. All subsequent application traffic is readable and modifiable by the attacker. This issue affects OTP from OTP 17.0 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to ssl from 5.3.4 before 11.7.4, 11.6.0.4 and 11.2.12.11.	N/A	More Details
CVE-2026-55737	Signed to Unsigned Conversion Error and Out-of-bounds Write vulnerability in Erlang OTP erts allows an attacker who can supply a crafted Erlang external term format (ETF) binary to binary_to_term/1 to corrupt the BEAM heap pointer and crash the virtual machine. When decoding a LARGE_TUPLE_EXT term, the validation pass decoded_size() in erts/emulator/beam/external.c reads the 32-bit arity field as unsigned (get_uint32()), while the decode pass dec_term() reads the same field as a signed 32-bit integer (get_int32()) into an int. An arity wire value of 0x80000000 passes validation as 2147483648 but decodes as -2147483648, so the subsequent hp += n moves the heap allocation pointer backward. Neither pass enforces the runtime tuple-arity limit MAX_ARITYVAL. The result is an out-of-bounds heap write; in practice the VM detects an impossible heap size and aborts, denying service. The required padding is large when uncompressed but the compressed-ETF envelope shrinks it to a small payload on the wire. This issue affects OTP from OTP 25.0 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to erts from 13.0 before 17.0.4, 16.4.0.4 and 15.2.7.11.	N/A	More Details
CVE-2026-54890	Integer Underflow (Wrap or Wraparound) vulnerability in erlang otp erlang/otp (erts modules), erlang otp erts (erts modules) allows Forced Integer Overflow, Excessive Allocation. This vulnerability is associated with program files erts/emulator/beam/external.c, emulator/beam/external.c. The BIT_BINARY_EXT tag (77) handler in the External Term Format (ETF) decoder accepts an encoding with both length and trailing-bits fields set to zero. The subsequent computation of the bitstring size underflows an unsigned integer, producing a value of roughly 2^64 that is then passed as a memory allocation size. The allocator aborts the entire node with a message such as "Cannot allocate 2305843009213693951 bytes of memory (of type binary)". The crash is a VM-level abort, not an Erlang-level exception. It cannot be intercepted by supervision trees, by try/catch, or by passing the [safe] option to binary_to_term/2 (which only restricts atom creation and does not perform structural validation of binary encodings). Any application that decodes ETF from untrusted sources via binary_to_term/1,2 or enif_binary_to_term() is exposed. The Erlang distribution protocol also decodes incoming terms through the same code path, but distribution is expected to run on trusted networks per the OTP Secure Coding Guidelines (DSG-011). This issue affects OTP from OTP 27.0 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to erts from 15.0 before 17.0.4, 16.4.0.4 and 15.2.7.11.	N/A	More Details
CVE-2026-12495	Denial-of-service (DoS) vulnerability due to a stack buffer overflow in the http_gdpr_decrypt function of the Mercusys MB115-4G device's web interface. An unauthenticated attacker could exploit this vulnerability by sending a specially crafted request to the /cgi/login endpoint, causing memory corruption and the httpd process to crash, resulting in a denial of service for the web administration service.	N/A	More Details
CVE-2026-14856	A stored Cross-Site Scripting (XSS) vulnerability in the file upload functionality of the Media Manager in TastyIgniter v4.3.0, caused by insufficient validation and sanitization of SVG files. An authenticated user with low privileges can upload a malicious SVG file containing JavaScript code. When an administrator views that file, the code executes in the context of their browser. By chaining this vulnerability with a Cross-Site Request Forgery (CSRF) attack, an attacker can extract the administrator’s CSRF token and perform unauthorized actions—such as modifying credentials—thereby gaining full control of the administrative account.	N/A	More Details
CVE-2026-57916	proCertum SmartSign opens Certificate Practice Statement (CPS) URI without schema validation. An attacker can prepare arbitrary certificate with CPS URI pointing to a local executable file or any URL, sign a document with it, and send it to the victim. When the victim opens the document in the application, the specified file will be executed (or webpage will be opened). This issue was fixed in version 9.4.3.90.	N/A	More Details
CVE-2026-47078	Relative Path Traversal vulnerability in Erlang OTP (stdlib zip module) allows writing files outside the intended extraction directory via a crafted zip archive. zip:unzip/1,2 and zip:extract/1,2 validate entry paths using zip:check_dir_level/2, which tracks directory depth as a running integer counter: .. decrements it, normal path components increment it. The caller rejects only paths where the final counter value is less than zero. A path such as ../x/y causes the counter to go negative mid-traversal then recover to zero, passing validation while resolving to a location outside the extraction directory when joined with the current working directory via add_cwd. This vulnerability is associated with program file lib/stdlib/src/zip.erl. This issue affects OTP from OTP 27.1 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to stdlib from 6.1 before 8.0.3, 7.3.0.1 and 6.2.2.4.	N/A	More Details
	Improper Handling of Exceptional Conditions vulnerability in Erlang OTP erts (epmd) allows an unauthenticated remote attacker to permanently terminate the Erlang Port Mapper Daemon (epmd) via connection slot exhaustion.		

CVE-2026-42792	The do_accept function in erts/epmd/src/epmd_srv.c calls epmd_cleanup_exit() when accept(2) returns EMFILE (per-process file descriptor limit reached) or ENFILE (system-wide file descriptor limit reached), rather than treating these as recoverable conditions. An attacker can exhaust epmd's file descriptor slots by holding many TCP connections open while periodically sending a single byte to reset the idle timeout, then causing accept(2) to return EMFILE, which kills the daemon. epmd has no per-source-IP connection cap, making the attack feasible from a single source. On Debian/Ubuntu default packaging the impact is amplified: the systemd unit inherits a low file descriptor soft limit, and repeated daemon deaths trigger systemd's start-rate-limit, permanently failing both epmd.service and epmd.socket and requiring manual operator intervention to recover. This issue affects OTP from OTP 17.0 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15.	N/A	More Details
CVE-2026-17574	HDF5 contains a NULL pointer dereference vulnerability. Processing a crafted HDF5 file containing an attribute with an invalid variable-length datatype type field may cause the application to crash when the attribute is read.	N/A	More Details
CVE-2026-17573	A double free vulnerability was discovered in the HDF5 library. Processing a crafted HDF5 file containing an oversized chunk size field via h5repack may cause the application to abort due to a double free.	N/A	More Details
CVE-2026-17572	Heap-based buffer overflow in the SOHM list-index deserialization code in HDF5 through 2.1.1 on all platforms allows attackers to cause a denial of service (crash) via a crafted HDF5 file whose shared-message list index declares a num_messages count exceeding list_max, triggering out-of-bounds heap reads and writes in H5SM__cache_list_deserialize and H5SM__cache_list_verify_chksum.	N/A	More Details
CVE-2026-57917	proCertum SmartSign parses external XML entities from arbitrary crafted signature files, enabling SSRF and potentially allowing the reading of local files, depending on the parser's configuration. The XML External Entity (XXE) vulnerability is triggered simply by previewing a file in the file selection window, before the victim clicks “Open”. This issue was fixed in version 9.4.3.90.	N/A	More Details
CVE-2026-12989	A lack of authentication in the mobile app (APK v5.5.0) for Ghost Robotics' Vision 60 robot allows an unauthenticated attacker connected to the device's internal Wi-Fi network to gain unrestricted access to the web administration interface and the HTTP API. Due to the lack of authorization mechanisms, the attacker can view real-time camera feeds, control the robot’s movements, manage sensors (GPS, RTK, SAM, LIDAR), and execute critical operational commands (Play, Pause, Stop, E-Stop). Successful exploitation completely compromises the confidentiality, integrity, and physical security of the system.	N/A	More Details
CVE-2026-12990	An access control vulnerability in the mobile app (APK v5.5.0) for Ghost Robotics' Vision 60 robot allows multiple simultaneous sessions to run without proper client validation or session integrity checks. An attacker with a modified version of the app can connect to the robot during an active, legitimate session. This allows the attacker to bypass control restrictions, intercept sensitive information (such as real-time video), and partially interact with the system unnoticed and without disconnecting the legitimate user, compromising confidentiality and operational security.	N/A	More Details
CVE-2026-12991	The lack of cryptographic mechanisms to ensure the integrity and authenticity of communications in Ghost Robotics' Vision 60 robot (APK v5.5.0) exposes the system to man-in-the-middle attacks. An attacker located on the local network can use ARP spoofing and selective traffic blocking techniques to intercept and manipulate packets between the legitimate operator and the robot. This allows the attacker to disconnect the original controller, establish unauthorized communications, and prevent the operator from regaining control of the device, seriously compromising the confidentiality, integrity, and availability (CIA) of operations.	N/A	More Details
CVE-2026-65766	Joomla Extension - joomshaper.com - Unauthenticated SQL injection in SP Page Builder < 6.7.1 - Improper validation of order parameters in the Dynamic Content endpoint leads to an SQL injection vector.	N/A	More Details
CVE-2026-65876	Joomla Extension - joomshaper.com - Unauthenticated SQL injection in SP Page Builder < 6.7.1 - Improper validation of catid parameters in the loadMoreArticles endpoint leads to an SQL injection vector.	N/A	More Details
CVE-2026-65877	Joomla Extension - joomshaper.com - Authenticated SQL injection in SP Page Builder < 6.7.1 - Improper validation of various parameters in the media manager search and date filters lead to an SQL injection vector.	N/A	More Details
CVE-2026-65878	Joomla Extension - joomshaper.com - Authenticated arbitrary file delete in SP Page Builder < 6.7.1- Improper path validation and ACL checks lead to a file deletion vector in the media manager.	N/A	More Details
CVE-2025-59172	Ericsson Packet Core Controller (PCC) versions prior to 1.38 contain an Improper Neutralization of Special Elements vulnerability allowing an attacker to execute arbitrary code as root.	N/A	More Details
CVE-2025-59177	Ericsson Packet Core Controller (PCC) versions prior to 1.39 contain a vulnerability in Configuration Management, allowing an attacker to execute specifically crafted commands to reveal system secret through error messages.	N/A	More Details
CVE-2025-59178	Ericsson Packet Core Controller (PCC) versions prior to 1.39 contain an Exposure of Sensitive System Information vulnerability in Configuration Management allowing an attacker to enumerate other users on the system.	N/A	More Details
CVE-	Ericsson Packet Core Controller (PCC) versions prior to 1.38 contain a hardcoded credential vulnerability in the alarm		

2025-59180	system. An attacker with access to the cluster with knowledge of the hardcoded credential can read alarm and alert information.	N/A	More Details
CVE-2026-59251	Allocation of resources without limits in Erlang/OTP public_key certificate path validation allows a remote unauthenticated attacker to cause denial of service by sending a crafted X.509 certificate chain during the TLS handshake. During RFC 5280 policy processing in public_key:pkix_path_validation/3, the certificate policy tree maintained by pubkey_policy_tree grows without an upper bound. When a certificate chain contains M policies per certificate and K certificates, the tree grows on the order of M^K nodes because pubkey_policy_tree:add_leaves/2 and pubkey_policy_tree:add_leaf_siblings/2 extend the tree per policy per certificate. A modest chain with many policies per certificate is enough to pin BEAM schedulers and exhaust the node's memory, taking down the entire VM. The attacker only needs to be able to present a certificate chain to the victim, which is the normal precondition for a TLS handshake, so exploitation succeeds against any incoming or outgoing TLS connection that validates the peer's chain (the default for SSL/TLS clients and mutual-TLS servers). This is the same vulnerability class as OpenSSL's X509_verify_cert policy tree DoS. This vulnerability is associated with program files lib/public_key/src/pubkey_policy_tree.erl and program routines pubkey_policy_tree:add_leaves/2 and pubkey_policy_tree:add_leaf_siblings/2. This issue affects OTP from OTP 26.2 before OTP 29.0.4, OTP 28.5.0.4 and OTP 27.3.4.15, corresponding to public_key from 1.15 before 1.21.4, 1.20.3.4 and 1.17.1.5.	N/A	More Details
CVE-2026-65764	Joomla Extension - phoca.cz - Reflected XSS vulnerability in Phoca Commander 5.0.0-6.1.1 - Improper validation of user inputs lead to a reflective XSS vulnerability.	N/A	More Details
CVE-2026-64647	Next.js is a React framework for building full-stack web applications. In versions 12.0.0 through 15.5.20 and 16.0.0 through 16.2.10, a server-side fetch with a request body may return a cached response body from a different request to the same URL but different body. Confidential data in the POST's response body would then leak to unauthorized requests. Though the request itself will not be deduped. This is only an issue when receiving request bodies with a content type charset other than UTF-8. For example, the UTF-16 byte sequences for and in the request body would share the same cache. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-66397	phpMyFAQ before 4.1.6 fails to validate path traversal sequences in the existing_image field during category updates, allowing authenticated attackers to delete arbitrary files by exploiting insufficient sanitization in Image::delete(). Attackers can delete the database.php configuration file to disable the installation gate and access the public setup wizard to create new superadmin accounts.	N/A	More Details
CVE-2026-64646	Next.js is a React framework for building full-stack web applications. In versions 13.0.0 through 15.5.20 and 16.0.0 through 16.2.10, requests targeting Next.js applications using App Router with at least one Server Action can lead to excessive memory consumption if that Server Actions uses the Edge runtime. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-64527	In the Linux kernel, the following vulnerability has been resolved: drm/hyperv: validate VMBus packet size in receive callback hyperv_receive_sub() reads msg->vid_hdr.type and dispatches into one of four message-type branches without knowing how many bytes the host wrote into hv->recv_buf. The completion path then runs memcpy(hv->init_buf, msg, VMBUS_MAX_PACKET_SIZE), so the consumer that wakes on wait_for_completion_timeout() can read up to 16 KiB of residue from a prior message as if it were the response payload. Pass bytes_recvd into hyperv_receive_sub() and reject any packet that does not cover the pipe + synthvid header. A single switch on msg->vid_hdr.type then computes the type-specific payload size: the three completion-driving types (SYNTHVID_VERSION_RESPONSE, SYNTHVID_RESOLUTION_RESPONSE, SYNTHVID_VRAM_LOCATION_ACK) fall through to a shared exit that requires that size before memcpy/complete, while SYNTHVID_FEATURE_CHANGE validates its own payload and returns before reading is_dirt_needed. Unknown types are dropped. SYNTHVID_RESOLUTION_RESPONSE is variable length: the host fills resolution_count entries, not the full SYNTHVID_MAX_RESOLUTION_COUNT array. Validate the fixed prefix first so resolution_count can be read, bound it against the array, then require only the count-sized array, so the shorter responses the host actually sends are accepted. Only run the sub-handler when vmbus_recvpacket() returned success. The memcpy length is bytes_recvd, which is bounded by VMBUS_MAX_PACKET_SIZE only on a successful receive; on -ENOBUFS vmbus_recvpacket() instead reports the required length, which can exceed hv->recv_buf, so copying bytes_recvd would read and write past the 16 KiB buffers. Gating on the success return keeps the copy bounded. The nonzero-return path is itself a malformed-message case and is now logged rather than silently skipped; channel recovery is not attempted. Rejected packets are reported via drm_err_ratelimited() rather than silently dropped, matching the CoCo-hardened pattern in hv_kvp_onchannelcallback().	N/A	More Details
CVE-2026-17612	Honeywell S35 Series 3M/5M/8M/PinHole Cameras, all versions prior to and including version HC5.26.1.14.20260207 contains an audit log disclosure Vulnerability that could allow an attacker to access audit logs without authentication, potentially resulting in the disclosure of sensitive information. Honeywell recommends updating to the latest available version (HC5.26.1.16.20260207) once available.	N/A	More Details
CVE-2026-16481	A Server-Side Request Forgery (SSRF) and credential exfiltration vulnerability exists in the cloud-healthcare-fhir-fetch-page tool of googleapis/mcp-toolbox. The tool takes an unvalidated pageURL parameter from the client and issues an HTTP GET request to it using an authenticated client. The underlying transport automatically attaches an Authorization: Bearer header to every outbound request regardless of the destination host. An attacker can supply an arbitrary external URL to the pageURL parameter (either directly via the tool execution payload or implicitly via data-driven pagination tracking loops), leading Toolbox into sending its OAuth/service-account access token to an attacker-controlled listener. Depending on the configuration, this leaks either the end-user's token or the broader service-account access token (ADC), potentially exposing Protected Health Information (PHI) and secondary Google Cloud Platform services.	N/A	More Details

CVE-2026-64528	In the Linux kernel, the following vulnerability has been resolved: tty: serial: samsung: Remove redundant port lock acquisition in rx helpers Sashiko identified a deadlock when the console flow is engaged [1]. When console flow control is enabled (UPF_CONS_FLOW), s3c24xx_serial_stop_tx() calls s3c24xx_serial_rx_enable() and s3c24xx_serial_start_tx() calls s3c24xx_serial_rx_disable(). The serial core framework invokes the .stop_tx() and .start_tx() callbacks with the port->lock spinlock already held. Furthermore, all internal driver paths that invoke stop_tx (such as the DMA TX completion handler s3c24xx_serial_tx_dma_complete() or the PIO TX IRQ handler s3c24xx_serial_tx_irq()) also acquire port->lock prior to calling it. (Note that s3c24xx_serial_start_tx() is only invoked by the serial core). However, s3c24xx_serial_rx_enable() and s3c24xx_serial_rx_disable() unconditionally attempt to acquire port->lock again using uart_port_lock_irqsave(). Since spinlocks are not recursive, this causes a deadlock on the same CPU when console flow control is engaged. Remove the redundant lock acquisition from both rx helper functions.	N/A	More Details
CVE-2026-66013	OpenRemote before 1.26.2 contains an authentication bypass vulnerability in the console registration API that allows unauthenticated attackers to update existing console assets by supplying a known asset identifier. Attackers can overwrite push notification tokens and console metadata without authentication or ownership validation, redirecting notifications or denying delivery to legitimate consoles.	N/A	More Details
CVE-2026-64645	Next.js is a React framework for building full-stack web applications. In versions 12.0.0 through 15.5.20 and 16.0.0 through 16.2.10, a rewrites() or redirects() rule that builds its external destination hostname from request-controlled input can be pointed at an arbitrary hostname, regardless of the rule's hostname suffix. For a rewrite, Next.js proxies the request to that arbitrary host and serves the response from the application's origin, leading to Server-Side Request forgery. A redirects() rule configured this way is vulnerable to an Open Redirect. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-64644	Next.js is a React framework for building full-stack web applications. In versions 15.5.0 through 15.5.20 and 16.0.0 through 16.2.10, when self-hosting Next.js with the default image loader, the Image Optimization API can optimize remotely hosted images if configured (not enabled by default). If those images contain malicious content, they can cause CPU exhaustion in /_next/image endpoints.Only config.images.remotePatterns is affected, and just the patterns in that array, whereas config.images.unoptimized: true, config.images.loader: 'custom', and Vercel are not impacted. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-64643	Next.js is a React framework for building full-stack web applications. In versions 12.0.0 through 15.5.20 and 16.0.0 through 16.2.10, Next.js applications using App Router, Server Actions (use server) or use cache endpoints can be disclosed bypassing any authentication on the pages where these endpoints are usually used. Server Action IDs can be disclosed to unauthenticated users via publicly served client artifacts (for example, static chunks containing action references). Affected users are applications using App Router and Server Actions. By itself, this disclosure is typically a recon/enumeration primitive; however, it can increase risk when combined with other weaknesses. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-64642	Next.js is a React framework for building full-stack web applications. In versions 16.0.0 through 16.2.10, crafted requests targeting Next.js applications using App Router built with Turbopack and a single entry in config.i18n.locales can bypass middleware/proxy based authentication. This issue has been fixed in version 16.2.11.	N/A	More Details
CVE-2026-64641	Next.js is a React framework for building full-stack web applications. In versions 13.0.0 through 15.5.20 and 16.0.0 through 16.2.10, crafted requests targeting Next.js applications using App Router with at least one Server Action can lead to excessive CPU usage blocking processing of further requests in the same process. This issue has been fixed in versions 15.5.21 and 16.2.11.	N/A	More Details
CVE-2026-59239	Stored Cross-site Scripting (CWE-79) in the email module in Roskus Prospero Flow CRM before 5.4.4 allows a remote, authenticated low-privileged user to execute arbitrary JavaScript in another user's browser, including administrators, leading to session compromise and account takeover, via a payload stored in an email body that is persisted without sanitization and rendered unescaped with {!! $email->body !!} when the recipient opens the message.	N/A	More Details
CVE-2026-15928	XMLRPC-C Library versions 1.07 through 1.67.01 are vulnerable to a reflected cross-site scripting (XSS) vulnerability in the error page component.	N/A	More Details
CVE-2026-54272	ip-address is a library for parsing and manipulating IPv4 and IPv6 addresses in JavaScript. Versions 10.1.1 through 10.2.0 are vulnerable to SSRF through misclassification of IPv4-mapped/NAT64 IPv6 addresses. Address6.getType() classifies an address by matching it against a table of known IPv6 special-use prefixes, returning Global unicast when nothing matches. That table had no entry for the IPv4-mapped range (::ffff:0:0/96), so every mapped address fell through to Global unicast; NAT64 addresses matched their own NAT64 … labels. The boolean checks isLoopback, isUnspecified, and isMulticast compared getType() against a fixed label and so returned false, while isLinkLocal and isULA checked only the native IPv6 ranges. The library already exposed isMapped4() and to4(), but did not apply them inside these checks, so a mapped or NAT64 address was never normalized to its embedded IPv4 address before classification. For IPv4-mapped addresses the host OS routes to the IPv4 stack, so the misclassification is reachable on any dual-stack host. For NAT64, the classification bypass is unconditional but end-to-end reachability additionally requires a NAT64/DNS64 gateway in the deployment network.This issue has been fixed in version 10.2.1.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: net: openvswitch: reject oversized nested action attrs Open vSwitch stores generated flow actions as nlattrs, whose nla_len field is u16. Commit a1e64addf3ff ("net: openvswitch: remove misbehaving actions length check") allowed the total sw_flow_actions stream to grow beyond 64 KiB, which is valid, but also removed the last guard preventing a generated nested action attribute from exceeding U16_MAX. An oversized generated container can thus be closed with a truncated nla_len. A later dump or teardown then walks a structurally different stream than the one that was validated. In particular, an oversized nested CLONE/CT action may cause subsequent bytes in the generated stream to be interpreted as independent actions.		More

2026-64531	Keep the larger total-action-stream behavior, but make nested action close reject generated containers that do not fit in nla_len, and return the error through all callers. For recursive SAMPLE, CLONE, DEC_TTL, and CHECK_PKT_LEN builders, trim resource-owning action-list tails in reverse construction order before discarding failed wrappers, so resources copied into the rejected tails are released before the wrappers are removed. Most failed outer wrappers are discarded by truncating actions_len after child resources have been released. CHECK_PKT_LEN also trims its parent after branch resources are gone. SET/TUNNEL close failures unwind their known tun_dst ownership directly, and SET_TO_MASKED has no external ownership and truncates on close failure.	N/A	Details
CVE-2026-64532	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: bound NTFS_DE view.data_off in UpdateRecordData{Root,Allocation} In do_action()'s UpdateRecordDataRoot (fslog.c:3489) and UpdateRecordDataAllocation (fslog.c:3697) cases, the memmove destination is `Add2Ptr(e, le16_to_cpu(e->view.data_off))`, where e->view.data_off comes from an on-disk NTFS_DE inside an INDEX_ROOT or INDEX_BUFFER. Neither case validates view.data_off + dlen against e->size; the existing check_if_index_root / check_if_alloc_index helpers walk the entry chain and validate the entry's offset, but not its internal view fields. The neighbouring read sites (e.g., fs/ntfs3/index.c when iterating view entries) check view.data_off + view.data_size <= e->size. Apply the same bound at the two memmove sites. Reproduced under UML+KASAN on mainline 8d90b09e6741 via pr_warn-only probe instrumentation: with view.data_off forced to 0xFFFC, the memmove writes 32 bytes past the end of the NTFS_DE. This is similar in shape to Pavitra Jha's 2026-05-02 patch "fs/ntfs3: prevent oob in case UpdateRecordDataRoot" (<20260502105008.21827-1-jhapavitra98@gmail.com>) which proposes calling ntfs3_bad_de_range(); that helper does not exist in mainline. This patch uses inline checks.	N/A	More Details
CVE-2026-64533	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: validate lcns_follow in log_replay conversion log_replay() converts DIR_PAGE_ENTRY_32 records into DIR_PAGE_ENTRY records when replaying version 0 restart tables. During this conversion, the memmove() length is derived directly from the on-disk lcns_follow field: memmove(&dp->vcn, &dp0->vcn_low, 2 * sizeof(u64) + le32_to_cpu(dp->lcns_follow) * sizeof(u64)); check_rstbl() validates restart table structure, but does not constrain per-entry lcns_follow values relative to the entry size. A malformed filesystem image can provide an oversized lcns_follow value, causing the conversion memmove() to access memory beyond the bounds of the allocated restart table buffer. The same field is later used to bound iteration over page_lcns[], so validating lcns_follow during conversion also prevents downstream out-of-bounds access from the same malformed metadata. Compute the maximum valid lcns_follow from the already-validated restart table entry size and reject entries that exceed this bound. Reuse the existing t16/t32 scratch variables already declared in log_replay() to avoid introducing new declarations. [almaz.alexandrovich@paragon-software.com: fixed the conflicts]	N/A	More Details
CVE-2026-64534	In the Linux kernel, the following vulnerability has been resolved: nvmet-tcp: check INIT_FAILED before nvmet_req_uninit in digest error path In nvmet_tcp_try_recv_ddgst(), when a data digest mismatch is detected, nvmet_req_uninit() is called unconditionally. However, if the command arrived via the nvmet_tcp_handle_req_failure() path, nvmet_req_init() had returned false and percpu_ref_tryget_live() was never executed. The unconditional percpu_ref_put() inside nvmet_req_uninit() then causes a refcount underflow, leading to a WARNING in percpu_ref_switch_to_atomic_rcu, a use-after-free diagnostic, and eventually a permanent workqueue deadlock. Check cmd->flags & NVMET_TCP_F_INIT_FAILED before calling nvmet_req_uninit(), matching the existing pattern in nvmet_tcp_execute_request().	N/A	More Details
CVE-2026-64535	In the Linux kernel, the following vulnerability has been resolved: nvmet-tcp: Fix potential UAF when ddgst mismatch Shivam Kumar found via vulnerability testing: When data digest is enabled on an NVMe/TCP connection and a digest mismatch occurs on a non-final H2C_DATA PDU during an R2T-based data transfer, the digest error handler in nvmet_tcp_try_recv_ddgst() calls nvmet_req_uninit() — which performs percpu_ref_put() on the submission queue — but does NOT mark the command as completed. It does not set cqe->status, does not modify rbytes_done, and does not clear any flag. When the subsequent fatal error triggers queue teardown, nvmet_tcp_uninit_data_in_cmds() iterates all commands, checks nvmet_tcp_need_data_in() for each one, and finds that the already-uninited command still appears to need data (because rbytes_done < transfer_len and cqe->status == 0). It therefore calls nvmet_req_uninit() a second time on the same command — a double percpu_ref_put against a single percpu_ref_get.	N/A	More Details
CVE-2026-64536	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8723bs: fix OOB reads in is_ap_in_tkip() IE loop The loop in is_ap_in_tkip() iterates over IEs without verifying that enough bytes remain before dereferencing the IE header or its payload: - pIE->element_id and pIE->length are read without checking that i + sizeof(*pIE) <= ie_length, so a truncated IE at the end of the buffer causes an OOB read. - For WLAN_EID_VENDOR_SPECIFIC the code compares pIE->data + 12, which requires pIE->length >= 16. For WLAN_EID_RSN it compares pIE->data + 8, requiring pIE->length >= 12. Neither requirement is checked. Add the missing IE header and payload bounds checks and guard each data access with an explicit pIE->length minimum, matching the pattern established in update_beacon_info().	N/A	More Details
CVE-2026-65893	This vulnerability exists in CP PLUS EZ-P21 IP Camera due to an insecure debug feature enabled in the firmware. An attacker with physical access could exploit this vulnerability by placing arbitrary code on removable media and triggering their execution through the debug mechanism. Successful exploitation of this vulnerability could allow an attacker to execute arbitrary code with elevated privileges on the targeted device.	N/A	More Details
CVE-2026-65894	This vulnerability exists in CP PLUS EZ-P21 IP Camera due to improper authentication of HTTP endpoints. A remote attacker could exploit this vulnerability by conducting brute-force attacks against HTTP endpoint on the targeted device. Successful exploitation of this vulnerability could allow an attacker to gain unauthorized access to live video snapshots from the targeted device.	N/A	More Details
CVE-2026-	Rejected reason: This is a duplicate.	N/A	More Details

15799			
CVE-2026-16554	cJSON library is vulnerable to an integer overflow in the print_string_ptr() function in cJSON.c on 32-bit platforms. The escape_characters counter, a 32-bit size_t, can wrap around when processing strings containing approximately 858,993,460 or more control characters, causing the output buffer to be allocated based on an underestimated length. When cJSON_PrintBuffered() is used with a pre-allocated buffer, the subsequent write loop overflows the heap allocation. An attacker supplying a crafted JSON string to an application using cJSON on a 32-bit platform can cause a heap buffer overflow, potentially leading to remote code execution, information disclosure, or denial of service. Because project creator contact attempts were unsuccessful, the vulnerability has only been confirmed in version 1.7.19 but may also affect other versions.	N/A	More Details
CVE-2026-66398	phpMyFAQ before v4.1.6 contains a remote code execution vulnerability in the configuration API that allows authenticated administrators with CONFIGURATION_EDIT and ATTACHMENT_ADD privileges to write arbitrary PHP files by manipulating the upgrade.lastDownloadedPackage setting. Attackers can upload a malicious ZIP file as an attachment, point the updater configuration to its stored path, and extract it into the application root to achieve code execution as the web server user.	N/A	More Details
CVE-2026-64451	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix NULL pointer dereference in func_set_flag() func_set_flag() dereferences tr->current_trace_flags before verifying that the current tracer is actually the function tracer. When the active tracer has been switched away from "function" (e.g., to "wakeup_rt"), tr->current_trace_flags can be NULL, leading to a NULL pointer dereference and kernel crash. The call chain that triggers this is: trace_options_write() -> __set_tracer_option() -> trace->set_flag() /* func_set_flag */ In func_set_flag(), the first operation is: if (!!set == !!(tr->current_trace_flags->val & bit)) This dereferences tr->current_trace_flags unconditionally. The safety check that guards against a non-function tracer: if (tr->current_trace != &function_trace) return 0; is placed *after* the dereference, which is too late. This was observed with the following crash dump: BUG: unable to handle page fault at 0000000000000000 RIP: func_set_flag+0xd Call Trace: __set_tracer_option+0x27 trace_options_write+0x75 vfs_write+0x12a ksys_write+0x66 do_syscall_64+0x5b RIP: ffffffff914c973d RSP: ff67ec88b01dfdf0 RFLAGS: 00010202 RAX: 0000000000000000 RBX: ff3a826e80354580 RCX: 0000000000000001 RDX: 0000000000000001 RSI: 0000000000000000 RDI: ffffffff93918080 The disassembly confirms the fault: func_set_flag+0: mov 0x1f08(%rdi), %rax ; RAX = tr->current_trace_flags = NULL func_set_flag+13: mov (%rax), %eax ; page fault: dereference NULL At the time of the crash: tr->current_trace_flags = 0x0 (NULL) tr->current_trace = wakeup_rt_tracer (not function_trace) The scenario is that a process opens a function tracer option file (such as "func_stack_trace"), then the current tracer is switched to another tracer (e.g., "wakeup_rt"), which sets current_trace_flags to NULL. When the process subsequently writes to the option file, func_set_flag() is invoked and crashes on the NULL dereference. Fix this by moving the current_trace check before the current_trace_flags dereference, so that func_set_flag() returns early when the function tracer is not active.	N/A	More Details
CVE-2026-64433	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: MGMT: Fix UAF of hci_conn_params in add_device_complete add_device_complete() runs from the hci_cmd_sync_work kworker, which holds only hci_req_sync_lock and *not* hci_dev_lock. It calls hci_conn_params_lookup() and then dereferences the returned object (params->flags) without taking hci_dev_lock: params = hci_conn_params_lookup(hdev, &cp->addr.bdaddr, le_addr_type(cp->addr.type)); ... device_flags_changed(NULL, hdev, &cp->addr.bdaddr, cp->addr.type, hdev->conn_flags, params ? params->flags : 0); hci_conn_params_lookup() walks hdev->le_conn_params and is documented to require hdev->lock. A concurrent MGMT_OP_REMOVE_DEVICE (remove_device()), which does run under hci_dev_lock, can call hci_conn_params_free() to list_del() and kfree() the very object the lookup returned, so the subsequent params->flags read touches freed memory [0]. Hold hci_dev_lock() across the hci_conn_params_lookup() and the read of params->flags (and the matching event emission) so the lookup result cannot be freed by a concurrent remove_device() before it is used, honouring the locking contract of hci_conn_params_lookup(). [0]: (trailing page/memory-state dump trimmed) BUG: KASAN: slab-use-after-free in add_device_complete+0x358/0x3d8 net/bluetooth/mgmt.c:7671 Read of size 1 at addr ffff000017ab26c1 by task kworker/u9:8/388 CPU: 1 UID: 0 PID: 388 Comm: kworker/u9:8 Not tainted 7.0.11 #20 PREEMPT Hardware name: linux,dummy-virt (DT) Workqueue: hci0 hci_cmd_sync_work Call trace: show_stack+0x2c/0x3c arch/arm64/kernel/stacktrace.c:499 (C) __dump_stack lib/dump_stack.c:94 [inline] dump_stack_lvl+0xb4/0xd4 lib/dump_stack.c:120 print_address_description mm/kasan/report.c:378 [inline] print_report+0x118/0x5d8 mm/kasan/report.c:482 kasan_report+0xb0/0xf4 mm/kasan/report.c:595 __asan_report_load1_noabort+0x20/0x2c mm/kasan/report_generic.c:378 add_device_complete+0x358/0x3d8 net/bluetooth/mgmt.c:7671 hci_cmd_sync_work+0x14c/0x240 net/bluetooth/hci_sync.c:334 process_one_work+0x628/0xd38 kernel/workqueue.c:3289 process_scheduled_works kernel/workqueue.c:3372 [inline] worker_thread+0x7a8/0xac0 kernel/workqueue.c:3453 kthread+0x39c/0x444 kernel/kthread.c:436 ret_from_fork+0x10/0x20 arch/arm64/kernel/entry.S:860 Allocated by task 3401: kasan_save_stack+0x3c/0x64 mm/kasan/common.c:57 kasan_save_track+0x20/0x3c mm/kasan/common.c:78 kasan_save_alloc_info+0x40/0x54 mm/kasan/generic.c:570 poison_kmalloc_redzone mm/kasan/common.c:398 [inline] __kasan_kmalloc+0xd4/0xd8 mm/kasan/common.c:415 kasan_kmalloc include/linux/kasan.h:263 [inline] __kmalloc_cache_noprof+0x1b0/0x458 mm/slub.c:5385 kmalloc_noprof include/linux/slab.h:950 [inline] kzalloc_noprof include/linux/slab.h:1188 [inline] hci_conn_params_add+0x10c/0x4b0 net/bluetooth/hci_core.c:2279 hci_conn_params_set net/bluetooth/mgmt.c:5162 [inline] add_device+0x5b4/0xa54 net/bluetooth/mgmt.c:7755 hci_mgmt_cmd net/bluetooth/hci_sock.c:1721 [inline] hci_sock_sendmsg+0x10b4/0x1dd0 net/bluetooth/hci_sock.c:1841 sock_sendmsg_nosec net/socket.c:727 [inline] __sock_sendmsg+0xe0/0x128 net/socket.c:742 sock_write_iter+0x250/0x390 net/socket.c:1195 new_sync_write fs/read_write.c:595 [inline] vfs_write+0x66c/0xab0 fs/read_write.c:688 ksys_write+0x1fc/0x24c fs/read_write.c:740 __do_sys_write fs/read_write.c:751 [inline] __se_sys_write fs/read_write.c:748 [inline] __arm64_sys_write+0x70/0xa4 fs/read_write.c:748 __invoke_syscall arch/arm64/kernel/syscall.c:35 [inline] invoke_syscall+0x84/0x2a8 arch/arm64/kernel/syscall.c:49 el0_svc_common.constprop.0+0xe4/0x294 arch/arm64/kernel/syscall.c:132 do_el0_svc+0x44/0x5c arch/arm64/kernel/syscall.c:151 el0_svc+0x38/0xac arch/arm64/kernel/entry-common.c:724 el0t_64_sync_handler+0xa0/0xe4 arch/arm64/kernel/entry-common.c:743 el0t_64_sync+0x198/0x19c	N/A	More Details

	arch/arm64/kernel/entry.S:596 Freed by task 3740: kasan_save_stack+0x3c/0x64 ---truncated---		
CVE-2026-64262	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: end fuse_req on io-uring cancel task work When io_uring delivers task work with tw.cancel set (PF_EXITING, PF_KTHREAD fallback, or percpu_ref_is_dying on the ring context), fuse_uring_send_in_task() takes the cancel branch, assigns -ECANCELED, and falls through to fuse_uring_send(). That path only flips the entry to FRRS_USERSPACE and completes the io_uring cmd; it never discharges the ring entry's owning reference to the fuse_req that fuse_uring_add_req_to_ring_ent() handed it at dispatch time. fuse_uring_send_in_task() tw.cancel == true err = -ECANCELED fuse_uring_send(ent, cmd, err, issue_flags) ent->state = FRRS_USERSPACE list_move(&ent->list, &queue->ent_in_userspace) ent->cmd = NULL io_uring_cmd_done(-ECANCELED) /* ent->fuse_req still set, req still hashed */ The fuse_req stays linked on fpq->processing[hash] and fuse_request_end() is never invoked. The originating syscall thread blocks in D-state in request_wait_answer() until fuse_abort_conn() runs, which can be the entire connection lifetime. For FR_BACKGROUND requests fc->num_background is never decremented either, so repeated cancels inflate the counter until max_background is hit and all later background ops stall. tw.cancel does not imply a connection abort (e.g. a single io_uring worker thread exits while the fuse connection stays up), so this cannot be left for fuse_abort_conn() to clean up. Ending the req but still routing the entry through fuse_uring_send() is not enough: that leaves a req-less entry on ent_in_userspace, and ent_list_request_expired() dereferences ent->fuse_req unconditionally on the head of that list, which would then NULL-deref. Fix the cancel branch to release the entry directly. Remove it from the queue, complete the io_uring cmd, end the fuse_req, free the entry, and drop its queue_refs (waking the teardown waiter if it was the last).	N/A	More Details
CVE-2026-64329	In the Linux kernel, the following vulnerability has been resolved: usb: typec: ucsi: ccg: Fix use-after-free of ucsi on remove The threaded IRQ handler ccg_irq_handler() calls ucsi_notify_common(), which on a connector-change event calls ucsi_connector_change() and schedules connector work. In ucsi_ccg_remove(), ucsi_destroy() frees uc->ucsi (kfree) before free_irq() is called, so a handler invocation already in flight may access the freed object after ucsi_destroy(). CPU 0 (remove) | CPU 1 (threaded IRQ) ucsi_destroy(uc->ucsi) | ccg_irq_handler() kfree(ucsi) // FREE | ucsi_notify_common(uc->ucsi) // USE Move free_irq() before ucsi_destroy() in the remove path. It is kept after ucsi_unregister(): ucsi_unregister() cancels connector work whose handler issues GET_CONNECTOR_STATUS through ucsi_send_command_common(), which waits for a completion that is signalled from the IRQ handler, so the IRQ must stay active until that work has been cancelled. The probe error path already orders free_irq() before ucsi_destroy(). This bug was found by static analysis.	N/A	More Details
CVE-2026-64306	In the Linux kernel, the following vulnerability has been resolved: crypto: drbg - Fix returning success on failure in CTR_DRBG drbg_ctr_generate() sometimes returns success when it fails, leaving the output buffer uninitialized. Fix it.	N/A	More Details
CVE-2026-64307	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Do not initialize SNP for ioctl(SNP_CONFIG) Sashiko notes: > if SEV initialization fails and KVM is actively running normal VMs, could a > userspace process trigger this code path via /dev/sev ioctls (e.g., > SEV_PDH_GEN) and zero out MSR_VM_HSAVE_PA globally? Would the next VMRUN > execution for an active VM trigger a general protection fault and crash the > host? Refuse to re-try initialization if SNP is not already initialized for SNP_CONFIG. This is technically an ABI break: before if SNP initialization failed it could be transparently retriggered by this ioctl, and if no VMs were running, everything worked fine. Hopefully this is enough of a corner case that nobody will notice, but someone does, there are a few options: * do something like symbol_get() for kvm and refuse to initialize if KVM is loaded * check each cpu's HSAVE_PA for non-zero data before re-initializing * once initialization has failed, continue to refuse to initialize until the ccp module is unloaded	N/A	More Details
CVE-2026-64308	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Do not initialize SNP for ioctl(SNP_VLEK_LOAD) Sashiko notes: > if SEV initialization fails and KVM is actively running normal VMs, could a > userspace process trigger this code path via /dev/sev ioctls (e.g., > SEV_PDH_GEN) and zero out MSR_VM_HSAVE_PA globally? Would the next VMRUN > execution for an active VM trigger a general protection fault and crash the > host? The SEV firmware docs for SNP_VLEK_LOAD note: > On SNP_SHUTDOWN, the VLEK is deleted. That is, the initialization/shutdown wrapper here is pointless, because the firmware immediately throws away the key anyway. Instead, refuse to do anything if SNP has not been previously initialized. This is an ABI break: before, this was a no-op and almost certainly a mistake by userspace, and now it returns -ENODEV. ABI compatibility could be maintained here by simply returning 0 in the check instead.	N/A	More Details
CVE-2026-64309	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Do not initialize SNP for ioctl(SNP_COMMIT) Sashiko notes: > if SEV initialization fails and KVM is actively running normal VMs, could a > userspace process trigger this code path via /dev/sev ioctls (e.g., > SEV_PDH_GEN) and zero out MSR_VM_HSAVE_PA globally? Would the next VMRUN > execution for an active VM trigger a general protection fault and crash the > host? The SNP_COMMIT command does not require the firmware to be in any particular state. Skip initializing it if it was previously uninitialized. The SEV-SNP firmware specification doc 56860 does not mention SNP_COMMIT in Table 5 as a command that is allowed in the UNINIT state, but it is in fact allowed and a future documentation update will reflect that.	N/A	More Details
CVE-2026-64310	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Do not initialize SNP for SEV ioctls Sashiko notes: > if SEV initialization fails and KVM is actively running normal VMs, could a > userspace process trigger this code path via /dev/sev ioctls (e.g., > SEV_PDH_GEN) and zero out MSR_VM_HSAVE_PA globally? Would the next VMRUN > execution for an active VM trigger a general protection fault and crash the > host? sev_move_to_init_state() is called for ioctls requiring only SEV firmware: SEV_PEK_GEN, SEV_PDH_GEN, SEV_PEK_CSR, SEV_PEK_CERT_IMPORT, and SEV_PDH_CERT_EXPORT. After the firmware command, it does SEV_SHUTDOWN on the SEV firmware. Since these commands do not require SNP to be initialized, skip it by calling __sev_platform_init_locked() which only initializes the SEV firmware. This way SNP is not Initialized at all, and HSAVE_PA is not cleared. The previous code saved any SEV initialization firmware error to init_args.error and then	N/A	More Details

	threw it away and hardcoded the return value of INVALID_PLATFORM_STATE regardless of the real firmware error. This patch changes it to surface the underlying error, which is hopefully both more useful and doesn't cause any problems. Note that it is still safe to call __sev_firmware_shutdown() directly: it calls __sev_snp_shutdown_locked(), which skips SNP shutdown if SNP was not initialized.		
CVE-2026-64314	In the Linux kernel, the following vulnerability has been resolved: crypto: chacha20poly1305 - validate poly1305 template argument chachapoly_create() still accepts the compatibility poly1305 parameter in the template name, but it assumes the second template argument is always present and immediately passes it to strcmp(). When the argument is missing, crypto_attr_alg_name() returns an error pointer. Check for that before comparing the name so malformed template instantiations fail with an error instead of dereferencing the error pointer in strcmp(). This matches the surrounding Crypto API template pattern where crypto_attr_alg_name() results are validated before string-specific use.	N/A	More Details
CVE-2026-64316	In the Linux kernel, the following vulnerability has been resolved: crypto: caam - use print_hex_dump_devel to guard key hex dumps Use print_hex_dump_devel() for dumping sensitive key material in *_setkey() and gen_split_key() to avoid leaking secrets at runtime when CONFIG_DYNAMIC_DEBUG is enabled.	N/A	More Details
CVE-2026-64321	In the Linux kernel, the following vulnerability has been resolved: nvme: target: rdma: fix ndev refcount leak on queue connect nvmet_rdma_queue_connect() calls nvmet_rdma_find_get_device() which acquires a reference on the returned ndev via kref_get(). On the path where the host queue backlog is exceeded and the function returns NVME_SC_CONNECT_CTRL_BUSY, reference of ndev is not released, leaking the kref. Fix this by adding a goto to the existing put_device label before the early return.	N/A	More Details
CVE-2026-64325	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7921/mt7925: fix NULL dereference in CSA beacon This patch is based on a BUG as reported by Bongani Hlope at https://lore.kernel.org/all/20260502125824.425d7159@bongani-mini.home.org.za/ When a channel-switch announcement (CSA) beacon is received, cfg80211 queues a wiphy work item that eventually calls mt7921_channel_switch_rx_beacon(). If the station disconnects (or the channel context is otherwise torn down) between the time the work is queued and the time it runs, the driver's dev->new_ctx pointer can already have been cleared to NULL. mt7921_channel_switch_rx_beacon() then dereferences new_ctx unconditionally, triggering a NULL pointer dereference at address 0x0: BUG: kernel NULL pointer dereference, address: 0000000000000000 RIP: 0010:mt7921_channel_switch_rx_beacon+0x1f/0x100 [mt7921_common] The same missing guard exists in mt7925_channel_switch_rx_beacon(), which shares the same code pattern introduced by the same commit. Add an early-return NULL check for dev->new_ctx in both mt7921_channel_switch_rx_beacon() and mt7925_channel_switch_rx_beacon(). When new_ctx is NULL there is no pending channel switch to process, so returning immediately is the correct and safe action. Oops-Analysis: http://oops.fenrus.org/reports/lkml/20260502125824.425d7159@bongani-mini.home.org.za/report.html	N/A	More Details
CVE-2026-64326	In the Linux kernel, the following vulnerability has been resolved: block: skip sync_blockdev() on surprise removal in bdev_mark_dead() bdev_mark_dead()'s @surprise == true means the device is already gone. The filesystem callback fs_bdev_mark_dead() honours this and skips sync_filesystem(), but the bare block device path (no ->mark_dead op) lost its !surprise guard when the holder ->mark_dead callback was wired up (see Fixes), and now calls sync_blockdev() unconditionally, which can hang forever waiting on writeback that can no longer complete. syzkaller hit this via nvme_reset_work()'s "I/O queues lost" path: nvme_mark_namespaces_dead() -> blk_mark_disk_dead() -> bdev_mark_dead(bdev, true) -> sync_blockdev() blocks in folio_wait_writeback(), wedging the reset worker and every task waiting on it. Skip the sync on surprise removal, matching fs_bdev_mark_dead(); invalidate_bdev() still runs. Orderly removal (surprise == false) is unchanged. Found by FuzzNvme(Syzkaller with FEMU fuzzing framework).	N/A	More Details
CVE-2026-64327	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: Initialize epfile->in early to fix endpoint direction checks When parsing endpoint descriptors, ffs_data_got_descs() generates the eps_addrmap which contains the endpoint direction. However, epfile->in was previously only populated in ffs_func_eps_enable() which executes upon USB host connection. As a result, early userspace ioctls like FUNCTIONFS_DMABUF_ATTACH that run before the host connects would see epfile->in as 0, leading to incorrect DMA directions. By moving the initialization to ffs_epfiles_create(), epfile->in is accurate before userspace opens the endpoint files.	N/A	More Details
CVE-2026-64328	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: Fix DMA fence leak In ffs_dmabuf_transfer(), a ffs_dma_fence object is kmalloc'd, with the underlying dma_fence later initialized by dma_fence_init(), which sets its kref counter to 1. Then, dma_resv_add_fence() gets a second reference, and a pointer to the ffs_dma_fence is passed as the usb_request's "context" field. The dma-resv mechanism will manage the second reference, but the first reference is never properly released; the ffs_dmabuf_cleanup() function decreases the reference count, but only to balance with the reference grab in ffs_dmabuf_signal_done(). The code will then slowly leak memory as more ffs_dma_fence objects are created without being ever freed. Address this issue by transferring ownership of the fence to the DMA reservation object, by calling dma_fence_put() right after dma_resv_add_fence(). The ffs_dma_fence then gets properly discarded after being signalled.	N/A	More Details
CVE-2026-64330	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpm: Validate SVID index in svdm_consume_modes() In svdm_consume_modes(), the SVID value is read from pmdata->svids using pmdata->svid_index as an array index without bounds validation: paltmode->svid = pmdata->svids[pmdata->svid_index]; If pmdata->svid_index is driven beyond SVID_DISCOVERY_MAX (16), it results in an out-of-bounds read of the pmdata->svids array. Because pd_mode_data is embedded inside struct tcpm_port, indexing past svids reads into adjacent fields. In particular: - At index 16, it reads the altmodes count. - At index 18 and beyond, it reads into altmode_desc[], which contains partner-supplied SVDM Discovery Modes VDOs. By injecting a chosen SVID into altmode_desc[0].vdo and driving svid_index to 20, the partner can force paltmode->svid to be loaded with an arbitrary, partner- chosen SVID, which is then registered via typec_partner_register_altmode(). Fix this by validating that pmdata->svid_index is non-negative and strictly less than pmdata->nsvids before accessing the pmdata->svids array inside	N/A	More Details

	svdm_consume_modes().		
CVE-2026-64429	In the Linux kernel, the following vulnerability has been resolved: gpio: eic-sprd: use raw_spinlock_t in the irq startup path sprd_eic_irq_unmask() enables the GPIO IRQ and then updates controller state through sprd_eic_update(), which takes sprd_eic->lock with spin_lock_irqsave(). The callback can be reached from irq_startup() while setting up a requested IRQ. That path is not sleepable, but on PREEMPT_RT a regular spinlock_t becomes a sleeping lock. This issue was found by our static analysis tool and then manually reviewed against the current tree. The grounded PoC kept the request_threaded_irq() -> __setup_irq() -> irq_startup() -> sprd_eic_irq_unmask() -> sprd_eic_update() carrier and used the original spin_lock_irqsave(&sprd_eic->lock) edge. Lockdep BUG: sleeping function called from invalid context hardirqs last disabled at ... __setup_irq.constprop.0 ... [vuln_msv] sprd_rt_spin_lock_irqsave+0x1c/0x30 [vuln_msv] sprd_eic_update.constprop.0+0x48/0x90 [vuln_msv] sprd_eic_irq_unmask.constprop.0+0x35/0x50 [vuln_msv] __setup_irq.constprop.0+0xd/0x30 [vuln_msv] Convert the Spreadtrum EIC controller lock to raw_spinlock_t. The locked section only serializes MMIO register updates and does not contain sleepable operations, so keeping it non-sleeping is appropriate for the irqchip callbacks.	N/A	More Details
CVE-2026-64331	In the Linux kernel, the following vulnerability has been resolved: usbip: vudc: fix NULL deref in vep_dequeue() vep_alloc_request() wasn't initializing vrequest->udc, so cancellations on the FunctionFS AIO path were arriving in vep_dequeue without a valid UDC reference. Since vrequest->udc is never actually properly used anywhere, we opt to remove it, and update vep_dequeue to obtain a reference to the udc with ep_to_vudc(), consistent with the other vep_ ops. AFAICT this bug has existed for ~10 years. Seems that nobody has really stressed the FunctionFS AIO path on usbip's vudc. I tested this fix in a QEMU aarch64 guest driving FunctionFS endpoints via AIO. Before the fix, running `usbip attach` from the host would cause the guest to oops with the following backtrace: Call trace: vep_dequeue+0x1c/0xe4 (P) usb_ep_dequeue+0x14/0x20 ffs_aio_cancel+0x24/0x34 __arm64_sys_io_cancel+0xb0/0x124 do_el0_svc+0x68/0x100 el0_svc+0x18/0x5c el0t_64_sync_handler+0x98/0xdc el0t_64_sync+0x154/0x158	N/A	More Details
CVE-2026-64332	In the Linux kernel, the following vulnerability has been resolved: USB: ulpi: fix memory leak on registration failure The allocated device name is never freed on early ULPI device registration failures. Fix this by initialising the device structure earlier and releasing the initial reference whenever registration fails.	N/A	More Details
CVE-2026-64334	In the Linux kernel, the following vulnerability has been resolved: USB: serial: digi_acceleport: fix hard lockup on disconnect If submitting the OOB write urb fails persistently (e.g if the device is being disconnected) the driver would loop indefinitely with interrupts disabled. Check for urb submission errors when sending OOB commands to avoid hanging if, for example, open(), set_termios() or close() races with a physical disconnect. This is issue was flagged by Sashiko when reviewing an unrelated change to the driver.	N/A	More Details
CVE-2026-64335	In the Linux kernel, the following vulnerability has been resolved: USB: serial: digi_acceleport: fix broken rx after throttle If the port is closed while throttled, the read urb is never resubmitted and the port will not receive any further data until the device is reconnected (or the driver is rebound). Clear the throttle flags and submit the urb if needed when opening the port.	N/A	More Details
CVE-2026-64336	In the Linux kernel, the following vulnerability has been resolved: USB: serial: keyspan_pda: fix information leak The write() callback is supposed to return the number of characters accepted or a negative errno. Since the addition of write fifo support the keyspan_pda implementation will however return the number characters submitted to the device if the write urb is not already in use. If this number is larger than the number of characters passed to write(), the line discipline continues writing data from beyond the tty write buffer. Fix the information leak by making sure that keyspan_pda_write_start() returns zero on success as intended.	N/A	More Details
CVE-2026-64337	In the Linux kernel, the following vulnerability has been resolved: usb: mtu3: unmap request DMA on queue failure mtu3_gadget_queue() maps the request before checking whether the QMU GPD ring can accept another transfer. the request is returned with -EAGAIN before it is linked on the endpoint request list if mtu3_prepare_transfer() fails. Normal completion and dequeue paths unmap requests from mtu3_req_complete(), but this error path never reaches that helper, so the DMA mapping is left active. Unmap the request before returning from the failed queue path.	N/A	More Details
CVE-2026-64338	In the Linux kernel, the following vulnerability has been resolved: USB: misc: uss720: unregister parport on probe failure uss720_probe() registers a parport before reading the 1284 register used to detect unsupported Belkin F5U002 adapters. If get_1284_register() fails, the error path drops the driver private data and the USB device reference, but leaves the parport device registered. Leaving the port registered is more than a private allocation leak: parport_register_port() has already reserved a parport number and registered the parport bus device, while pp->private_data still points at the private data that the common error path is about to release. Undo the pre-announce registration in the get_1284_register() failure branch before jumping to the common private-data cleanup path. Clear priv->pp first, matching the disconnect path and avoiding a stale pointer in the private data. This issue was identified during our ongoing static-analysis research while reviewing kernel code.	N/A	More Details
CVE-2026-64339	In the Linux kernel, the following vulnerability has been resolved: usb: misc: usbio: bound bulk IN response length to the received transfer usbio_bulk_msg() copies bpkt_len = le16_to_cpu(bpkt->len) bytes out of the bulk IN buffer (usbio->rxbuf, allocated with size usbio->rxbuf_len) into the caller's buffer. bpkt_len is fully controlled by the device and is only checked against ibuf_len; ibuf_len in turn is checked against usbio->txbuf_len, not against rxbuf_len: if ((obuf_len > (usbio->txbuf_len - sizeof(*bpkt))) || (ibuf_len > (usbio->txbuf_len - sizeof(*bpkt)))) return -EMSGSIZE; txbuf_len and rxbuf_len are taken independently from the bulk OUT and bulk IN endpoint wMaxPacketSize in usbio_probe(). A malicious or malfunctioning device that advertises a large bulk OUT endpoint and a small bulk IN endpoint (e.g. by claiming one of the quirk-free IDs such as the Lattice NX33U, 0x2ac1:0x20cb) therefore makes ibuf_len, and hence the device-supplied bpkt_len, exceed rxbuf_len. memcpy() then reads up to txbuf_len - rxbuf_len bytes past the end of the rxbuf slab object. The over-read bytes are handed back to the i2c layer and on to user space through i2c-dev, disclosing adjacent slab memory; with KASAN this is reported as a slab-out-of-bounds read. The number of bytes actually received is already known: act equals the URB actual_length and is bounded by	N/A	More Details

	rxbuf_len. Reject any response that claims more payload than was received, mirroring the existing "act < sizeof(*bpkt)" check just above. The control path (usbio_ctrl_msg()) is not affected: it uses a single buffer (ctrlbuf) for both directions, so its analogous copy can never leave the allocation. Found by code review. The out-of-bounds read was confirmed under AddressSanitizer with a faithful userspace model of usbio_bulk_msg()'s receive path (an rxbuf_len-sized buffer, the same act/ibuf_len/bpkt_len checks and the memcpy). A USB raw-gadget + dummy_hcd reproducer is also available.		
CVE-2026-64340	In the Linux kernel, the following vulnerability has been resolved: USB: legousbtower: fix use-after-free on disconnect race mutex_unlock() may access the mutex structure after releasing the lock and therefore cannot be used to manage lifetime of objects directly (unlike spinlocks and refcounts). [1][2] Use a kref to release the driver data to avoid use-after-free in mutex_unlock() when release() races with disconnect(). [1] a51749ab34d9 ("locking/mutex: Document that mutex_unlock() is non-atomic") [2] 2b9d9e0a9ba0 ("locking/mutex: Clarify that mutex_unlock(), and most other sleeping locks, can still use the lock object after it's unlocked")	N/A	More Details
CVE-2026-64341	In the Linux kernel, the following vulnerability has been resolved: USB: iowarrior: fix use-after-free on disconnect race mutex_unlock() may access the mutex structure after releasing the lock and therefore cannot be used to manage lifetime of objects directly (unlike spinlocks and refcounts). [1][2] Use a kref to release the driver data to avoid use-after-free in mutex_unlock() when release() races with disconnect(). [1] a51749ab34d9 ("locking/mutex: Document that mutex_unlock() is non-atomic") [2] 2b9d9e0a9ba0 ("locking/mutex: Clarify that mutex_unlock(), and most other sleeping locks, can still use the lock object after it's unlocked")	N/A	More Details
CVE-2026-64342	In the Linux kernel, the following vulnerability has been resolved: USB: iowarrior: fix use-after-free on disconnect Submitted write URBs are not stopped on close() and therefore need to be stopped unconditionally on disconnect() to avoid use-after-free in the completion handler.	N/A	More Details
CVE-2026-64343	In the Linux kernel, the following vulnerability has been resolved: USB: ldusb: fix use-after-free on disconnect race mutex_unlock() may access the mutex structure after releasing the lock and therefore cannot be used to manage lifetime of objects directly (unlike spinlocks and refcounts). [1][2] Use a kref to release the driver data to avoid use-after-free in mutex_unlock() when release() races with disconnect(). [1] a51749ab34d9 ("locking/mutex: Document that mutex_unlock() is non-atomic") [2] 2b9d9e0a9ba0 ("locking/mutex: Clarify that mutex_unlock(), and most other sleeping locks, can still use the lock object after it's unlocked")	N/A	More Details
CVE-2026-66824	A stored cross-site scripting vulnerability existed in the capture tree visualization page. The application embedded the serialized capture tree directly into an inline JavaScript block using the Jinja safe filter. Because the tree data can contain values derived from captured and potentially attacker-controlled web content, a specially crafted value could prematurely terminate the surrounding <script> element and inject arbitrary HTML or JavaScript. The malicious code would execute in the browser of a user viewing the affected capture tree. Successful exploitation could allow an attacker to perform actions using the victim’s authenticated session, access information available to the victim, or modify application data within the permissions of the affected user. The patch removes the JSON data from the HTML document and retrieves it through a dedicated API endpoint. The client then processes the response using response.json(), preventing capture data from being interpreted as executable content within the original page’s HTML or JavaScript context.	N/A	More Details
CVE-2026-66825	Pivotick contains a cross-site scripting vulnerability in the sidebar property-list component. Values associated with link-like properties, such as url, uri, href, link, website, or homepage, were rendered as hyperlinks without validating their URL scheme. An attacker able to supply or influence node or edge property data could provide a malicious value using the javascript: scheme, including variants obfuscated with whitespace or control characters. If a user clicked the generated property link, attacker-controlled JavaScript could execute in the context of the Pivotick application. Successful exploitation could allow the attacker to access information available to the victim’s browser session or perform actions with the victim’s privileges. The vulnerability was addressed by normalizing property values and preventing URLs with non-allowlisted schemes from being rendered as clickable links.	N/A	More Details
CVE-2026-64305	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - protect service table iterations with service_lock The service_table list is protected by service_lock when entries are added or removed (in adf_service_add() and adf_service_remove()), but several functions iterate over the list without holding this lock. A concurrent adf_service_register() or adf_service_unregister() call could modify the list during traversal, leading to list corruption or a use-after-free. Fix this by holding service_lock across all list_for_each_entry() iterations of service_table in adf_dev_init(), adf_dev_start(), adf_dev_stop(), adf_dev_shutdown(), adf_dev_restarting_notify(), adf_dev_restarted_notify(), and adf_error_notifier(). The lock ordering is safe: callers of the static helpers (adf_dev_up() and adf_dev_down()) acquire state_lock before service_lock, and no event_hld callback or service_lock holder ever acquires state_lock in the reverse order.	N/A	More Details
CVE-2026-64302	In the Linux kernel, the following vulnerability has been resolved: x86/mm: Fix freeing of PMD-sized vmemmap pages Commit bf9e4e30f353 ("x86/mm: use pagetable_free()"), switched from freeing non-boot page tables through __free_pages() to pagetable_free(). However, the function is also called to free vmemmap pages. Given that vmemmap pages are not page tables, already the page_ptdesc(page) is wrong. But worse, pagetable_free() calls: __free_pages(page, compound_order(page)); Since vmemmap pages are not compound pages (see vmemmap_alloc_block()) -- except for HVO, which doesn't apply here -- only first page of a PMD-sized vmemmap page is freed, leaking the other ones. Fix it by properly decoupling pagetable and vmemmap freeing. free_pagetable() no longer has to mess with SECTION_INFO, as only the vmemmap is marked like that in register_page_bootmem_memmap(). The indentation in remove_pmd_table() is messed up. Fix that while touching it. Bootmem info handling will soon be fixed up. For now, handle it similar to free_pagetable(), just avoiding the ifdef. [dhansen: changelog munging. More imperative voice]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: fix moving cancelled entry to ent_in_userspace list fuse_uring_cancel() moves entries that are available (these have no reqs attached) to the		

CVE-2026-64263	ent_in_userspace list. ent_list_request_expired() checks the first entry on ent_in_userspace and dereferences ent->fuse_req unconditionally, which will crash on a cancelled entry that was moved to this list. Fix this by freeing the entry and dropping queue_refs directly in fuse_uring_cancel(). This is safe because cancel is the cancel handler itself - after io_uring_cmd_done(), no more cancels will be dispatched for this command, and teardown serializes with cancel via queue->lock. Since cancel now decrements queue_refs, fuse_uring_abort() must no longer gate fuse_uring_abort_end_requests() on queue_refs > 0, as cancelled entries may have already dropped queue_refs while requests are still queued. Remove the gate so abort always flushes requests and stops queues.	N/A	More Details
CVE-2026-64264	In the Linux kernel, the following vulnerability has been resolved: fuse-uring: fix EFAULT clobber in fuse_uring_commit copy_from_user() returns the number of bytes not copied as an unsigned residual on failure (1..sizeof(struct fuse_out_header)). fuse_uring_commit stores that residual in ssize_t err, sets req->out.h.error to -EFAULT, then jumps to out: with err still holding the positive residual. err = copy_from_user(&req->out.h, &ent->headers->in_out, sizeof(req->out.h)); if (err) { req->out.h.error = -EFAULT; goto out; /* err is the positive residual */ } ... out: fuse_uring_req_end(ent, req, err); fuse_uring_req_end() then runs if (error) req->out.h.error = error; which overwrites the just-assigned -EFAULT with the positive residual. FUSE callers such as fuse_simple_request() test err < 0 to detect failure, so the positive value is interpreted as success and the caller proceeds with an uninitialised or partial req->out.args. Fix by assigning err = -EFAULT in the failure branch before jumping to out, so fuse_uring_req_end() receives a negative errno and sets req->out.h.error to -EFAULT.	N/A	More Details
CVE-2026-64267	In the Linux kernel, the following vulnerability has been resolved: fuse: avoid 32-bit prune notification count wrap FUSE_NOTIFY_PRUNE validates the nodeid payload length with: size - sizeof(outarg) != outarg.count * sizeof(u64) On 32-bit kernels, size_t is also 32 bits, so the daemon-controlled count multiplication can wrap. A prune notification with count 0x20000000 and no nodeid payload passes the check, enters the copy loop, and asks the device copy path to read nodeids that are not present in the userspace write buffer. In QEMU this reaches the fuse_copy_fill() BUG_ON(!err) path. Validate the payload length with array_size() instead. That accepts exactly the same valid messages, but avoids wrapping arithmetic before the copy loop consumes the count.	N/A	More Details
CVE-2026-64270	In the Linux kernel, the following vulnerability has been resolved: Input: mms114 - reject an oversized device packet size mms114_interrupt() reads a packet of touch data from the device into a fixed-size on-stack buffer struct mms114_touch touch[MMS114_MAX_TOUCH]; which holds MMS114_MAX_TOUCH (10) events of MMS114_EVENT_SIZE (8) bytes, i.e. 80 bytes. The length of the I2C read into it is taken verbatim from the device: packet_size = mms114_read_reg(data, MMS114_PACKET_SIZE); if (packet_size <= 0) goto out; ... error = __mms114_read_reg(data, MMS114_INFORMATION, packet_size, (u8 *)touch); packet_size is a single device register byte (0x0F) and the only check is the lower bound packet_size <= 0; it is never bounded against the size of touch[]. A malfunctioning, malicious or counterfeit controller (or an attacker tampering with the I2C bus) can report a packet_size of up to 255, so __mms114_read_reg() writes up to 175 bytes past the end of touch[] on the IRQ-thread stack: a stack out-of-bounds write that can overwrite the stack canary, saved registers and the return address. A well-formed device never reports more than the buffer holds, so reject an oversized packet and drop the report, consistent with the handler's other error paths, rather than reading past the buffer.	N/A	More Details
CVE-2026-64271	In the Linux kernel, the following vulnerability has been resolved: Input: touchwin - reset the packet index on every complete packet tw_interrupt() accumulates each non-zero serial byte into a fixed three-byte buffer with a running index that is only reset once a full packet has been received *and* the device's two Y bytes agree: tw->data[tw->idx++] = data; if (tw->idx == TW_LENGTH && tw->data[1] == tw->data[2]) { ... tw->idx = 0; } The reset is gated on tw->data[1] == tw->data[2], a value the device controls. A malicious, malfunctioning or counterfeit Touchwindow peripheral can stream non-zero bytes whose 2nd and 3rd bytes differ: the index reaches TW_LENGTH without the equality holding, is never reset, and keeps growing, so tw->data[tw->idx++] walks off the end of the three-byte array and the rest of the heap-allocated struct tw, one attacker-chosen byte at a time -- an unbounded, device-driven heap out-of-bounds write. Reset the index on every completed packet and report an event only when the two Y bytes match, like the other serio touchscreen drivers do.	N/A	More Details
CVE-2026-64272	In the Linux kernel, the following vulnerability has been resolved: Input: mms114 - fix touch indexing for MMS134S and MMS136 The MMS134S and MMS136 touch controllers have an event size of 6 bytes rather than 8 bytes. When __mms114_read_reg() reads the touch data packet from the device into the touch buffer, the events are packed tightly at 6-byte intervals. However, the driver iterates through the events using standard C array indexing (touch[index]), where each element is sizeof(struct mms114_touch) (8 bytes) apart. As a result, any touch events beyond the first one are read from incorrect offsets and parsed improperly. Fix this by explicitly calculating the byte offset for each touch event based on the device's specific event size.	N/A	More Details
CVE-2026-64273	In the Linux kernel, the following vulnerability has been resolved: Input: iforce - bound the device-reported force-feedback effect index iforce_process_packet() handles a status report (packet id 0x02) by taking a force-feedback effect index straight from the device wire and using it to address the per-effect state array: i = data[1] & 0x7f; if (data[1] & 0x80) { if (!test_and_set_bit(FF_CORE_IS_PLAYED, iforce->core_effects[i].flags)) ... } else if (test_and_clear_bit(FF_CORE_IS_PLAYED, iforce->core_effects[i].flags)) { ... } The index is masked only with 0x7f, so it ranges 0..127, but core_effects[] holds only IFORCE_EFFECTS_MAX (32) entries. For an index of 32..127 the test_and_set_bit()/test_and_clear_bit() is an out-of-bounds single-bit read-modify-write past the array. core_effects[] is the second-to-last member of struct iforce, so the write lands in the trailing members and beyond the embedding kzalloc()'d iforce_serio / iforce_usb object. data[1] is unvalidated device payload on both transports (the USB interrupt endpoint and serio), and the status path is not gated on force feedback being present, so a malicious or counterfeit device can set or clear a bit at an attacker-chosen offset past the object. Reject an out-of-range index instead of indexing with it. Bound against the array dimension IFORCE_EFFECTS_MAX rather than dev->ff->max_effects so the check guarantees memory safety regardless of how many effects the device registered. A legitimate "effect started/stopped" status always carries an index below IFORCE_EFFECTS_MAX, so well-formed devices are unaffected; the neighbouring mark_core_as_ready() loop is already bounded and is left untouched.	N/A	More Details

CVE-2026-64274	In the Linux kernel, the following vulnerability has been resolved: Input: goodix - clamp the device-reported contact count goodix_ts_read_input_report() copies the number of touch points reported by the device into an on-stack buffer u8 point_data[2 + GOODIX_MAX_CONTACT_SIZE * GOODIX_MAX_CONTACTS]; which is sized for at most GOODIX_MAX_CONTACTS (10) contacts. The only runtime check bounds the per-interrupt count against ts->max_touch_num, but that value is taken verbatim from a 4-bit field of the device configuration block and is never clamped: ts->max_touch_num = ts->config[MAX_CONTACTS_LOC] & 0x0f; The nibble can be 0..15, so a malfunctioning, malicious or counterfeit controller (or an attacker tampering with the I2C bus) can advertise up to 15 contacts. goodix_ts_read_input_report() then accepts a touch_num of up to 15 and the second goodix_i2c_read() writes ts->contact_size * (touch_num - 1) bytes past the one-contact header into point_data - up to 30 bytes (45 with the 9-byte report format) beyond the 92-byte buffer: a stack out-of-bounds write. Clamp max_touch_num to GOODIX_MAX_CONTACTS, the number of contacts point_data[] is sized for, when reading it from the configuration.	N/A	More Details
CVE-2026-64275	In the Linux kernel, the following vulnerability has been resolved: Input: elan_i2c - prevent division by zero and arithmetic underflow The Elan I2C touchpad driver queries the device for its physical dimensions and trace counts to calculate the device resolution and width. However, if the device firmware or device tree provides invalid zero values for x_traces or y_traces, it results in a fatal division-by-zero exception leading to a kernel panic during device probe. Add checks to ensure these parameters are non-zero before performing the division. If invalid trace values are detected, fall back to a safe default of 1. Additionally, prevent an arithmetic underflow in the touch reporting logic. Previously, if the calculated or fallback width was smaller than ETP_FWIDTH_REDUCE (90), the subtraction would underflow, resulting in a massive unsigned integer being reported to userspace. Clamp the adjusted width to a minimum of 0 to safely handle small physical dimensions and fallback scenarios. Completing the probe with safe fallback values ensures the sysfs nodes are created, keeping the firmware update path intact so a recovery firmware can be flashed to the device.	N/A	More Details
CVE-2026-64278	In the Linux kernel, the following vulnerability has been resolved: i2c: imx-lpi2c: mark I2C adapter when hardware is powered down On some i.MX platforms, certain I2C client drivers keep a periodic workqueue which continues to trigger I2C transfers. During system suspend/resume, there exists a time window between: - suspend_noirq and the system entering suspend - the system starting to resume and resume_noirq In this window, the I2C controller resources such as clock and pinctrl may already be disabled or not yet restored. If a workqueue triggers an I2C transfer in this period, the driver attempts to access I2C registers while the hardware resources are unavailable, which may lead to system hang. Mark the I2C adapter as suspended during noirq suspend and block new transfers until resume, ensuring that I2C transfers are only issued when hardware resources are available.	N/A	More Details
CVE-2026-64282	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Don't leak PFN when kvm_translate_vncr() races MMU notifier In the case that kvm_translate_vncr() races with an MMU notifier the early return does not release a reference on the faulted in PFN. Add the necessary call to kvm_release_faultin_page() for the unused PFN.	N/A	More Details
CVE-2026-64283	In the Linux kernel, the following vulnerability has been resolved: KVM: guest_memfd: Treat memslot binding offset+size as unsigned values When binding a memslot to a guest_memfd file, treat the offset and size as unsigned values to fix a bug where the sum of the two can result in a false negative when checking for overflow against the size of the file. Passing unsigned values also avoids relying on somewhat obscure checks in other flows for safety, and tracks the offset and size as they are intended to be tracked, as unsigned values. On 64-bit kernels, the number of pages a memslot contains and thus the size (and offset) of its guest_memfd binding are unsigned 64-bit values. Taking the offset+size as an loff_t instead of a uoff_t inadvertently converts the unsigned value to a signed value if the offset and/or size is massive. Locally storing the offset and size as signed values is benign in and of itself (though even that is *extremely* difficult to discern), but operating on their sum is not. For the offset, KVM explicitly checks against a negative value, which might seem like a bug as KVM could incorrectly reject a legitimate binding, but that's not actually the case as KVM_CREATE_GUEST_MEMFD takes a signed value for its size, i.e. a would-be-negative offset is also greater than the maximum possible size of any guest_memfd file. Regarding the size, while KVM lacks an explicit check for a negative value, i.e. seemingly has a flawed overflow check, KVM restricts the number of pages in a single memslot to the largest positive signed 32-bit value: if (id < KVM_USER_MEM_SLOTS && (mem->memory_size >> PAGE_SHIFT) > KVM_MEM_MAX_NR_PAGES) return -EINVAL; and so that maximum "size" will ever be is 0x7fffffff000. The sum of the two is, however, problematic. While the size is restricted by KVM's memslot logic, the offset is not, i.e. the offset is completely unchecked until the "offset + size > i_size_read(inode)" check. If the offset is the (nearly) largest possible _positive_ value, then adding size to the offset can result in a signed, negative 64-bit value. When compared against the size of the file (guaranteed to be positive), the negative sum is always smaller, and KVM incorrectly allows the absurd offset. Opportunistically add missing includes in kvm_mm.h (instead of relying on its parents).	N/A	More Details
CVE-2025-59181	Ericsson Packet Core Controller (PCC) versions prior to 1.39 contain a directory traversal vulnerability in Configuration Management that could allow an attacker to change directory permissions, denying access to legitimate users.	N/A	More Details
CVE-2026-64288	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: nv: Avoid dereferencing NULL VNCR pseudo-TLB VNCR TLB invalidation occurs from MMU notifiers or TLBI instructions, and either can race against a vcpu not being onlined yet (no pseudo-TLB allocated). Similarly, the TLB might be invalid, and the invalidation should be skipped in this case. Both kvm_invalidate_vncr_ipa() and kvm_invalidate_vncr_va() are expected to perform the same checks, except that the latter doesn't check for the allocation and blindly dereferences the pointer. Solve this by introducing a new iterator built on top of the usual kvm_for_each_vcpu() that checks for both of the above conditions, and convert the two users to it.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: iommufd: Set upper bounds on cache invalidation entry_num and entry_len iommufd_hwpt_invalidate() takes a user-controlled entry_num and entry_len, each bounded only by U32_MAX. An entry_len beyond the kernel's struct size makes the copy helper verify the extra bytes are zero,		

CVE-2026-64289	scanning that excess in one uninterruptible pass; a multi-gigabyte value over zeroed user memory trips the soft-lockup watchdog. A large entry_num is the other half, driving the backend invalidation loop with no reschedule. The VT-d nested handler, for one, copies each entry and flushes caches per iteration, pinning the CPU on a non-preemptible kernel. Cap both in the ioctl. entry_len is held under PAGE_SIZE, above any request struct, and entry_num under 1 << 19, the order of a hardware invalidation queue and well beyond any real batch, bounding the per-call loop length.	N/A	More Details
CVE-2026-64290	In the Linux kernel, the following vulnerability has been resolved: iommufd: Break the loop on failure in iommufd_fault_fops_read() On a copy_to_user() failure inside the inner list_for_each_entry, only the inner loop breaks; the outer while re-fetches the just-restored fault group and retries the failing copy_to_user() forever, spinning the reader at 100% CPU with fault->mutex held. Check rc after the inner loop and break the outer while as well.	N/A	More Details
CVE-2026-64291	In the Linux kernel, the following vulnerability has been resolved: iommufd: Set veventq_depth upper bound iommufd_veventq_alloc() accepts any !0 veventq_depth from userspace, with an upper bound at U32_MAX. This leaves a vulnerability where userspace can allocate excessively large queues to exhaust kernel memory reserves. Cap the veventq_depth (maximum number of entries) to 1 << 19, matching the maximum number of entries in the SMMUv3 EVTQ (the largest use case today).	N/A	More Details
CVE-2026-59240	The vulnerability involves an Insecure Direct Object Reference (IDOR) in the `DeleteNotificationController::delete()` method at endpoint `GET /notification/delete/{id}`. The flaw allows any authenticated user, regardless of company or permissions, to delete notifications belonging to any other user in the system. The controller retrieves the target record with `Notification::findOrFail($id)` and deletes it without validating `user_id` or `company_id` ownership, unlike the sibling `SetNotificationReadAjaxController`, which correctly scopes lookups by `Auth::id()`. Because notification identifiers are sequential, an attacker can iterate over IDs to systematically delete notifications belonging to any user, denying them visibility of ticket alerts, task assignments, and other system events.	N/A	More Details
CVE-2026-55685	React Router is a router for React. In versions 7.0.0 through 7.17.0, the manifest endpoint could be accessed via unauthenticated targeted requests that would put heavy load on the server and slow down response times. This issue is a follow up to CVE-2026-42342, and does not does not impact React Router applications using Declarative Mode (<BrowserRouter>) or Data Mode (createBrowserRouter/<RouterProvider>). This issue has been fixed in version 7.18.0.	N/A	More Details
CVE-2026-53669	React Router is a router for React. Versions 6.0.0 through 7.17.0 are vulnerable to Open Redirtect through use of backslashes in <Link> and useNavigate. This issue is a follow up to CVE-2025-68470 and has been fixed in version 7.18.0.	N/A	More Details
CVE-2026-64292	In the Linux kernel, the following vulnerability has been resolved: iommufd: Move vevent memory allocation outside spinlock The veventq memory allocation happens inside the spinlock. Given its depth is decided by the user space, this leaves a vulnerability, where userspace can allocate large queues to exhaust atomic memory reserves. Move the allocation outside the spinlock and use GFP_NOWAIT, which can fail fast under memory pressure without dipping into the GFP_ATOMIC reserves or direct-reclaiming from the threaded IRQ handler. On allocation failure, queue the lost_events_header (so userspace learns of the drop) and return -ENOMEM so the caller learns of the kernel-side memory pressure. This is intentionally distinct from the queue-overflow path, which also queues the lost_events_header but returns 0: a full queue is an expected userspace-pacing condition rather than a kernel error. A subsequent change will cap the upper bound of the veventq_depth.	N/A	More Details
CVE-2026-64294	In the Linux kernel, the following vulnerability has been resolved: mm: do file ownership checks with the proper mount idmap Ever since idmapped mounts were introduced, inode ownership checks (for side-channel protection) in mincore() and madvise(MADV_PAGEOUT) were done against the nop_mnt_idmap, which completely ignores the file's mount's idmap. This results in odd edgecases like: 1) mount/bind-mount with an idmap userA:userB:1 2) userB runs an owner_or_capable() check on file that is owned by userA on-disk/in-memory, but owned by userB after idmap translation 3) owner_or_capable() mysteriously fails as the correct idmap wasn't supplied In the case of mincore/madvise MADV_PAGEOUT, this is usually benign, because file_permission(file, MAY_WRITE) will probably succeed, as it uses the proper idmap internally, but it does not need to be the case on e.g a 0444 file where even the owner itself doesn't have permissions to write to it. Since this is clearly not trivial to get right, introduce a file_owner_or_capable() that can carry the correct semantics, and switch the various users in mm to it. The issue was found by manual code inspection & an off-list discussion with Jan Kara.	N/A	More Details
CVE-2026-64295	In the Linux kernel, the following vulnerability has been resolved: mm: page_ext: add count limit to page_ext_iter_next to prevent invalid PFN access The page_ext iteration API does not validate if the PFN still belongs to a valid section while advancing the iterator. When dynamically adding memory in the hotplug path, it can lead to a NULL pointer dereference during page_ext_lookup at the boundary of the last valid section when iterator count equals __pgcount. The for_each_page_ext() macro calls page_ext_iter_next() as its loop increment. for_each_page_ext() does a "__page_ext = page_ext_iter_next(&__iter)" at the end. This causes page_ext_iter_next() to increment iter->index past __pgcount and call page_ext_lookup(start_pfn + __pgcount). During memory hotplug (online), the PFN at start_pfn + __pgcount may belong to a section that has not yet been initialized, causing page_ext_lookup() to trigger a NULL pointer dereference. [14.555124][T846] Call trace: [14.555125][T846] lookup_page_ext+0x6c/0x108 (P) [14.555127][T846] page_ext_lookup+0x30/0x3c [14.555129][T846] __reset_page_owner+0x11c/0x260 [14.571201] [T846] __free_pages_ok+0x5e8/0x8e0 [14.571204][T846] __free_pages_core+0x78/0xf0 [14.571206][T846] generic_online_page+0x14/0x24 [14.597782][T846] online_pages+0x178/0x30c [14.597784][T846] memory_block_change_state+0x284/0x32c [14.597787][T846] memory_subsys_online+0x4c/0x64 [14.597789][T846] device_online+0x88/0xb0 [14.597791][T846] online_memory_block+0x30/0x40 [14.597793][T846] walk_memory_blocks+0xac/0xe8 [14.597794][T846] add_memory_resource+0x280/0x298 [14.656161][T846] add_memory+0x60/0x98 Move the iteration boundary enforcement inside the iterator functions, so callers cannot inadvertently access beyond the requested range.	N/A	More Details

CVE-2026-64297	In the Linux kernel, the following vulnerability has been resolved: module: decompress: check return value of module_extend_max_pages() module_extend_max_pages() calls kvrealloc() internally and returns -ENOMEM on allocation failure. The return value is never checked. If the initial allocation fails, info->pages remains NULL and info->max_pages remains 0. Subsequent calls to module_get_next_page() will attempt to dynamically grow the array by calling module_extend_max_pages(info, 0) since info->used_pages is 0. This results in kvrealloc(NULL, 0) returning ZERO_SIZE_PTR, which is treated as a success, leading to a dereference of ZERO_SIZE_PTR and a kernel oops. Fix: add the missing error check after module_extend_max_pages() and return immediately on failure. This matches the pattern used by every other kvrealloc() caller in the module loading path. [Sami: Corrected the analysis in the commit message.]	N/A	More Details
CVE-2026-64301	In the Linux kernel, the following vulnerability has been resolved: regulator: scmi: fix of_node refcount leak in scmi_regulator_probe() scmi_regulator_probe() calls of_find_node_by_name() which takes a reference on the returned device node. On the error path where process_scmi_regulator_of_node() fails, the function returns without calling of_node_put() on the child node, leaking the reference. Add of_node_put(np) on the error path to properly release the reference.	N/A	More Details
CVE-2026-64344	In the Linux kernel, the following vulnerability has been resolved: USB: idmouse: fix use-after-free on disconnect race mutex_unlock() may access the mutex structure after releasing the lock and therefore cannot be used to manage lifetime of objects directly (unlike spinlocks and refcounts). [1][2] Use a kref to release the driver data to avoid use-after-free in mutex_unlock() when release() races with disconnect(). [1] a51749ab34d9 ("locking/mutex: Document that mutex_unlock() is non-atomic") [2] 2b9d9e0a9ba0 ("locking/mutex: Clarify that mutex_unlock(), and most other sleeping locks, can still use the lock object after it's unlocked")	N/A	More Details
CVE-2026-64345	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_printer: take kref only for successful open printer_open() returns -EBUSY when the character device is already open, but it increments dev->kref regardless of the return value. VFS does not call ->release() for a failed open, so every rejected second open permanently leaks one reference. Move kref_get() into the successful-open branch.	N/A	More Details
CVE-2026-64346	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: udc: Fix use-after-free in gadget_match_driver The udc structure acts as the management structure for the gadget, but their lifecycles are decoupled. A race condition exists where usb_del_gadget() frees the udc memory (e.g., via mode-switch work) while gadget_match_driver() concurrently accesses the freed udc memory (e.g., via configfs), causing a Use-After-Free (UAF) that triggers a NULL pointer dereference when the freed memory is zeroed: [39430.908615][T1171] Unable to handle kernel NULL pointer dereference at virtual address 0000000000000000 [39430.911397][T1171] pc : __pi_strcmp+0x20/0x140 [39430.911441][T1171] lr : gadget_match_driver+0x34/0x60 ... [39430.911890][T1171] usb_gadget_register_driver_owner+0x50/0xf8 [39430.911910][T1171] gadget_dev_desc_UDC_store+0xf4/0x140 [39430.931308][T1171] configfs_write_iter+0xec/0x134 [39430.957058][T1171] Workqueue: events_freezable __dwc3_set_mode [39430.957287][T1171] dwc3_gadget_exit+0x34/0x8c [39430.957304][T1171] __dwc3_set_mode+0xc0/0x664 Fix this by ensuring the udc structure remains allocated until the gadget is released. To achieve this, introduce a new usb_gadget_release() routine to the core. When the gadget is added, usb_add_gadget() stores the gadget's release routine in the udc structure and takes a reference to the udc. When the gadget is released, usb_gadget_release() drops the reference to the udc and then calls the gadget's release routine.	N/A	More Details
CVE-2026-64538	In the Linux kernel, the following vulnerability has been resolved: ipv6: Fix null-ptr-deref in fib6_nh_mtu_change(). fib6_nh_mtu_change() re-fetches idev via __in6_dev_get(arg->dev) and dereferences idev->cnf.mtu6 without a NULL check. addrconf_ifdown() clears dev->ip6_ptr with RCU_INIT_POINTER() after rt6_disable_ip() has released tb6_lock, so the RA-driven MTU walk can observe a NULL idev and oops. The caller rt6_mtu_change_route() guards its own __in6_dev_get(), but this re-fetch is unguarded; nexthop-backed routes survive addrconf_ifdown()'s flush, so the walk still reaches it after ip6_ptr is nulled. Return 0 when idev is NULL, matching rt6_mtu_change_route() and the fib6_mtu() fix in commit 5ad509c1fdad ("ipv6: Fix null-ptr-deref in fib6_mtu()."). Oops: general protection fault, ... KASAN: null-ptr-deref in range [0x00000000000002a8-0x00000000000002af] RIP: 0010:fib6_nh_mtu_change+0x203/0x990 rt6_mtu_change_route+0x141/0x1d0 __fib6_clean_all+0xd0/0x160 rt6_mtu_change+0xb4/0x100 ndisc_router_discovery+0x24b5/0x2cb0 icmpv6_rcv+0x12e9/0x1710 ipv6_rcv+0x39b/0x410	N/A	More Details
CVE-2026-59730	Astro is a web framework for content-driven websites. In versions 8.1.0 through 11.0.1, when trailingSlash: 'always' is configured, the @astrojs/node standalone server's static file handler appends a trailing slash to request paths and issues a 301 redirect. Paths beginning with /\ (slash-backslash) were not recognized as internal paths, so the handler would echo the raw path back in the Location header. Because browsers treat \ as / per the WHATWG URL specification, the resulting redirect could resolve to an external host. Preconditions for exploitation: trailingSlash: 'always' must be set (non-default; the default is 'ignore'), the request path must not have a file extension in its final segment, and an attacker must deliver the crafted link to a user. This issue has been fixed in version 11.0.2.	N/A	More Details
CVE-2026-64362	In the Linux kernel, the following vulnerability has been resolved: HID: lg-g15: cancel pending work on remove to fix a use-after-free lg_g15_data is allocated with devm and holds a work item. The report handlers schedule that work straight from device input. lg_g15_event() and lg_g15_v2_event() do it on the backlight cycle key, and lg_g510_leds_event() does it too. The worker dereferences the lg_g15_data back through container_of. The driver had no remove callback and never cancelled the work. So if a report scheduled the work and the keyboard was then unplugged, devres freed lg_g15_data while the work was still pending or running, and the worker touched freed memory. This is a use-after-free. It is reachable as a race on device unplug. Add a remove callback that cancels the work before devres frees the state. g15->work is only initialized for the models that schedule it (G15, G15 v2, G510). The G13 and Z-10 leave it zeroed, so guard the cancel on g15->work.func to avoid cancelling a work that was never set up. The g15 NULL test mirrors the one already in lg_g15_raw_event().	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: HID: appleir: fix UAF on pending key_up_timer in		

CVE-2026-64363	remove() appleir_remove() runs hid_hw_stop() before timer_delete_sync(). hid_hw_stop() synchronously unregisters the HID input device via hid_disconnect() -> hidinput_disconnect() -> input_unregister_device(), which drops the last reference and frees the underlying input_dev when no userspace handle holds it open. key_up_tick() reads appleir->input_dev and calls input_report_key() / input_sync() on it. The timer is armed from appleir_raw_event() with a HZ/8 (~125 ms) timeout on every keydown and key-repeat report. If a key was pressed shortly before the device is disconnected, the timer can fire after hid_hw_stop() has freed input_dev but before the teardown drains it. A simple reorder is not sufficient. Putting the timer drain first still leaves a window where a USB URB completion (raw_event) running during hid_hw_stop() can call mod_timer() and re-arm the timer, which then fires after hidinput_disconnect() has freed input_dev. The same URB-completion window also lets raw_event() reach key_up(), key_down() and battery_flat() directly, all of which dereference appleir->input_dev. Introduce a 'removing' flag on struct appleir, gated by the existing spinlock. appleir_remove() sets the flag under the lock and then shuts down the timer with timer_shutdown_sync(), which both drains any in-flight callback and permanently disables further mod_timer() calls. appleir_raw_event() and key_up_tick() bail out early if the flag is set, so no path can arm or run the timer, or dereference appleir->input_dev, after remove() has started tearing down. The keyrepeat and flatbattery branches of appleir_raw_event() previously called into the input layer without holding the spinlock; take it now so the flag check is well-defined. This incidentally closes a pre-existing read-side race on appleir->current_key in the keyrepeat branch. This bug is structurally a sibling of commit 4db2af929279 ("HID: appletb-kbd: fix UAF in inactivity-timer cleanup path") and has been present since the driver was introduced.	N/A	More Details
CVE-2026-64365	In the Linux kernel, the following vulnerability has been resolved: HID: letsketch: fix UAF on inrange_timer at driver unbind letsketch_driver does not provide a .remove callback, but letsketch_probe() arms a per-device timer: timer_setup(&data->inrange_timer, letsketch_inrange_timeout, 0); The timer is re-armed from letsketch_raw_event() with a 100 ms timeout on every pen-in-range report, and its callback dereferences data->input_tablet to deliver a synthetic BTN_TOOL_PEN release. letsketch_data is allocated with devm_kzalloc(), and its input_dev fields are devm-allocated via letsketch_setup_input_tablet(). On device unbind (USB unplug or rmmod), the HID core runs its default teardown and devm cleanup frees both letsketch_data and the input devices. Because no .remove callback exists, nothing drains the timer first: if raw_event armed it within ~100 ms of the unbind, the pending timer fires on freed memory. This is a UAF read of data and of data->input_tablet, followed by input_report_key() / input_sync() into the freed input_dev. The same problem can occur on the probe error path: if hid_hw_start() enabled I/O on an always-poll-quirk device and then failed, raw_event may have armed the timer before devm releases data. Fix by adding a .remove callback that calls hid_hw_stop() first. hid_hw_stop() synchronously kills the URBs that deliver raw_event(), so once it returns no path can re-arm the timer. timer_shutdown_sync() then drains any in-flight callback and permanently disables further mod_timer() calls. Apply the same timer_shutdown_sync() in the probe error path so the timer is guaranteed not to outlive data.	N/A	More Details
CVE-2026-64369	In the Linux kernel, the following vulnerability has been resolved: s390: Revert support for DCACHE_WORD_ACCESS load_unaligned_zeropad() reads eight bytes from unaligned addresses and may cross page boundaries. It handles exceptions which may happen if reading from the second page results in an exception. For pages which are donated to the Ultravisor for secure execution purposes the do_secure_storage_access() exception handler however does not handle such exceptions correctly. Such an exception may result in an endless exception loop which will never be resolved. An attempt to fix this [1] turned out to be not sufficient. For now revert load_unaligned_zeropad() until this problem has been resolved in a proper way. Note that the implementation of load_unaligned_zeropad() itself is correct. The revert is just a temporary workaround until there is complete fix for secure storage access exceptions. [1] commit b00be77302d7 ("s390/mm: Add missing secure storage access fixups for donated memory")	N/A	More Details
CVE-2026-64370	In the Linux kernel, the following vulnerability has been resolved: posix-cpu-timers: Fix pid refcount leak in do_cpu_nanosleep() error path In do_cpu_nanosleep(), posix_cpu_timer_create() takes a pid reference via get_pid() and stores it in timer.it.cpu.pid. If the subsequent posix_cpu_timer_set() call fails, the function returns immediately without calling posix_cpu_timer_del() to release the pid reference, causing a leak. Fix it by calling posix_cpu_timer_del() before the unlock-and-return on the error path, consistent with the other exit paths in the same function.	N/A	More Details
CVE-2026-64371	In the Linux kernel, the following vulnerability has been resolved: proc: protect ptrace_may_access() with exec_update_lock (part 1) Fix the easy cases where procfs currently calls ptrace_may_access() without exec_update_lock protection, where the fix is to simply add the extra lock or use mm_access(): - do_task_stat(): grab exec_update_lock - proc_pid_wchan(): grab exec_update_lock - proc_map_files_lookup(): use mm_access() instead of get_task_mm() - proc_map_files_readdir(): use mm_access() instead of get_task_mm() - proc_ns_get_link(): grab exec_update_lock - proc_ns_readlink(): grab exec_update_lock	N/A	More Details
CVE-2026-64373	In the Linux kernel, the following vulnerability has been resolved: cpufreq: Fix hotplug-suspend race during reboot During system reboot, cpufreq_suspend() is called via the kernel_restart() -> device_shutdown() path. Unlike the normal system suspend path, the reboot path does not call freeze_processes(), so userspace processes and kernel threads remain active. This allows CPU hotplug operations to run concurrently with cpufreq_suspend(). The original code has no synchronization with CPU hotplug, leading to a race condition where governor_data can be freed by the hotplug path while cpufreq_suspend() is still accessing it, resulting in a null pointer dereference: Unable to handle kernel NULL pointer dereference Call Trace: do_kernel_fault+0x28/0x3c cpufreq_suspend+0xdc/0x160 device_shutdown+0x18/0x200 kernel_restart+0x40/0x80 arm64_sys_reboot+0x1b0/0x200 Fix this by adding cpus_read_lock()/cpus_read_unlock() to cpufreq_suspend() to block CPU hotplug operations while suspend is in progress. [rjw: Changelog edits]	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: firmware_loader: fix device reference leak in firmware_upload_register() firmware_upload_register() -> fw_create_instance() -> device_initialize() After fw_create_instance() succeeds, the lifetime of the embedded struct device is expected to be managed through the device core reference counting, since fw_create_instance() has already called device_initialize(). In firmware_upload_register(), if alloc_lookup_fw_priv() fails after fw_create_instance() succeeds, the code reaches		More

2026-64376	free_fw_sysfs and frees fw_sysfs directly instead of releasing the device reference with put_device(). This may leave the reference count of the embedded struct device unbalanced, resulting in a refcount leak. The issue was identified by a static analysis tool I developed and confirmed by manual review. Fix this by using put_device(fw_dev) in the failure path and letting fw_dev_release() handle the final cleanup, instead of freeing the instance directly from the error path.	N/A	Details
CVE-2026-64377	In the Linux kernel, the following vulnerability has been resolved: cpufreq: qcom-cpufreq-hw: Fix possible double free qcom_cpufreq.data is allocated with devm_kzalloc() in probe() as an array of per-domain data. qcom_cpufreq_hw_cpu_init() stores a pointer to one element of this array in policy->driver_data. qcom_cpufreq_hw_cpu_exit() currently calls kfree() on policy->driver_data. This is not valid because the memory is devm-managed. For the first domain, this can free the devm-managed allocation while the devres entry is still active, leading to a possible double free when the platform device is later detached. For other domains, the pointer may refer to an element inside the array rather than the allocation base. Remove the kfree(data) call and let devres release qcom_cpufreq.data. This issue was found by a static analysis tool I am developing.	N/A	More Details
CVE-2026-64381	In the Linux kernel, the following vulnerability has been resolved: smb: client: Fix next buffer leak in receive_encrypted_standard() receive_encrypted_standard() allocates next_buffer before checking whether the number of compound PDUs already reached MAX_COMPOUND. If the limit check fails, the function returns immediately and the newly allocated next_buffer is not assigned to server->smallbuf/server->bigbuf, making it leaked. Move the MAX_COMPOUND check before allocating next_buffer.	N/A	More Details
CVE-2026-64404	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: avoid NULL deref of conn in iso_conn_big_sync() iso_conn_big_sync() drops the socket lock to call hci_get_route() and then re-acquires it, but dereferences iso_pi(sk)->conn->hcon afterwards without re-checking that conn is still valid. While the lock is dropped, the connection can be torn down under the same socket lock: iso_disconn_cfm() -> iso_conn_del() -> iso_chan_del() sets iso_pi(sk)->conn to NULL (and the broadcast teardown path can also clear conn->hcon on its own). When iso_conn_big_sync() re-acquires the lock and reads conn->hcon, conn may be NULL, causing a NULL pointer dereference (hcon is the first member of struct iso_conn). This path is reached from iso_sock_recvmsg() for a PA-sync broadcast sink socket (BT_SK_DEFER_SETUP | BT_SK_PA_SYNC), so the dropped-lock window can race with connection teardown driven by controller events. Re-validate iso_pi(sk)->conn and its hcon after re-acquiring the socket lock and bail out if the connection went away, as already done in the sibling iso_sock_rebind_bc().	N/A	More Details
CVE-2026-64405	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_conn: Fix null ptr deref in hci_abort_conn() hci_abort_conn() read hci_skb_event(hdev->sent_cmd) when a connection was pending, but hdev->sent_cmd can be NULL while req_status is still HCI_REQ_PEND, leading to a NULL pointer dereference and a general protection fault from the hci_rx_work() receive path. Instead of inspecting hdev->sent_cmd, track the in-flight create connection command with a new per-connection HCI_CONN_CREATE flag and route all cancellation through hci_cancel_connect_sync(), which dispatches to a dedicated per-type cancel function. The create command is in exactly one of two states: still queued, or in flight. The cancel function holds cmd_sync_work_lock across the whole decision: the worker takes this lock to dequeue every entry, so while it is held a queued command cannot start running and an in-flight command cannot complete and let the next command become pending. This keeps the flag test and hci_cmd_sync_cancel() atomic with respect to the worker, so a queued command is simply dequeued, and an in-flight command owned by this connection is cancelled without the risk of cancelling an unrelated command that became pending in the meantime. CIS uses the same flag mechanism via HCI_CONN_CREATE_CIS but cannot be dequeued per-connection. hci_acl_create_conn_sync() and hci_le_create_conn_sync() clear HCI_CONN_CREATE after the create command completes, but the command status handler can free conn via hci_conn_del() (for example when the controller rejects the connection) while the worker is still blocked on the connection complete event. Hold a reference on conn across the create command so the flag can be cleared without a use-after-free.	N/A	More Details
CVE-2026-64407	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btnxpuart: Fix out-of-bounds firmware read in nxp_recv_fw_req_v3() During the v3 firmware download the controller sends a v3_data_req with a 32 bit offset and a 16 bit len. nxp_recv_fw_req_v3() checks only the lower bound of the offset and then sends firmware from that offset. nxpdev->fw_dnld_v3_offset = offset - nxpdev->fw_v3_offset_correction; serdev_device_write_buf(nxpdev->serdev, nxpdev->fw->data + nxpdev->fw_dnld_v3_offset, len); Nothing checks that fw_dnld_v3_offset + len stays within nxpdev->fw->size, so a controller that asks for an offset or length past the firmware image makes the driver read past the end of nxpdev->fw->data and send that memory back over UART. nxp_recv_fw_req_v1() already bounds the same write. Add the equivalent check to the v3 path, reject the request when it falls outside the firmware image, and zero len on the error path so the fw_v3_prev_sent bookkeeping at free_skb stays consistent.	N/A	More Details
CVE-2026-64409	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: btmtksdio: fix infinite loop in btmtksdio_txrx_work() Every once in a while we see a hung btmtksdio_flush() task: INFO: task kworker/u17:0:189 blocked for more than 122 seconds. __cancel_work_timer+0x3f4/0x460 cancel_work_sync+0x1c/0x2c btmtksdio_flush+0x2c/0x40 hci_dev_open_sync+0x10c4/0x2190 [..] It all boils down to incorrect time_is_before_jiffies() usage in btmtksdio_txrx_work(). The btmtksdio_txrx_work() loop is expected to be terminated if running for longer than 5*HZ. However the timeout check is twisted: time_is_before_jiffies(old_jiffies + 5*HZ) evaluates to true when old_jiffies + 5*HZ is in the past i.e. when a timeout has occurred. Using OR with time_is_before_jiffies(txrx_timeout) means that: - before the 5-second timeout: the condition is `int_status || false`, so it loops as long as there are pending interrupts. - after the 5-second timeout: the condition becomes `int_status || true`, which is always true. When the loop becomes infinite btmtksdio_txrx_work() loop never terminates and never releases the SDIO host. Fix loop termination condition to actually enforce a 5*HZ timeout.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mm/swap: add cond_resched() in swap_reclaim_full_clusters to prevent softlockup We hit a real softlockup in an internal stress test environment. The workload was LTP memory/swap stress on a large arm64 machine, with 320 CPUs, about 1TB memory and an 8.6GB swap device. The system was under heavy load and the swap device had a large number of full clusters. The		

CVE-2026-64415	softlockup was triggered during a stress test after about 3 days. So, add periodic cond_resched() calls during large full_clusters reclaim operations to prevent softlockup issues. Detailed call trace as follow: PID: 3817773 TASK: ffff0883bb28b780 CPU: 48 COMMAND: "kworker/48:7" #0 [ffff800080183d10] __crash_kexec at ffffa4c1361e5de4 #1 [ffff800080183d90] panic at ffffa4c1360d5e9c #2 [ffff800080183e20] watchdog_timer_fn at ffffa4c136231fa8 ... #16 [ffff8000c4ad3cb0] swap_cache_del_folio at ffffa4c1363e1614 #17 [ffff8000c4ad3ce0] __try_to_reclaim_swap at ffffa4c1363e4bfc #18 [ffff8000c4ad3d40] swap_reclaim_full_clusters at ffffa4c1363e5474 #19 [ffff8000c4ad3da0] swap_reclaim_work at ffffa4c1363e550c #20 [ffff8000c4ad3dc0] process_one_work at ffffa4c136102edc #21 [ffff8000c4ad3e10] worker_thread at ffffa4c136103398 #22 [ffff8000c4ad3e70] kthread at ffffa4c13610d95c	N/A	More Details
CVE-2026-64416	In the Linux kernel, the following vulnerability has been resolved: mm: swap_cgroup: fix NULL deref in lookup_swap_cgroup_id on swapless host lookup_swap_cgroup_id() passes swap_cgroup_ctrl[type].map to __swap_cgroup_id_lookup() without checking that the type was ever registered via swap_cgroup_swapon(). On a swapless host every ctrl->map is NULL, so __swap_cgroup_id_lookup() dereferences NULL + a scaled swp_offset(). Since commit bea67dcc5eea ("mm: attempt to batch free swap entries for zap_pte_range()"), zap_pte_range() -> swap_pte_batch() calls lookup_swap_cgroup_id() on any non-present, non-none PTE that decodes as a real swap entry, without first validating it against swap_info[]. A single PTE corrupted into a type-0 swap entry takes the host down at process exit. We hit this in production on a swapless 6.12.58 host: ~1s of "get_swap_device: Bad swap file entry 3f800204222bb" (do_swap_page() being correctly defensive about the same entry) followed by BUG: unable to handle page fault for address: 000003f800204220 RIP: 0010:lookup_swap_cgroup_id+0x2b/0x60 Call Trace: swap_pte_batch+0xbf/0x230 zap_pte_range+0x4c8/0x780 unmap_page_range+0x190/0x3e0 exit_mmap+0xd9/0x3c0 do_exit+0x20c/0x4b0 syzbot has reported the identical stack. The source of the PTE corruption is a separate bug; this change makes the teardown path as robust as the fault path already is. Every other caller of lookup_swap_cgroup_id() is downstream of a get_swap_device() that has already validated the entry, so the new branch is cold.	N/A	More Details
CVE-2026-64417	In the Linux kernel, the following vulnerability has been resolved: mm: shrinker: fix NULL pointer dereference in debugfs shrinker_debugfs_add() creates both "count" and "scan" debugfs files unconditionally. That assumes every shrinker implements both count_objects() and scan_objects(), which is not guaranteed. For example, the xen-backend shrinker sets count_objects() but leaves scan_objects() NULL, so writing to its scan file calls through a NULL function pointer and panics the kernel: BUG: kernel NULL pointer dereference, address: 0000000000000000 RIP: 0010:0x0 Code: Unable to access opcode bytes at 0xffffffffffffffd6. Call Trace: <TASK> shrinker_debugfs_scan_write+0x12e/0x270 full_proxy_write+0x5f/0x90 vfs_write+0xde/0x420 ? filp_flush+0x75/0x90 ? filp_close+0x1d/0x30 ? do_dup2+0xb8/0x120 ksys_write+0x68/0xf0 ? filp_flush+0x75/0x90 do_syscall_64+0xb3/0x5b0 entry_SYSCALL_64_after_hwframe+0x76/0x7e The count path has the same issue in principle if a shrinker omits count_objects(). To fix it, only create "count" and "scan" debugfs files when the corresponding callbacks are present.	N/A	More Details
CVE-2026-64419	In the Linux kernel, the following vulnerability has been resolved: mm/shrinker: do not hold RCU lock in shrinker_debugfs_count_show() Reading the debugfs "count" file of a memcg-aware shrinker can sleep inside an RCU read-side critical section: BUG: sleeping function called from invalid context at kernel/cgroup/rstat.c:421 RCU nest depth: 1, expected: 0 css_rstat_flush mem_cgroup_flush_stats zswap_shrinker_count shrinker_debugfs_count_show shrinker_debugfs_count_show() invokes the ->count_objects() callback under rcu_read_lock(). The zswap callback flushes memcg stats via css_rstat_flush(), which may sleep, so it must not run under RCU. The RCU lock is not needed here. mem_cgroup_iter() takes RCU internally and returns a memcg holding a css reference (dropped on the next iteration or by mem_cgroup_iter_break()), so the memcg stays alive without it. The shrinker is kept alive by the open debugfs file: shrinker_free() removes the debugfs entries via debugfs_remove_recursive(), which waits for in-flight readers to drain, before call_rcu(..., shrinker_free_rcu_cb). The sibling "scan" handler already invokes the sleeping ->scan_objects() callback with no RCU section. Drop the rcu_read_lock()/rcu_read_unlock().	N/A	More Details
CVE-2026-64421	In the Linux kernel, the following vulnerability has been resolved: media: nxp: imx8-isi: Fix use-after-free on remove KASAN reports a slab-use-after-free in __media_entity_remove_link() during rmmod of imx8_isi: BUG: KASAN: slab-use-after-free in __media_entity_remove_link+0x608/0x650 Read of size 2 at addr ffff0000d47cb02a by task rmmod/724 Call trace: __media_entity_remove_link+0x608/0x650 __media_entity_remove_links+0x78/0x144 __media_device_unregister_entity+0x150/0x280 media_device_unregister_entity+0x48/0x68 v4l2_device_unregister_subdev+0x158/0x300 v4l2_async_unbind_subdev_one+0x22c/0x358 v4l2_async_nf_unbind_all_subdevs+0xfc/0x1c0 v4l2_async_nf_unregister+0x5c/0x14c mxc_isi_remove+0x124/0x2a0 [imx8_isi] Allocated by task 249: __kmalloc_noprof+0x27c/0x690 mxc_isi_crossbar_init+0x22c/0x560 [imx8_isi] Freed by task 724: kfree+0x1e4/0x5b0 mxc_isi_crossbar_cleanup+0x34/0x80 [imx8_isi] mxc_isi_remove+0x11c/0x2a0 [imx8_isi] The problem is that mxc_isi_remove() calls mxc_isi_crossbar_cleanup() before mxc_isi_v4l2_cleanup(). The crossbar cleanup frees the media entity pads, but the subsequent v4l2 cleanup still tries to remove media links that reference those pads. Fix this by calling mxc_isi_v4l2_cleanup() before mxc_isi_crossbar_cleanup() to ensure all media entities are properly unregistered while the pads are still valid.	N/A	More Details
CVE-2026-64424	In the Linux kernel, the following vulnerability has been resolved: netpoll: fix a use-after-free on shutdown path There is a use-after-free error on netpoll, which is clearly detected by KASAN. BUG: KASAN: slab-use-after-free in _raw_spin_lock_irqsave+0x3b/0x80 Read of size 1 at addr ... by task kworker/9:1 Workqueue: events queue_process Call Trace: skb_dequeue+0x1e/0xb0 queue_process+0x2c/0x600 process_scheduled_works+0x4b6/0x850 worker_thread+0x414/0x5a0 Allocated by task 242: __netpoll_setup+0x201/0x4a0 netpoll_setup+0x249/0x550 enabled_store+0x32f/0x380 Freed by task 0: kfree+0x1b7/0x540 rcu_core+0x3f8/0x7a0 The problem happens when there is a pending TX worker running in parallel with the cleanup path. This is what happens on netpoll shutdown path: 1) __netpoll_cleanup() is called 2) set dev->npinfo to NULL 3) call_rcu() with rcu_cleanup_netpoll_info() 3.1) rcu_cleanup_netpoll_info() tries to cancel all workers with cancel_delayed_work(), but doesn't wait for the worker to finish 4) and kfree(npinfo); Because 3.1) doesn't really cancel the work, as the comment says "we can't call cancel_delayed_work_sync here, as we are in softirq", the TX worker can run after 4). TI;DR: queue_process() is not an	N/A	More Details

	RCU reader, it reaches npinfo through the work item via container_of(). Use disable_delayed_work_sync() to ensure the worker is completely stopped and prevent any future re-arming attempts. Once npinfo is set to NULL, senders will bail out and not queue new work. The disable flag ensures any in-flight re-arming attempts also fail silently. In the future, we can do the cleanup inline here without needing the npinfo->rcu rcu_head, but that is net-next material.		
CVE-2026-64425	In the Linux kernel, the following vulnerability has been resolved: io_uring/io-wq: re-check IO_WQ_BIT_EXIT for each linked work item commit 10dc95939817 ("io_uring/io-wq: check IO_WQ_BIT_EXIT inside work run loop") fixed the obvious case where io_worker_handle_work() took one exit-bit snapshot before draining pending work, but the fix stops one level too early. io_worker_handle_work() now re-checks IO_WQ_BIT_EXIT in its outer work run loop, yet it still snapshots that bit once before processing a whole dependent linked-work chain. If io_wq_exit_start() sets IO_WQ_BIT_EXIT after the first linked item has started, the remaining linked items can still reuse stale do_kill = false, skip IO_WQ_WORK_CANCEL, and continue running after exit has begun. Move the check further inside, so it covers linked items too. Note: this is a syzbot special as it loves setting up tons of slow linked work on weird devices like msr that take forever to read, and immediately close the ring. Exit then takes a long time.	N/A	More Details
CVE-2026-64426	In the Linux kernel, the following vulnerability has been resolved: io_uring/nop: fix file reference leak with IOSQE_FIXED_FILE NOP file-acquisition support choses between a fixed (registered) file and a normal fget()'d file based on its own IORING_NOP_FIXED_FILE flag in sqe->nop_flags. However, a request's REQ_F_FIXED_FILE is set independently from the generic IOSQE_FIXED_FILE sqe flag during request init, before the issue handler runs. If a NOP is submitted with IOSQE_FIXED_FILE set (so REQ_F_FIXED_FILE is set) but without IORING_NOP_FIXED_FILE, io_nop() takes the normal path and grabs a real reference via io_file_get_normal(). On completion, io_put_file() only drops the reference when REQ_F_FIXED_FILE is clear, so the fget()'d file is never released and leaks: BUG: memory leak unreferenced object 0xffff88800f42c240 (size 176): kmem_cache_alloc_noprof+0x358/0x440 alloc_empty_file+0x57/0x180 path_openat+0x44/0x1e50 do_file_open+0x121/0x200 do_sys_openat2+0xa7/0x150 __x64_sys_openat+0x82/0xf0 Decide between fixed and normal file acquisition from REQ_F_FIXED_FILE, the same way io_assign_file() does for every other opcode, and fold IORING_NOP_FIXED_FILE into REQ_F_FIXED_FILE at prep time.	N/A	More Details
CVE-2026-64427	In the Linux kernel, the following vulnerability has been resolved: HID: logitech-dj: Fix maxfield check in DJ short report validation Commit b6a57912854e ("HID: logitech-dj: Prevent REPORT_ID_DJ_SHORT related user initiated OOB write") added validation for the DJ short output report, but the error path dereferences rep->field[0] even when rep->maxfield is zero. Commit 8b9a097eb2fc ("HID: logitech-dj: fix wrong detection of bad DJ_SHORT output report") made the check conditional on rep being present, but a crafted descriptor can still create report ID 0x20 with only padding output items. hid-core registers the report, ignores the padding field, and leaves rep->maxfield as zero. In that case the validation enters the rep->maxfield < 1 branch and then dereferences rep->field[0]->report_count while printing the error message, causing a NULL pointer dereference during probe. This is reproducible with uhid by emulating a Logitech receiver with a padding-only DJ short output report: BUG: KASAN: null-ptr-deref in logi_dj_probe+0xb1/0x754 [hid_logitech_dj] Read of size 4 at addr 0000000000000028 by task kworker/4:1/129 ... Call Trace: logi_dj_probe+0xb1/0x754 [hid_logitech_dj] hid_device_probe+0x329/0x3f0 [hid] really_probe+0x162/0x570 __device_attach+0x137/0x2c0 bus_probe_device+0x38/0xc0 device_add+0xa56/0xce0 hid_add_device+0x19c/0x280 [hid] uhid_device_add_worker+0x2c/0xb0 [uhid] Reject the zero-field report before printing the field report_count.	N/A	More Details
CVE-2026-64428	In the Linux kernel, the following vulnerability has been resolved: gpio: sch: use raw_spinlock_t in the irq startup path sch_irq_unmask() enables the GPIO IRQ and then updates the controller state through sch_irq_mask_unmask(), which takes sch->lock with spin_lock_irqsave(). The callback can be reached from irq_startup() while setting up a requested IRQ. That path is not sleepable, but on PREEMPT_RT a regular spinlock_t becomes a sleeping lock. This issue was found by our static analysis tool and then manually reviewed against the current tree. The grounded PoC kept the request_threaded_irq() -> __setup_irq() -> irq_startup() -> sch_irq_unmask() -> sch_irq_mask_unmask() carrier and used the original spin_lock_irqsave(&sch->lock) edge. Lockdep reported: BUG: sleeping function called from invalid context hardirqs last disabled at ... __setup_irq.constprop.0 ... [vuln_msv] sch_rt_spin_lock_irqsave+0x1c/0x30 [vuln_msv] sch_irq_mask_unmask.constprop.0+0x31/0x70 [vuln_msv] __setup_irq.constprop.0+0xd/0x30 [vuln_msv] Convert the SCH controller lock to raw_spinlock_t. The same lock is also used by the GPIO direction and value callbacks, but those critical sections only update MMIO-backed GPIO registers and do not contain sleepable operations. Keeping this register lock non-sleeping is therefore appropriate for the irqchip callbacks and does not change the GPIO-side locking contract.	N/A	More Details
CVE-2026-64537	In the Linux kernel, the following vulnerability has been resolved: bridge: cfm: reject invalid CCM interval at configuration time ccm_tx_work_expired() re-arms itself via queue_delayed_work() using the configured exp_interval converted by interval_to_us(). When exp_interval is BR_CFM_CCM_INTERVAL_NONE or out of range, interval_to_us() returns 0, causing the worker to fire immediately in a tight loop that allocates skbs until OOM. Fix this by validating exp_interval at configuration time: - Constrain IFLA_BRIDGE_CFM_CC_CONFIG_EXP_INTERVAL to the valid range [BR_CFM_CCM_INTERVAL_3_3_MS, BR_CFM_CCM_INTERVAL_10_MIN] in the netlink policy so userspace cannot set an invalid value. - Reject starting CCM TX in br_cfm_cc_ccm_tx() when exp_interval has not yet been configured (defaults to 0 from kzalloc).	N/A	More Details
CVE-2026-64539	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: eir: Fix stack OOB write when prepending the Flags AD eir_create_adv_data() builds the advertising data into a fixed-size buffer ("size", 31 for the legacy path). It may prepend a 3-byte "Flags" AD structure (LE_AD_NO_BREDR on an LE-only controller) and then copies the per-instance data without checking that it still fits: memcpy(ptr, adv->adv_data, adv->adv_data_len); tlv_data_max_len() only reserves those 3 bytes when the user-supplied flags carry a managed-flags bit, so an instance added with flags == 0 is accepted with adv_data_len up to the full buffer. At advertise time the flags are still prepended, and the memcpy() writes 3 + adv_data_len bytes into the size-byte buffer: BUG: KASAN: stack-out-of-bounds in eir_create_adv_data (net/bluetooth/eir.c:301) Write of size 31 at addr ffff88800a547bdc by task kworker/u9:0/65 Workqueue: hci0 hci_cmd_sync_work __asan_memcpy (mm/kasan/shadow.c:106) eir_create_adv_data	N/A	More Details

	(net/bluetooth/eir.c:301) hci_update_adv_data_sync (net/bluetooth/hci_sync.c:1310) hci_schedule_adv_instance_sync (net/bluetooth/hci_sync.c:1817) hci_cmd_sync_work (net/bluetooth/hci_sync.c:332) This frame has 1 object: [32, 64) 'cp' The "Flags" structure is added by the kernel, not requested by userspace, so only prepend it when it fits together with the instance advertising data; when there is no room for both, drop the flags rather than the user-provided data. Reachable by a local user with CAP_NET_ADMIN owning an LE-only controller on the legacy advertising path.		
CVE-2026-64347	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: composite: fix dead empty check in the USB_DT_OTG handler The OTG branch of composite_setup() falls back to the first configuration when none is selected: if (cdev->config) config = cdev->config; else config = list_first_entry(&cdev->configs, struct usb_configuration, list); if (!config) goto done; ... memcpy(req->buf, config->descriptors[0], value); list_first_entry() never returns NULL. On an empty list it returns container_of() of the list head. So the "if (!config)" check is dead. When cdev->configs is empty, config points at the head inside struct usb_composite_dev. config->descriptors[0] reads whatever sits at that offset. The memcpy copies up to w_length bytes of it into the response buffer. cdev->configs can be empty in two cases. One is a teardown race on gadget unbind with a control transfer in flight. The other is a driver that sets is_otg before it adds a config. A reproducer that holds cdev->configs empty triggers a KASAN fault in this branch. Use list_first_entry_or_null() so the existing check does its job.	N/A	More Details
CVE-2026-64540	In the Linux kernel, the following vulnerability has been resolved: usbnet: gl620a: fix out-of-bounds read in genelink_rx_fixup() genelink_rx_fixup() splits an aggregated RX frame into its individual packets, using a per-packet length taken from device-supplied data. That length is only bounded by GL_MAX_PACKET_LEN (1514); it is never compared against how many bytes were actually received. A malicious GeneLink (GL620A) device can therefore send a short URB whose header claims packet_count > 1 and a first packet of up to 1514 bytes. skb_put_data(gl_skb, packet->packet_data, size); then copies past the end of the receive buffer and hands the adjacent slab contents up the network stack, an out-of-bounds read that leaks kernel heap. No privilege is required: the path runs in the usbnet RX softirq as soon as the interface is up. BUG: KASAN: slab-out-of-bounds in genelink_rx_fixup (drivers/net/usb/gl620a.c:112) Read of size 1514 at addr ffff888011309708 by task ksoftirqd/0/14 Call Trace: ... __asan_memcpy (mm/kasan/shadow.c:105) genelink_rx_fixup (include/linux/skbuff.h:2814 drivers/net/usb/gl620a.c:112) usbnet_bh (drivers/net/usb/usbnet.c:572 drivers/net/usb/usbnet.c:1589) process_one_work (kernel/workqueue.c:3322) bh_worker (kernel/workqueue.c:3405) tasklet_action (kernel/softirq.c:965) handle_softirqs (kernel/softirq.c:622) run_ksoftirqd (kernel/softirq.c:1076) ... skb_pull() already verifies that the requested length fits the buffer and returns NULL otherwise. Move it ahead of the copy and check its result, so a packet that overruns the received data is rejected before it is read. Well-formed frames, whose packets are fully present, are unaffected.	N/A	More Details
CVE-2026-64348	In the Linux kernel, the following vulnerability has been resolved: usb: free iso schedules on failed submit EHCI and FOTG210 isochronous submits build an ehci_iso_sched before linking the URB to the endpoint queue, and keep the staged schedule in urb->hcpriv until iso_stream_schedule() and the link helpers consume it. If the controller is no longer accessible, or usb_hcd_link_urb_to_ep() fails, submit jumps to done_not_linked before that handoff happens and leaks the staged schedule still attached to urb->hcpriv. Free the staged schedule from done_not_linked when submit fails before the URB is linked and clear urb->hcpriv after the free. The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1.1. An x86_64 allyesconfig build showed no new warnings. As we do not have an EHCI host controller with a USB isochronous device to test with, no runtime testing was able to be performed.	N/A	More Details
CVE-2026-64349	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: fix dwc3_readl() and dwc3_writel() calls in dwc3_ulpi_setup() The dwc3_ulpi_setup() calls the register read and write calls with dwc3->regs when both these calls take the dwc3 structure directly. Chnage these two calls to fix the following sparse warning, and possibly a nasty bug in the dwc3_ulpi_setup() code: drivers/usb/dwc3/core.c:796:45: warning: incorrect type in argument 1 (different address spaces) drivers/usb/dwc3/core.c:796:45: expected struct dwc3 *dwc drivers/usb/dwc3/core.c:796:45: got void [noderef] __iomem *regs drivers/usb/dwc3/core.c:798:40: warning: incorrect type in argument 1 (different address spaces) drivers/usb/dwc3/core.c:798:40: expected struct dwc3 *dwc drivers/usb/dwc3/core.c:798:40: got void [noderef] __iomem *regs	N/A	More Details
CVE-2026-64350	In the Linux kernel, the following vulnerability has been resolved: usb: cdnsp: fix stream context array leak in cdnsp_alloc_stream_info() cdnsp_alloc_stream_info() allocates stream_info->stream_ctx_array with cdnsp_alloc_stream_ctx(). If a later stream ring allocation or stream mapping update fails, the error path frees the allocated stream rings and stream_rings array, but leaves stream_ctx_array allocated. Free the stream context array before falling through to the stream_rings cleanup path.	N/A	More Details
CVE-2026-64352	In the Linux kernel, the following vulnerability has been resolved: bpf: Allow LPM map access from sleepable BPF programs trie_lookup_elem() annotates its rcu_dereference_check() walks with only rcu_read_lock_bh_held(). Because rcu_dereference_check(p, c) resolves to "c || rcu_read_lock_held()", this passes for XDP/NAPI and classic RCU readers but fails for sleepable BPF programs, which enter via __bpf_prog_enter_sleepable() and hold only rcu_read_lock_trace(). trie_update_elem() and trie_delete_elem() have the same problem in a different form: they walk the trie with plain rcu_dereference(), which asserts rcu_read_lock_held() unconditionally. Both are reachable from sleepable BPF programs via the bpf_map_update_elem / bpf_map_delete_elem helpers, and from the syscall path under classic rcu_read_lock(). In the writer paths the trie is actually protected by trie->lock (an rqspinlock taken across the walk); we never relied on the RCU read-side lock to keep nodes alive there. A sleepable LSM hook that ends up touching an LPM trie therefore triggers lockdep on debug kernels: ============================= WARNING: suspicious RCU usage 7.1.0-... Tainted: G E ----------------------------- kernel/bpf/lpm_trie.c:249 suspicious rcu_dereference_check() usage! 1 lock held by net_tests/540: #0: (rcu_tasks_trace_srcu_struct){....}-{0:0}, at: __bpf_prog_enter_sleepable+0x26/0x280 Call Trace: dump_stack_lvl lockdep_rcu_suspicious trie_lookup_elem bpf_prog_..._enforce_security_socket_connect bpf_trampoline_...	N/A	More Details

	security_socket_connect __sys_connect do_syscall_64 This is lockdep-only -- no UAF, since Tasks Trace RCU does serialize against the trie's reclaim path -- but it spams the console once per distinct callsite on every debug kernel running a sleepable BPF LSM that touches an LPM trie, which is increasingly common. For the lookup path, switch the rcu_dereference_check() annotation from rcu_read_lock_bh_held() to bpf_rcu_lock_held(), which accepts all three contexts (classic, BH, Tasks Trace). Other map types already follow this convention. For trie_update_elem() and trie_delete_elem(), annotate the walks as rcu_dereference_protected(*p, 1) -- matching trie_free() in the same file -- since trie->lock is held across the walk. rqspinlock has no lockdep_map, so the predicate degenerates to '1' rather than lockdep_is_held(&trie->lock); the protection is real but not machine-verifiable. trie_get_next_key() also uses bare rcu_dereference() but is reachable only from the BPF syscall, which holds classic rcu_read_lock() before dispatching, so it is left untouched.		
CVE-2026-64353	In the Linux kernel, the following vulnerability has been resolved: bpf: Keep dynamic inner array lookups nullable An ARRAY_OF_MAPS can use an array created with BPF_F_INNER_MAP as its inner map template. A concrete inner array with a different max_entries value can then replace the template. After a successful outer map lookup, the verifier represents the resulting map pointer using the inner map template. Const-key lookup nullness elision consequently uses the template max_entries even though the runtime helper uses the concrete inner map max_entries. Do not elide lookup result nullness for maps marked with BPF_F_INNER_MAP, because the template max_entries does not prove that the key is in bounds for the concrete runtime map.	N/A	More Details
CVE-2026-64356	In the Linux kernel, the following vulnerability has been resolved: xfs: fix memory leak in xfs_dqinode_metadir_create() If xfs_metadir_create() fails in xfs_dqinode_metadir_create(), the current code returns directly, leaking the allocated update and transaction state. If the subsequent commit fails, the caller-owned inode reference is left behind. Fix this memory leak by routing the create failure path through xfs_metadir_cancel(). For both create and commit failures, finish and release any inode returned to the caller, mirroring the unwind pattern in xfs_metadir_mkdir(). The bug was first flagged by an experimental analysis tool we are developing for kernel memory-management bugs while analyzing v6.13-rc1. The tool is still under development and is not yet publicly available. Manual inspection confirms that the bug is still present in v7.1.1. An x86_64 allyesconfig build showed no new warnings. Runtime validation used kprobe fault injection during `mount -o uquota` on a metadir XFS image. Injecting xfs_metadir_create() reproduced the old active-update path that left mount stuck later in mount setup; after this change, the same injection reported cancel_hits=1 and irele_hits=1. Injecting xfs_metadir_commit() exercised the old inode-reference leak path; after this change, it reported irele_hits=1.	N/A	More Details
CVE-2026-64357	In the Linux kernel, the following vulnerability has been resolved: xfs: fix exchmaps reservation limit check xfs_exchmaps_estimate_overhead() adds the bmbt and rmapbt overhead to a local resblks variable, but the final UINT_MAX check still tests req->resblks. That is the reservation value from before the overhead was added. The computed value is stored back in req->resblks and later passed to xfs_trans_alloc(), whose block reservation argument is unsigned int. Check the computed reservation so the existing limit applies to the value that will be used.	N/A	More Details
CVE-2026-64358	In the Linux kernel, the following vulnerability has been resolved: media: mtk-jpeg: cancel workqueue on release for supported platforms only Since a recent fix the mtk_jpeg_release function cancels any pending or running work present in the driver workqueue using cancel_work_sync function. Currently, only the multicore based variants use this workqueue and they have the jpeg_worker platform data field initialized with a workqueue callback function. For the others, this field value remain NULL by default. The cancel_work_sync function is unconditionally called in mtk_jpeg_release function, even for the variants that do not use the workqueue. This call generates a WARN_ON print in __flush_work because the workqueue callback function presence check fails in __flush_work function (used by cancel_work_sync). So, to avoid these warnings, call cancel_work_sync only if a workqueue callback is defined in platform data.	N/A	More Details
CVE-2026-64359	In the Linux kernel, the following vulnerability has been resolved: nilfs2: reject CLEAN_SEGMENTS ioctl with out-of-range segment numbers Syzbot reported a hung task in nilfs_transaction_begin() where multiple tasks performing chmod() on a nilfs2 mount blocked for over 143 seconds waiting to acquire ns_segctor_sem for read: INFO: task syz.0.17:5918 blocked for more than 143 seconds. Call Trace: schedule+0x164/0x360 rwsem_down_read_slowpath+0x6d9/0x940 down_read+0x99/0x2e0 nilfs_transaction_begin+0x364/0x710 fs/nilfs2/segment.c:221 nilfs_setattr+0x124/0x2c0 fs/nilfs2/inode.c:921 notify_change+0xc1a/0xf40 chmod_common+0x273/0x4a0 do_fchmodat+0x12d/0x230 The writer holding ns_segctor_sem was a concurrent NILFS_IOCTL_CLEAN_SEGMENTS caller, stuck inside printk while emitting per-element warnings from nilfs_sufile_updatev(): __nilfs_msg+0x373/0x450 fs/nilfs2/super.c:78 nilfs_sufile_updatev+0x21c/0x6d0 fs/nilfs2/sufile.c:186 nilfs_sufile_freev fs/nilfs2/sufile.h:93 [inline] nilfs_free_segments fs/nilfs2/segment.c:1140 [inline] nilfs_segctor_collect_blocks fs/nilfs2/segment.c:1261 [inline] nilfs_segctor_do_construct+0x1f55/0x76c0 nilfs_clean_segments+0x3bd/0xa50 nilfs_ioctl_clean_segments fs/nilfs2/ioctl.c:922 [inline] nilfs_ioctl+0x261f/0x2780 The root cause is that user-supplied segment numbers are not validated before nilfs_clean_segments() begins doing work; the range check on each segnum is performed deep inside the call chain by nilfs_sufile_updatev(), which emits a nilfs_warn() per invalid entry while still holding the segctor lock and the sufile mi_sem. Under load (repeated invocations across multiple mounts saturating the global printk path), the cumulative printk latency keeps ns_segctor_sem held long enough to trip the hung_task watchdog, blocking concurrent operations such as chmod() that need ns_segctor_sem for read. Fix by validating the contents of kbufs[4] in nilfs_clean_segments() immediately after acquiring ns_segctor_sem via nilfs_transaction_lock(). Holding ns_segctor_sem serializes the check against nilfs_ioctl_resize(), which can modify ns_nsegments, so the validation uses a consistent value. Out-of-range segment numbers are rejected with -EINVAL before any segment-cleaning work begins, so the bad entries never reach the per-element diagnostic path inside nilfs_sufile_updatev().	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: hfs/hfsplus: zero-initialize buffer in hfs_bnode_read hfs_bnode_read() can return early without writing to the output buffer when is_bnode_offset_valid() fails or when check_and_correct_requested_ length() corrects the length to zero. Callers such as hfs_bnode_read_ u16() and hfs_bnode_read_u8() pass stack-allocated buffers and use the result unconditionally, leading to KMSAN uninit-value	N/A	More

64360	reports. Rather than initializing at each individual call site, zero the buffer at the start of hfs_bnode_read() before any validation checks. This ensures all callers in both hfs and hfsplus get a deterministic zero value regardless of which early-return path is taken.		Details
CVE-2026-64555	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: nv: Fix SPSR_EL2 restore in kvm_hyp_handle_mops() kvm_hyp_handle_mops() resets the single-step state machine as part of rewinding state for a MOPS exception by modifying vcpu_cpsr() and writing the result directly into hardware. In the case of nested virtualization, vcpu_cpsr() is a synthetic value such that the rest of KVM can deal with vEL2 cleanly. That means the value requires translation before being written into hardware, which is unfortunately missing from the MOPS handler. Fix it by directly modifying SPSR_EL2 and avoiding the synthetic state altogether, which will be resynchronized on the next 'full' exit back to KVM.	N/A	More Details
CVE-2026-64554	In the Linux kernel, the following vulnerability has been resolved: netfilter: bridge: fix stale prevhdr pointer in br_ip6_fragment() br_ip6_fragment() gets prevhdr, a pointer into the skb head, from ip6_find_1stfragopt(), then calls skb_checksum_help(). For a cloned skb skb_checksum_help() reallocates the head via pskb_expand_head(), leaving prevhdr dangling. It is later dereferenced in ip6_frag_next(), causing a use-after-free write. Save prevhdr's offset before skb_checksum_help() and recompute it after, like commit ef0efcd3bd3f ("ipv6: Fix dangling pointer when ipv6 fragment"). BUG: KASAN: slab-use-after-free in ip6_frag_next (net/ipv6/ip6_output.c:857) Write of size 1 at addr ffff888013ff5016 by task exploit/141 Call Trace: ... kasan_report (mm/kasan/report.c:595) ip6_frag_next (net/ipv6/ip6_output.c:857) br_ip6_fragment (net/ipv6/netfilter.c:212) nf_ct_bridge_post (net/bridge/netfilter/nf_conntrack_bridge.c:407) nf_hook_slow (net/netfilter/core.c:619) br_forward_finish (net/bridge/br_forward.c:66) __br_forward (net/bridge/br_forward.c:115) maybe_deliver (net/bridge/br_forward.c:191) br_flood (net/bridge/br_forward.c:245) br_handle_frame_finish (net/bridge/br_input.c:229) br_handle_frame (net/bridge/br_input.c:442) ... packet_sendmsg (net/packet/af_packet.c:3114) ... do_syscall_64 (arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121) Kernel panic - not syncing: Fatal exception in interrupt	N/A	More Details
CVE-2026-64553	In the Linux kernel, the following vulnerability has been resolved: net: psample: fix info leak in PSAMPLE_ATTR_DATA psample open codes nla_put() presumably to avoid wiping the data with 0s just to override it with packet data. This open coding is missing clearing the pad, however, each netlink attr is padded to 4B and data_len may not be divisible by 4B.	N/A	More Details
CVE-2026-64552	In the Linux kernel, the following vulnerability has been resolved: virtio-net: fix len check in receive_big() receive_big() bounds the device-announced length by (big_packets_num_skbfrags + 1) * PAGE_SIZE. That is still too loose: add_recvbuf_big() sets sg[1] to start at offset sizeof(struct padded_vnet_hdr) into the first page, so the chain actually carries hdr_len + (PAGE_SIZE - sizeof(padded_vnet_hdr)) + big_packets_num_skbfrags * PAGE_SIZE bytes -- 20 bytes less than the check allows for the common hdr_len == 12 case. A malicious virtio backend can announce a len in that gap. page_to_skb() then walks one frag past the page chain, storing a NULL page->private into skb_shinfo()->frags[MAX_SKB_FRAGS], which is both an out-of-bounds write past the static frag array and a NULL frag handed up the rx path. Bound len by the size add_recvbuf_big() actually advertised.	N/A	More Details
CVE-2026-64551	In the Linux kernel, the following vulnerability has been resolved: sctp: validate STALE_COOKIE cause length before reading staleness When an ERROR chunk with a STALE_COOKIE cause is received in the COOKIE_ECHOED state, sctp_sf_do_5_2_6_stale() reads the 4-byte Measure of Staleness that follows the cause header: err = (struct sctp_errhdr *)(chunk->skb->data); stale = ntohl(*(__be32 *)((u8 *)err + sizeof(*err))); err is the first cause in the chunk, not the STALE_COOKIE cause that caused the dispatch, and nothing guarantees the staleness field is present. sctp_walk_errors() only requires a cause to be as long as the 4-byte header, so for a STALE_COOKIE cause of length 4 the read runs past the cause, and for a minimal ERROR chunk past skb->tail. The value is echoed to the peer in the Cookie Preservative of the reply INIT, leaking uninitialized memory. sctp_sf_cookie_echoed_err() already walks to the STALE_COOKIE cause, so check its length there and pass it to sctp_sf_do_5_2_6_stale(), which reads that cause instead of the first one. A STALE_COOKIE cause too short to hold the staleness field is discarded. The read is reachable by any peer that can drive an association into COOKIE_ECHOED, including an unprivileged process using a raw SCTP socket in a user and network namespace.	N/A	More Details
CVE-2026-64550	In the Linux kernel, the following vulnerability has been resolved: net: qualcomm: rmnet: validate MAP frame length before ingress parsing When ingress deaggregation is disabled, rmnet_map_ingress_handler() passes the skb straight to __rmnet_map_ingress_handler(), skipping the length validation that rmnet_map_deaggregate() performs on the aggregated path. The parser then dereferences the MAP header and csum header/trailer based on the on-wire pkt_len without checking skb->len, so a short frame is read out of bounds: BUG: KASAN: slab-out-of-bounds in rmnet_map_checksum_downlink_packet Read of size 1 at addr ffff88801118ed00 by task exploit/147 Call Trace: ... rmnet_map_checksum_downlink_packet (drivers/net/ethernet/qualcomm/rmnet/rmnet_map_data.c:413) __rmnet_map_ingress_handler (drivers/net/ethernet/qualcomm/rmnet/rmnet_handlers.c:96) rmnet_rx_handler (drivers/net/ethernet/qualcomm/rmnet/rmnet_handlers.c:129) __netif_receive_skb_core.constprop.0 (net/core/dev.c:6089) netif_receive_skb (net/core/dev.c:6460) tun_get_user (drivers/net/tun.c:1955) tun_chr_write_iter (drivers/net/tun.c:2001) vfs_write (fs/read_write.c:688) ksys_write (fs/read_write.c:740) do_syscall_64 (arch/x86/entry/syscall_64.c:94) ... Factor that validation out of rmnet_map_deaggregate() into rmnet_map_validate_packet_len() and run it on the no-aggregation path too. The MAP header is bounds-checked first, since this path can receive a frame shorter than the header.	N/A	More Details
CVE-2026-	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: bpa10x: avoid OOB read of revision string in bpa10x_setup() bpa10x_setup() sends the vendor command 0xfc0e and passes the response to bt_dev_info() and hci_set_fw_info() as a "%s" string starting at skb->data + 1, without checking the length: bt_dev_info(hdev, "%s", (char *)(skb->data + 1)); hci_set_fw_info(hdev, "%s", skb->data + 1); A device that returns a one-byte response (status only) leaves skb->data + 1 past the end of the data, and the %s walk reads adjacent slab memory until it	N/A	More Details

64549	meets a NUL. The same happens when the payload is not NUL-terminated within skb->len. The out-of-bounds bytes end up in the kernel log and the firmware-info debugfs file. Print the revision string with a bounded "%.*s" limited to skb->len - 1 instead. This keeps the string readable for well-behaved devices while never reading past the received data, and does not fail setup, so a device returning a short or unterminated response keeps working.		
CVE-2026-64548	In the Linux kernel, the following vulnerability has been resolved: bpf, sockmap: reject overflowing copy + len in bpf_msg_push_data() When the scatterlist ring is full or nearly full, bpf_msg_push_data() enters a copy fallback path and computes copy + len for the page allocation size. Since len comes from BPF with arg3_type = ARG_ANYTHING and both are u32, a crafted len can wrap the sum to a small value, causing an undersized allocation followed by an out-of-bounds memcpy. BUG: unable to handle page fault for address: ffffed104089a402 Oops: Oops: 0000 [#1] SMP KASAN NOPTI Call Trace: __asan_memcpy (mm/kasan/shadow.c:105) bpf_msg_push_data (net/core/filter.c:2852 net/core/filter.c:2788) bpf_prog_9ed8b5711920a7d7+0x2e/0x36 sk_psock_msg_verdict (net/core/skmsg.c:934) tcp_bpf_sendmsg (net/ipv4/tcp_bpf.c:421 net/ipv4/tcp_bpf.c:584) __sys_sendto (net/socket.c:2206) do_syscall_64 (arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:130) Add an overflow check before the allocation.	N/A	More Details
CVE-2026-64547	In the Linux kernel, the following vulnerability has been resolved: net: usb: net1080: validate packet_len before pad-byte access in rx_fixup For an even packet_len, net1080_rx_fixup() reads the pad byte at skb->data[packet_len] before the skb->len != packet_len check further down, and packet_len is only bounded against NC_MAX_PACKET. A malicious NetChip 1080 device can send a short frame advertising a large even packet_len (e.g. 0x4000), so the pad-byte read lands past the end of the skb: BUG: KASAN: slab-out-of-bounds in net1080_rx_fixup Read of size 1 at addr ffff8880106c83c6 by task ksoftirqd/0/14 ... net1080_rx_fixup (drivers/net/usb/net1080.c:384) usbnet_bh (drivers/net/usb/usbnet.c:1589) process_one_work (kernel/workqueue.c:3322) bh_worker (kernel/workqueue.c:3708) tasklet_action (kernel/softirq.c:965) handle_softirqs (kernel/softirq.c:622) ... Reject the frame when packet_len >= skb->len before reading.	N/A	More Details
CVE-2026-64546	In the Linux kernel, the following vulnerability has been resolved: drm/edid: fix OOB read in drm_parse_tiled_block() drm_parse_tiled_block() casts the DisplayID block to a struct displayid_tiled_block and reads the full fixed layout up to tile->topology_id[7] without checking block->num_bytes. The DisplayID iterator only validates the declared payload length, so a crafted EDID can advertise a tiled-display block (tag DATA_BLOCK_TILED_DISPLAY, or DATA_BLOCK_2_TILED_DISPLAY_TOPOLOGY for v2.0) with a small num_bytes at the end of a DisplayID extension. The read then runs past the end of the exact-sized kmemdup()'d EDID allocation, a heap out-of-bounds read. Reject blocks shorter than the spec's 22-byte tiled payload before reading the fixed struct, as drm_parse_vesa_mso_data() already does. BUG: KASAN: slab-out-of-bounds in drm_edid_connector_update Read of size 2 at addr ffff888010077700 by task exploit/147 dump_stack_lvl (lib/dump_stack.c:94 ...) print_report (mm/kasan/report.c:378 ...) kasan_report (mm/kasan/report.c:595) drm_edid_connector_update (drivers/gpu/drm/drm_edid.c:7581) bochs_connector_helper_get_modes (drivers/gpu/drm/tiny/bochs.c:574) drm_helper_probe_single_connector_modes (drivers/gpu/drm/drm_probe_helper.c:426) status_store (drivers/gpu/drm/drm_sysfs.c:219) ... vfs_write (fs/read_write.c:595 fs/read_write.c:688) ksys_write (fs/read_write.c:740)	N/A	More Details
CVE-2026-64545	In the Linux kernel, the following vulnerability has been resolved: net, bpf: check master for NULL in xdp_master_redirect() xdp_master_redirect() dereferences the result of netdev_master_upper_dev_get_rcu() without a NULL check, but that helper returns NULL when the receiving device has no upper-master adjacency. The reach guard only checks netif_is_bond_slave(). On bond slave release bond_upper_dev_unlink() drops the upper-master adjacency before clearing IFF_SLAVE, so an XDP_TX reaching xdp_master_redirect() in that window still passes netif_is_bond_slave() while master is already NULL, and faults on master->flags at offset 0xb0: BUG: kernel NULL pointer dereference, address: 00000000000000b0 RIP: 0010:xdp_master_redirect (net/core/filter.c:4432) Call Trace: xdp_master_redirect (net/core/filter.c:4432) bpf_prog_run_generic_xdp (include/net/xdp.h:700) do_xdp_generic (net/core/dev.c:5608) __netif_receive_skb_one_core (net/core/dev.c:6204) process_backlog (net/core/dev.c:6319) __napi_poll (net/core/dev.c:7729) net_rx_action (net/core/dev.c:7792) handle_softirqs (kernel/softirq.c:622) __dev_queue_xmit (include/linux/bottom_half.h:33) packet_sendmsg (net/packet/af_packet.c:3082) __sys_sendto (net/socket.c:2252) Kernel panic - not syncing: Fatal exception in interrupt The missing check dates back to the original code; commit 1921f91298d1 ("net, bpf: fix null-ptr-deref in xdp_master_redirect() for down master") later added the master->flags read where the fault now lands but kept the unconditional deref. Check master for NULL before use; a NULL master is treated the same as one that is not up.	N/A	More Details
CVE-2026-64544	In the Linux kernel, the following vulnerability has been resolved: crypto: asymmetric_keys - fix OOB read in pefile_digest_pe_contents pefile_digest_pe_contents() computes the trailing-data hash length as pelen - (hashed_bytes + certs_size). A crafted PE can make the addition exceed pelen, causing the unsigned subtraction to underflow to ~4 GiB. This is passed to crypto_shash_update() which reads out of bounds and panics on unmapped vmalloc guard pages. BUG: unable to handle page fault for address: ffffc900038d8000 Oops: Oops: 0000 [#1] SMP KASAN NOPTI RIP: 0010:sha256_blocks_generic (lib/crypto/sha256.c:152) Call Trace: <TASK> __sha256_update (lib/crypto/sha256.c:208) crypto_sha256_update (crypto/sha256.c:142) verify_pefile_signature (crypto/asymmetric_keys/verify_pefile.c:436) kexec_kernel_verify_pe_sig (kernel/kexec_file.c:151) __do_sys_kexec_file_load (kernel/kexec_file.c:406) do_syscall_64 (arch/x86/entry/syscall_64.c:94) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:121) </TASK> Kernel panic - not syncing: Fatal exception Validate that the addition does not overflow and the result does not exceed pelen before the subtraction. Return -ELIBBAD on failure.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: tipc: fix use-after-free of the discoverer in tipc_disc_rcv() bearer_disable() frees b->disc with tipc_disc_delete()'s plain kfree(), but tipc_disc_rcv() still dereferences b->disc in RX softirq under rcu_read_lock() (tipc_udp_recv -> tipc_rcv -> tipc_disc_rcv). L2 bearers are safe thanks to the synchronize_net() in tipc_disable_l2_media(), but the UDP bearer defers that call to the cleanup_bearer() workqueue, so the discoverer is freed with no grace period: BUG: KASAN: slab-use-after-free in		

CVE-2026-64543	tipc_disc_rcv (net/tipc/discover.c:149) Read of size 8 at addr ffff88802348b728 by task poc_tipc/184 <IRQ> tipc_disc_rcv (net/tipc/discover.c:149) tipc_rcv (net/tipc/node.c:2126) tipc_udp_recv (net/tipc/udp_media.c:391) udp_rcv (net/ipv4/udp.c:2643) ip_local_deliver_finish (net/ipv4/ip_input.c:241) </IRQ> Freed by task 181: kfree (mm/slub.c:6565) bearer_disable (net/tipc/bearer.c:418) tipc_nl_bearer_disable (net/tipc/bearer.c:1001) The bearer is freed with kfree_rcu(); free the discoverer the same way. Add an rcu_head to struct tipc_discoverer and free it and its skb from an RCU callback. Because the RCU callback (tipc_disc_free_rcu) lives in module text, a call_rcu() that is still pending when the tipc module is unloaded would invoke a freed function. Add an rcu_barrier() to tipc_exit() after the bearer subsystem has been torn down, so all pending discoverer callbacks have run before the module text goes away. Reachable from an unprivileged user namespace: the TIPCv2 genl family is netnsok and its bearer commands have no GENL_ADMIN_PERM. Needs CONFIG_TIPC and CONFIG_TIPC_MEDIA_UDP.	N/A	More Details
CVE-2026-64542	In the Linux kernel, the following vulnerability has been resolved: ipv6: ndisc: fix NULL deref in accept_untracked_na() accept_untracked_na() re-fetches the inet6_dev with __in6_dev_get(dev) and dereferences idev->cnf.accept_untracked_na without a NULL check, even though its only caller ndisc_recv_na() already fetched and NULL-checked idev for the same device. Both reads of dev->ip6_ptr run in the same RCU read-side critical section, but a concurrent addrconf_ifdown() can clear dev->ip6_ptr between them: lowering the MTU below IPV6_MIN_MTU calls addrconf_ifdown() without the synchronize_net() that orders the unregister path, so the re-fetch returns NULL and oopses: BUG: KASAN: null-ptr-deref in ndisc_recv_na (net/ipv6/ndisc.c:974) Read of size 4 at addr 0000000000000364 Call Trace: <IRQ> ndisc_recv_na (net/ipv6/ndisc.c:974) icmpv6_rcv (net/ipv6/icmp.c:1193) ip6_protocol_deliver_rcu (net/ipv6/ip6_input.c:479) ip6_input_finish (net/ipv6/ip6_input.c:534) ip6_input (net/ipv6/ip6_input.c:545) ip6_mc_input (net/ipv6/ip6_input.c:635) ipv6_rcv (net/ipv6/ip6_input.c:351) </IRQ> It is reachable by an unprivileged user via a network namespace. Pass the caller's already validated idev instead of re-fetching it; the idev stays alive for the whole RCU critical section, so it is safe even after dev->ip6_ptr has been cleared.	N/A	More Details
CVE-2026-64541	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix UAF in smc_cdc_rx_handler() by pinning the socket smc_cdc_rx_handler() looks up the connection by token under the link group's conns_lock, drops the lock, and then dereferences conn and the smc_sock derived from it, ending in sock_hold(&smc->sk) inside smc_cdc_msg_recv(). No reference is held across the lock release. The only reference pinning the socket while the connection is discoverable in the link group is taken in smc_lgr_register_conn() (sock_hold) and dropped in __smc_lgr_unregister_conn() (sock_put), both under conns_lock. Once the handler drops conns_lock, a concurrent close() -> smc_release() -> smc_conn_free() -> smc_lgr_unregister_conn() can drop that reference and free the smc_sock, so the handler's later sock_hold() runs on freed memory: WARNING: lib/refcount.c:25 at refcount_warn_saturate Workqueue: rxe_wq do_work refcount_warn_saturate (lib/refcount.c:25) smc_cdc_msg_recv (net/smc/smc_cdc.c:430) smc_cdc_rx_handler (net/smc/smc_cdc.c:502) smc_wr_rx_tasklet_fn (net/smc/smc_wr.c:445) tasklet_action_common (kernel/softirq.c:938) handle_softirqs (kernel/softirq.c:622) Kernel panic - not syncing: panic_on_warn set Only SMC-R is affected. The SMC-D receive tasklet is stopped by tasklet_kill(&conn->rx_tsklet) in smc_conn_free() before the connection is unregistered, so it cannot run concurrently with the free. Take the socket reference while still holding conns_lock, so the registration reference can no longer be the last one, and drop it once the handler is done.	N/A	More Details
CVE-2026-64351	In the Linux kernel, the following vulnerability has been resolved: net: usb: kalmia: bound RX frame length in kalmia_rx_fixup() kalmia_rx_fixup() computes usb_packet_length = skb->len - (2 * KALMIA_HEADER_LENGTH) as a u16, guarded only by a pre-loop check that skb->len is at least KALMIA_HEADER_LENGTH, which is 6. A device can deliver a short bulk-IN frame with skb->len in the 6 to 11 range, or leave a short trailing remainder on a later loop iteration. Either case underflows usb_packet_length to about 65530. That bypasses the usb_packet_length < ether_packet_length truncation path. The device-supplied ether_packet_length, a le16 up to 65535 read from header_start[2], then drives a memcmp() and the following skb_trim() and skb_pull() past the end of the rx buffer. The rx buffer is hard_mtu * 10, which is 14000 bytes. That is an out of bounds read. Require both the start and end framing headers to be present before subtracting them, on every loop iteration.	N/A	More Details