SingCERT
Singapore Cyber Emergency Response Team

Security Bulletin 31 December 2025

Generated on 31 December 2025

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-2025-54322	Xspeeder SXZOS through 2025-12-26 allows root remote code execution via base64-encoded Python code in the chkid parameter to vLogin.py. The title and oIP parameters are also used.	10.0	More Details
CVE-2025-52691	Successful exploitation of the vulnerability could allow an unauthenticated attacker to upload arbitrary files to any location on the mail server, potentially enabling remote code execution.	10.0	More Details
CVE-2025-68668	n8n is an open source workflow automation platform. From version 1.0.0 to before 2.0.0, a sandbox bypass vulnerability exists in the Python Code Node that uses Pyodide. An authenticated user with permission to create or modify workflows can exploit this vulnerability to execute arbitrary commands on the host system running n8n, using the same privileges as the n8n process. This issue has been patched in version 2.0.0. Workarounds for this issue involve disabling the Code Node by setting the environment variable NODES_EXCLUDE: "[\"n8n-nodes-base.code\"]", disabling Python support in the Code node by setting the environment variable N8N_PYTHON_ENABLED=false, which was introduced in n8n version 1.104.0, and configuring n8n to use the task runner based Python sandbox via the N8N_RUNNERS_ENABLED and N8N_NATIVE_PYTHON_RUNNER environment variables.	9.9	More Details
CVE-2025-66203	StreamVault is a video download integration solution. Prior to version 251126, a Remote Code Execution (RCE) vulnerability exists in the stream-vault application (SpiritApplication). The application allows administrators to configure yt-dlp arguments via the /admin/api/saveConfig endpoint without sufficient validation. These arguments are stored globally and subsequently used in YtDlpUtil.java when constructing the command line to execute yt-dlp. This issue has been patched in version 251126.	9.9	More Details
CVE-2025-68897	Improper Control of Generation of Code ('Code Injection') vulnerability in Mohammad I. Okfie IF AS Shortcode allows Code Injection.This issue affects IF AS Shortcode: from n/a through 1.2.	9.9	More Details
CVE-2025-68562	Unrestricted Upload of File with Dangerous Type vulnerability in RomanCode MapSVG allows Upload a Web Shell to a Web Server.This issue affects MapSVG: from n/a through 8.7.3.	9.9	More Details
CVE-2025-13915	IBM API Connect 10.0.8.0 through 10.0.8.5, and 10.0.11.0 could allow a remote attacker to bypass authentication mechanisms and gain unauthorized access to the application.	9.8	More Details
CVE-2025-15226	WMPro developed by Sunnet has a Arbitrary File Upload vulnerability, allowing unauthenticated remote attackers to upload and execute web shell backdoors, thereby enabling arbitrary code execution on the server.	9.8	More Details
CVE-2025-15228	BPMFlowWebkit developed by WELLTEND TECHNOLOGY has a Arbitrary File Upload vulnerability, allowing unauthenticated remote attackers to upload and execute web shell backdoors, thereby enabling arbitrary code execution on the server.	9.8	More Details
CVE-2025-15194	A vulnerability was found in D-Link DIR-600 up to 2.15WWb02. Affected by this vulnerability is an unknown functionality of the file hedwig.cgi of the component HTTP Header Handler. The manipulation of the argument Cookie results in stack-based buffer overflow. It is possible to launch the attack remotely. The exploit has been made public and could be used. This vulnerability only affects products that are no longer supported by the maintainer.	9.8	More Details
CVE-2025-68860	Authentication Bypass Using an Alternate Path or Channel vulnerability in Mobile Builder Mobile builder allows Authentication Abuse.This issue affects Mobile builder: from n/a through 1.4.2.	9.8	More Details
CVE-2025-68974	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in miniOrange WordPress Social Login and Register miniorange-login-openid allows PHP Local File Inclusion.This issue affects WordPress Social Login and Register: from n/a through <= 7.7.0.	9.8	More Details

CVE-2025-68983	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in thembay Greenmart greenmart allows PHP Local File Inclusion.This issue affects Greenmart: from n/a through <= 4.2.11.	9.8	More Details
CVE-2025-68984	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in thembay Puca puca allows PHP Local File Inclusion.This issue affects Puca: from n/a through <= 2.6.39.	9.8	More Details
CVE-2025-68985	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in thembay Aora aora allows PHP Local File Inclusion.This issue affects Aora: from n/a through <= 1.3.15.	9.8	More Details
CVE-2025-68987	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in Edge-Themes Cinerama - A WordPress Theme for Movie Studios and Filmmakers cinerama allows PHP Local File Inclusion.This issue affects Cinerama - A WordPress Theme for Movie Studios and Filmmakers: from n/a through <= 2.4.	9.8	More Details
CVE-2025-68990	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in xenioushk BWL Pro Voting Manager bwl-pro-voting-manager allows Blind SQL Injection.This issue affects BWL Pro Voting Manager: from n/a through <= 1.4.9.	9.8	More Details
CVE-2025-15255	A vulnerability was determined in Tenda W6-S 1.0.0.4(510). This impacts an unknown function of the file /bin/httpd of the component R7websSsecurityHandler. Executing manipulation of the argument Cookie can lead to stack-based buffer overflow. The attack may be launched remotely. The exploit has been publicly disclosed and may be utilized.	9.8	More Details
CVE-2025-68926	RustFS is a distributed object storage system built in Rust. In versions prior to 1.0.0-alpha.77, RustFS implements gRPC authentication using a hardcoded static token `"rustfs rpc"` that is publicly exposed in the source code repository, hardcoded on both client and server sides, non-configurable with no mechanism for token rotation, and universally valid across all RustFS deployments. Any attacker with network access to the gRPC port can authenticate using this publicly known token and execute privileged operations including data destruction, policy manipulation, and cluster configuration changes. Version 1.0.0-alpha.77 contains a fix for the issue.	9.8	More Details
CVE-2022-50691	MiniDVBLinux 5.4 contains a remote command execution vulnerability that allows unauthenticated attackers to execute arbitrary commands as root through the 'command' GET parameter. Attackers can exploit the /tpl/commands.sh endpoint by sending malicious command values to gain root-level system access.	9.8	More Details
CVE-2022-50695	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x contains a network vulnerability that allows unauthenticated attackers to send ICMP signals to arbitrary hosts through network command scripts. Attackers can abuse ping.php, traceroute.php, and dns.php to generate network flooding attacks targeting external hosts.	9.8	More Details
CVE-2022-50790	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x and below contain an unauthenticated vulnerability that allows remote attackers to access live radio stream information through webplay or ffmpeg scripts. Attackers can exploit the vulnerability by calling specific web scripts to disclose radio stream details without requiring authentication.	9.8	More Details
CVE-2022-50792	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x and below contain an unauthenticated file disclosure vulnerability that allows remote attackers to access sensitive system files. Attackers can exploit the vulnerability by manipulating the 'file' GET parameter to disclose arbitrary files on the affected device.	9.8	More Details
CVE-2022-50794	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x and below contain an unauthenticated command injection vulnerability in the username parameter. Attackers can exploit index.php and login.php scripts by injecting arbitrary shell commands through the HTTP POST 'username' parameter to execute system commands.	9.8	More Details
CVE-2022-50803	JM-DATA ONU JF511-TV version 1.0.67 uses default credentials that allow attackers to gain unauthorized access to the device with administrative privileges.	9.8	More Details
CVE-2024-58336	Akuvox Smart Intercom S539 contains an unauthenticated vulnerability that allows remote attackers to access live video streams by requesting the video.cgi endpoint on port 8080. Attackers can retrieve video stream data without authentication by directly accessing the specified endpoint on affected Akuvox doorphone and intercom devices.	9.8	More Details
CVE-2024-58338	Anevia Flamingo XL 3.2.9 contains a restricted shell vulnerability that allows remote attackers to escape the sandboxed environment through the traceroute command. Attackers can exploit the traceroute command to inject shell commands and gain full root access to the device by bypassing the restricted login environment.	9.8	More Details
CVE-2024-44065	Time-based blind SQL Injection vulnerability in Cloudlog v2.6.15 at the endpoint /index.php/logbookadvanced/search in the qsoresults parameter.	9.8	More Details
CVE-2025-13773	The Print Invoice & Delivery Notes for WooCommerce plugin for WordPress is vulnerable to Remote Code Execution in all versions up to, and including, 5.8.0 via the 'WooCommerce_Delivery_Notes::update' function. This is due to missing capability check in the 'WooCommerce_Delivery_Notes::update' function, PHP enabled in Dompdf, and missing escape in the 'template.php' file. This makes it possible for unauthenticated attackers to execute code on the server.	9.8	More Details
CVE-2025-8769	Telenium Online Web Application is vulnerable due to a Perl script that is called to load the login page. Due to improper input validation, an attacker can inject arbitrary Perl code through a crafted HTTP request, leading to remote code execution on the server.	9.8	More Details
CVE-2018-25134	Synaccess netBooter NP-02x/NP-08x 6.8 contains an authentication bypass vulnerability in the webNewAcct.cgi script that allows unauthenticated attackers to create admin user accounts. Attackers can exploit the missing control check by sending crafted POST requests to create administrative accounts and gain unauthorized control over power supply management.	9.8	More Details
CVE-2025-68038	Deserialization of Untrusted Data vulnerability in Icegram Icegram Express Pro email-subscribers-premium allows Object Injection.This issue affects Icegram Express Pro: from n/a through <= 5.9.11.	9.8	More Details
CVE-2025-68496	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Syed Balkhi User Feedback userfeedback-lite allows Blind SQL Injection.This issue affects User Feedback: from n/a through <= 1.10.1.	9.8	More Details
CVE-2025-68506	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in Nawawi Jamili Docket Cache docket-cache allows PHP Local File Inclusion.This issue affects Docket Cache: from n/a through <= 24.07.03.	9.8	More Details

CVE-2025-68519	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in BeRocket Brands for WooCommerce brands-for-woocommerce allows Blind SQL Injection.This issue affects Brands for WooCommerce: from n/a through <= 3.8.6.3.	9.8	More Details
CVE-2025-68530	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in pavothemes Bookory bookory allows PHP Local File Inclusion.This issue affects Bookory: from n/a through <= 2.2.7.	9.8	More Details
CVE-2025-68537	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in thembay Zota zota allows PHP Local File Inclusion.This issue affects Zota: from n/a through <= 1.3.14.	9.8	More Details
CVE-2025-68540	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in thembay Fana fana allows PHP Local File Inclusion.This issue affects Fana: from n/a through <= 1.1.35.	9.8	More Details
CVE-2025-68563	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in WP Shuffle Subscribe to Unlock Lite subscribe-to-unlock-lite allows PHP Local File Inclusion.This issue affects Subscribe to Unlock Lite: from n/a through <= 1.3.0.	9.8	More Details
CVE-2025-68565	Missing Authorization vulnerability in JayBee Twitch Player ttv-easy-embed-player allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Twitch Player: from n/a through <= 2.1.3.	9.8	More Details
CVE-2025-68570	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in captivateaudio Captivate Sync captivatesync-trade allows Blind SQL Injection.This issue affects Captivate Sync: from n/a through <= 3.2.2.	9.8	More Details
CVE-2025-68590	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in CRM Perks Integration for Contact Form 7 HubSpot cf7-hubspot allows Blind SQL Injection.This issue affects Integration for Contact Form 7 HubSpot: from n/a through <= 1.4.2.	9.8	More Details
CVE-2018-25135	Anviz AIM CrossChex Standard 4.3.6.0 contains a CSV injection vulnerability that allows attackers to execute commands by inserting malicious formulas in user import fields. Attackers can craft payloads in fields like 'Name', 'Gender', or 'Position' to trigger Excel macro execution when importing user data.	9.8	More Details
CVE-2018-25154	GNU Barcode 0.99 contains a buffer overflow vulnerability in its code 93 encoding process that allows attackers to trigger memory corruption. Attackers can exploit boundary errors during input file processing to potentially execute arbitrary code on the affected system.	9.8	More Details
CVE-2019-25249	devolo dLAN 500 AV Wireless+ 3.1.0-1 contains an authentication bypass vulnerability that allows attackers to enable hidden services through the htmlmgr CGI script. Attackers can enable telnet and remote shell services, reboot the device, and gain root access without a password by manipulating system configuration parameters.	9.8	More Details
CVE-2019-25240	Rifatron 5brid DVR contains an unauthenticated vulnerability in the animate.cgi script that allows unauthorized access to live video streams. Attackers can exploit the Mobile Web Viewer module by specifying channel numbers to retrieve sequential video snapshots without authentication.	9.8	More Details
CVE-2019-25237	V-SOL GPON/EPON OLT Platform v2.03 contains a privilege escalation vulnerability that allows normal users to gain administrative access by manipulating the user role parameter. Attackers can send a crafted HTTP POST request to the user management endpoint with 'user_role_mod' set to integer value '1' to elevate their privileges.	9.8	More Details
CVE-2019-25236	iSeeQ Hybrid DVR WH-H4 1.03R contains an unauthenticated vulnerability in the get_jpeg script that allows unauthorized access to live video streams. Attackers can retrieve video snapshots from specific camera channels by sending requests to the /cgi-bin/get_jpeg endpoint without authentication.	9.8	More Details
CVE-2019-25235	Smartwares HOME easy 1.0.9 contains an authentication bypass vulnerability that allows unauthenticated attackers to access administrative web pages by disabling JavaScript. Attackers can navigate to multiple administrative endpoints and to bypass client-side validation and access sensitive system information.	9.8	More Details
CVE-2025-15114	Ksenia Security Lares 4.0 Home Automation version 1.6 contains a critical security flaw that exposes the alarm system PIN in the 'basisInfo' XML file after authentication. Attackers can retrieve the PIN from the server response to bypass security measures and disable the alarm system without additional authentication.	9.8	More Details
CVE-2018-25142	NovaRad NovaPACS Diagnostics Viewer 8.5.19.75 contains an unauthenticated XML External Entity (XXE) injection vulnerability in XML preference import settings. Attackers can craft malicious XML files with DTD parameter entities to retrieve arbitrary system files through an out-of-band channel attack.	9.8	More Details
CVE-2025-52835	Cross-Site Request Forgery (CSRF) vulnerability in ConoHa by GMO WING WordPress Migrator allows Upload a Web Shell to a Web Server.This issue affects WING WordPress Migrator: from n/a through 1.1.9.	9.6	More Details
CVE-2025-15102	DVP-12SE11T - Password Protection Bypass	9.1	More Details
CVE-2025-68535	Missing Authorization vulnerability in sunshinephotocart Sunshine Photo Cart sunshine-photo-cart allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Sunshine Photo Cart: from n/a through <= 3.5.7.1.	9.1	More Details
CVE-2024-25181	A critical vulnerability has been identified in givanz VvvebJs 1.7.2, which allows both Server-Side Request Forgery (SSRF) and arbitrary file reading. The vulnerability stems from improper handling of user-supplied URLs in the "file_get_contents" function within the "save.php" file.	9.1	More Details
CVE-2025-68916	Riello UPS NetMan 208 Application before 1.12 allows cgi-bin/certsupload.cgi /../ directory traversal for file upload with resultant code execution.	9.1	More Details
CVE-2025-68511	Missing Authorization vulnerability in Jegstudio Gutenverse Form gutenverse-form allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Gutenverse Form: from n/a through <= 2.3.1.	9.1	More Details
CVE-2025-68508	Missing Authorization vulnerability in Brave Brave brave-popup-builder allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Brave: from n/a through <= 0.8.3.	9.1	More Details

CVE-2025-68600	Server-Side Request Forgery (SSRF) vulnerability in Yannick Lefebvre Link Library link-library allows Server Side Request Forgery.This issue affects Link Library: from n/a through <= 7.8.4.	9.1	More Details
CVE-2025-68500	Server-Side Request Forgery (SSRF) vulnerability in bdthemes Prime Slider – Addons For Elementor bdthemes-prime-slider-lite allows Server Side Request Forgery.This issue affects Prime Slider – Addons For Elementor: from n/a through <= 4.0.10.	9.1	More Details
CVE-2025-15359	DVP-12SE11T - Out-of-bound memory write Vulnerability	9.1	More Details
CVE-2025-67623	Server-Side Request Forgery (SSRF) vulnerability in 6Storage 6Storage Rentals 6storage-rentals allows Server Side Request Forgery.This issue affects 6Storage Rentals: from n/a through <= 2.19.9.	9.1	More Details
CVE-2025-68929	Frappe is a full-stack web application framework. Prior to versions 14.99.6 and 15.88.1, an authenticated user with specific permissions could be tricked into accessing a specially crafted link. This could lead to a malicious template being executed on the server, resulting in remote code execution. Versions 14.99.6 and 15.88.1 fix the issue. No known workarounds are available.	9.0	More Details

OTHER VULNERABILITIES

CVE Number	Description	Base Score	Reference
CVE-2025-68920	C-Kermit (aka ckermit) through 10.0 Beta.12 (aka 416-beta12) before 244644d allows a remote Kermit system to overwrite files on the local system, or retrieve arbitrary files from the local system.	8.9	More Details
CVE-2025-67622	Cross-Site Request Forgery (CSRF) vulnerability in titopandub Evergreen Post Tweeter evergreen-post-tweeter allows Stored XSS.This issue affects Evergreen Post Tweeter: from n/a through <= 1.8.9.	8.8	More Details
CVE-2025-15216	A vulnerability was identified in Tenda AC23 16.03.07.52. This impacts the function fromSetIpMacBind of the file /goform/SetIpMacBind. Such manipulation of the argument bindnum leads to stack-based buffer overflow. It is possible to launch the attack remotely. The exploit is publicly available and might be used.	8.8	More Details
CVE-2025-15218	A weakness has been identified in Tenda AC10U 15.03.06.48/15.03.06.49. Affected by this vulnerability is the function fromadvsetlanip of the file /goform/AdvSetLanip of the component POST Request Parameter Handler. Executing manipulation of the argument lanMask can lead to buffer overflow. The attack can be launched remotely. The exploit has been made available to the public and could be exploited.	8.8	More Details
CVE-2025-15230	A vulnerability was found in Tenda M3 1.0.0.13(4903). Affected by this issue is the function formSetVlanPolicy of the file /goform/setVlanPolicyData. Performing manipulation of the argument qvlan_truck_port results in heap-based buffer overflow. Remote exploitation of the attack is possible. The exploit has been made public and could be used.	8.8	More Details
CVE-2025-15231	A vulnerability was determined in Tenda M3 1.0.0.13(4903). This affects the function formSetRemoteVlanInfo of the file /goform/setVlanInfo. Executing manipulation of the argument ID/vlan/port can lead to stack-based buffer overflow. The attack can be executed remotely. The exploit has been publicly disclosed and may be utilized.	8.8	More Details
CVE-2025-15232	A vulnerability was identified in Tenda M3 1.0.0.13(4903). This vulnerability affects the function formSetAdPushInfo of the file /goform/setAdPushInfo. The manipulation of the argument mac/terminal leads to stack-based buffer overflow. The attack is possible to be carried out remotely. The exploit is publicly available and might be used.	8.8	More Details
CVE-2025-15233	A security flaw has been discovered in Tenda M3 1.0.0.13(4903). This issue affects the function formSetAdInfoDetails of the file /goform/setAdInfoDetail. The manipulation of the argument adName/smsPassword/smsAccount/weixinAccount/weixinName/smsSignature/adRedirectUrl/adCopyRight/smsContent/adItemUID results in heap-based buffer overflow. The attack may be performed from remote. The exploit has been released to the public and may be exploited.	8.8	More Details
CVE-2025-68976	Missing Authorization vulnerability in Eagle-Themes Eagle Booking eagle-booking allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Eagle Booking: from n/a through <= 1.3.4.3.	8.8	More Details
CVE-2025-68981	Missing Authorization vulnerability in designthemes HomeFix Elementor Portfolio homefix-ele-portfolio allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects HomeFix Elementor Portfolio: from n/a through <= 1.0.1.	8.8	More Details
CVE-2025-2155	Unrestricted Upload of File with Dangerous Type vulnerability in Echo Call Center Services Trade and Industry Inc. Specto CM allows Remote Code Inclusion.This issue affects Specto CM: before 17032025.	8.8	More Details
CVE-2025-68608	Missing Authorization vulnerability in DeluxeThemes Userpro userpro allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Userpro: from n/a through <= 5.1.9.	8.8	More Details
CVE-2025-68601	Cross-Site Request Forgery (CSRF) vulnerability in Rustaurius Five Star Restaurant Reservations restaurant-reservations allows Cross Site Request Forgery.This issue affects Five Star Restaurant Reservations: from n/a through <= 2.7.7.	8.8	More Details
CVE-2025-68596	Missing Authorization vulnerability in Bit Apps Bit Assist bit-assist allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Bit Assist: from n/a through <= 1.5.11.	8.8	More Details
CVE-2025-68595	Missing Authorization vulnerability in Trustindex Widgets for Social Photo Feed social-photo-feed-widget allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Widgets for Social Photo Feed: from n/a through <= 1.7.7.	8.8	More Details

CVE-2025-68593	Missing Authorization vulnerability in Liton Arefin WP Adminify adminify allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Adminify: from n/a through <= 4.0.6.1.	8.8	More Details
CVE-2025-68592	Missing Authorization vulnerability in Liton Arefin WP Adminify adminify allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Adminify: from n/a through <= 4.0.6.1.	8.8	More Details
CVE-2025-15252	A flaw has been found in Tenda M3 1.0.0.13(4903). The affected element is the function formSetRemoteDhcpForAp of the file /goform/setDhcpAP. This manipulation of the argument startip/endip/leasetime/gateway/dns1/dns2 causes stack-based buffer overflow. The attack can be initiated remotely. The exploit has been published and may be used.	8.8	More Details
CVE-2025-15253	A vulnerability has been found in Tenda M3 1.0.0.13(4903). The impacted element is an unknown function of the file /goform/exeCommand. Such manipulation of the argument cmdinput leads to stack-based buffer overflow. The attack can be launched remotely. The exploit has been disclosed to the public and may be used.	8.8	More Details
CVE-2025-15217	A security flaw has been discovered in Tenda AC23 16.03.07.52. Affected is the function formSetPPTPUserList of the component HTTP POST Request Handler. Performing manipulation of the argument list results in buffer overflow. The attack can be initiated remotely.	8.8	More Details
CVE-2025-15215	A vulnerability was determined in Tenda AC10U 15.03.06.48/15.03.06.49. This affects the function formSetPPTPUserList of the file /goform/setPptpUserList of the component HTTP POST Request Handler. This manipulation of the argument list causes buffer overflow. It is possible to initiate the attack remotely. The exploit has been publicly disclosed and may be utilized.	8.8	More Details
CVE-2025-68585	Missing Authorization vulnerability in Ben Balter WP Document Revisions wp-document-revisions allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Document Revisions: from n/a through <= 3.7.2.	8.8	More Details
CVE-2025-15092	A vulnerability was identified in UTT 进取 512W up to 1.7.7-171114. Impacted is the function strcpy of the file /goform/ConfigExceptMSN. Such manipulation of the argument remark leads to buffer overflow. It is possible to launch the attack remotely. The exploit is publicly available and might be used.	8.8	More Details
CVE-2025-15137	A vulnerability was detected in TRENDnet TEW-800MB 1.0.1.0. Affected by this vulnerability is the function sub_F934 of the file NTPSyncWithHost.cgi. The manipulation results in command injection. The attack may be launched remotely. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	8.8	More Details
CVE-2025-15136	A security vulnerability has been detected in TRENDnet TEW-800MB 1.0.1.0. Affected is the function do_setWizard_asp of the file /goform/wizardset of the component Management Interface. The manipulation of the argument WizardConfigured leads to command injection. The attack may be initiated remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	8.8	More Details
CVE-2025-15189	A vulnerability was identified in D-Link DWR-M920 up to 1.1.50. This issue affects the function sub_464794 of the file /boafrm/formDefRoute. The manipulation of the argument submit-url leads to buffer overflow. The attack may be initiated remotely. The exploit is publicly available and might be used.	8.8	More Details
CVE-2025-15190	A security flaw has been discovered in D-Link DWR-M920 up to 1.1.50. Impacted is the function sub_42261C of the file /boafrm/formFilter. The manipulation of the argument ip6addr results in stack-based buffer overflow. The attack may be launched remotely. The exploit has been released to the public and may be exploited.	8.8	More Details
CVE-2025-67729	LMDeploy is a toolkit for compressing, deploying, and serving LLMs. Prior to version 0.11.1, an insecure deserialization vulnerability exists in lmdeploy where torch.load() is called without the weights_only=True parameter when loading model checkpoint files. This allows an attacker to execute arbitrary code on the victim's machine when they load a malicious .bin or .pt model file. This issue has been patched in version 0.11.1.	8.8	More Details
CVE-2025-15193	A vulnerability was detected in D-Link DWR-M920 up to 1.1.50. This affects the function sub_423848 of the file /boafrm/formParentControl. Performing manipulation of the argument submit-url results in buffer overflow. The attack is possible to be carried out remotely. The exploit is now public and may be used.	8.8	More Details
CVE-2025-66738	An issue in Yealink T21P_E2 Phone 52.84.0.15 allows a remote normal privileged attacker to execute arbitrary code via a crafted request the ping function of the diagnostic component.	8.8	More Details
CVE-2025-15091	A vulnerability was determined in UTT 进取 512W up to 1.7.7-171114. This issue affects the function strcpy of the file /goform/formPictureUrl. This manipulation of the argument importpictureurl causes buffer overflow. It is possible to initiate the attack remotely. The exploit has been publicly disclosed and may be utilized.	8.8	More Details
CVE-2018-25143	Microhard Systems IPn4G 1.1.0 contains a service vulnerability that allows authenticated users to enable a restricted SSH shell with a default 'msshc' user. Attackers can exploit a custom 'ping' command in the NcFTP environment to escape the restricted shell and execute commands with root privileges.	8.8	More Details
CVE-2025-15090	A vulnerability was found in UTT 进取 512W up to 1.7.7-171114. This vulnerability affects the function strcpy of the file /goform/formConfigNoticeConfig. The manipulation of the argument timestart results in buffer overflow. The attack may be performed from remote. The exploit has been made public and could be used.	8.8	More Details
CVE-2025-15089	A vulnerability has been found in UTT 进取 512W up to 1.7.7-171114. This affects the function strcpy of the file /goform/APSecurity. The manipulation of the argument wepkey1 leads to buffer overflow. The attack is possible to be carried out remotely. The exploit has been disclosed to the public and may be used.	8.8	More Details
CVE-2025-55061	CWE-434 Unrestricted Upload of File with Dangerous Type	8.8	More Details
CVE-2019-25246	Beward N100 H.264 VGA IP Camera M2.1.6 contains an authenticated file disclosure vulnerability that allows attackers to read arbitrary system files via the 'READ.filePath' parameter. Attackers can exploit the fileread script or SendCGICMD API to access sensitive files like /etc/passwd and /etc/issue by supplying absolute file paths.	8.8	More Details

CVE-2019-25245	Ross Video DashBoard 8.5.1 contains an elevation of privileges vulnerability that allows authenticated users to modify executable files due to improper permission settings. Attackers can exploit the 'M' or 'C' flags for 'Authenticated Users' group to replace the DashBoard.exe binary with a malicious executable.	8.8	More Details
CVE-2019-25243	FaceSentry 6.4.8 contains an authenticated remote command injection vulnerability in pingTest.php and tcpPortTest.php scripts. Attackers can exploit unsanitized input parameters to inject and execute arbitrary shell commands with root privileges by manipulating the 'strInIP' and 'strInPort' parameters.	8.8	More Details
CVE-2018-25148	Microhard Systems IPn4G 1.1.0 contains multiple authenticated remote code execution vulnerabilities in the admin interface that allow attackers to create crontab jobs and modify system startup scripts. Attackers can exploit hidden admin features to execute arbitrary commands with root privileges, including starting services, disabling firewalls, and writing files to the system.	8.8	More Details
CVE-2025-68586	Missing Authorization vulnerability in Gora Tech Cooked cooked allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Cooked: from n/a through <= 1.11.2.	8.8	More Details
CVE-2025-15234	A weakness has been identified in Tenda M3 1.0.0.13(4903). Impacted is the function formSetRemoteInternetLanInfo of the file /goform/setInternetLanInfo. This manipulation of the argument portIp/portMask/portGateWay/portDns/portSecDns causes heap-based buffer overflow. It is possible to initiate the attack remotely. The exploit has been made available to the public and could be exploited.	8.8	More Details
CVE-2025-68572	Missing Authorization vulnerability in Spider Themes BBP Core bbp-core allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects BBP Core: from n/a through <= 1.4.1.	8.8	More Details
CVE-2025-68529	Cross-Site Request Forgery (CSRF) vulnerability in Rhys Wynne WP Email Capture wp-email-capture allows Cross Site Request Forgery.This issue affects WP Email Capture: from n/a through <= 3.12.5.	8.8	More Details
CVE-2025-67625	Cross-Site Request Forgery (CSRF) vulnerability in tmtraderunner Trade Runner traderunner allows Cross Site Request Forgery.This issue affects Trade Runner: from n/a through <= 3.14.	8.8	More Details
CVE-2022-50793	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains an authenticated command injection vulnerability in the www-data-handler.php script that allows attackers to inject system commands through the 'services' POST parameter. Attackers can exploit this vulnerability by crafting malicious 'services' parameter values to execute arbitrary system commands with www-data user privileges.	8.8	More Details
CVE-2025-68575	Missing Authorization vulnerability in Wappointment team Wappointment wappointment allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Wappointment: from n/a through <=2.7.2.	8.8	More Details
CVE-2025-68505	Missing Authorization vulnerability in icc0rz H5P h5p allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects H5P: from n/a through <= 1.16.1.	8.8	More Details
CVE-2025-68569	Missing Authorization vulnerability in codepeople WP Time Slots Booking Form wp-time-slots-booking-form allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Time Slots Booking Form: from n/a through <= 1.2.38.	8.8	More Details
CVE-2025-15356	A vulnerability has been found in Tenda AC20 up to 16.03.08.12. The impacted element is the function sscanf of the file /goform/PowerSaveSet. The manipulation of the argument powerSavingEn/time/powerSaveDelay/ledCloseType leads to buffer overflow. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used.	8.8	More Details
CVE-2025-68577	Missing Authorization vulnerability in Virusdie Virusdie virusdie allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Virusdie: from n/a through <= 1.1.6.	8.8	More Details
CVE-2025-68573	Cross-Site Request Forgery (CSRF) vulnerability in Alessandro Piconi Simple Keyword to Link simple-keyword-to-link allows Cross Site Request Forgery.This issue affects Simple Keyword to Link: from n/a through <= 1.5.	8.8	More Details
CVE-2025-68567	Cross-Site Request Forgery (CSRF) vulnerability in wphocus My auctions allegro my-auctions-allegro-free-edition allows Cross Site Request Forgery.This issue affects My auctions allegro: from n/a through <= 3.6.32.	8.8	More Details
CVE-2025-68571	Missing Authorization vulnerability in SALESmanago SALESmanago salesmanago allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects SALESmanago: from n/a through <= 3.9.0.	8.8	More Details
CVE-2025-68580	Cross-Site Request Forgery (CSRF) vulnerability in pluginsware Advanced Classifieds & Directory Pro advanced-classifieds-and-directory-pro allows Cross Site Request Forgery.This issue affects Advanced Classifieds & Directory Pro: from n/a through <= 3.2.9.	8.8	More Details
CVE-2025-68582	Missing Authorization vulnerability in Funnelforms Funnelforms Free funnelforms-free allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Funnelforms Free: from n/a through <= 3.8.	8.8	More Details
CVE-2025-68583	Cross-Site Request Forgery (CSRF) vulnerability in Tikweb Management Fast User Switching fast-user-switching allows Cross Site Request Forgery.This issue affects Fast User Switching: from n/a through <= 1.4.10.	8.8	More Details
CVE-2025-68584	Cross-Site Request Forgery (CSRF) vulnerability in Constantin Boiangiu Vimeotheque codeflavors-vimeo-video-post-lite allows Cross Site Request Forgery.This issue affects Vimeotheque: from n/a through <= 2.3.5.2.	8.8	More Details

CVE-2025-68522	Missing Authorization vulnerability in wpstream WpStream wpstream allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WpStream: from n/a through <= 4.9.5.	8.8	More Details
CVE-2025-68521	Missing Authorization vulnerability in wpstream WpStream wpstream allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WpStream: from n/a through <= 4.9.5.	8.8	More Details
CVE-2025-59887	Improper authentication of library files in the Eaton UPS Companion software installer could lead to arbitrary code execution of an attacker with the access to the software package. This security issue has been fixed in the latest version of EUC which is available on the Eaton download center.	8.6	More Details
CVE-2023-36525	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in WPJobBoard allows Blind SQL Injection.This issue affects WPJobBoard: from n/a through 5.9.0.	8.6	More Details
CVE-2025-13417	The Plugin Organizer WordPress plugin before 10.2.4 does not sanitize and escape a parameter before using it in a SQL statement, allowing subscribers to perform SQL injection attacks.	8.6	More Details
CVE-2024-58315	Tosibox Key Service 3.3.0 contains an unquoted service path vulnerability that allows local non-privileged users to potentially execute code with elevated system privileges. Attackers can exploit the service startup process by inserting malicious code in the system root path, enabling unauthorized code execution during application startup or system reboot.	8.4	More Details
CVE-2022-50795	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains a conditional command injection vulnerability that allows local authenticated users to create malicious files in the /tmp directory. Unauthenticated attackers can execute commands by making a single HTTP POST request to the traceroute.php script, which triggers the malicious file and then deletes it after execution.	8.4	More Details
CVE-2022-50791	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains a conditional command injection vulnerability that allows local authenticated users to create malicious files in the /tmp directory. Unauthenticated attackers can execute commands by making a single HTTP POST request to the vulnerable ping.php script, which triggers the malicious file and then deletes it.	8.4	More Details
CVE-2022-50789	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains a command injection vulnerability that allows local authenticated users to create malicious files in the /tmp directory with .dns.pid extension. Unauthenticated attackers can execute the malicious commands by making a single HTTP POST request to the vulnerable dns.php script, which triggers command execution and then deletes the file.	8.4	More Details
CVE-2025-68939	Gitea before 1.23.0 allows attackers to add attachments with forbidden file extensions by editing an attachment name via an attachment API.	8.2	More Details
CVE-2023-54163	NLB mKlik Macedonia 3.3.12 contains a SQL injection vulnerability in international transfer parameters that allows attackers to manipulate database queries. Attackers can inject arbitrary SQL code through unsanitized input to potentially disclose sensitive information from the mobile banking application.	8.2	More Details
CVE-2025-66444	Cross-site Scripting vulnerability in Hitachi Infrastructure Analytics Advisor (Data Center Analytics component) and Hitachi Ops Center Analyzer (Hitachi Ops Center Analyzer detail view component).This issue affects Hitachi Infrastructure Analytics Advisor:; Hitachi Ops Center Analyzer: from 10.0.0-00 before 11.0.5-00.	8.2	More Details
CVE-2025-59683	Pexip Infinity 15.0 through 38.0 before 38.1 has Improper Access Control in the Secure Scheduler for Exchange service, when used with Office 365 Legacy Exchange Tokens. This allows a remote attacker to read potentially sensitive data and excessively consume resources, leading to a denial of service.	8.2	More Details
CVE-2022-50694	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains an SQL injection vulnerability in the 'username' POST parameter of index.php that allows attackers to manipulate database queries. Attackers can inject arbitrary SQL code through the username parameter to bypass authentication and potentially access unauthorized database information.	8.2	More Details
CVE-2018-25128	SOCA Access Control System 180612 contains multiple SQL injection vulnerabilities that allow attackers to manipulate database queries through unvalidated POST parameters. Attackers can bypass authentication, retrieve password hashes, and gain administrative access with full system privileges by exploiting injection flaws in Login.php and Card_Edit_GetJson.php.	8.2	More Details
CVE-2025-68523	Missing Authorization vulnerability in Spiffy Plugins Spiffy Calendar spiffy-calendar allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Spiffy Calendar: from n/a through <= 5.0.7.	8.1	More Details
CVE-2025-68603	Missing Authorization vulnerability in Marketing Fire Editorial Calendar editorial-calendar allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Editorial Calendar: from n/a through <= 3.8.8.	8.1	More Details
CVE-2025-68587	Missing Authorization vulnerability in Bob Watu Quiz watu allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Watu Quiz: from n/a through <= 3.4.5.	8.1	More Details
CVE-2025-68579	Missing Authorization vulnerability in FolioVision FV Simpler SEO fv-all-in-one-seo-pack allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects FV Simpler SEO: from n/a through <= 1.9.6.	8.1	More Details
CVE-2025-68594	Missing Authorization vulnerability in Assaf Parag Poll, Survey & Quiz Maker Plugin by Opinion Stage social-polls-by-opinionstage allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Poll, Survey & Quiz Maker Plugin by Opinion Stage: from n/a through <= 19.12.1.	8.1	More Details
CVE-2025-68982	Missing Authorization vulnerability in designthemes DesignThemes LMS Addon designthemes-lms-addon allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects DesignThemes LMS Addon: from n/a through <= 2.6.	8.1	More Details
CVE-	Missing Authorization vulnerability in Essekia Tablesome tablesome allows Exploiting Incorrectly Configured Access Control Security		More

2025-68517	Levels.This issue affects Tablesome: from n/a through <= 1.1.35.1.	8.1	Details
CVE-2025-68581	Missing Authorization vulnerability in YITHEMES YITH Slider for page builders yith-slider-for-page-builders allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects YITH Slider for page builders: from n/a through <= 1.0.11.	8.1	More Details
CVE-2025-68980	Missing Authorization vulnerability in designthemes WeDesignTech Portfolio wedesigntech-portfolio allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WeDesignTech Portfolio: from n/a through <= 1.0.2.	8.1	More Details
CVE-2025-68591	Missing Authorization vulnerability in Mitchell Bennis Simple File List simple-file-list allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Simple File List: from n/a through <= 6.1.15.	8.1	More Details
CVE-2025-68589	Missing Authorization vulnerability in WP Socio WP Telegram Widget and Join Link wptelegram-widget allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Telegram Widget and Join Link: from n/a through <= 2.2.11.	8.1	More Details
CVE-2025-68979	Authorization Bypass Through User-Controlled Key vulnerability in SimpleCalendar Google Calendar Events google-calendar-events allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Google Calendar Events: from n/a through <= 3.5.9.	8.1	More Details
CVE-2025-68975	Authorization Bypass Through User-Controlled Key vulnerability in Eagle-Themes Eagle Booking eagle-booking allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Eagle Booking: from n/a through <= 1.3.4.3.	8.1	More Details
CVE-2025-68588	Missing Authorization vulnerability in totalsoft TS Poll poll-wp allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects TS Poll: from n/a through <= 2.5.3.	8.1	More Details
CVE-2025-67909	Authorization Bypass Through User-Controlled Key vulnerability in WP Swings Membership For WooCommerce membership-for-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Membership For WooCommerce: from n/a through <= 3.0.3.	8.1	More Details
CVE-2025-15103	DVP-12SE11T - Authentication Bypass via Partial Password Disclosure	8.1	More Details
CVE-2025-68578	Missing Authorization vulnerability in Addonify Addonify addonify-quick-view allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Addonify: from n/a through <= 2.0.4.	8.1	More Details
CVE-2025-15112	Ksenia Security Lares 4.0 version 1.6 contains a URL redirection vulnerability in the 'cmdOk.xml' script that allows attackers to manipulate the 'redirectPage' GET parameter. Attackers can craft malicious links that redirect authenticated users to arbitrary websites when clicking on a specially constructed link hosted on a trusted domain.	8.0	More Details
CVE-2025-67450	Due to insecure library loading in the Eaton UPS Companion software executable, an attacker with access to the software package could perform arbitrary code execution . This security issue has been fixed in the latest version of EUC which is available on the Eaton download center.	7.8	More Details
CVE-2025-12771	IBM Concert 1.0.0 through 2.1.0 is vulnerable to a stack-based buffer overflow, caused by improper bounds checking. A local user could overflow the buffer and execute arbitrary code on the system.	7.8	More Details
CVE-2025-15113	Ksenia Security Lares 4.0 Home Automation version 1.6 contains an unprotected endpoint vulnerability that allows authenticated attackers to upload MPFS File System binary images. Attackers can exploit this vulnerability to overwrite flash program memory and potentially execute arbitrary code on the home automation system's web server.	7.8	More Details
CVE-2025-68973	In GnuPG through 2.4.8, armor_filter in g10/armor.c has two increments of an index variable where one is intended, leading to an out-of-bounds write for crafted input. (For ExtendedLTS, 2.2.51 and later are fixed versions.)	7.8	More Details
CVE-2025-64645	IBM Concert 1.0.0 through 2.1.0 could allow a local user to escalate their privileges due to a race condition of a symbolic link.	7.7	More Details
CVE-2025-15067	Unrestricted Upload of File with Dangerous Type vulnerability in Innorix Innorix WP allows Upload a Web Shell to a Web Server.This issue affects Innorix WP from All versions If the "exam" directory exists under the directory where the product is installed (ex: innorix/exam)	7.7	More Details
CVE-2025-69217	coturn is a free open source implementation of TURN and STUN Server. Versions 4.6.2r5 through 4.7.0-r4 have a bad random number generator for nonces and port randomization after refactoring. Additionally, random numbers aren't generated with openssl's RAND_bytes but libc's random() (if it's not running on Windows). When fetching about 50 sequential nonces (i.e., through sending 50 unauthenticated allocations requests) it is possible to completely reconstruct the current state of the random number generator, thereby predicting the next nonce. This allows authentication while spoofing IPs. An attacker can send authenticated messages without ever receiving the responses, including the nonce (requires knowledge of the credentials, which is e.g., often the case in IoT settings). Since the port randomization is deterministic given the pseudorandom seed, an attacker can exactly reconstruct the ports and, hence predict the randomization of the ports. If an attacker allocates a relay port, they know the current port, and they are able to predict the next relay port (at least if it is not used before). Commit 11fc465f4bba70bb0ad8aae17d6c4a63a29917d9 contains a fix.	7.7	More Details
CVE-2025-15068	Missing Authorization vulnerability in Gmission Web Fax allows Privilege Abuse, Session Credential Falsification through Manipulation.This issue affects Web Fax: from 3.0 before 4.0.	7.7	More Details

CVE-2025-2307	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Verisay Communication and Information Technology Industry and Trade Ltd. Co. Aidango allows Cross-Site Scripting (XSS).This issue affects Aidango: before 2.144.4.	7.6	More Details
CVE-2025-2406	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Verisay Communication and Information Technology Industry and Trade Ltd. Co. Trizbi allows Cross-Site Scripting (XSS).This issue affects Trizbi: before 2.144.4.	7.6	More Details
CVE-2025-2405	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Verisay Communication and Information Technology Industry and Trade Ltd. Co. Titarus allows Cross-Site Scripting (XSS).This issue affects Titarus: before 2.144.4.	7.6	More Details
CVE-2025-59129	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection') vulnerability in Appointify allows Blind SQL Injection.This issue affects Appointify: from n/a through 1.0.8.	7.6	More Details
CVE-2025-66379	Pexip Infinity before 39.0 has Improper Input Validation in the media implementation, allowing a remote attacker to trigger a software abort via a crafted media stream, resulting in a denial of service.	7.5	More Details
CVE-2025-66443	Pexip Infinity 35.0 through 38.1 before 39.0, in non-default configurations that use Direct Media for WebRTC, has Improper Input Validation in signalling that allows an attacker to trigger a software abort, resulting in a temporary denial of service.	7.5	More Details
CVE-2025-68606	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in WPXPO PostX ultimate-post allows Retrieve Embedded Sensitive Data.This issue affects PostX: from n/a through <= 5.0.3.	7.5	More Details
CVE-2025-67621	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in 10up Eight Day Week Print Workflow eight-day-week-print-workflow allows Retrieve Embedded Sensitive Data.This issue affects Eight Day Week Print Workflow: from n/a through <= 1.2.5.	7.5	More Details
CVE-2025-68576	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in Virusdie Virusdie virusdie allows Retrieve Embedded Sensitive Data.This issue affects Virusdie: from n/a through <= 1.1.6.	7.5	More Details
CVE-2018-25129	SOCA Access Control System 180612 contains multiple insecure direct object reference vulnerabilities that allow attackers to access sensitive user credentials. Attackers can retrieve authenticated and unauthenticated user password hashes and pins through unprotected endpoints like Get_Permissions_From_DB.php and Ac10_ReadSortCard.	7.5	More Details
CVE-2018-25136	FLIR Brickstream 3D+ 2.1.742.1842 contains an unauthenticated vulnerability that allows remote attackers to access live video streams without credentials. Attackers can retrieve video stream images by directly accessing multiple image endpoints like middleImage.jpg, rightimage.jpg, and leftimage.jpg.	7.5	More Details
CVE-2025-15225	WMPro developed by Sunnet has an Arbitrary File Read vulnerability, allowing unauthenticated remote attackers to exploit Relative Path Traversal to read arbitrary system files.	7.5	More Details
CVE-2018-25137	FLIR Brickstream 3D+ 2.1.742.1842 contains an unauthenticated vulnerability in the ExportConfig REST API that allows attackers to download sensitive configuration files. Attackers can exploit the getConfigExportFile.cgi endpoint to retrieve system configurations, potentially enabling authentication bypass and privilege escalation.	7.5	More Details
CVE-2018-25138	FLIR AX8 Thermal Camera 1.32.16 contains hard-coded SSH and web panel credentials that cannot be changed through normal camera operations. Attackers can exploit these persistent credentials to gain unauthorized shell access and login to multiple camera interfaces using predefined username and password combinations.	7.5	More Details
CVE-2025-59946	NanoMQ MQTT Broker (NanoMQ) is an Edge Messaging Platform. Prior to version 0.24.2, there is a classical data racing issue about sub info list which could result in heap use after free crash. This issue has been patched in version 0.24.2.	7.5	More Details
CVE-2025-68494	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in Leap13 Premium Addons for Elementor premium-addons-for-elementor allows Retrieve Embedded Sensitive Data.This issue affects Premium Addons for Elementor: from n/a through <= 4.11.53.	7.5	More Details
CVE-2018-25139	FLIR AX8 Thermal Camera 1.32.16 contains an unauthenticated vulnerability that allows remote attackers to access live video streams without credentials. Attackers can directly connect to the RTSP stream using tools like VLC or FFmpeg to view and record thermal camera footage.	7.5	More Details
CVE-2025-67015	Incorrect access control in Comtech EF Data CDM-625 / CDM-625A Advanced Satellite Modem with firmware v2.5.1 allows attackers to change the Administrator password and escalate privileges via sending a crafted POST request to /Forms/admin_access_1.	7.5	More Details
CVE-2025-67014	Incorrect access control in DEV Systemtechnik GmbH DEV 7113 RF over Fiber Distribution System 32-0078 H.01 allows unauthenticated attackers to access an administrative endpoint.	7.5	More Details
CVE-2018-25140	FLIR thermal traffic cameras contain an unauthenticated device manipulation vulnerability in their WebSocket implementation that allows attackers to bypass authentication and authorization controls. Attackers can directly modify device configurations, access system information, and potentially initiate denial of service by sending crafted WebSocket messages without authentication.	7.5	More Details
CVE-2025-57403	Cola Dnslog v1.3.2 is vulnerable to Directory Traversal. When a DNS query for a TXT record is processed, the application concatenates the requested URL (or a portion of it) directly with a base path using os.path.join. This bypass allows directory traversal or absolute path injection, leading to the potential exposure of sensitive information.	7.5	More Details
CVE-	FLIR thermal traffic cameras contain an unauthenticated vulnerability that allows remote attackers to access live video streams		More

2018-25141	without credentials. Attackers can directly retrieve video streams by accessing specific endpoints like /live.mjpeg, /snapshot.jpg, and RTSP streaming URLs without authentication.	7.5	Details
CVE-2025-25341	A vulnerability exists in the libxmljs 1.0.11 when parsing a specially crafted XML document. Accessing the internal _ref property on entity_ref and entity_decl nodes causes a segmentation fault, potentially leading to a denial-of-service (DoS).	7.5	More Details
CVE-2018-25147	Microhard Systems IPn4G 1.1.0 contains hardcoded default credentials that cannot be changed through normal gateway operations. Attackers can exploit these default credentials to gain unauthorized root-level access to the device by logging in with predefined username and password combinations.	7.5	More Details
CVE-2025-68568	Missing Authorization vulnerability in integrationclaspo Popup Builder: Exit-Intent pop-up, Spin the Wheel, Newsletter signup, Email Capture & Lead Generation forms maker claspo allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Popup Builder: Exit-Intent pop-up, Spin the Wheel, Newsletter signup, Email Capture & Lead Generation forms maker: from n/a through <= 1.0.5.	7.5	More Details
CVE-2019-25239	V-SOL GPON/EPON OLT Platform 2.03 contains an unauthenticated information disclosure vulnerability that allows attackers to download configuration files via direct object reference. Attackers can retrieve sensitive configuration data by sending HTTP GET requests to the usrcfg.conf endpoint, potentially enabling authentication bypass and system access.	7.5	More Details
CVE-2019-25248	Beward N100 M2.1.6.04C014 contains an unauthenticated vulnerability that allows remote attackers to access live video streams without credentials. Attackers can directly retrieve the camera's RTSP stream by exploiting the lack of authentication in the video access mechanism.	7.5	More Details
CVE-2025-68516	Insertion of Sensitive Information Into Sent Data vulnerability in Essekia Tablesome tablesome allows Retrieve Embedded Sensitive Data.This issue affects Tablesome: from n/a through <= 1.1.35.1.	7.5	More Details
CVE-2019-25253	KYOCERA Net Admin 3.4.0906 contains an XML External Entity (XXE) injection vulnerability in the Multi-Set Template Editor that allows unauthenticated attackers to read arbitrary system files. Attackers can craft a malicious XML file with external entity references to retrieve sensitive configuration data like database credentials through an out-of-band channel attack.	7.5	More Details
CVE-2019-25258	LogicalDOC Enterprise 7.7.4 contains multiple post-authentication file disclosure vulnerabilities that allow attackers to read arbitrary files through unverified 'suffix' and 'fileVersion' parameters. Attackers can exploit directory traversal techniques in /thumbnail and /convertpdf endpoints to access sensitive system files like win.ini and /etc/passwd by manipulating path traversal sequences.	7.5	More Details
CVE-2025-3232	A remote unauthenticated attacker may be able to bypass authentication by utilizing a specific API route to execute arbitrary OS commands.	7.5	More Details
CVE-2025-32095	Pexip Infinity before 37.0 has improper input validation in signalling that allows a remote attacker to trigger a software abort via a crafted signalling message, resulting in a denial of service.	7.5	More Details
CVE-2025-32096	Pexip Infinity 33.0 through 37.0 before 37.1 has improper input validation in signaling that allows an attacker to trigger a software abort, resulting in a denial of service.	7.5	More Details
CVE-2025-48704	Pexip Infinity 35.0 through 37.2 before 38.0 has Improper Input Validation in signalling that allows an attacker to trigger a software abort, resulting in a denial of service.	7.5	More Details
CVE-2025-66377	Pexip Infinity before 39.0 has Missing Authentication for a Critical Function in a product-internal API, allowing an attacker (who already has access to execute code on one node within a Pexip Infinity installation) to impact the operation of other nodes within the installation.	7.5	More Details
CVE-2019-25241	FaceSentry Access Control System 6.4.8 contains a critical authentication vulnerability with hard-coded SSH credentials for the wwwuser account. Attackers can leverage the insecure sudoers configuration to escalate privileges and gain root access by executing sudo commands without authentication.	7.5	More Details
CVE-2025-62753	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in MadrasThemes MAS Videos allows PHP Local File Inclusion.This issue affects MAS Videos: from n/a through 1.3.2.	7.5	More Details
CVE-2025-68870	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in reDim GmbH CookieHint WP allows PHP Local File Inclusion.This issue affects CookieHint WP: from n/a through 1.0.0.	7.5	More Details
CVE-2025-69256	The Serverless Framework is a framework for using AWS Lambda and other managed cloud services to build applications. Starting in version 4.29.0 and prior to version 4.29.3, a command injection vulnerability exists in the Serverless Framework's built-in MCP server package (@serverless/mcp). This vulnerability only affects users of the experimental MCP server feature (serverless mcp), which represents less than 0.1% of Serverless Framework users. The core Serverless Framework CLI and deployment functionality are not affected. The vulnerability is caused by the unsanitized use of input parameters within a call to `child_process.exec`, enabling an attacker to inject arbitrary system commands. Successful exploitation can lead to remote code execution under the server process's privileges. The server constructs and executes shell commands using unvalidated user input directly within command-line strings. This introduces the possibility of shell metacharacter injection (`|`, `>`, `&&`, etc.). Version 4.29.3 fixes the issue.	7.5	More Details
CVE-2022-50788	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains an information disclosure vulnerability that allows unauthenticated attackers to access sensitive log files. Attackers can directly browse the /log directory to retrieve system and sensitive information without authentication.	7.5	More Details
CVE-2025-66869	Buffer overflow vulnerability in function strcat in asan_interceptors.cpp in libming 0.4.8.	7.5	More Details

CVE-2025-66865	An issue was discovered in function d_print_comp_inner in file cp-demangle.c in BinUtils 2.26 allows attackers to cause a denial of service via crafted PE file.	7.5	More Details
CVE-2022-50692	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x and below contain an insufficient session expiration vulnerability that allows attackers to reuse old session credentials. Attackers can exploit weak session management to potentially hijack active user sessions and gain unauthorized access to the application.	7.5	More Details
CVE-2025-15284	Improper Input Validation vulnerability in qs (parse modules) allows HTTP DoS.This issue affects qs: < 6.14.1. SummaryThe arrayLimit option in qs does not enforce limits for bracket notation (a[]=1&a[]=2), allowing attackers to cause denial-of-service via memory exhaustion. Applications using arrayLimit for DoS protection are vulnerable. DetailsThe arrayLimit option only checks limits for indexed notation (a[0]=1&a[1]=2) but completely bypasses it for bracket notation (a[]=1&a[]=2). Vulnerable code (lib/parse.js:159-162): if (root === '[]' && options.parseArrays) { obj = utils.combine([], leaf); // No arrayLimit check } Working code (lib/parse.js:175): else if (index <= options.arrayLimit) { // Limit checked here obj = []; obj[index] = leaf; } The bracket notation handler at line 159 uses utils.combine([], leaf) without validating against options.arrayLimit, while indexed notation at line 175 checks index <= options.arrayLimit before creating arrays. PoCTest 1 - Basic bypass: npm install qs const qs = require('qs'); const result = qs.parse('a[]=1&a[]=2&a[]=3&a[]=4&a[]=5&a[]=6', { arrayLimit: 5 }); console.log(result.a.length); // Output: 6 (should be max 5) Test 2 - DoS demonstration: const qs = require('qs'); const attack = 'a[]=' + Array(10000).fill('x').join('&a[]='); const result = qs.parse(attack, { arrayLimit: 100 }); console.log(result.a.length); // Output: 10000 (should be max 100) Configuration: * arrayLimit: 5 (test 1) or arrayLimit: 100 (test 2) * Use bracket notation: a[]=value (not indexed a[0]=value) ImpactDenial of Service via memory exhaustion. Affects applications using qs.parse() with user-controlled input and arrayLimit for protection. Attack scenario: * Attacker sends HTTP request: GET /api/search?filters[]=x&filters[]=x&...&filters[]=x (100,000+ times) * Application parses with qs.parse(query, { arrayLimit: 100 }) * qs ignores limit, parses all 100,000 elements into array * Server memory exhausted → application crashes or becomes unresponsive * Service unavailable for all users Real-world impact: * Single malicious request can crash server * No authentication required * Easy to automate and scale * Affects any endpoint parsing query strings with bracket notation	7.5	More Details
CVE-2025-66863	An issue was discovered in function d_discriminator in file cp-demangle.c in BinUtils 2.26 allows attackers to cause a denial of service via crafted PE file.	7.5	More Details
CVE-2025-66862	A buffer overflow vulnerability in function gnu_special in file cplus-dem.c in BinUtils 2.26 allows attackers to cause a denial of service via crafted PE file.	7.5	More Details
CVE-2022-50796	SOUND4 IMPACT/FIRST/PULSE/Eco <=2.x contains an unauthenticated remote code execution vulnerability in the firmware upload functionality with path traversal flaw. Attackers can exploit the upload.cgi script to write malicious files to the system with www-data permissions, enabling unauthorized access and code execution.	7.5	More Details
CVE-2025-68036	Missing Authorization vulnerability in Emraan Cheema CubeWP allows Accessing Functionality Not Properly Constrained by ACLs.This issue affects CubeWP: from n/a through 1.1.27.	7.5	More Details
CVE-2025-69200	phpMyFAQ is an open source FAQ web application. In versions prior to 4.0.16, an unauthenticated remote attacker can trigger generation of a configuration backup ZIP via `POST /api/setup/backup` and then download the generated ZIP from a web-accessible location. The ZIP contains sensitive configuration files (e.g., `database.php` with database credentials), leading to high-impact information disclosure and potential follow-on compromise. Version 4.0.16 fixes the issue.	7.5	More Details
CVE-2024-25183	givanz VvvebJs 1.7.2 is vulnerable to Directory Traversal via scan.php.	7.5	More Details
CVE-2025-15358	DVP-12SE11T - Denial of Service Vulnerability	7.5	More Details
CVE-2025-68988	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in o2oe E-Invoice App Malaysia einvoiceapp-malaysia allows Retrieve Embedded Sensitive Data.This issue affects E-Invoice App Malaysia: from n/a through <= 1.1.0.	7.5	More Details
CVE-2025-68996	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in WebCodingPlace Responsive Posts Carousel Pro responsive-posts-carousel-pro allows PHP Local File Inclusion.This issue affects Responsive Posts Carousel Pro: from n/a through <= 15.1.	7.5	More Details
CVE-2025-68989	Insertion of Sensitive Information Into Sent Data vulnerability in Renzo Johnson Contact Form 7 Extension For Mailchimp contact-form-7-mailchimp-extension allows Retrieve Embedded Sensitive Data.This issue affects Contact Form 7 Extension For Mailchimp: from n/a through <= 0.9.49.	7.5	More Details
CVE-2025-68877	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in CedCommerce CedCommerce Integration for Good Market allows PHP Local File Inclusion.This issue affects CedCommerce Integration for Good Market: from n/a through 1.0.6.	7.5	More Details
CVE-2022-50798	SoX 14.4.2 contains a division by zero vulnerability when handling WAV files that can cause program crashes. Attackers can trigger a floating point exception by providing a specially crafted WAV file that causes arithmetic errors during sound file processing.	7.5	More Details
CVE-2022-50799	Fetch FTP Client 5.8.2 contains a denial of service vulnerability that allows attackers to trigger 100% CPU consumption by sending long server responses. Attackers can send specially crafted FTP server responses exceeding 2K bytes to cause excessive resource utilization and potentially crash the application.	7.5	More Details
CVE-2025-15227	BPMFlowWebkit developed by WELLTEND TECHNOLOGY has a Arbitrary File Read vulnerability, allowing unauthenticated remote attackers to exploit Absolute Path Traversal to download arbitrary system files.	7.5	More Details
CVE-	H3C SSL VPN contains a user enumeration vulnerability that allows attackers to identify valid usernames through the 'txtUsrName'		

2022-50800	POST parameter. Attackers can submit different usernames to the login_submit.cgi endpoint and analyze response messages to distinguish between existing and non-existing accounts.	7.5	More Details
CVE-2023-53983	Anevia Flamingo XL/XS 3.6.20 contains a critical vulnerability with weak default administrative credentials that can be easily guessed. Attackers can leverage these hard-coded credentials to gain full remote system control without complex authentication mechanisms.	7.5	More Details
CVE-2023-54327	Tinycontrol LAN Controller 1.58a contains an authentication bypass vulnerability that allows unauthenticated attackers to change admin passwords through a crafted API request. Attackers can exploit the /stm.cgi endpoint with a specially crafted authentication parameter to disable access controls and modify administrative credentials.	7.5	More Details
CVE-2025-15111	Ksenia Security Lares 4.0 Home Automation version 1.6 contains a default credentials vulnerability that allows unauthorized attackers to gain administrative access. Attackers can exploit the weak default administrative credentials to obtain full control of the home automation system.	7.5	More Details
CVE-2024-58337	Akuvox Smart Intercom S539 contains an improper access control vulnerability that allows users with 'User' privileges to modify API access settings and configurations. Attackers can exploit this vulnerability to escalate privileges and gain unauthorized access to administrative functionalities.	7.5	More Details
CVE-2025-66877	Buffer overflow vulnerability in function dcputchar in decompile.c in libming 0.4.8.	7.5	More Details
CVE-2025-68922	OpenOps before 0.6.11 allows remote code execution in the Terraform block.	7.4	More Details
CVE-2025-15196	A vulnerability was identified in code-projects Assessment Management 1.0. This affects an unknown part of the file login.php. Such manipulation of the argument userid leads to sql injection. The attack can be launched remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2025-15077	A security vulnerability has been detected in itsourcecode Student Management System 1.0. The affected element is an unknown function of the file /form137.php. The manipulation of the argument ID leads to sql injection. The attack may be initiated remotely. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2025-15243	A flaw has been found in code-projects Simple Stock System 1.0. This affects an unknown function of the file /market/login.php. Executing manipulation of the argument Username can lead to sql injection. The attack can be launched remotely. The exploit has been published and may be used.	7.3	More Details
CVE-2025-15076	A weakness has been identified in Tenda CH22 1.0.0.1. Impacted is an unknown function of the file /public/. Executing manipulation can lead to path traversal. The attack can be launched remotely. The exploit has been made available to the public and could be exploited.	7.3	More Details
CVE-2025-15198	A weakness has been identified in code-projects College Notes Uploading System 1.0. This issue affects some unknown processing of the file /login.php. Executing manipulation of the argument User can lead to sql injection. The attack may be launched remotely. The exploit has been made available to the public and could be exploited.	7.3	More Details
CVE-2025-15075	A security flaw has been discovered in itsourcecode Student Management System 1.0. This issue affects some unknown processing of the file /student_p.php. Performing manipulation of the argument ID results in sql injection. The attack can be initiated remotely. The exploit has been released to the public and may be exploited.	7.3	More Details
CVE-2025-15099	A vulnerability was identified in simstudioai sim up to 0.5.27. This vulnerability affects unknown code of the file apps/sim/lib/auth/internal.ts of the component CRON Secret Handler. The manipulation of the argument INTERNAL_API_SECRET leads to improper authentication. It is possible to initiate the attack remotely. The exploit is publicly available and might be used. The identifier of the patch is e359dc2946b12ed5e45a0ec9c95ecf91bd18502a. Applying a patch is the recommended action to fix this issue.	7.3	More Details
CVE-2025-15207	A vulnerability has been found in Campcodes Supplier Management System 1.0. Affected is an unknown function of the file /admin/view_products.php. The manipulation of the argument chkId[] leads to sql injection. It is possible to initiate the attack remotely. The exploit has been disclosed to the public and may be used.	7.3	More Details
CVE-2025-15206	A flaw has been found in Campcodes Supplier Management System 1.0. This impacts an unknown function of the file /admin/add_area.php. Executing manipulation of the argument txtAreaCode can lead to sql injection. The attack may be performed from remote. The exploit has been published and may be used.	7.3	More Details
CVE-2025-15097	A vulnerability was found in Alteryx Server. Affected by this issue is some unknown functionality of the file /gallery/api/status/. Performing manipulation results in improper authentication. The attack is possible to be carried out remotely. The exploit has been made public and could be used. Upgrading to version 2023.1.1.13.486, 2023.2.1.10.293, 2024.1.1.9.236, 2024.2.1.6.125 and 2025.1.1.1.31 can resolve this issue. Upgrading the affected component is recommended.	7.3	More Details
CVE-2025-15073	A vulnerability was determined in itsourcecode Online Frozen Foods Ordering System 1.0. This affects an unknown part of the file /contact_us.php. This manipulation of the argument Name causes sql injection. It is possible to initiate the attack remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2025-15208	A security flaw has been discovered in code-projects Refugee Food Management System 1.0. Affected by this issue is some unknown functionality of the file /home/editrefugee.php. The manipulation of the argument rfid results in sql injection. The attack can be launched remotely. The exploit has been released to the public and may be exploited.	7.3	More Details
CVE-2025-15074	A vulnerability was identified in itsourcecode Online Frozen Foods Ordering System 1.0. This vulnerability affects unknown code of the file /customer_details.php. Such manipulation leads to sql injection. It is possible to launch the attack remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-	A vulnerability was identified in Edimax BR-6208AC 1.02/1.03. Affected is the function formStaDrvSetup of the file /goform/formStaDrvSetup of the component Web-based Configuration Interface. The manipulation of the argument rootAPmac leads to command injection. Remote exploitation of the attack is possible. The exploit is publicly available and might be used. Edimax		

2025-15256	confirms this issue: "The product mentioned, EDIMAX BR-6208AC V2, has reached its End of Life (EOL) status. It is no longer supported or maintained by Edimax, and it is no longer available for purchase in the market. Consequently, there will be no further firmware updates or patches for this device. We recommend users upgrade to newer models for better security." This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2025-15186	A vulnerability has been found in code-projects Refugee Food Management System 1.0. Affected by this issue is some unknown functionality of the file /home/addusers.php. Such manipulation of the argument a leads to sql injection. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used.	7.3	More Details
CVE-2025-15195	A vulnerability was determined in code-projects Assessment Management 1.0. Affected by this issue is some unknown functionality of the file /admin/add-module.php. This manipulation of the argument linked[] causes sql injection. The attack can be initiated remotely. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2025-15053	A flaw has been found in code-projects Student Information System 1.0. This issue affects some unknown processing of the file /searchresults.php. Executing manipulation of the argument searchbox can lead to sql injection. The attack may be performed from remote. The exploit has been published and may be used.	7.3	More Details
CVE-2025-15181	A security flaw has been discovered in code-projects Refugee Food Management System 1.0. The impacted element is an unknown function of the file /home/pagenateRefugeesList.php. Performing manipulation of the argument rfid results in sql injection. Remote exploitation of the attack is possible. The exploit has been released to the public and may be exploited.	7.3	More Details
CVE-2025-15168	A vulnerability was identified in itsourcecode Student Management System 1.0. Affected is an unknown function of the file /statistical.php. Such manipulation of the argument ID leads to sql injection. The attack can be executed remotely. The exploit is publicly available and might be used.	7.3	More Details
CVE-2025-15167	A vulnerability was determined in itsourcecode Online Cake Ordering System 1.0. This impacts an unknown function of the file /detailtransac.php. This manipulation of the argument ID causes sql injection. Remote exploitation of the attack is possible. The exploit has been publicly disclosed and may be utilized.	7.3	More Details
CVE-2025-15166	A vulnerability was found in itsourcecode Online Cake Ordering System 1.0. This affects an unknown function of the file /updatesupplier.php?action=edit. The manipulation of the argument ID results in sql injection. The attack may be launched remotely. The exploit has been made public and could be used.	7.3	More Details
CVE-2025-15165	A vulnerability has been found in itsourcecode Online Cake Ordering System 1.0. The impacted element is an unknown function of the file /updatecustomer.php?action=edit. The manipulation of the argument ID leads to sql injection. The attack may be initiated remotely. The exploit has been disclosed to the public and may be used.	7.3	More Details
CVE-2025-15182	A weakness has been identified in code-projects Refugee Food Management System 1.0. This affects an unknown function of the file /home/served.php. Executing manipulation of the argument refNo can lead to sql injection. The attack can be executed remotely. The exploit has been made available to the public and could be exploited.	7.3	More Details
CVE-2025-15183	A security vulnerability has been detected in code-projects Refugee Food Management System 1.0. This impacts an unknown function of the file /home/viewtakenfd.php. The manipulation of the argument tfid leads to sql injection. The attack is possible to be carried out remotely. The exploit has been disclosed publicly and may be used.	7.3	More Details
CVE-2025-15184	A vulnerability was detected in code-projects Refugee Food Management System 1.0. Affected is an unknown function of the file /home/refugeesreport2.php. The manipulation of the argument a results in sql injection. The attack may be performed from remote. The exploit is now public and may be used.	7.3	More Details
CVE-2025-15142	A vulnerability was identified in 9786 phpok3w up to 901d96a06809fb28b17f3a4362c59e70411c933c. Impacted is an unknown function of the file show.php. The manipulation of the argument ID leads to sql injection. It is possible to initiate the attack remotely. The exploit is publicly available and might be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2025-15140	A vulnerability was found in saiftheboss7 onlinemcqexam up to 0e56806132971e49721db3ef01868098c7b42ada. This vulnerability affects unknown code of the file /admin/quesadd.php. Performing manipulation of the argument ans1/ans2 results in sql injection. The attack is possible to be carried out remotely. The exploit has been made public and could be used. This product adopts a rolling release strategy to maintain continuous delivery The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2025-15185	A flaw has been found in code-projects Refugee Food Management System 1.0. Affected by this vulnerability is an unknown functionality of the file /home/refugeesreport.php. This manipulation of the argument a causes sql injection. It is possible to initiate the attack remotely. The exploit has been published and may be used.	7.3	More Details
CVE-2025-15078	A vulnerability was detected in itsourcecode Student Management System 1.0. The impacted element is an unknown function of the file /list_report.php. The manipulation of the argument sy results in sql injection. The attack may be launched remotely. The exploit is now public and may be used.	7.3	More Details
CVE-2025-15354	A flaw has been found in itsourcecode Society Management System 1.0. The affected element is an unknown function of the file /admin/add_admin.php. Executing manipulation of the argument Username can lead to sql injection. It is possible to launch the attack remotely. The exploit has been published and may be used.	7.3	More Details
CVE-2025-15353	A vulnerability was detected in itsourcecode Society Management System 1.0. Impacted is the function edit_admin_query of the file /admin/edit_admin_query.php. Performing manipulation of the argument Username results in sql injection. It is possible to initiate the attack remotely. The exploit is now public and may be used.	7.3	More Details
CVE-2025-15127	A security vulnerability has been detected in FantasticLBP Hotels_Server up to 67b44df162fab26df209bd5d5d542875fcbec1d0. Affected by this issue is some unknown functionality of the file /controller/api/Room.php. Such manipulation of the argument hotelId leads to sql injection. The attack may be launched remotely. The exploit has been disclosed publicly and may be used. This product operates on a rolling release basis, ensuring continuous delivery. Consequently, there are no version details for either affected or updated releases. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-	A vulnerability was determined in FeehiCMS up to 2.1.1. Impacted is an unknown function of the file frontend/web/timthumb.php of		

2025-15264	the component TimThumb. Executing manipulation of the argument src can lead to server-side request forgery. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	7.3	More Details
CVE-2025-15109	A flaw has been found in jackq XCMS up to 3fab5342cc509945a7ce1b8ec39d19f701b89261. This impacts an unknown function of the file Public/javascripts/admin/plupload-2.1.2/examples/upload.php. This manipulation causes unrestricted upload. It is possible to initiate the attack remotely. The exploit has been published and may be used. This product is using a rolling release to provide continious delivery. Therefore, no version details for affected nor updated releases are available. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2025-15263	A weakness has been identified in BiggiDroid Simple PHP CMS 1.0. Affected is an unknown function of the file /admin/login.php of the component Admin Login. Executing manipulation of the argument Username can lead to sql injection. The attack can be executed remotely. The exploit has been made available to the public and could be exploited.	7.3	More Details
CVE-2025-15257	A security flaw has been discovered in Edimax BR-6208AC 1.02/1.03. Affected by this vulnerability is the function formRoute of the file /gogorm/formRoute of the component Web-based Configuration Interface. The manipulation of the argument strIp/strMask/strGateway results in command injection. The attack can be executed remotely. The exploit has been released to the public and may be exploited. Edimax confirms this issue: "The product mentioned, EDIMAX BR-6208AC V2, has reached its End of Life (EOL) status. It is no longer supported or maintained by Edimax, and it is no longer available for purchase in the market. Consequently, there will be no further firmware updates or patches for this device. We recommend users upgrade to newer models for better security." This vulnerability only affects products that are no longer supported by the maintainer.	7.3	More Details
CVE-2025-15247	A vulnerability was identified in gmg137 snap7-rs up to 153d3e8c16decd7271e2a5b2e3da4d6f68589424. Affected by this issue is the function snap7_rs::client::S7Client::download of the file client.rs. Such manipulation leads to heap-based buffer overflow. The attack can be executed remotely. The exploit is publicly available and might be used. This product implements a rolling release for ongoing delivery, which means version information for affected or updated releases is unavailable. The project was informed of the problem early through an issue report but has not responded yet.	7.3	More Details
CVE-2025-61914	n8n is an open source workflow automation platform. Prior to version 1.114.0, a stored Cross-Site Scripting (XSS) vulnerability may occur in n8n when using the “Respond to Webhook” node. When this node responds with HTML content containing executable scripts, the payload may execute directly in the top-level window, rather than within the expected sandbox introduced in version 1.103.0. This behavior can enable a malicious actor with workflow creation permissions to execute arbitrary JavaScript in the context of the n8n editor interface. This issue has been patched in version 1.114.0. Workarounds for this issue involve restricting workflow creation and modification privileges to trusted users only, avoiding use of untrusted HTML responses in the “Respond to Webhook” node, and using an external reverse proxy or HTML sanitizer to filter responses that include executable scripts.	7.3	More Details
CVE-2022-50787	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x contains an unauthenticated stored cross-site scripting vulnerability in the username parameter that allows attackers to inject malicious scripts. Attackers can exploit the unvalidated username input to execute arbitrary HTML and JavaScript code in victim browser sessions without authentication.	7.2	More Details
CVE-2025-15178	A vulnerability was found in Tenda WH450 1.0.0.18. This issue affects some unknown processing of the file /goform/VirtualSer of the component HTTP Request Handler. The manipulation of the argument page results in stack-based buffer overflow. The attack can be launched remotely. The exploit has been made public and could be used.	7.2	More Details
CVE-2018-25131	Leica Geosystems GR10/GR25/GR30/GR50 GNSS 4.30.063 contains a stored cross-site scripting vulnerability in the configuration file upload functionality. Attackers can upload a malicious HTML file to that executes arbitrary JavaScript in a user's browser session when viewed.	7.2	More Details
CVE-2025-14509	The Lucky Wheel for WooCommerce – Spin a Sale plugin for WordPress is vulnerable to PHP Code Injection in all versions up to, and including, 1.1.13. This is due to the plugin using eval() to execute user-supplied input from the 'Conditional Tags' setting without proper validation or sanitization. This makes it possible for authenticated attackers, with Administrator-level access and above, to execute arbitrary PHP code on the server. In WordPress multisite installations, this allows Site Administrators to execute arbitrary code, a capability they should not have since plugin/theme file editing is disabled for non-Super Admins in multisite environments.	7.2	More Details
CVE-2025-2515	A vulnerability was found in BlueChi, a multi-node systemd service controller used in RHIVOS. This flaw allows a user with root privileges on a managed node (qm) to create or override systemd service unit files that affect the host node. This issue can lead to privilege escalation, unauthorized service execution, and potential system compromise.	7.2	More Details
CVE-2025-15179	A vulnerability was determined in Tenda WH450 1.0.0.18. Impacted is an unknown function of the file /goform/qossetting. This manipulation of the argument page causes stack-based buffer overflow. The attack may be initiated remotely. The exploit has been publicly disclosed and may be utilized.	7.2	More Details
CVE-2025-15164	A security flaw has been discovered in Tenda WH450 1.0.0.18. This affects an unknown part of the file /goform/SafeMacFilter. The manipulation of the argument page results in stack-based buffer overflow. The attack may be performed from remote. The exploit has been released to the public and may be exploited.	7.2	More Details
CVE-2025-15180	A vulnerability was identified in Tenda WH450 1.0.0.18. The affected element is an unknown function of the file /goform/webExcptypemanFilte of the component HTTP Request Handler. Such manipulation of the argument page leads to stack-based buffer overflow. The attack may be launched remotely. The exploit is publicly available and might be used.	7.2	More Details
CVE-2025-15161	A vulnerability was found in Tenda WH450 1.0.0.18. Affected is an unknown function of the file /goform/PPTPUserSetting. Performing manipulation of the argument delno results in stack-based buffer overflow. Remote exploitation of the attack is possible. The exploit has been made public and could be used.	7.2	More Details
CVE-2025-15162	A vulnerability was determined in Tenda WH450 1.0.0.18. Affected by this vulnerability is an unknown functionality of the file /goform/RouteStatic. Executing manipulation of the argument page can lead to stack-based buffer overflow. The attack can be executed remotely. The exploit has been publicly disclosed and may be utilized.	7.2	More Details
CVE-2025-15163	A vulnerability was identified in Tenda WH450 1.0.0.18. Affected by this issue is some unknown functionality of the file /goform/SafeEmailFilter. The manipulation of the argument page leads to stack-based buffer overflow. The attack is possible to be carried out remotely. The exploit is publicly available and might be used.	7.2	More Details
CVE-2025-	A vulnerability has been found in Tenda WH450 1.0.0.18. This impacts an unknown function of the file /goform/PPTPServer. Such manipulation of the argument ip1 leads to stack-based buffer overflow. The attack may be launched remotely. The exploit has been	7.2	More Details

15160	disclosed to the public and may be used.		
CVE-2025-13592	The Advanced Ads plugin for WordPress is vulnerable to Remote Code Execution in versions up to, and including, 2.0.14 via the 'change-ad__content' shortcode parameter. This allows authenticated attackers with editor-level permissions or above, to execute code on the server.	7.2	More Details
CVE-2025-15177	A vulnerability has been found in Tenda WH450 1.0.0.18. This vulnerability affects unknown code of the file /goform/SetIpBind of the component HTTP Request Handler. The manipulation of the argument page leads to stack-based buffer overflow. The attack can be initiated remotely. The exploit has been disclosed to the public and may be used.	7.2	More Details
CVE-2025-68861	Missing Authorization vulnerability in Plugin Optimizer allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Plugin Optimizer: from n/a through 1.3.7.	7.1	More Details
CVE-2025-15069	Improper Authentication vulnerability in Gmission Web Fax allows Privilege Escalation.This issue affects Web Fax: from 3.0 before 4.0.	7.1	More Details
CVE-2025-23554	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Jakub Glos Off Page SEO allows Reflected XSS.This issue affects Off Page SEO: from n/a through 3.0.3.	7.1	More Details
CVE-2025-23550	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Kemal YAZICI Product Puller allows Reflected XSS.This issue affects Product Puller: from n/a through 1.5.1.	7.1	More Details
CVE-2025-23469	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Sleekplan allows Reflected XSS.This issue affects Sleekplan: from n/a through 0.2.0.	7.1	More Details
CVE-2025-23458	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Rakessh Ads24 Lite allows Reflected XSS.This issue affects Ads24 Lite: from n/a through 1.0.	7.1	More Details
CVE-2025-68697	n8n is an open source workflow automation platform. Prior to version 2.0.0, in self-hosted n8n instances where the Code node runs in legacy (non-task-runner) JavaScript execution mode, authenticated users with workflow editing access can invoke internal helper functions from within the Code node. This allows a workflow editor to perform actions on the n8n host with the same privileges as the n8n process, including: reading files from the host filesystem (subject to any file-access restrictions configured on the instance and OS/container permissions), and writing files to the host filesystem (subject to the same restrictions). This issue has been patched in version 2.0.0. Workarounds for this issue involve limiting file operations by setting N8N_RESTRICT_FILE_ACCESS_TO to a dedicated directory (e.g., ~/.n8n-files) and ensure it contains no sensitive data, keeping N8N_BLOCK_FILE_ACCESS_TO_N8N_FILES=true (default) to block access to .n8n and user-defined config files, and disabling high-risk nodes (including the Code node) using NODES_EXCLUDE if workflow editors are not fully trusted.	7.1	More Details
CVE-2025-68879	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Councilsoft Content Grid Slider allows Reflected XSS.This issue affects Content Grid Slider: from n/a through 1.5.	7.1	More Details
CVE-2025-68878	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Prasadkirpekar Advanced Custom CSS allows Reflected XSS.This issue affects Advanced Custom CSS: from n/a through 1.1.0.	7.1	More Details
CVE-2025-59131	Cross-Site Request Forgery (CSRF) vulnerability in Hoernerfranz WP-CalDav2ICS allows Stored XSS.This issue affects WP-CalDav2ICS: from n/a through 1.3.4.	7.1	More Details
CVE-2025-68876	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in INVELITY Invelity SPS connect allows Reflected XSS.This issue affects Invelity SPS connect: from n/a through 1.0.8.	7.1	More Details
CVE-2025-66445	Authorization bypass vulnerability in Hitachi Infrastructure Analytics Advisor (Data Center Analytics component) and Hitachi Ops Center Analyzer (Hitachi Ops Center Analyzer detail view component).This issue affects Hitachi Infrastructure Analytics Advisor:; Hitachi Ops Center Analyzer: from 10.0.0-00 before 11.0.5-00.	7.1	More Details
CVE-2025-13407	The Gravity Forms WordPress plugin before 2.9.23.1 does not properly prevent users from uploading dangerous files through its chunked upload functionality, allowing attackers to upload PHP files to affected sites and achieve Remote Code Execution, granted they can discover or enumerate the upload path.	6.8	More Details
CVE-2025-14728	Rapid7 Velociraptor versions before 0.75.6 contain a directory traversal issue on Linux servers that allows a rogue client to upload a file which is written outside the datastore directory. Velociraptor is normally only allowed to write in the datastore directory. The issue occurs due to insufficient sanitization of directory names which end with a ".", only encoding the final "." AS "%2E". Although files can be written to incorrect locations, the containing directory must end with "%2E". This limits the impact of this vulnerability, and prevents it from overwriting critical files.	6.8	More Details
CVE-2025-69257	theshit is a command-line utility that automatically detects and fixes common mistakes in shell commands. Prior to version 0.1.1, the application loads custom Python rules and configuration files from user-writable locations (e.g., `~/.config/theshit/`) without validating ownership or permissions when executed with elevated privileges. If the tool is invoked with `sudo` or otherwise runs with an effective UID of root, it continues to trust configuration files originating from the unprivileged user's environment. This allows a local attacker to inject arbitrary Python code via a malicious rule or configuration file, which is then executed with root privileges. Any system where this tool is executed with elevated privileges is affected. In environments where the tool is permitted to run via `sudo` without a password (`NOPASSWD`), a local unprivileged user can escalate privileges to root without additional interaction. The issue has been fixed in version 0.1.1. The patch introduces strict ownership and permission checks for all configuration files and custom rules. The application now enforces that rules are only loaded if they are owned by the effective user executing the tool. When executed with elevated privileges (`EUID=0`), the application refuses to load any files that are not owned by root or that are writable by non-root users. When executed as a non-root user, it similarly refuses to load rules owned by other users. This prevents	6.7	More Details

	both vertical and horizontal privilege escalation via execution of untrusted code. If upgrading is not possible, users should avoid executing the application with `sudo` or as the root user. As a temporary mitigation, ensure that directories containing custom rules and configuration files are owned by root and are not writable by non-root users. Administrators may also audit existing custom rules before running the tool with elevated privileges.		
CVE-2025-36192	IBM DS8A00(R10.1) 10.10.106.0 and IBM DS8A00 (R10.0) 10.1.3.010.2.45.0 and IBM DS8900F (R9.4) 89.40.83.089.42.18.089.44.5.0 IBM System Storage DS8000 could allow a local user with authorized CCW update permissions to delete or corrupt backups due to missing authorization in IBM Safeguarded Copy / GDPS Logical corruption protection mechanisms.	6.7	More Details
CVE-2025-59888	Improper quotation in search paths in the Eaton UPS Companion software installer could lead to arbitrary code execution of an attacker with the access to the file system. This security issue has been fixed in the latest version of EUC which is available on the Eaton download center.	6.7	More Details
CVE-2025-60458	UxPlay 1.72 contains a double free vulnerability in its RTSP request handling. A specially crafted RTSP TEARDOWN request can trigger multiple calls to free() on the same memory address, potentially causing a Denial of Service.	6.5	More Details
CVE-2025-69020	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Tribulant Software Newsletters newsletters-lite allows Stored XSS.This issue affects Newsletters: from n/a through <= 4.12.	6.5	More Details
CVE-2025-69017	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Magnigenie RestroPress restropress allows Stored XSS.This issue affects RestroPress: from n/a through <= 3.2.4.2.	6.5	More Details
CVE-2025-69018	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Shamalli Web Directory Free web-directory-free allows DOM-Based XSS.This issue affects Web Directory Free: from n/a through <= 1.7.12.	6.5	More Details
CVE-2025-69019	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in FlippingBook FlippingBook flippingbook allows DOM-Based XSS.This issue affects FlippingBook: from n/a through <= 2.0.1.	6.5	More Details
CVE-2025-66094	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Yada Wiki yada-wiki allows Stored XSS.This issue affects Yada Wiki: from n/a through 3.5.	6.5	More Details
CVE-2024-42718	A path traversal vulnerability in Croogo CMS 4.0.7 allows remote attackers to read arbitrary files via a specially crafted path in the 'edit-file' parameter.	6.5	More Details
CVE-2025-62746	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in CodeFlavors Featured Video for WordPress & VideographyWP allows Stored XSS.This issue affects Featured Video for WordPress & VideographyWP: from n/a through 1.0.18.	6.5	More Details
CVE-2025-63027	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Webcreations907 WBC907 Core allows Stored XSS.This issue affects WBC907 Core: from n/a through 3.4.1.	6.5	More Details
CVE-2025-64190	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in 8theme.Com XStore Core allows DOM-Based XSS.This issue affects XStore Core: from n/a before 5.6.	6.5	More Details
CVE-2025-60935	An open redirect vulnerability in the login endpoint of Blitz Panel v1.17.0 allows attackers to redirect users to malicious domains via a crafted URL. This issue affects the next_url parameter in the login endpoint and could lead to phishing or token theft after successful authentication.	6.5	More Details
CVE-2025-66103	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Revmakx WPCal.Io allows DOM-Based XSS.This issue affects WPCal.Io: from n/a through 0.9.5.9.	6.5	More Details
CVE-2025-14178	In PHP versions:8.1.* before 8.1.34, 8.2.* before 8.2.30, 8.3.* before 8.3.29, 8.4.* before 8.4.16, 8.5.* before 8.5.1, a heap buffer overflow occurs in array_merge() when the total element count of packed arrays exceeds 32-bit limits or HT_MAX_SIZE, due to an integer overflow in the precomputation of element counts using zend_hash_num_elements(). This may lead to memory corruption or crashes and affect the integrity and availability of the target server.	6.5	More Details
CVE-2022-50696	SOUND4 IMPACT/FIRST/PULSE/Eco versions 2.x and below contain hardcoded credentials embedded in server binaries that cannot be modified through normal device operations. Attackers can leverage these static credentials to gain unauthorized access to the device across Linux and Windows distributions without requiring user interaction.	6.5	More Details
CVE-2022-50804	JM-DATA ONU JF511-TV version 1.0.67 is vulnerable to cross-site request forgery (CSRF) attacks, allowing attackers to perform administrative actions on behalf of authenticated users without their knowledge or consent.	6.5	More Details
CVE-2025-68992	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in xenioushk BWL Knowledge Base Manager bwl-kb-manager allows Stored XSS.This issue affects BWL Knowledge Base Manager: from n/a through <= 1.6.3.	6.5	More Details
CVE-2025-67013	The web management interface in ETL Systems Ltd DEXTRA Series ' Digital L-Band Distribution System v1.8 does not implement Cross-Site Request Forgery (CSRF) protection mechanisms (no tokens, no Origin/Referer validation) on critical configuration endpoints.	6.5	More Details
CVE-2019-25257	LogicalDOC Enterprise 7.7.4 contains multiple authenticated OS command execution vulnerabilities that allow attackers to manipulate binary paths when changing system settings. Attackers can exploit these vulnerabilities by modifying configuration parameters like antivirus.command, ocr.Tesseract.path, and other system paths to execute arbitrary system commands with	6.5	More Details

	elevated privileges.		
CVE-2024-39037	MyNET up to v26.08.316 was discovered to contain an Unauthenticated SQL Injection vulnerability via the intmenu parameter.	6.5	More Details
CVE-2025-68607	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Hiroaki Miyashita Custom Field Template allows Stored XSS.This issue affects Custom Field Template: from n/a through 2.7.5.	6.5	More Details
CVE-2025-68914	Riello UPS NetMan 208 Application before 1.12 allows cgi-bin/login.cgi username SQL Injection. For example, an attacker can delete the LOGINFAILEDTABLE table.	6.5	More Details
CVE-2025-68868	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Codeaffairs Wp Text Slider Widget allows Stored XSS.This issue affects Wp Text Slider Widget: from n/a through 1.0.	6.5	More Details
CVE-2025-68504	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Crocoblock JetSearch allows DOM-Based XSS.This issue affects JetSearch: from n/a through 3.5.16.	6.5	More Details
CVE-2019-25256	VideoFlow Digital Video Protection DVP 2.10 contains an authenticated directory traversal vulnerability that allows attackers to access arbitrary system files through unvalidated 'ID' parameters. Attackers can exploit multiple Perl scripts like downloadsys.pl to read sensitive files by manipulating directory path traversal in download requests.	6.5	More Details
CVE-2018-25146	Microhard Systems IPn4G 1.1.0 contains an undocumented vulnerability that allows authenticated attackers to list and manipulate running system processes. Attackers can send arbitrary signals to kill background processes and system services through a hidden feature, potentially causing service disruption and requiring device restart.	6.5	More Details
CVE-2018-25145	Microhard Systems IPn4G 1.1.0 contains a configuration file disclosure vulnerability that allows authenticated attackers to download sensitive system configuration files. Attackers can retrieve configuration files from multiple directories including '/www', '/etc/m_cli/', and '/tmp' to access system passwords and network settings.	6.5	More Details
CVE-2025-68503	Missing Authorization vulnerability in Crocoblock JetBlog allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects JetBlog: from n/a through 2.4.7.	6.5	More Details
CVE-2023-40679	Missing Authorization vulnerability in Jewel Theme Master Addons for Elementor allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Master Addons for Elementor: from n/a through 2.0.5.3.	6.5	More Details
CVE-2025-68040	Insertion of Sensitive Information Into Sent Data vulnerability in weDevs WP Project Manager wedevs-project-manager allows Retrieve Embedded Sensitive Data.This issue affects WP Project Manager: from n/a through 3.0.1.	6.5	More Details
CVE-2025-68498	Missing Authorization vulnerability in Crocoblock JetTabs allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects JetTabs: from n/a through 2.2.12.	6.5	More Details
CVE-2025-68499	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Crocoblock JetTabs allows DOM-Based XSS.This issue affects JetTabs: from n/a through 2.2.12.	6.5	More Details
CVE-2025-68431	libheif is an HEIF and AVIF file format decoder and encoder. Prior to version 1.21.0, a crafted HEIF that exercises the overlay image item path triggers a heap buffer over-read in `HeifPixelImage::overlay()`. The function computes a negative row length (likely from an unclipped overlay rectangle or invalid offsets), which then underflows when converted to `size_t` and is passed to `memcpy`, causing a very large read past the end of the source plane and a crash. Version 1.21.0 contains a patch. As a workaround, avoid decoding images using `iovl` overlay boxes.	6.5	More Details
CVE-2025-66947	SQL injection vulnerability in krishanmuraiji SMS v.1.0, within the /studentms/admin/edit-class-detail.php via the editid GET parameter. An attacker can trigger controlled delays using SQL SLEEP() to infer database contents. Successful exploitation may lead to full database compromise, especially within an administrative module.	6.5	More Details
CVE-2025-68936	ONLYOFFICE Docs before 9.2.1 allows XSS via the Color theme name. This is related to DocumentServer.	6.4	More Details
CVE-2025-68917	ONLYOFFICE Docs before 9.2.1 allows XSS in the textarea of the comment editing form. This is related to DocumentServer.	6.4	More Details
CVE-2025-68935	ONLYOFFICE Docs before 9.2.1 allows XSS via the Font field for the Multilevel list settings window. This is related to DocumentServer.	6.4	More Details
CVE-2025-15152	A vulnerability was identified in h-moses moga-mall up to 392d631a5ef15962a9bddeeb9f1269b9085473fa. This vulnerability affects the function addProduct of the file src/main/java/com/ms/product/controller/PmsProductController.java. Such manipulation of the argument objectName leads to unrestricted upload. The attack may be performed from remote. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed.	6.3	More Details
CVE-2025-15139	A vulnerability has been found in TRENDnet TEW-822DRE 1.00B21/1.01B06. This affects the function sub_43ACF4 of the file /boafrm/formWsc. Such manipulation of the argument peerPin leads to command injection. The attack can be executed remotely. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details

CVE-2025-69205	Micro Registration Utility (µURU) is a telephone self registration utility based on asterisk. In versions up to and including commit 88db9a953f38a3026bcd6816d51c7f3b93c55893, an attacker can crafts a special federation name and characters treated special by asterisk can be injected into the `Dial( )` application due to improper input validation. This allows an attacker to redirect calls on both of the federating instances. If the attack succeeds, the impact is very high. However, the requires that an admin accept the federation requests. As of time of publication, a known patched version of µURU is not available.	6.3	More Details
CVE-2025-15088	A vulnerability was detected in ketr JEPaaS up to 7.2.8. Affected by this vulnerability is the function postilService.loadPostils of the file /je/postil/postil/loadPostil. Performing manipulation of the argument keyWord results in sql injection. Remote exploitation of the attack is possible. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2025-15135	A weakness has been identified in joey-zhou xiaozhi-esp32-server-java up to 3.0.0. This impacts the function tryAuthenticateWithCookies of the file AuthenticationInterceptor.java of the component Cookie Handler. Executing manipulation can lead to improper authentication. The attack can be launched remotely. The exploit has been made available to the public and could be exploited. Upgrading to version 4.0.0 will fix this issue. It is recommended to upgrade the affected component.	6.3	More Details
CVE-2025-15081	A vulnerability has been found in JD Cloud BE6500 4.4.1.r4308. This issue affects the function sub_4780 of the file /jdcapi. Such manipulation of the argument ddns_name leads to command injection. The attack may be performed from remote. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2025-15199	A security vulnerability has been detected in code-projects College Notes Uploading System 1.0. Impacted is an unknown function of the file /dashboard/userprofile.php. The manipulation of the argument image leads to unrestricted upload. Remote exploitation of the attack is possible. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2025-15065	Exposure of Sensitive Information to an Unauthorized Actor, Missing Encryption of Sensitive Data, Files or Directories Accessible to External Parties vulnerability in Kings Information & Network Co. KESS Enterprise on Windows allows Privilege Escalation, Modify Existing Service, Modify Shared File.This issue affects KESS Enterprise: before *.25.9.19.Exe.	6.3	More Details
CVE-2025-15131	A vulnerability was found in ZSPACE Z4Pro+ 1.0.0440024. Impacted is the function zfilev2_api_SafeStatus of the file /v2/file/safe/status of the component HTTP POST Request Handler. The manipulation results in command injection. The attack may be performed from remote. The exploit has been made public and could be used. The vendor was contacted early about this disclosure.	6.3	More Details
CVE-2025-15133	A vulnerability was identified in ZSPACE Z4Pro+ 1.0.0440024. The impacted element is the function zfilev2_api_CloseSafe of the file /v2/file/safe/close of the component HTTP POST Request Handler. Such manipulation leads to command injection. It is possible to launch the attack remotely. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure.	6.3	More Details
CVE-2025-15357	A vulnerability was found in D-Link DI-7400G+ 19.12.25A1. This affects an unknown function of the file /msp_info.htm?flag=cmd. The manipulation of the argument cmd results in command injection. The attack can be launched remotely. The exploit has been made public and could be used.	6.3	More Details
CVE-2025-15132	A vulnerability was determined in ZSPACE Z4Pro+ 1.0.0440024. The affected element is the function zfilev2_api_open of the file /v2/file/safe/open of the component HTTP POST Request Handler. This manipulation causes command injection. It is possible to initiate the attack remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure.	6.3	More Details
CVE-2025-15129	A flaw has been found in ChenJinchuang Lin-CMS-TP5 up to 0.3.3. This vulnerability affects the function Upload of the file application/lib/file/LocalUploader.php of the component File Upload Handler. Executing manipulation of the argument File can lead to code injection. The attack can be executed remotely. The exploit has been published and may be used. The project was informed of the problem early through an issue report but has not responded yet.	6.3	More Details
CVE-2025-15106	A weakness has been identified in getmaxun maxun up to 0.0.28. The affected element is the function router.get of the file server/src/routes/auth.ts of the component Authentication Endpoint. Executing manipulation can lead to improper authorization. The attack can be executed remotely. The exploit has been made available to the public and could be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2025-15209	A weakness has been identified in code-projects Refugee Food Management System 1.0. This affects an unknown part of the file /home/editfood.php. This manipulation of the argument a/b/c/d causes sql injection. The attack may be initiated remotely. The exploit has been made available to the public and could be exploited.	6.3	More Details
CVE-2025-15191	A weakness has been identified in D-Link DWR-M920 up to 1.1.50. The affected element is the function sub_4155B4 of the file /boafrm/formLtefotaUpgradeFibocom. This manipulation of the argument fota_url causes command injection. Remote exploitation of the attack is possible. The exploit has been made available to the public and could be exploited.	6.3	More Details
CVE-2025-15098	A vulnerability was determined in YunaiV yudao-cloud up to 2025.11. This affects the function BpmHttpCallbackTrigger/BpmSyncHttpRequestTrigger of the component Business Process Management. Executing manipulation of the argument url/header/body can lead to server-side request forgery. The attack may be performed from remote. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	6.3	More Details
CVE-2025-15210	A security vulnerability has been detected in code-projects Refugee Food Management System 1.0. This vulnerability affects unknown code of the file /home/editrefugee.php. Such manipulation of the argument a/b/c/sex/d/e/nationality_nid leads to sql injection. The attack may be launched remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2025-15254	A vulnerability was found in Tenda W6-S 1.0.0.4(510). This affects the function TendaAte of the file /goform/ate of the component ATE Service. Performing manipulation results in os command injection. The attack may be initiated remotely. The exploit has been made public and could be used.	6.3	More Details
CVE-2025-15192	A security vulnerability has been detected in D-Link DWR-M920 up to 1.1.50. The impacted element is the function sub_415328 of the file /boafrm/formLtefotaUpgradeQuectel. Such manipulation of the argument fota_url leads to command injection. The attack can be executed remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2025-	A vulnerability was determined in aizuda snail-job up to 1.7.0 on macOS. Affected by this vulnerability is the function FurySerializer.deserialize of the component API. This manipulation of the argument argsStr causes deserialization. Remote	6.3	More

15246	exploitation of the attack is possible. The exploit has been publicly disclosed and may be utilized.		Details
CVE-2025-15211	A flaw has been found in code-projects Refugee Food Management System 1.0. Impacted is an unknown function of the file /home/refugee.php. Executing manipulation of the argument refNo/Fname/Lname/sex/age/contact/nationality_nid can lead to sql injection. The attack can be executed remotely. The exploit has been published and may be used.	6.3	More Details
CVE-2025-15212	A vulnerability was detected in code-projects Refugee Food Management System 1.0. This issue affects some unknown processing of the file /home/regfood.php. Performing manipulation of the argument a results in sql injection. Remote exploitation of the attack is possible. The exploit is now public and may be used.	6.3	More Details
CVE-2025-15205	A vulnerability was identified in code-projects Student File Management System 1.0. Affected by this vulnerability is an unknown functionality of the file /download.php. The manipulation of the argument istore_id leads to sql injection. The attack can be initiated remotely. The exploit is publicly available and might be used.	6.3	More Details
CVE-2025-15050	A security vulnerability has been detected in code-projects Student File Management System 1.0. This affects an unknown part of the file /save_file.php. Such manipulation of the argument File leads to unrestricted upload. The attack can be executed remotely. The exploit has been disclosed publicly and may be used.	6.3	More Details
CVE-2025-15066	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal'), Missing Authorization vulnerability in Innorix WP allows Path Traversal.This issue affects Innorix WP from All versions If the "exam" directory exists under the directory where the product is installed (ex: innorix/exam)	6.2	More Details
CVE-2025-36154	IBM Concert 1.0.0 through 2.1.0 stores sensitive information in cleartext during recursive docker builds which could be obtained by a local user.	6.2	More Details
CVE-2018-25130	Beward Intercom 2.3.1 contains a credentials disclosure vulnerability that allows local attackers to access plain-text authentication credentials stored in an unencrypted database file. Attackers can read the BEWARD.INTERCOM.FDB file to extract usernames and passwords, enabling unauthorized access to IP cameras and door stations.	6.2	More Details
CVE-2024-29720	An issue in Terra Informatica Software, Inc Sciter v.4.4.7.0 allows a local attacker to obtain sensitive information via the adopt component of the Sciter video rendering function.	6.2	More Details
CVE-2025-68509	URL Redirection to Untrusted Site ('Open Redirect') vulnerability in Jeff Starr User Submitted Posts user-submitted-posts allows Phishing.This issue affects User Submitted Posts: from n/a through <= 20251121.	6.1	More Details
CVE-2025-68602	URL Redirection to Untrusted Site ('Open Redirect') vulnerability in Scott Paterson Accept Donations with PayPal easy-paypal-donation allows Phishing.This issue affects Accept Donations with PayPal: from n/a through <= 1.5.1.	6.1	More Details
CVE-2024-35322	MyNET up to v26.08 was discovered to contain a reflected cross-site scripting (XSS) vulnerability via the ficheiro parameter.	6.1	More Details
CVE-2024-40317	A reflected cross-site scripting (XSS) vulnerability in MyNET up to v26.08 allows attackers to execute arbitrary code in the context of a user's browser via injecting a crafted payload into the parameter HTTP.	6.1	More Details
CVE-2025-57462	Stored cross-site scripting (xss) in machsol machpanel 8.0.32 allows attackers to execute arbitrary web scripts or HTML via a crafted PDF file.	6.1	More Details
CVE-2025-65442	DOM-based Cross-Site Scripting (XSS) vulnerability in 201206030 novel V3.5.0 allows remote attackers to execute arbitrary JavaScript code or disclose sensitive information (e.g., user session cookies) via a crafted "wvstest" parameter in the URL or malicious script injection into window.localStorage. The vulnerability arises from insufficient validation and encoding of user-controllable data in the book comment module: unfiltered user input is stored in the backend database (book_comment table, commentContent field) and returned via API, then rendered directly into the page DOM via Vue 3's v-html directive without sanitization. Even if modern browsers' built-in XSS filters block pop-up alerts, attackers can use concealed payloads to bypass interception and achieve actual harm.	6.1	More Details
CVE-2025-67349	A cross-site scripting (XSS) vulnerability was identified in FluentCMS 1.2.3. After logging in as an admin and navigating to the "Add Page" function, the application fails to properly sanitize input in the <head> section, allowing remote attackers to inject arbitrary script tags.	6.1	More Details
CVE-2025-68991	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in xenioushk BWL Pro Voting Manager bwl-pro-voting-manager allows DOM-Based XSS.This issue affects BWL Pro Voting Manager: from n/a through <= 1.4.9.	6.1	More Details
CVE-2025-68978	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in designthemes DesignThemes Core designthemes-core allows DOM-Based XSS.This issue affects DesignThemes Core: from n/a through <= 1.6.	6.1	More Details
CVE-2025-68977	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in designthemes DesignThemes Portfolio Addon designthemes-portfolio-addon allows DOM-Based XSS.This issue affects DesignThemes Portfolio Addon: from n/a through <= 1.5.	6.1	More Details
CVE-2025-55060	CWE-601 URL Redirection to Untrusted Site ('Open Redirect')	6.1	More Details
CVE-2025-68574	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in voidcoders WPBakery Visual Composer WHMCS Elements void-visual-whmcs-element allows DOM-Based XSS.This issue affects WPBakery Visual Composer WHMCS Elements: from n/a through <= 1.0.4.3.	6.1	More Details

CVE-2025-15355	ISOinsight developed by NetVision Information has a Reflected Cross-site Scripting vulnerability, allowing unauthenticated remote attackers to execute arbitrary JavaScript codes in user's browser through phishing attacks.	6.1	More Details
CVE-2025-67632	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in The Plugin Factory Google AdSense for Responsive Design – GARD google-adsense-for-responsive-design-gard allows DOM-Based XSS.This issue affects Google AdSense for Responsive Design – GARD: from n/a through <= 2.23.	6.1	More Details
CVE-2025-67633	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in brownbagmarketing Greenhouse Job Board greenhouse-job-board allows DOM-Based XSS.This issue affects Greenhouse Job Board: from n/a through <= 2.7.3.	6.1	More Details
CVE-2025-14313	The Advance WP Query Search Filter WordPress plugin through 1.0.10 does not sanitise and escape a parameter before outputting it back in the page, leading to a Reflected Cross-Site Scripting which could be used against high privilege users such as admin	6.1	More Details
CVE-2022-50802	ETAP Safety Manager 1.0.0.32 contains a cross-site scripting vulnerability in the 'action' GET parameter that allows unauthenticated attackers to inject malicious HTML and JavaScript. Attackers can craft specially formed requests to execute arbitrary scripts in victim browser sessions, potentially stealing credentials or performing unauthorized actions.	6.1	More Details
CVE-2025-14312	The Advance WP Query Search Filter WordPress plugin through 1.0.10 does not sanitise and escape a parameter before outputting it back in the page, leading to a Reflected Cross-Site Scripting which could be used against high privilege users such as admin	6.1	More Details
CVE-2025-13958	The YaMaps for WordPress Plugin WordPress plugin before 0.6.40 does not validate and escape some of its shortcode attributes before outputting them back in a page/post where the shortcode is embed, which could allow users with the contributor role and above to perform Stored Cross-Site Scripting attacks.	5.9	More Details
CVE-2025-69006	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Atte Moisio AM Events am-events allows Stored XSS.This issue affects AM Events: from n/a through <= 1.13.1.	5.9	More Details
CVE-2025-69007	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in OTWthemes Popping Sidebars and Widgets Light popping-sidebars-and-widgets-light allows Stored XSS.This issue affects Popping Sidebars and Widgets Light: from n/a through <= 1.27.	5.9	More Details
CVE-2023-32120	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Bob Hostel allows DOM-Based XSS.This issue affects Hostel: from n/a through 1.1.5.1.	5.9	More Details
CVE-2025-69008	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Inboxify Inboxify Sign Up Form inboxify-sign-up-form allows Stored XSS.This issue affects Inboxify Sign Up Form: from n/a through <= 1.0.4.	5.9	More Details
CVE-2025-68972	In GnuPG through 2.4.8, if a signed message has \f at the end of a plaintext line, an adversary can construct a modified message that places additional text after the signed material, such that signature verification of the modified message succeeds (although an "invalid armor" message is printed during verification). This is related to use of \f as a marker to denote truncation of a long plaintext line.	5.9	More Details
CVE-2025-49088	Pexip Infinity 32.0 through 37.1 before 37.2, in certain configurations of OTJ (One Touch Join) for Teams SIP Guest Join, has Improper Input Validation in the OTJ service, allowing a remote attacker to trigger a software abort via a crafted calendar invite, leading to a denial of service.	5.9	More Details
CVE-2025-66378	Pexip Infinity 38.0 and 38.1 before 39.0 has insufficient access control in the RTMP implementation, allowing an attacker to disconnect RTMP streams traversing a Proxy Node.	5.9	More Details
CVE-2025-1721	IBM Concert 1.0.0 through 2.1.0 could allow a remote attacker to obtain sensitive information from allocated memory due to improper clearing of heap memory.	5.9	More Details
CVE-2025-68945	In Gitea before 1.21.2, an anonymous user can visit a private user's project.	5.8	More Details
CVE-2025-15251	A vulnerability was detected in beecue FastBee up to 2.1. Impacted is the function getRootElement of the file springboot/fastbee-server/sip-server/src/main/java/com/fastbee/sip/handler/req/ReqAbstractHandler.java of the component SIP Message Handler. The manipulation results in xml external entity reference. It is possible to launch the attack remotely. A high complexity level is associated with this attack. The exploitability is considered difficult. The project owner replied to the issue report: "Okay, we'll handle it as soon as possible."	5.6	More Details
CVE-2025-68919	Fujitsu / Fsas Technologies ETERNUS SF ACM/SC/Express (DX / AF Management Software) before 16.8-16.9.1 PA 2025-12, when collected maintenance data is accessible by a principal/authority other than ETERNUS SF Admin, allows an attacker to potentially affect system confidentiality, integrity, and availability.	5.6	More Details
CVE-2025-15070	Exposure of Sensitive Information to an Unauthorized Actor, Missing Authorization vulnerability in Gmission Web Fax allows Authentication Abuse.This issue affects Web Fax: from 3.0 before 4.0.	5.5	More Details
CVE-2018-25144	Microhard Systems IPn4G 1.1.0 contains an authentication bypass vulnerability in the hidden system-editor.sh script that allows authenticated attackers to read, modify, or delete arbitrary files. Attackers can exploit unsanitized 'path', 'savefile', 'edit', and 'delfile' parameters to perform unauthorized file system modifications through GET and POST requests.	5.5	More Details
CVE-2025-	Riello UPS NetMan 208 Application before 1.12 allows cgi-bin/loginbanner_w.cgi XSS via a crafted banner.	5.5	More

68915			Details
CVE-2023-32238	Vulnerability in CodexThemes TheGem (Elementor), CodexThemes TheGem (WPBakery).This issue affects TheGem (Elementor): from n/a before 5.8.1.1; TheGem (WPBakery): from n/a before 5.8.1.1.	5.4	More Details
CVE-2025-68512	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in creativeinteractivemedia Real 3D FlipBook real3d-flipbook-lite allows Stored XSS.This issue affects Real 3D FlipBook: from n/a through <= 4.11.4.	5.4	More Details
CVE-2025-68513	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in boldthemes Bold Timeline Lite bold-timeline-lite allows Stored XSS.This issue affects Bold Timeline Lite: from n/a through <= 1.2.7.	5.4	More Details
CVE-2025-68599	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Embeds For YouTube Plugin Support YouTube Embed youtube-embed allows Stored XSS.This issue affects YouTube Embed: from n/a through <= 5.4.	5.4	More Details
CVE-2025-68120	To prevent unexpected untrusted code execution, the Visual Studio Code Go extension is now disabled in Restricted Mode.	5.4	More Details
CVE-2025-68525	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in pixelgrade Category Icon category-icon allows Stored XSS.This issue affects Category Icon: from n/a through <= 1.0.2.	5.4	More Details
CVE-2025-68527	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Kodezen LLC Academy LMS academy allows Stored XSS.This issue affects Academy LMS: from n/a through <= 3.4.0.	5.4	More Details
CVE-2025-68528	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WPFactory Free Shipping Bar: Amount Left for Free Shipping for WooCommerce amount-left-free-shipping-woocommerce allows Stored XSS.This issue affects Free Shipping Bar: Amount Left for Free Shipping for WooCommerce: from n/a through <= 2.4.9.	5.4	More Details
CVE-2025-68533	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in HasThemes WC Builder wc-builder allows Stored XSS.This issue affects WC Builder: from n/a through <= 1.2.0.	5.4	More Details
CVE-2025-68532	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in modeltheme ModelTheme Addons for WPBakery and Elementor modeltheme-addons-for-wpbakery allows Stored XSS.This issue affects ModelTheme Addons for WPBakery and Elementor: from n/a through < 1.5.6.	5.4	More Details
CVE-2025-2154	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Echo Call Center Services Trade and Industry Inc. Specto CM allows Stored XSS.This issue affects Specto CM: before 17032025.	5.4	More Details
CVE-2025-68942	Gitea before 1.22.2 allows XSS because the search input box (for creating tags and branches) is v-html instead of v-text.	5.4	More Details
CVE-2025-68566	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in wphocus My auctions allegro my-auctions-allegro-free-edition allows Stored XSS.This issue affects My auctions allegro: from n/a through <= 3.6.32.	5.4	More Details
CVE-2025-68598	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in LiveComposer Page Builder: Live Composer live-composer-page-builder allows Stored XSS.This issue affects Page Builder: Live Composer: from n/a through <= 2.0.5.	5.4	More Details
CVE-2025-68946	In Gitea before 1.20.1, a forbidden URL scheme such as javascript: can be used for a link, aka XSS.	5.4	More Details
CVE-2025-68951	phpMyFAQ is an open source FAQ web application. Versions 4.0.14 and 4.0.15 have a stored cross-site scripting (XSS) vulnerability that allows an attacker to execute arbitrary JavaScript in an administrator’s browser by registering a user whose display name contains HTML entities. When an administrator views the admin user list, the payload is decoded server-side and rendered without escaping, resulting in script execution in the admin context. Version 4.0.16 contains a patch for the issue.	5.4	More Details
CVE-2025-68605	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in PickPlugins Post Grid and Gutenberg Blocks post-grid allows Stored XSS.This issue affects Post Grid and Gutenberg Blocks: from n/a through <= 2.3.18.	5.4	More Details
CVE-2025-36230	IBM Aspera Faspex 5 5.0.0 through 5.0.14.1 is vulnerable to HTML injection. A remote attacker could inject malicious HTML code, which when viewed, would be executed in the victim's Web browser within the security context of the hosting site.	5.4	More Details
CVE-2025-68928	Frappe CRM is an open-source customer relationship management tool. Prior to version 1.56.2, authenticated users could set crafted URLs in a website field, which were not sanitized, causing cross-site scripting. Version 1.56.2 fixes the issue. No known workarounds are available.	5.4	More Details
CVE-2025-68597	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in BlueGlass Interactive AG Jobs for WordPress job-postings allows Stored XSS.This issue affects Jobs for WordPress: from n/a through <= 2.7.17.	5.4	More Details
CVE-	Missing Authorization vulnerability in wpdive Better Elementor Addons allows Exploiting Incorrectly Configured Access Control		More

2023-41656	Security Levels.This issue affects Better Elementor Addons: from n/a through 1.3.7.	5.4	Details
CVE-2025-68497	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Brainstorm Force Astra Widgets astra-widgets allows Stored XSS.This issue affects Astra Widgets: from n/a through <= 1.2.16.	5.4	More Details
CVE-2025-67631	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Ecommerce Platforms Gift Hunt gift-hunt allows Stored XSS.This issue affects Gift Hunt: from n/a through <= 2.0.2.	5.4	More Details
CVE-2025-67627	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in TouchOfTech Draft Notify draft-notify allows Stored XSS.This issue affects Draft Notify: from n/a through <= 1.5.	5.4	More Details
CVE-2025-67628	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in AMP-MODE Review Disclaimer review-disclaimer allows Stored XSS.This issue affects Review Disclaimer: from n/a through <= 2.0.3.	5.4	More Details
CVE-2025-67630	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in webheadcoder WH Tweaks wh-tweaks allows Stored XSS.This issue affects WH Tweaks: from n/a through <= 1.0.2.	5.4	More Details
CVE-2025-67629	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Basticom Basticom Framework basticom-framework allows Stored XSS.This issue affects Basticom Framework: from n/a through <= 1.5.2.	5.4	More Details
CVE-2025-68998	Cross-Site Request Forgery (CSRF) vulnerability in Heateor Support Heateor Social Login heateor-social-login allows Cross Site Request Forgery.This issue affects Heateor Social Login: from n/a through <= 1.1.39.	5.4	More Details
CVE-2019-25250	Devolo dLAN 500 AV Wireless+ 3.1.0-1 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without proper request validation. Attackers can craft malicious web pages that trigger unauthorized configuration changes by exploiting predictable URL actions when a logged-in user visits the site.	5.3	More Details
CVE-2025-68943	Gitea before 1.21.8 inadvertently discloses users' login times by allowing (for example) the lastlogintime explore/users sort order.	5.3	More Details
CVE-2025-68997	Authorization Bypass Through User-Controlled Key vulnerability in AdvancedCoding wpDiscuz wpdiscuz allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects wpDiscuz: from n/a through <= 7.6.40.	5.3	More Details
CVE-2019-25251	Teradek VidiU Pro 3.0.3 contains a server-side request forgery vulnerability in the management interface that allows attackers to manipulate GET parameters 'url' and 'xml_url'. Attackers can exploit this flaw to bypass firewalls, initiate network enumeration, and potentially trigger external HTTP requests to arbitrary destinations.	5.3	More Details
CVE-2019-25247	Beward N100 H.264 VGA IP Camera M2.1.6 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without proper request validation. Attackers can craft a malicious web page with a hidden form to add an admin user by tricking a logged-in user into submitting the form.	5.3	More Details
CVE-2019-25244	Legrand BTicino Driver Manager F454 1.0.51 contains multiple web vulnerabilities that allow attackers to perform administrative actions without proper request validation. Attackers can exploit cross-site request forgery to change passwords and inject stored cross-site scripting payloads through unvalidated GET parameters.	5.3	More Details
CVE-2025-66080	Missing Authorization vulnerability in WP Legal Pages WP Cookie Notice for GDPR, CCPA & ePrivacy Consent allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects WP Cookie Notice for GDPR, CCPA & ePrivacy Consent: from n/a through 4.0.3.	5.3	More Details
CVE-2025-69009	Missing Authorization vulnerability in kamleshyadav Medicalequipment medicalequipment allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Medicalequipment: from n/a through <= 1.0.9.	5.3	More Details
CVE-2025-69010	Missing Authorization vulnerability in themebeez Themebeez Toolkit themebeez-toolkit allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Themebeez Toolkit: from n/a through <= 1.3.5.	5.3	More Details
CVE-2025-15155	A vulnerability was detected in floooh sokol up to 16cbcc864012898793cd2bc57f802499a264ea40. The impacted element is the function _sg_pipeline_desc_defaults in the library sokol_gfx.h. The manipulation results in stack-based buffer overflow. The attack requires a local approach. The exploit is now public and may be used. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The patch is identified as 5d11344150973f15e16d3ec4ee7550a73fb995e0. It is advisable to implement a patch to correct this issue.	5.3	More Details
CVE-2025-15150	A vulnerability was found in PX4 PX4-Autopilot up to 1.16.0. Affected by this issue is the function MavlinkLogHandler::state_listing/MavlinkLogHandler::log_entry_from_id of the file src/modules/mavlink/mavlink_log_handler.cpp. The manipulation results in stack-based buffer overflow. The attack is only possible with local access. The patch is identified as 338595edd1d235efd885fd5e9f45e7f9dcf4013d. It is best practice to apply a patch to resolve this issue.	5.3	More Details
CVE-2025-15154	A security vulnerability has been detected in PbootCMS up to 3.2.12. The affected element is the function get_user_ip of the file core/function/handle.php of the component Header Handler. The manipulation of the argument X-Forwarded-For leads to use of less trusted source. The attack can be initiated remotely. The exploit has been disclosed publicly and may be used.	5.3	More Details
CVE-2019-25234	SmartHouse Webapp 6.5.33 contains multiple cross-site request forgery and cross-site scripting vulnerabilities that allow attackers to perform unauthorized actions. Attackers can exploit these vulnerabilities by tricking logged-in users into visiting malicious websites or injecting malicious scripts into various application parameters.	5.3	More Details

CVE-2025-15176	A flaw has been found in Open5GS up to 2.7.5. This affects the function decode_ipv6_header/ogs_pfcp_pdr_rule_find_by_packet of the file lib/pfcp/rule-match.c of the component PFCP Session Establishment Request Handler. Executing manipulation can lead to reachable assertion. It is possible to launch the attack remotely. The exploit has been published and may be used. This patch is called b72d8349980076e2c033c8324f07747a86eea4f8. Applying a patch is advised to resolve this issue.	5.3	More Details
CVE-2019-25233	AVE DOMINAplus 1.10.x contains cross-site request forgery and cross-site scripting vulnerabilities that allow attackers to perform administrative actions without user consent. Attackers can craft malicious web pages to exploit login.php parameters and execute arbitrary scripts in user browser sessions.	5.3	More Details
CVE-2018-25156	Teradek Cube 7.3.6 contains a cross-site request forgery vulnerability that allows attackers to change administrative passwords without proper request validation. Attackers can craft a malicious web page with a hidden form to submit password change requests to the device's system configuration interface.	5.3	More Details
CVE-2018-25155	Teradek Slice 7.3.15 contains a cross-site request forgery vulnerability that allows attackers to change administrative passwords without proper request validation. Attackers can craft a malicious web page that automatically submits password change requests to the device when a logged-in user visits the page.	5.3	More Details
CVE-2025-15128	A vulnerability was detected in ZKTeco BioTime up to 9.0.3/9.0.4/9.5.2. This affects an unknown part of the file /base/safe_setting/ of the component Endpoint. Performing manipulation of the argument backup_encryption_password_decrypt/export_encryption_password_decrypt results in unprotected storage of credentials. Remote exploitation of the attack is possible. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2018-25152	Ecessa Edge EV150 10.7.4 contains a cross-site request forgery vulnerability that allows attackers to create administrative user accounts without authentication. Attackers can craft a malicious web page with a form that submits requests to the /cgi-bin/pl_web.cgi/util_configlogin_act endpoint to add superuser accounts with arbitrary credentials.	5.3	More Details
CVE-2018-25150	Ecessa ShieldLink SL175EHQ 10.7.4 contains a cross-site request forgery vulnerability that allows attackers to create administrative user accounts without authentication. Attackers can craft a malicious web page with a hidden form to add a superuser account by tricking a logged-in administrator into loading the page.	5.3	More Details
CVE-2025-69204	ImageMagick is free and open-source software used for editing and manipulating digital images. Prior to version 7.1.2-12, in the WriteSVGImage function, using an int variable to store number_attributes caused an integer overflow. This, in turn, triggered a buffer overflow and caused a DoS attack. Version 7.1.2-12 fixes the issue.	5.3	More Details
CVE-2025-68618	ImageMagick is free and open-source software used for editing and manipulating digital images. Prior to version 7.1.2-12, using Magick to read a malicious SVG file resulted in a DoS attack. Version 7.1.2-12 fixes the issue.	5.3	More Details
CVE-2019-25252	Teradek VidiU Pro 3.0.3 contains a cross-site request forgery vulnerability that allows attackers to change administrative passwords without proper request validation. Attackers can craft malicious web pages that automatically submit password change requests to the device when a logged-in administrator visits the page.	5.3	More Details
CVE-2018-25127	SOCA Access Control System 180612 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without proper request validation. Attackers can craft malicious web pages that submit forged requests to create admin accounts by tricking logged-in users into visiting a malicious site.	5.3	More Details
CVE-2019-25254	KYOCERA Net Admin 3.4.0906 contains a cross-site request forgery vulnerability that allows attackers to create administrative users without proper request validation. Attackers can craft malicious web pages that automatically submit forms to add new admin accounts with predefined credentials when a logged-in user visits the page.	5.3	More Details
CVE-2025-15229	A vulnerability has been found in Tenda CH22 up to 1.0.0.1. Affected by this vulnerability is the function fromDhcpListClient of the file /goform/DhcpListClient. Such manipulation of the argument LISTLEN leads to denial of service. The attack may be launched remotely. The exploit has been disclosed to the public and may be used.	5.3	More Details
CVE-2025-15082	A vulnerability was found in TOZED ZLT M30s up to 1.47. Impacted is an unknown function of the file /reqproc/proc_post of the component Web Management Interface. Performing manipulation of the argument goformId results in information disclosure. It is possible to initiate the attack remotely. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way.	5.3	More Details
CVE-2025-14280	The PixelYourSite plugin for WordPress is vulnerable to Sensitive Information Exposure in all versions up to, and including, 11.1.5 through publicly exposed log files. This makes it possible for unauthenticated attackers to view potentially sensitive information contained in the exposed log files, when the "Meta API logs" setting is enabled (disabled by default). The vulnerability was partially patched in version 11.1.5 and fully patched in version 11.1.5.1.	5.3	More Details
CVE-2025-14913	The Frontend Post Submission Manager Lite – Frontend Posting WordPress Plugin plugin for WordPress is vulnerable to unauthorized loss of data due to an incorrect authorization check on the 'media_delete_action' function in all versions up to, and including, 1.2.6. This makes it possible for unauthenticated attackers to delete arbitrary attachments.	5.3	More Details
CVE-2025-53627	Meshtastic is an open source mesh networking solution. The Meshtastic firmware (starting from version 2.5) introduces asymmetric encryption (PKI) for direct messages, but when the `pki_encrypted` flag is missing, the firmware silently falls back to legacy AES-256-CTR channel encryption. This was an intentional decision to maintain backwards compatibility. However, the end-user applications, like Web app, iOS/Android app, and applications built on top of Meshtastic using the SDK, did not have a way to differentiate between end-to-end encrypted DMs and the legacy DMs. This creates a downgrade attack path where adversaries who know a shared channel key can craft and inject spoofed direct messages that are displayed as if they were PKC encrypted. Users are not given any feedback of whether a direct message was decrypted with PKI or with legacy symmetric encryption, undermining the expected security guarantees of the PKI rollout. Version 2.7.15 fixes this issue.	5.3	More Details
CVE-2025-68993	Missing Authorization vulnerability in XforWooCommerce Share, Print and PDF Products for WooCommerce share-print-pdf-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Share, Print and PDF Products for WooCommerce: from n/a through <= 3.1.2.	5.3	More Details
CVE-2025-	Missing Authorization vulnerability in XforWooCommerce Product Loops for WooCommerce product-loops allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Product Loops for WooCommerce: from n/a through <=	5.3	More Details

68994	2.1.2.		
CVE-2025-65885	An issue was discovered in the Delight Custom Firmware (CFW) for Nokia Symbian Belle devices on Nokia 808 (Delight v1.8), Nokia N8 (Delight v6.7), Nokia E7 (Delight v1.3), Nokia C7 (Delight v6.7), Nokia 700 (Delight v1.2), Nokia 701 (Delight v1.1), Nokia 603 (Delight v1.0), Nokia 500 (Delight v1.2), Nokia E6 (Delight v1.0), Nokia Oro (Delight v1.0), and Vertu Constellation T (Delight v1.0) allowing local attackers to inject startup scripts via crafted .txt files in the :\Data directory.	5.1	More Details
CVE-2024-58335	OpenXRechnungToolbox through 2024-10-05-3.0.0 before 6c50e89 allows XXE because the disallow-doctype-decl feature is not enabled in visualization/VisualizerImpl.java.	5.0	More Details
CVE-2025-15222	A vulnerability has been found in Dromara Sa-Token up to 1.44.0. This issue affects the function ObjectInputStream.readObject of the file SaSerializerTemplateForJdkUseBase64.java. Such manipulation leads to deserialization. The attack can be executed remotely. This attack is characterized by high complexity. The exploitability is assessed as difficult. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	5.0	More Details
CVE-2025-68944	Gitea before 1.22.2 sometimes mishandles the propagation of token scope for access control within one of its own package registries.	5.0	More Details
CVE-2025-68893	Server-Side Request Forgery (SSRF) vulnerability in HETWORKS WordPress Image shrinker allows Server Side Request Forgery.This issue affects WordPress Image shrinker: from n/a through 1.1.0.	4.9	More Details
CVE-2025-68941	Gitea before 1.22.3 mishandles access to a private resource upon receiving an API token with scope limited to public resources.	4.9	More Details
CVE-2025-69014	Server-Side Request Forgery (SSRF) vulnerability in Youzify Youzify youzify allows Server Side Request Forgery.This issue affects Youzify: from n/a through <= 1.3.5.	4.9	More Details
CVE-2025-55064	CWE-79 Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting')	4.8	More Details
CVE-2025-55062	CWE-79 Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting')	4.8	More Details
CVE-2025-55063	CWE-79 Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting')	4.8	More Details
CVE-2025-15130	A vulnerability has been found in shanyu SyCms up to a242ef2d194e8bb249dc175e7c49f2c1673ec921. This issue affects the function addPost of the file Application/Admin/Controller/FileManageController.class.php of the component Administrative Panel. The manipulation leads to code injection. The attack is possible to be carried out remotely. The exploit has been disclosed to the public and may be used. This product adopts a rolling release strategy to maintain continuous delivery The project was informed of the problem early through an issue report but has not responded yet. This vulnerability only affects products that are no longer supported by the maintainer.	4.7	More Details
CVE-2025-15143	A security flaw has been discovered in EyouCMS up to 1.7.6. The affected element is an unknown function of the file /application/admin/logic/FilemanagerLogic.php of the component Backend Template Management. The manipulation of the argument content results in sql injection. It is possible to launch the attack remotely. The exploit has been released to the public and may be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	4.7	More Details
CVE-2025-15138	A flaw has been found in prasathmani TinyFileManager up to 2.6. Affected by this issue is some unknown functionality of the file tinyfilemanager.php. This manipulation of the argument fullpath causes path traversal. Remote exploitation of the attack is possible. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.7	More Details
CVE-2025-15110	A vulnerability has been found in jackq XCMS up to 3fab5342cc509945a7ce1b8ec39d19f701b89261. Affected is the function Upload of the file Admin/Home/Controller/ProductImageController.class.php of the component Backend. Such manipulation of the argument File leads to unrestricted upload. It is possible to launch the attack remotely. The exploit has been disclosed to the public and may be used. This product takes the approach of rolling releases to provide continious delivery. Therefore, version details for affected and updated releases are not available. The project was informed of the problem early through an issue report but has not responded yet.	4.7	More Details
CVE-2025-15250	A security vulnerability has been detected in 08CMS Novel System up to 3.4. This issue affects some unknown processing of the file admina/mtpls.inc.php of the component Template Handler. The manipulation leads to code injection. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used.	4.7	More Details
CVE-2025-15262	A security flaw has been discovered in BiggiDroid Simple PHP CMS 1.0. This impacts an unknown function of the file /admin/edit.php of the component Site Logo Handler. Performing manipulation of the argument image results in unrestricted upload. Remote exploitation of the attack is possible. The exploit has been released to the public and may be exploited.	4.7	More Details
CVE-2025-15169	A weakness has been identified in BiggiDroid Simple PHP CMS 1.0. Affected by this issue is some unknown functionality of the file /admin/editsite.php. Executing manipulation of the argument ID can lead to sql injection. The attack may be performed from remote. The exploit has been made available to the public and could be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	4.7	More Details
CVE-2025-15197	A security flaw has been discovered in code-projects/anirbandutta9 Content Management System and News-Buzz 1.0. This vulnerability affects unknown code of the file /admin/editposts.php. Performing manipulation of the argument image results in unrestricted upload. The attack may be initiated remotely. The exploit has been released to the public and may be exploited.	4.7	More Details

CVE-2025-15148	A flaw has been found in CmsEasy up to 7.7.7. Affected is the function savetemp_action in the library /lib/admin/template_admin.php of the component Backend Template Management Page. Executing manipulation of the argument content/tempdata can lead to code injection. The attack may be launched remotely. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.7	More Details
CVE-2025-15360	A vulnerability was determined in newbee-mall-plus 2.0.0. This impacts the function Upload of the file src/main/java/ltd/newbee/mall/controller/common/UploadController.java of the component Product Information Edit Page. This manipulation of the argument File causes unrestricted upload. The attack may be initiated remotely. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	4.7	More Details
CVE-2025-15085	A security flaw has been discovered in youlaitech youlai-mall 1.0.0/2.0.0. This affects the function deductBalance of the file mall-ums/ums-boot/src/main/java/com/youlai/mall/ums/controller/app/MemberController.java of the component Balance Handler. The manipulation results in improper authorization. The attack can be launched remotely. The exploit has been released to the public and may be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2025-68938	Gitea before 1.25.2 mishandles authorization for deletion of releases.	4.3	More Details
CVE-2025-68148	FreshRSS is a free, self-hostable RSS aggregator. From version 1.27.0 to before 1.28.0, An attacker could globally deny access to feeds via proxy modifying to 429 Retry-After for a large list of feeds on given instance, making it unusable for majority of users. This issue has been patched in version 1.28.0.	4.3	More Details
CVE-2019-25255	VideoFlow Digital Video Protection DVP 2.10 contains an authenticated remote code execution vulnerability that allows attackers to execute system commands with root privileges. Attackers can exploit the vulnerability through a cross-site request forgery (CSRF) mechanism to gain unauthorized system access.	4.3	More Details
CVE-2018-25133	Synaccess netBooter NP-0801DU 7.4 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without proper request validation. Attackers can craft malicious web pages with hidden form submissions to add admin users by tricking authenticated administrators into loading a malicious page.	4.3	More Details
CVE-2025-66737	Yealink T21P_E2 Phone 52.84.0.15 is vulnerable to Directory Traversal. A remote normal privileged attacker can read arbitrary files via a crafted request result read function of the diagnostic component.	4.3	More Details
CVE-2025-62112	Cross-Site Request Forgery (CSRF) vulnerability in Merv Barrett Import into Easy Property Listings allows Cross Site Request Forgery.This issue affects Import into Easy Property Listings: from n/a through 2.2.1.	4.3	More Details
CVE-2025-62128	Missing Authorization vulnerability in SiteLock SiteLock Security allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects SiteLock Security: from n/a through 5.0.1.	4.3	More Details
CVE-2025-15093	A security flaw has been discovered in sunkaifei FlyCMS up to abbaa5a8daefb146ad4d61027035026b052cb414. The affected element is an unknown function of the file src/main/java/com/flycms/web/system/IndexAdminController.java of the component Admin Login. Performing manipulation of the argument redirectUrl results in cross site scripting. The attack can be initiated remotely. The exploit has been released to the public and may be exploited. Continious delivery with rolling releases is used by this product. Therefore, no version details of affected nor updated releases are available. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2025-14687	IBM Db2 Intelligence Center 1.1.0, 1.1.1, 1.1.2 could allow an authenticated user to perform unauthorized actions due to client-side enforcement of sever side security mechanisms.	4.3	More Details
CVE-2025-13767	Mattermost versions 11.1.x <= 11.1.0, 11.0.x <= 11.0.5, 10.12.x <= 10.12.3, 10.11.x <= 10.11.7 fails to validate user channel membership when attaching Mattermost posts as comments to Jira issues, which allows an authenticated attacker with access to the Jira plugin to read post content and attachments from channels they do not have access to.	4.3	More Details
CVE-2022-50801	JM-DATA ONU JF511-TV version 1.0.67 is vulnerable to authenticated stored cross-site scripting (XSS) attacks, allowing attackers with authenticated access to inject malicious scripts that will be executed in other users' browsers when they view the affected content.	4.3	More Details
CVE-2019-25242	FaceSentry Access Control System 6.4.8 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without user consent. Attackers can craft malicious web pages to change administrator passwords, add new admin users, or open access control doors by tricking authenticated users into loading a specially crafted webpage.	4.3	More Details
CVE-2018-25149	Microhard Systems IPn4G 1.1.0 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without user consent. Attackers can craft malicious web pages to change admin passwords, add new users, and modify system settings by tricking authenticated users into loading a specially crafted page.	4.3	More Details
CVE-2018-25151	Ecessa WANWorx WVR-30 versions before 10.7.4 contain a cross-site request forgery vulnerability that allows attackers to perform administrative actions without request validation. Attackers can craft a malicious web page with a hidden form to create a new superuser account by tricking an authenticated administrator into loading the page.	4.3	More Details
CVE-2025-15087	A security vulnerability has been detected in youlaitech youlai-mall 1.0.0/2.0.0. Affected is the function submitOrderPayment of the file mall-oms/oms-boot/src/main/java/com/youlai/mall/oms/controller/app/OrderController.java. Such manipulation of the argument orderSn leads to improper authorization. The attack may be launched remotely. The exploit has been disclosed publicly and may be used. The real existence of this vulnerability is still doubted at the moment. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2019-25238	V-SOL GPON/EPON OLT Platform 2.03 contains a cross-site request forgery vulnerability that allows attackers to perform administrative actions without user consent. Attackers can craft malicious web pages to create admin users, enable SSH, or modify system settings by tricking authenticated administrators into loading a specially crafted page.	4.3	More Details

CVE-2025-15086	A weakness has been identified in youlaitech youlai-mall 1.0.0/2.0.0. This impacts the function getMemberByMobile of the file mall-ums/ums-boot/src/main/java/com/youlai/mall/ums/controller/app/MemberController.java. This manipulation causes improper access controls. The attack may be initiated remotely. The exploit has been made available to the public and could be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2025-15094	A weakness has been identified in sunkaifei FlyCMS up to abbaa5a8daefb146ad4d61027035026b052cb414. The impacted element is the function userLogin of the file src/main/java/com/flycms/web/front/UserController.java of the component User Login. Executing manipulation of the argument redirectUrl can lead to cross site scripting. The attack can be launched remotely. The exploit has been made available to the public and could be exploited. This product does not use versioning. This is why information about affected and unaffected releases are unavailable. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2025-69016	Missing Authorization vulnerability in averta Shortcodes and extra features for Phlox theme auxin-elements allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Shortcodes and extra features for Phlox theme: from n/a through <= 2.17.12.	4.3	More Details
CVE-2025-15170	A security vulnerability has been detected in Advaya Softech GEMS ERP Portal up to 2.1. This affects an unknown part of the file /home.jsp?isError=true of the component Error Message Handler. The manipulation of the argument Message leads to cross site scripting. It is possible to initiate the attack remotely. The exploit has been disclosed publicly and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2025-15220	A vulnerability was detected in SohuTV CacheCloud up to 3.2.0. This affects the function init of the file src/main/java/com/sohu/cache/web/controller/LoginController.java. The manipulation results in cross site scripting. The attack may be launched 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.	4.3	More Details
CVE-2025-15213	A vulnerability has been found in code-projects Student File Management System 1.0. The affected element is an unknown function of the file /download.php of the component File Download Handler. The manipulation of the argument store_id leads to improper authorization. The attack is possible to be carried out remotely. The exploit has been disclosed to the public and may be used.	4.3	More Details
CVE-2025-68502	Authorization Bypass Through User-Controlled Key vulnerability in Crocoblock JetPopup allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects JetPopup: from n/a through 2.0.20.1.	4.3	More Details
CVE-2025-15156	A flaw has been found in omec-project UPF up to 2.1.3-dev. This affects the function handleSessionEstablishmentRequest of the file /pfcpiface/pfcpiface/messages_session.go of the component PFCP Session Establishment Request Handler. This manipulation causes null pointer dereference. The attack may be initiated remotely. The exploit has been published and may be used. The project was informed of the problem early through an issue report but has not responded yet.	4.3	More Details
CVE-2025-14426	The Strong Testimonials plugin for WordPress is vulnerable to unauthorized modification of data due to a missing capability check in the 'edit_rating' function in all versions up to, and including, 3.2.18. This makes it possible for authenticated attackers with Contributor-level access and above to modify or delete the rating meta on any testimonial post, including those created by other users, by reusing a valid nonce obtained from their own testimonial edit screen.	4.3	More Details
CVE-2025-15118	A security vulnerability has been detected in macrozheng mall up to 1.0.3. This vulnerability affects unknown code of the file /member/address/update/ of the component Member Endpoint. The manipulation leads to improper authorization. Remote exploitation of the attack is possible. The exploit has been disclosed publicly and may be used.	4.3	More Details
CVE-2025-69206	Hemmelig is a messing app with with client-side encryption and self-destructing messages. Prior to version 7.3.3, a Server-Side Request Forgery (SSRF) filter bypass vulnerability exists in the webhook URL validation of the Secret Requests feature. The application attempts to block internal/private IP addresses but can be bypassed using DNS rebinding or open redirect services. This allows an authenticated user to make the server initiate HTTP requests to internal network resources. Version 7.3.3 contains a patch for the issue.	4.3	More Details
CVE-2025-69013	Missing Authorization vulnerability in jetmonsters Stratum stratum allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Stratum: from n/a through <= 1.6.1.	4.3	More Details
CVE-2025-15144	A weakness has been identified in dayrui XunRuiCMS up to 4.7.1. The impacted element is the function dr_show_error/dr_exit_msg of the file /dayrui/Fcms/Init.php of the component JSONP Callback Handler. This manipulation of the argument callback causes cross site scripting. The attack can be initiated remotely. The exploit has been made available to the public and could be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	4.3	More Details
CVE-2025-68995	Missing Authorization vulnerability in Gal Dubinski My Sticky Elements mystickyelements allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects My Sticky Elements: from n/a through <= 2.3.3.	4.3	More Details
CVE-2025-69012	Missing Authorization vulnerability in Stephen Harris Event Organiser event-organiser allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Event Organiser: from n/a through <= 3.12.8.	4.3	More Details
CVE-2023-28619	Missing Authorization vulnerability in bnayawpguy Resoto allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Resoto: from n/a through 1.0.8.	4.3	More Details
CVE-2025-64641	Mattermost versions 11.1.x <= 11.1.0, 11.0.x <= 11.0.5, 10.12.x <= 10.12.3, 10.11.x <= 10.11.7 fail to verify that post actions invoking /share-issue-publicly were created by the Jira plugin which allowed a malicious Mattermost user to exfiltrate Jira tickets when victim users interacted with affected posts	4.1	More Details
CVE-2025-68950	ImageMagick is free and open-source software used for editing and manipulating digital images. Prior to version 7.1.2-12, Magick fails to check for circular references between two MVGs, leading to a stack overflow. This is a DoS vulnerability, and any situation that allows reading the mvg file will be affected. Version 7.1.2-12 fixes the issue.	4.0	More Details
CVE-2025-	Missing Authorization vulnerability in Automattic Crowdsignal Forms crowdsignal-forms allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Crowdsignal Forms: from n/a through <= 1.7.2.	3.8	More Details

69015			
CVE-2025-15187	A vulnerability was found in GreenCMS up to 2.3. This affects an unknown part of the file /DataController.class.php of the component File Handler. Performing manipulation of the argument sqlFiles/zipFiles results in path traversal. The attack can be initiated remotely. The exploit has been made public and could be used. This vulnerability only affects products that are no longer supported by the maintainer.	3.8	More Details
CVE-2025-36228	IBM Aspera Faspex 5 5.0.0 through 5.0.14.1 may allow inconsistent permissions between the user interface and backend API allowed users to access features that appeared disabled, potentially leading to misuse.	3.8	More Details
CVE-2025-15244	A vulnerability has been found in PHPEMS up to 11.0. This impacts an unknown function of the component Purchase Request Handler. The manipulation leads to race condition. The attack may be initiated remotely. A high degree of complexity is needed for the attack. The exploitability is said to be difficult. The exploit has been disclosed to the public and may be used.	3.7	More Details
CVE-2025-15107	A security vulnerability has been detected in actiontech sqle up to 4.2511.0. The impacted element is an unknown function of the file sqle/utils/jwt.go of the component JWT Secret Handler. The manipulation of the argument JWTSecretKey leads to use of hard-coded cryptographic key . The attack is possible to be carried out remotely. The attack's complexity is rated as high. The exploitability is regarded as difficult. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report and is planning to fix this flaw in an upcoming release.	3.7	More Details
CVE-2025-15153	A weakness has been identified in PbootCMS up to 3.2.12. Impacted is an unknown function of the file /data/pbootcms.db of the component SQLite Database. Executing manipulation can lead to files or directories accessible. It is possible to launch the attack remotely. Attacks of this nature are highly complex. The exploitability is considered difficult. The exploit has been made available to the public and could be exploited. Modifying the configuration settings is advised.	3.7	More Details
CVE-2025-15151	A vulnerability was determined in TaleLin Lin-CMS up to 0.6.0. This affects an unknown part of the file /tests/config.py of the component Tests Folder. This manipulation of the argument username/password causes password in configuration file. The attack is possible to be carried out remotely. The complexity of an attack is rather high. It is indicated that the exploitability is difficult. The exploit has been publicly disclosed and may be utilized.	3.7	More Details
CVE-2025-15105	A security flaw has been discovered in getmaxun maxun up to 0.0.28. Impacted is an unknown function of the file /getmaxun/maxun/blob/develop/server/src/routes/auth.ts. Performing manipulation of the argument api_key results in use of hard-coded cryptographic key . Remote exploitation of the attack is possible. The attack is considered to have high complexity. The exploitability is considered difficult. The exploit has been released to the public and may be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	3.7	More Details
CVE-2025-15108	A vulnerability was detected in PandaXGO PandaX up to fb8ff40f7ce5dfebdf66306c6d85625061faf7e5. This affects an unknown function of the file config.yml of the component JWT Secret Handler. The manipulation of the argument key results in use of hard-coded cryptographic key . The attack may be performed from remote. This attack is characterized by high complexity. The exploitability is reported as difficult. The exploit is now public and may be used. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	3.7	More Details
CVE-2025-15116	A security flaw has been discovered in OpenCart up to 4.1.0.3. Affected by this issue is some unknown functionality of the component Single-Use Coupon Handler. Performing manipulation results in race condition. The attack may be initiated remotely. The attack's complexity is rated as high. The exploitation is known to be difficult. The exploit has been released to the public and may be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	3.7	More Details
CVE-2025-15221	A flaw has been found in SohuTV CacheCloud up to 3.2.0. This vulnerability affects the function index of the file src/main/java/com/sohu/cache/web/controller/AppDataMigrateController.java. This manipulation causes cross site scripting. Remote exploitation of the attack is possible. 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.	3.5	More Details
CVE-2025-15174	A security vulnerability has been detected in SohuTV CacheCloud up to 3.2.0. Affected by this vulnerability is the function doAppAuditList of the file src/main/java/com/sohu/cache/web/controller/AppManageController.java. Such manipulation leads to cross site scripting. The attack may be performed from remote. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15219	A security vulnerability has been detected in SohuTV CacheCloud up to 3.2.0. Affected by this issue is the function doMachineList/doPodList of the file src/main/java/com/sohu/cache/web/controller/MachineManageController.java. The manipulation leads to cross site scripting. The attack may be initiated remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15249	A weakness has been identified in zhujunliang3 work_platform up to 6bc5a50bb527ce27f7906d11ea6ec139beb79c31. This vulnerability affects unknown code of the component Content Handler. Executing manipulation can lead to cross site scripting. The attack may be performed from remote. This product utilizes a rolling release system for continuous delivery, and as such, version information for affected or updated releases is not disclosed. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15241	A security vulnerability has been detected in CloudPanel Community Edition up to 2.5.1. The affected element is an unknown function of the file /admin/users of the component HTTP Header Handler. Such manipulation of the argument Referer leads to open redirect. It is possible to launch the attack remotely. The exploit has been disclosed publicly and may be used. Upgrading to version 2.5.2 is sufficient to fix this issue. Upgrading the affected component is recommended.	3.5	More Details
CVE-2025-15248	A security flaw has been discovered in sunhailin12315 product-review 商品评价系统 up to 91ead6890b4065bb45b7602d0d73348e75cb4639. This affects an unknown part of the component Write a Review. Performing manipulation of the argument content results in cross site scripting. The attack is possible to be carried out remotely. The exploit has been released to the public and may be exploited. This product adopts a rolling release strategy to maintain continuous delivery The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-	A flaw has been found in SohuTV CacheCloud up to 3.2.0. The impacted element is the function redirectNoPower of the file src/main/java/com/sohu/cache/web/controller/WebResourceController.java. This manipulation causes cross site scripting. The attack is possible to be carried out remotely. The exploit has been published and may be used. The project was informed of the problem	3.5	More Details

15201	early through an issue report but has not responded yet.		
CVE-2025-15245	A vulnerability was found in D-Link DCS-850L 1.02.09. Affected is the function uploadfirmware of the component Firmware Update Service. The manipulation of the argument DownloadFile results in path traversal. The attack must originate from the local network. The exploit has been made public and could be used. This vulnerability only affects products that are no longer supported by the maintainer.	3.5	More Details
CVE-2025-15134	A security flaw has been discovered in yourmaileyes MOOC up to 1.17. This affects the function subreview of the file mooc/controller/MainController.java of the component Submission Handler. Performing manipulation of the argument review results in cross site scripting. The attack can be initiated remotely. The exploit has been released to the public and may be exploited. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15175	A vulnerability was detected in SohuTV CacheCloud up to 3.2.0. Affected by this issue is the function doAppList/appCommandAnalysis of the file src/main/java/com/sohu/cache/web/controller/AppController.java. Performing manipulation results in cross site scripting. 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.	3.5	More Details
CVE-2025-15171	A vulnerability was identified in SohuTV CacheCloud up to 3.2.0. This affects the function index of the file src/main/java/com/sohu/cache/web/controller/ServerController.java. The manipulation leads to cross site scripting. Remote exploitation of the attack is possible. The exploit is publicly available and might be used. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15173	A weakness has been identified in SohuTV CacheCloud up to 3.2.0. Affected is the function advancedAnalysis of the file src/main/java/com/sohu/cache/web/controller/InstanceController.java. This manipulation causes cross site scripting. The attack is possible to be carried out remotely. The exploit has been made available to the public and could be exploited. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15258	A weakness has been identified in Edimax BR-6208AC 1.02/1.03. Affected by this issue is the function formALGSetup of the file /goform/formALGSetup of the component Web-based Configuration Interface. This manipulation of the argument wlan-url causes open redirect. The attack is possible to be carried out remotely. The exploit has been made available to the public and could be exploited. Edimax confirms this issue: "The product mentioned, EDIMAX BR-6208AC V2, has reached its End of Life (EOL) status. It is no longer supported or maintained by Edimax, and it is no longer available for purchase in the market. Consequently, there will be no further firmware updates or patches for this device. We recommend users upgrade to newer models for better security." This vulnerability only affects products that are no longer supported by the maintainer.	3.5	More Details
CVE-2025-15095	A security vulnerability has been detected in postmanlabs httpbin up to 0.6.1. This affects an unknown function of the file httpbin-master/httpbin/core.py. The manipulation leads to cross site scripting. The attack may be initiated remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15052	A vulnerability was detected in code-projects Student Information System 1.0. This vulnerability affects unknown code of the file /profile.php. Performing manipulation of the argument firstname/lastname results in cross site scripting. The attack is possible to be carried out remotely. The exploit is now public and may be used.	3.5	More Details
CVE-2025-15172	A security flaw has been discovered in SohuTV CacheCloud up to 3.2.0. This impacts the function preview of the file src/main/java/com/sohu/cache/web/controller/RedisConfigTemplateController.java. The manipulation results in cross site scripting. The attack can be executed remotely. The exploit has been released to the public and may be exploited. The project was informed of the problem early through an issue report but has not responded yet.	3.5	More Details
CVE-2025-15122	A vulnerability was found in JeecgBoot up to 3.9.0. The impacted element is the function loadDatarule of the file /sys/sysDepartRole/datarule/. Performing manipulation of the argument departId/roleId results in improper authorization. It is possible to initiate the attack remotely. The attack is considered to have high complexity. The exploitability is regarded as difficult. The exploit has been made public and could be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15126	A weakness has been identified in JeecgBoot up to 3.9.0. Affected by this vulnerability is the function getPositionUserList of the file /sys/position/getPositionUserList. This manipulation of the argument positionId causes improper authorization. The attack may be initiated remotely. The complexity of an attack is rather high. The exploitation appears to be difficult. The exploit has been made available to the public and could be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15125	A security flaw has been discovered in JeecgBoot up to 3.9.0. Affected is the function queryDepartPermission of the file /sys/permission/queryDepartPermission. The manipulation of the argument departId results in improper authorization. The attack can be launched remotely. This attack is characterized by high complexity. The exploitability is told to be difficult. The exploit has been released to the public and may be exploited. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15123	A vulnerability was determined in JeecgBoot up to 3.9.0. This affects an unknown function of the file /sys/sysDepartPermission/datarule/. Executing manipulation can lead to improper authorization. It is possible to launch the attack remotely. The attack requires a high level of complexity. The exploitability is reported as difficult. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15084	A vulnerability was identified in youlaitech youlai-mall 1.0.0/2.0.0. The impacted element is the function orderService.payOrder of the file mall-oms/oms-boot/src/main/java/com/youlai/mall/oms/controller/app/OrderController.java of the component Order Payment Handler. The manipulation leads to improper access controls. The attack can be initiated remotely. The attack is considered to have high complexity. The exploitability is regarded as difficult. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-68940	In Gitea before 1.22.5, branch deletion permissions are not adequately enforced after merging a pull request.	3.1	More Details
CVE-	A flaw has been found in JeecgBoot up to 3.9.0. Impacted is the function getDeptRoleList of the file /sys/sysDepartRole/getDeptRoleList. This manipulation of the argument departId causes improper authorization. The attack is		More

2025-15120	possible to be carried out remotely. A high degree of complexity is needed for the attack. The exploitability is considered difficult. The exploit has been published and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	Details
CVE-2025-15242	A vulnerability was detected in PHPEMS up to 11.0. The impacted element is an unknown function of the component Coupon Handler. Performing manipulation results in race condition. The attack can be initiated remotely. The complexity of an attack is rather high. The exploitability is regarded as difficult. The exploit is now public and may be used.	3.1	More Details
CVE-2025-15117	A weakness has been identified in Dromara Sa-Token up to 1.44.0. This affects the function ObjectInputStream.readObject of the file SaJdkSerializer.java. Executing manipulation can lead to deserialization. The attack may be launched remotely. This attack is characterized by high complexity. It is indicated that the exploitability is difficult. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-36229	IBM Aspera Faspex 5 5.0.0 through 5.0.14.1 could allow authenticated users to enumerate sensitive information of data due by enumerating package identifiers.	3.1	More Details
CVE-2025-15119	A vulnerability was detected in JeecgBoot up to 3.9.0. This issue affects the function queryPageList of the file /sys/sysDepartRole/list. The manipulation of the argument deptId results in improper authorization. The attack can be executed remotely. A high complexity level is associated with this attack. The exploitability is assessed as difficult. The exploit is now public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15141	A vulnerability was determined in Halo up to 2.21.10. This issue affects some unknown processing of the file /actuator of the component Configuration Handler. Executing manipulation can lead to information disclosure. The attack may be performed from remote. This attack is characterized by high complexity. The exploitability is assessed as difficult. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15124	A vulnerability was identified in JeecgBoot up to 3.9.0. This impacts the function getParameterMap of the file /sys/sysDepartPermission/list. The manipulation of the argument departId leads to improper authorization. The attack can be initiated remotely. The attack's complexity is rated as high. The exploitability is said to be difficult. The exploit is publicly available and might be used. The vendor was contacted early about this disclosure but did not respond in any way.	3.1	More Details
CVE-2025-15200	A vulnerability was detected in SohuTV CacheCloud up to 3.2.0. The affected element is the function getExceptionStatisticsByClient/getCommandStatisticsByClient/doIndex of the file src/main/java/com/sohu/cache/web/controller/AppClientDataShowController.java. The manipulation results in cross site scripting. The attack can be executed 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.	2.4	More Details
CVE-2025-15204	A vulnerability was determined in SohuTV CacheCloud up to 3.2.0. Affected is the function doQuartzList of the file src/main/java/com/sohu/cache/web/controller/QuartzManageController.java. Executing manipulation can lead to cross site scripting. 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.	2.4	More Details
CVE-2025-15121	A vulnerability has been found in JeecgBoot up to 3.9.0. The affected element is the function getDeptRoleByUserId of the file /sys/sysDepartRole/getDeptRoleByUserId. Such manipulation of the argument departId leads to information disclosure. The vendor was contacted early about this disclosure but did not respond in any way.	2.4	More Details
CVE-2025-15203	A vulnerability was found in SohuTV CacheCloud up to 3.2.0. This impacts the function index of the file src/main/java/com/sohu/cache/web/controller/ResourceController.java. Performing manipulation results in cross site scripting. 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.	2.4	More Details
CVE-2025-15145	A security vulnerability has been detected in SohuTV CacheCloud up to 3.2.0. This affects the function doTotalList of the file src/main/java/com/sohu/cache/web/controller/TotalManageController.java. Such manipulation leads to cross site scripting. The attack can be launched remotely. The exploit has been disclosed publicly and may be used. The project was informed of the problem early through an issue report but has not responded yet.	2.4	More Details
CVE-2025-15146	A vulnerability was detected in SohuTV CacheCloud up to 3.2.0. This impacts the function doUserList of the file src/main/java/com/sohu/cache/web/controller/UserManageController.java. Performing manipulation results in cross site scripting. The attack may be initiated 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.	2.4	More Details
CVE-2025-15202	A vulnerability has been found in SohuTV CacheCloud up to 3.2.0. This affects the function taskQueueList of the file src/main/java/com/sohu/cache/web/controller/TaskController.java. Such manipulation leads to cross site scripting. 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.	2.4	More Details
CVE-2025-15149	A vulnerability has been found in rawchen ecms up to b59d7feaa9094234e8aa6c8c6b290621ca575ded. Affected by this vulnerability is the function updateProductServlet of the file src/servlet/product/updateProductServlet.java of the component Add New Product Page. The manipulation of the argument productName leads to cross site scripting. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used. This product follows a rolling release approach for continuous delivery, so version details for affected or updated releases are not provided. The vendor was contacted early about this disclosure but did not respond in any way.	2.4	More Details
CVE-2025-15214	A vulnerability was found in Campcodes Park Ticketing System 1.0. The impacted element is the function save_pricing of the file admin_class.php. The manipulation of the argument Name results in cross site scripting. The attack may be performed from remote. The exploit has been made public and could be used.	2.4	More Details
CVE-2025-15188	A vulnerability was determined in Campcodes Complete Online Beauty Parlor Management System 1.0. This vulnerability affects unknown code of the file /admin/search-invoices.php. Executing manipulation of the argument searchdata can lead to cross site scripting. The attack can be launched remotely. The exploit has been publicly disclosed and may be utilized.	2.4	More Details
CVE-2025-57840	ADB(Android Debug Bridge) is affected by type privilege bypass, successful exploitation of this vulnerability may affect service availability.	2.2	More Details

CVE-2025-15083	A vulnerability was determined in TOZED ZLT M30s up to 1.47. The affected element is an unknown function of the component UART Interface. Executing manipulation can lead to on-chip debug and test interface with improper access control. The physical device can be targeted for the attack. Attacks of this nature are highly complex. The exploitability is described as difficult. The exploit has been publicly disclosed and may be utilized. The vendor was contacted early about this disclosure but did not respond in any way.	2.0	More Details
CVE-2023-54182	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to check readonly condition correctly With below case, it can mount multi-device image w/ rw option, however one of secondary device is set as ro, later update will cause panic, so let's introduce f2fs_dev_is_readonly(), and check multi-devices rw status in f2fs_remount() w/ it in order to avoid such inconsistent mount status. mkfs.f2fs -c /dev/zram1 /dev/zram0 -f blockdev --setro /dev/zram1 mount -t f2fs dev/zram0 /mnt/f2fs mount: /mnt/f2fs: WARNING: source write-protected, mounted read-only. mount -t f2fs -o remount,rw mnt/f2fs dd if=/dev/zero of=/mnt/f2fs/file bs=1M count=8192 kernel BUG at fs/f2fs/inline.c:258! RIP: 0010:f2fs_write_inline_data+0x23e/0x2d0 [f2fs] Call Trace: f2fs_write_single_data_page+0x26b/0x9f0 [f2fs] f2fs_write_cache_pages+0x389/0xa60 [f2fs] __f2fs_write_data_pages+0x26b/0x2d0 [f2fs] f2fs_write_data_pages+0x2e/0x40 [f2fs] do_writepages+0xd3/0x1b0 __writeback_single_inode+0x5b/0x420 writeback_sb_inodes+0x236/0x5a0 __writeback_inodes_wb+0x56/0xf0 wb_writeback+0x2a3/0x490 wb_do_writeback+0x2b2/0x330 wb_workfn+0x6a/0x260 process_one_work+0x270/0x5e0 worker_thread+0x52/0x3e0 kthread+0xf4/0x120 ret_from_fork+0x29/0x50	N/A	More Details
CVE-2023-54183	In the Linux kernel, the following vulnerability has been resolved: media: v4l2-core: Fix a potential resource leak in v4l2_fwnode_parse_link() If fwnode_graph_get_remote_endpoint() fails, 'fwnode' is known to be NULL, so fwnode_handle_put() is a no-op. Release the reference taken from a previous fwnode_graph_get_port_parent() call instead. Also handle fwnode_graph_get_port_parent() failures. In order to fix these issues, add an error handling path to the function and the needed gotos.	N/A	More Details
CVE-2023-54184	In the Linux kernel, the following vulnerability has been resolved: scsi: target: iscsit: Free cmds before session free Commands from recovery entries are freed after session has been closed. That leads to use-after-free at command free or NPE with such call trace: Time2Retain timer expired for SID: 1, cleaning up iSCSI session. BUG: kernel NULL pointer dereference, address: 0000000000000140 RIP: 0010:sbitmap_queue_clear+0x3a/0xa0 Call Trace: target_release_cmd_kref+0xd1/0x1f0 [target_core_mod] transport_generic_free_cmd+0xd1/0x180 [target_core_mod] iscsit_free_cmd+0x53/0xd0 [iscsi_target_mod] iscsit_free_connection_recovery_entries+0x29d/0x320 [iscsi_target_mod] iscsit_close_session+0x13a/0x140 [iscsi_target_mod] iscsit_check_post_dataout+0x440/0x440 [iscsi_target_mod] call_timer_fn+0x24/0x140 Move cleanup of recovery enrties to before session freeing.	N/A	More Details
CVE-2023-54180	In the Linux kernel, the following vulnerability has been resolved: btrfs: handle case when repair happens with dev-replace [BUG] There is a bug report that a BUG_ON() in btrfs_repair_io_failure() (originally repair_io_failure() in v6.0 kernel) got triggered when replacing a unreliable disk: BTRFS warning (device sda1): csum failed root 257 ino 2397453 off 39624704 csum 0xb0d18c75 expected csum 0x4dae9c5e mirror 3 kernel BUG at fs/btrfs/extent_io.c:2380! invalid opcode: 0000 [#1] PREEMPT SMP NOPTI CPU: 9 PID: 3614331 Comm: kworker/u257:2 Tainted: G OE 6.0.0-5-amd64 #1 Debian 6.0.10-2 Hardware name: Micro-Star International Co., Ltd. MS-7C60/TRX40 PRO WIFI (MS-7C60), BIOS 2.70 07/01/2021 Workqueue: btrfs-endio btrfs_end_bio_work [btrfs] RIP: 0010:repair_io_failure+0x24a/0x260 [btrfs] Call Trace: <TASK> clean_io_failure+0x14d/0x180 [btrfs] end_bio_extent_readpage+0x412/0x6e0 [btrfs] ? __switch_to+0x106/0x420 process_one_work+0x1c7/0x380 worker_thread+0x4d/0x380 ? rescuer_thread+0x3a0/0x3a0 kthread+0xe9/0x110 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork+0x22/0x30 [CAUSE] Before the BUG_ON(), we got some read errors from the replace target first, note the mirror number (3, which is beyond RAID1 duplication, thus it's read from the replace target device). Then at the BUG_ON() location, we are trying to writeback the repaired sectors back the failed device. The check looks like this: ret = btrfs_map_block(fs_info, BTRFS_MAP_WRITE, logical, &map_length, &bioc, mirror_num); if (ret) goto out_counter_dec; BUG_ON(mirror_num != bioc->mirror_num); But inside btrfs_map_block(), we can modify bioc->mirror_num especially for dev-replace: if (dev_replace_is_ongoing && mirror_num == map->num_stripes + 1 && !need_full_stripe(op) && dev_replace->tgtdev != NULL) { ret = get_extra_mirror_from_replace(fs_info, logical, *length, dev_replace->srcdev->devid, &mirror_num, &physical_to_patch_in_first_stripe); patch_the_first_stripe_for_dev_replace = 1; } Thus if we're repairing the replace target device, we're going to trigger that BUG_ON(). But in reality, the read failure from the replace target device may be that, our replace hasn't reached the range we're reading, thus we're reading garbage, but with replace running, the range would be properly filled later. Thus in that case, we don't need to do anything but let the replace routine to handle it. [FIX] Instead of a BUG_ON(), just skip the repair if we're repairing the device replace target device.	N/A	More Details
CVE-2023-54181	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix issue in verifying allow_ptr_leaks After we converted the capabilities of our networking-bpf program from cap_sys_admin to cap_net_admin+cap_bpf, our networking-bpf program failed to start. Because it failed the bpf verifier, and the error log is "R3 pointer comparison prohibited". A simple reproducer as follows, SEC("cls-ingress") int ingress(struct __sk_buff *skb) { struct iphdr *iph = (void *)(long)skb->data + sizeof(struct ethhdr); if ((long)(iph + 1) > (long)skb->data_end) return TC_ACT_STOLEN; return TC_ACT_OK; } Per discussion with Yonghong and Alexei [1], comparison of two packet pointers is not a pointer leak. This patch fixes it. Our local kernel is 6.1.y and we expect this fix to be backported to 6.1.y, so stable is CCed. [1]. https://lore.kernel.org/bpf/CAADnVQ+Nmspr7Si+pxWn8zkE7hX-7s93ugwC+94aXSy4uQ9vBg@mail.gmail.com/	N/A	More Details
CVE-2023-54185	In the Linux kernel, the following vulnerability has been resolved: btrfs: remove BUG_ON()'s in add_new_free_space() At add_new_free_space() we have these BUG_ON()'s that are there to deal with any failure to add free space to the in memory free space cache. Such failures are mostly -ENOMEM that should be very rare. However there's no need to have these BUG_ON()'s, we can just return any error to the caller and all callers and their upper call chain are already dealing with errors. So just make add_new_free_space() return any errors, while removing the BUG_ON()'s, and returning the total amount of added free space to an optional u64 pointer argument.	N/A	More Details
CVE-2023-54178	In the Linux kernel, the following vulnerability has been resolved: of: unittest: fix null pointer dereferencing in of_unittest_find_node_by_name() when kmalloc() fail to allocate memory in kasprintf(), name or full_name will be NULL, strcmp() will cause null pointer dereference.	N/A	More Details
CVE-2023-54186	In the Linux kernel, the following vulnerability has been resolved: usb: typec: altmodes/displayport: fix pin_assignment_show This patch fixes negative indexing of buf array in pin_assignment_show when get_current_pin_assignments returns 0 i.e. no compatible pin assignments are found. BUG: KASAN: use-after-free in pin_assignment_show+0x26c/0x33c ... Call trace: dump_backtrace+0x110/0x204 dump_stack_lvl+0x84/0xbc print_report+0x358/0x974 kasan_report+0x9c/0xfc __do_kernel_fault+0xd4/0x2d4 do_bad_area+0x48/0x168 do_tag_check_fault+0x24/0x38 do_mem_abort+0x6c/0x14c el1_abort+0x44/0x68 el1h_64_sync_handler+0x64/0xa4 el1h_64_sync+0x78/0x7c pin_assignment_show+0x26c/0x33c dev_attr_show+0x50/0xc0	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix potential corruption when moving a directory F2FS has the		

2023-54187	same issue in ext4_rename causing crash revealed by xfstests/generic/707. See also commit 0813299c586b ("ext4: Fix possible corruption when moving a directory")	N/A	More Details
CVE-2023-54188	In the Linux kernel, the following vulnerability has been resolved: dmaengine: apple-admac: Fix 'current_tx' not getting freed In terminate_all we should queue up all submitted descriptors to be freed. We do that for the content of the 'issued' and 'submitted' lists, but the 'current_tx' descriptor falls through the cracks as it's removed from the 'issued' list once it gets assigned to be the current descriptor. Explicitly queue up freeing of the 'current_tx' descriptor to address a memory leak that is otherwise present.	N/A	More Details
CVE-2023-54189	In the Linux kernel, the following vulnerability has been resolved: pstore/ram: Add check for kstrdup Add check for the return value of kstrdup() and return the error if it fails in order to avoid NULL pointer dereference.	N/A	More Details
CVE-2023-54190	In the Linux kernel, the following vulnerability has been resolved: leds: led-core: Fix refcount leak in of_led_get() class_find_device_by_of_node() calls class_find_device(), it will take the reference, use the put_device() to drop the reference when not need anymore.	N/A	More Details
CVE-2023-54191	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: fix memory leak in mt7996_mcu_exit Always purge mcu skb queues in mt7996_mcu_exit routine even if mt7996_firmware_state fails.	N/A	More Details
CVE-2023-54179	In the Linux kernel, the following vulnerability has been resolved: scsi: qla2xxx: Array index may go out of bound Klocwork reports array 'vha->host_str' of size 16 may use index value(s) 16..19. Use snprintf() instead of sprintf().	N/A	More Details
CVE-2022-50783	In the Linux kernel, the following vulnerability has been resolved: mptcp: use proper req destructor for IPv6 Before, only the destructor from TCP request sock in IPv4 was called even if the subflow was IPv6. It is important to use the right destructor to avoid memory leaks with some advanced IPv6 features, e.g. when the request socks contain specific IPv6 options.	N/A	More Details
CVE-2023-54177	In the Linux kernel, the following vulnerability has been resolved: quota: fix warning in dqgrab() There's issue as follows when do fault injection: WARNING: CPU: 1 PID: 14870 at include/linux/quotaops.h:51 dquot_disable+0x13b7/0x18c0 Modules linked in: CPU: 1 PID: 14870 Comm: fsconfig Not tainted 6.3.0-next-20230505-00006-g5107a9c821af-dirty #541 RIP: 0010:dquot_disable+0x13b7/0x18c0 RSP: 0018:ffffc9000acc79e0 EFLAGS: 00010246 RAX: 0000000000000000 RBX: 0000000000000000 RCX: ffff88825e41b980 RDX: 0000000000000000 RSI: ffff88825e41b980 RDI: 0000000000000002 RBP: ffff888179f68000 R08: ffffffff82087ca7 R09: 0000000000000000 R10: 0000000000000001 R11: ffffed102f3ed026 R12: ffff888179f68130 R13: ffff888179f68110 R14: dffffc0000000000 R15: ffff888179f68118 FS: 00007f450a073740(0000) GS:ffff88882fc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007ffe96f2efd8 CR3: 000000025c8ad000 CR4: 00000000000006e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> dquot_load_quota_sb+0xd53/0x1060 dquot_resume+0x172/0x230 ext4_reconfigure+0x1dc6/0x27b0 reconfigure_super+0x515/0xa90 __x64_sys_fsconfig+0xb19/0xd20 do_syscall_64+0x39/0xb0 entry_SYSCALL_64_after_hwframe+0x63/0xcd Above issue may happens as follows: ProcessA ProcessB ProcessC sys_fsconfig vfs_fsconfig_locked reconfigure_super ext4_remount dquot_suspend -> suspend all type quota sys_fsconfig vfs_fsconfig_locked reconfigure_super ext4_remount dquot_resume ret = dquot_load_quota_sb add_dquot_ref do_open -> open file O_RDWR vfs_open do_dentry_open get_write_access atomic_inc_unless_negative(&inode->i_writecount) ext4_file_open dquot_file_open dquot_initialize __dquot_initialize dqget atomic_inc(&dquot->dq_count); __dquot_initialize __dquot_initialize dqget if (!test_bit(DQ_ACTIVE_B, &dquot->dq_flags)) ext4_acquire_dquot -> Return error DQ_ACTIVE_B flag isn't set dquot_disable invalidate_dquots if (atomic_read(&dquot->dq_count)) dqgrab WARN_ON_ONCE(!test_bit(DQ_ACTIVE_B, &dquot->dq_flags)) -> Trigger warning In the above scenario, 'dquot->dq_flags' has no DQ_ACTIVE_B is normal when dqgrab(). To solve above issue just replace the dqgrab() use in invalidate_dquots() with atomic_inc(&dquot->dq_count).	N/A	More Details
CVE-2023-54176	In the Linux kernel, the following vulnerability has been resolved: mptcp: stricter state check in mptcp_worker As reported by Christoph, the mptcp protocol can run the worker when the relevant msk socket is in an unexpected state: connect() // incoming reset + fastclose // the mptcp worker is scheduled mptcp_disconnect() // msk is now CLOSED listen() mptcp_worker() Leading to the following splat: divide error: 0000 [#1] PREEMPT SMP CPU: 1 PID: 21 Comm: kworker/1:0 Not tainted 6.3.0-rc1-gde5e8fd0123c #11 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.11.0-2.el7 04/01/2014 Workqueue: events mptcp_worker RIP: 0010:__tcp_select_window+0x22c/0x4b0 net/ipv4/tcp_output.c:3018 RSP: 0018:ffffc900000b3c98 EFLAGS: 00010293 RAX: 000000000000ffd7 RBX: 000000000000ffd7 RCX: 0000000000000000 RDX: 0000000000000000 RSI: ffffffff8214ce97 RDI: 0000000000000004 RBP: 000000000000ffd7 R08: 0000000000000004 R09: 0000000000010000 R10: 000000000000ffd7 R11: ffff888005afa148 R12: 000000000000ffd7 R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000000 FS: 0000000000000000(0000) GS:ffff88803ed00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000405270 CR3: 000000003011e006 CR4: 0000000000370ee0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> tcp_select_window net/ipv4/tcp_output.c:262 [inline] __tcp_transmit_skb+0x356/0x1280 net/ipv4/tcp_output.c:1345 tcp_transmit_skb net/ipv4/tcp_output.c:1417 [inline] tcp_send_active_reset+0x13e/0x320 net/ipv4/tcp_output.c:3459 mptcp_check_fastclose net/mptcp/protocol.c:2530 [inline] mptcp_worker+0x6c7/0x800 net/mptcp/protocol.c:2705 process_one_work+0x3bd/0x950 kernel/workqueue.c:2390 worker_thread+0x5b/0x610 kernel/workqueue.c:2537 kthread+0x138/0x170 kernel/kthread.c:376 ret_from_fork+0x2c/0x50 arch/x86/entry/entry_64.S:308 </TASK> This change addresses the issue explicitly checking for bad states before running the mptcp worker.	N/A	More Details
CVE-2022-50885	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix NULL-ptr-deref in rxe_qp_do_cleanup() when socket create failed There is a null-ptr-deref when mount.cifs over rdma: BUG: KASAN: null-ptr-deref in rxe_qp_do_cleanup+0x2f3/0x360 [rdma_rxe] Read of size 8 at addr 0000000000000018 by task mount.cifs/3046 CPU: 2 PID: 3046 Comm: mount.cifs Not tainted 6.1.0-rc5+ #62 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.14.0-1.fc3 Call Trace: <TASK> dump_stack_lvl+0x34/0x44 kasan_report+0xad/0x130 rxe_qp_do_cleanup+0x2f3/0x360 [rdma_rxe] execute_in_process_context+0x25/0x90 __rxe_cleanup+0x101/0x1d0 [rdma_rxe] rxe_create_qp+0x16a/0x180 [rdma_rxe] create_qp.part.0+0x27d/0x340 ib_create_qp_kernel+0x73/0x160 rdma_create_qp+0x100/0x230 _smbd_get_connection+0x752/0x20f0 smbd_get_connection+0x21/0x40 cifs_get_tcp_session+0x8ef/0xda0 mount_get_conns+0x60/0x750 cifs_mount+0x103/0xd00 cifs_smb3_do_mount+0x1dd/0xcb0 smb3_get_tree+0x1d5/0x300 vfs_get_tree+0x41/0xf0 path_mount+0x9b3/0xdd0 __x64_sys_mount+0x190/0x1d0 do_syscall_64+0x35/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0 The root cause of the issue is the socket create failed in rxe_qp_init_req(). So move the reset rxe_qp_do_cleanup() after the NULL ptr check.	N/A	More Details

CVE-2022-50886	In the Linux kernel, the following vulnerability has been resolved: mmc: toshsd: fix return value check of mmc_add_host() mmc_add_host() may return error, if we ignore its return value, the memory that allocated in mmc_alloc_host() will be leaked and it will lead a kernel crash because of deleting not added device in the remove path. So fix this by checking the return value and goto error path which will call mmc_free_host(), besides, free_irq() also needs be called.	N/A	More Details
CVE-2022-50887	In the Linux kernel, the following vulnerability has been resolved: regulator: core: fix unbalanced of node refcount in regulator_dev_lookup() I got the the following report: OF: ERROR: memory leak, expected refcount 1 instead of 2, of_node_get()/of_node_put() unbalanced - destroy cset entry: attach overlay node /i2c/pmic@62/regulators/exten In of_get_regulator(), the node is returned from of_parse_phandle() with refcount incremented, after using it, of_node_put() need be called.	N/A	More Details
CVE-2022-50888	In the Linux kernel, the following vulnerability has been resolved: remoteproc: qcom: q6v5: Fix potential null-ptr-deref in q6v5_wcss_init_mmio() q6v5_wcss_init_mmio() will call platform_get_resource_byname() that may fail and return NULL. devm_ioremap() will use res->start as input, which may causes null-ptr-deref. Check the ret value of platform_get_resource_byname() to avoid the null-ptr-deref.	N/A	More Details
CVE-2022-50889	In the Linux kernel, the following vulnerability has been resolved: dm integrity: Fix UAF in dm_integrity_dtr() Dm_integrity also has the same UAF problem when dm_resume() and dm_destroy() are concurrent. Therefore, cancelling timer again in dm_integrity_dtr().	N/A	More Details
CVE-2023-54164	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: fix iso_conn related locking and validity issues sk->sk_state indicates whether iso_pi(sk)->conn is valid. Operations that check/update sk_state and access conn should hold lock_sock, otherwise they can race. The order of taking locks is hci_dev_lock > lock_sock > iso_conn_lock, which is how it is in connect/disconnect_cfm -> iso_conn_del -> iso_chan_del. Fix locking in iso_connect_cis/bis and sendmsg/recvmsg to take lock_sock around updating sk_state and conn. iso_conn_del must not occur during iso_connect_cis/bis, as it frees the iso_conn. Hold hdev->lock longer to prevent that. This should not reintroduce the issue fixed in commit 241f51931c35 ("Bluetooth: ISO: Avoid circular locking dependency"), since the we acquire locks in order. We retain the fix in iso_sock_connect to release lock_sock before iso_connect_* acquires hdev->lock. Similarly for commit 6a5ad251b7cd ("Bluetooth: ISO: Fix possible circular locking dependency"). We retain the fix in iso_conn_ready to not acquire iso_conn_lock before lock_sock. iso_conn_add shall return iso_conn with valid hcon. Make it so also when reusing an old CIS connection waiting for disconnect timeout (see __iso_sock_close where conn->hcon is set to NULL). Trace with iso_conn_del after iso_chan_add in iso_connect_cis: ========== iso_sock_create:771: sock 00000000be9b69b7 iso_sock_init:693: sk 000000004dff667e iso_sock_bind:827: sk 000000004dff667e 70:1a:b8:98:ff:a2 type 1 iso_sock_setsockopt:1289: sk 000000004dff667e iso_sock_setsockopt:1289: sk 000000004dff667e iso_sock_setsockopt:1289: sk 000000004dff667e iso_sock_connect:875: sk 000000004dff667e iso_connect_cis:353: 70:1a:b8:98:ff:a2 -> 28:3d:c2:4a:7e:da hci_get_route:1199: 70:1a:b8:98:ff:a2 -> 28:3d:c2:4a:7e:da hci_conn_add:1005: hci0 dst 28:3d:c2:4a:7e:da iso_conn_add:140: hcon 000000007b65d182 conn 00000000daf8625e __iso_chan_add:214: conn 00000000daf8625e iso_connect_cfm:1700: hcon 000000007b65d182 bdaddr 28:3d:c2:4a:7e:da status 12 iso_conn_del:187: hcon 000000007b65d182 conn 00000000daf8625e, err 16 iso_sock_clear_timer:117: sock 000000004dff667e state 3 <Note: sk_state is BT_BOUND (3), so iso_connect_cis is still running at this point> iso_chan_del:153: sk 000000004dff667e, conn 00000000daf8625e, err 16 hci_conn_del:1151: hci0 hcon 000000007b65d182 handle 65535 hci_conn_unlink:1102: hci0: hcon 000000007b65d182 hci_chan_list_flush:2780: hcon 000000007b65d182 iso_sock_getsockopt:1376: sk 000000004dff667e iso_sock_getname:1070: sock 00000000be9b69b7, sk 000000004dff667e iso_sock_getname:1070: sock 00000000be9b69b7, sk 000000004dff667e iso_sock_getsockopt:1376: sk 000000004dff667e iso_sock_getname:1070: sock 00000000be9b69b7, sk 000000004dff667e iso_sock_getname:1070: sock 00000000be9b69b7, sk 000000004dff667e iso_sock_shutdown:1434: sock 00000000be9b69b7, sk 000000004dff667e, how 1 __iso_sock_close:632: sk 000000004dff667e state 5 socket 00000000be9b69b7 <Note: sk_state is BT_CONNECT (5), even though iso_chan_del sets BT_CLOSED (6). Only iso_connect_cis sets it to BT_CONNECT, so it must be that iso_chan_del occurred between iso_chan_add and end of iso_connect_cis.> BUG: kernel NULL pointer dereference, address: 0000000000000000 PGD 8000000006467067 P4D 8000000006467067 PUD 3f5f067 PMD 0 Oops: 0000 [#1] PREEMPT SMP PTI Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.16.2-1.fc38 04/01/2014 RIP: 0010:__iso_sock_close (net/bluetooth/iso.c:664) bluetooth ========== Trace with iso_conn_del before iso_chan_add in iso_connect_cis: ========== iso_connect_cis:356: 70:1a:b8:98:ff:a2 -> 28:3d:c2:4a:7e:da ... iso_conn_add:140: hcon 0000000093bc551f conn 00000000768ae504 hci_dev_put:1487: hci0 orig refcnt 21 hci_event_packet:7607: hci0: e ---truncated---	N/A	More Details
CVE-2023-54165	In the Linux kernel, the following vulnerability has been resolved: zsmalloc: move LRU update from zs_map_object() to zs_malloc() Under memory pressure, we sometimes observe the following crash: [5694.832838] ------------[cut here]------------ [5694.842093] list_del corruption, ffff888014b6a448->next is LIST_POISON1 (dead000000000100) [5694.858677] WARNING: CPU: 33 PID: 418824 at lib/list_debug.c:47 __list_del_entry_valid+0x42/0x80 [5694.961820] CPU: 33 PID: 418824 Comm: fuse_counters.s Kdump: loaded Tainted: G S 5.19.0-0_fbk3_rc3_hoangnhatpzsdynshrv41_10870_g85a9558a25de #1 [5694.990194] Hardware name: Wiwynn Twin Lakes MP/Twin Lakes Passive MP, BIOS YMM16 05/24/2021 [5695.007072] RIP: 0010:__list_del_entry_valid+0x42/0x80 [5695.017351] Code: 08 48 83 c2 22 48 39 d0 74 24 48 8b 10 48 39 f2 75 2c 48 8b 51 08 b0 01 48 39 f2 75 34 c3 48 c7 c7 55 d7 78 82 e8 4e 45 3b 00 <0f> 0b eb 31 48 c7 c7 27 a8 70 82 e8 3e 45 3b 00 0f 0b eb 21 48 c7 [5695.054919] RSP: 0018:ffffc90027aef4f0 EFLAGS: 00010246 [5695.065366] RAX: 41fe484987275300 RBX: ffff888008988180 RCX: 0000000000000000 [5695.079636] RDX: ffff88886006c280 RSI: ffff888860060480 RDI: ffff888860060480 [5695.093904] RBP: 0000000000000002 R08: 0000000000000000 R09: ffffc90027aef370 [5695.108175] R10: 0000000000000000 R11: ffffffff82fdf1c0 R12: 0000000010000002 [5695.122447] R13: ffff888014b6a448 R14: ffff888014b6a420 R15: 00000000138dc240 [5695.136717] FS: 00007f23a7d3f740(0000) GS:ffff888860040000(0000) knlGS:0000000000000000 [5695.152899] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [5695.164388] CR2: 0000560ceaab6ac0 CR3: 000000001c06c001 CR4: 00000000007706e0 [5695.178659] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [5695.192927] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [5695.207197] PKRU: 55555554 [5695.212602] Call Trace: [5695.217486] <TASK> [5695.221674] zs_map_object+0x91/0x270 [5695.229000] zswap_frontswap_store+0x33d/0x870 [5695.237885] ? do_raw_spin_lock+0x5d/0xa0 [5695.245899] __frontswap_store+0x51/0xb0 [5695.253742] swap_writepage+0x3c/0x60 [5695.261063] shrink_page_list+0x738/0x1230 [5695.269255] shrink_lruvec+0x5ec/0xcd0 [5695.276749] ? shrink_slab+0x187/0x5f0 [5695.284240] ? mem_cgroup_iter+0x6e/0x120 [5695.292255] shrink_node+0x293/0x7b0 [5695.299402] do_try_to_free_pages+0xea/0x550 [5695.307940] try_to_free_pages+0x19a/0x490 [5695.316126] __folio_alloc+0x19ff/0x3e40 [5695.323971] ? __filemap_get_folio+0x8a/0x4e0 [5695.332681] ? walk_component+0x2a8/0xb50 [5695.340697] ? generic_permission+0xda/0x2a0 [5695.349231] ? __filemap_get_folio+0x8a/0x4e0 [5695.357940] ? walk_component+0x2a8/0xb50 [5695.365955] vma_alloc_folio+0x10e/0x570 [5695.373796] ? walk_component+0x52/0xb50 [5695.381634] wp_page_copy+0x38c/0xc10 [5695.388953] ? filename_lookup+0x378/0xbc0 [5695.397140] handle_mm_fault+0x87f/0x1800 [5695.405157] do_user_addr_fault+0x1bd/0x570 [	N/A	More Details

	5695.413520] exc_page_fault+0x5d/0x110 [5695.421017] asm_exc_page_fault+0x22/0x30 After some investigation, I have found the following issue: unlike other zswap backends, zsmalloc performs the LRU list update at the object mapping time, rather than when the slot for the object is allocated. This deviation was discussed and agreed upon during the review process of the zsmalloc writeback patch series: https://lore.kernel.org/lkml/Y3flcAXNxxrvy3ZH@cmpxchg.org/ Unfortunately, this introduces a subtle bug that occurs when there is a concurrent store and reclaim, which interleave as follows: zswap_frontswap_store() shrink_worker() zs_malloc() zs_zpool_shrink() spin_lock(&pool->lock) zs_reclaim_page() zspage = find_get_zspage() spin_unlock(&pool->lock) spin_lock(&pool->lock) zspage = list_first_entry(&pool->lru) ---truncated---		
CVE-2023-54166	In the Linux kernel, the following vulnerability has been resolved: igc: Fix Kernel Panic during ndo_tx_timeout callback The Xeon validation group has been carrying out some loaded tests with various HW configurations, and they have seen some transmit queue time out happening during the test. This will cause the reset adapter function to be called by igc_tx_timeout(). Similar race conditions may arise when the interface is being brought down and up in igc_reinit_locked(), an interrupt being generated, and igc_clean_tx_irq() being called to complete the TX. When the igc_tx_timeout() function is invoked, this patch will turn off all TX ring HW queues during igc_down() process. TX ring HW queues will be activated again during the igc_configure_tx_ring() process when performing the igc_up() procedure later. This patch also moved existing igc_disable_tx_ring_hw() to avoid using forward declaration. Kernel trace: [7678.747813] ------------[cut here]------------ [7678.757914] NETDEV WATCHDOG: enp1s0 (igc): transmit queue 2 timed out [7678.770117] WARNING: CPU: 0 PID: 13 at net/sched/sch_generic.c:525 dev_watchdog+0x1ae/0x1f0 [7678.784459] Modules linked in: xt_conntrack nft_chain_nat xt_MASQUERADE xt_addrtype nft_compat nf_tables nfnetlink br_netfilter bridge stp llc overlay dm_mod emrcha(PO) emriio(PO) rktpm(PO) cegbuf_mod(PO) patch_update(PO) se(PO) sgx_tgts(PO) mktme(PO) keylocker(PO) svtdx(PO) svfs_pci_hotplug(PO) vtd_mod(PO) davemem(PO) svmabort(PO) svindexio(PO) usbx2(PO) ehci_sched(PO) svheartbeat(PO) ioapic(PO) sv8259(PO) svintr(PO) lt(PO) pcierootport(PO) enginefw_mod(PO) ata(PO) smbus(PO) spiflash_cdf(PO) arden(PO) dsa_iax(PO) oobmsm_punit(PO) cpm(PO) svkdb(PO) ebg_pch(PO) pch(PO) sviotargets(PO) svbdf(PO) svmem(PO) svbios(PO) dram(PO) svtsc(PO) targets(PO) superio(PO) svkernel(PO) cswitch(PO) mcf(PO) pentiumIII_mod(PO) fs_svfs(PO) mdevdefdb(PO) svfs_os_services(O) ixgbe mdio mdio_devres libphy emeraldrapids_svdefs(PO) regsupport(O) libnvdimm nls_cp437 snd_hda_codec_realtek snd_hda_codec_generic ledtrig_audio snd_hda_intel snd_intel_dspcfg snd_hda_codec snd_hwdep x86_pkg_temp_thermal snd_hda_core snd_pcm snd_timer isst_if_mbox_pci [7678.784496] input_leds isst_if_mmio sg snd isst_if_common soundcore wmi button sad9(O) drm fuse backlight configfs efivarfs ip_tables x_tables vmd sdhci led_class rtl8150 r8152 hid_generic pegasus mmc_block usbhid mmc_core hid megaraid_sas ixgb igb i2c_algo_bit ice i40e hpsa scsi_transport_sas e1000e e1000 e100 ax88179_178a usbnet xhci_pci sd_mod xhci_hcd t10_pi crc32c_intel crc64_rocksoft igc crc64 crc_t10dif usbcore crct10dif_generic ptp crct10dif_common usb_common pps_core [7679.200403] RIP: 0010:dev_watchdog+0x1ae/0x1f0 [7679.210201] Code: 28 e9 53 ff ff ff 4c 89 e7 c6 05 06 42 b9 00 01 e8 17 d1 fb ff 44 89 e9 4c 89 e6 48 c7 c7 40 ad fb 81 48 89 c2 e8 52 62 82 ff <0f> 0b e9 72 ff ff ff 65 8b 05 80 7d 7c 7e 89 c0 48 0f a3 05 0a c1 [7679.245438] RSP: 0018:ffa00000001f7d90 EFLAGS: 00010282 [7679.256021] RAX: 0000000000000000 RBX: ff11000109938440 RCX: 0000000000000000 [7679.268710] RDX: ff11000361e26cd8 RSI: ff11000361e1b880 RDI: ff11000361e1b880 [7679.281314] RBP: ffa00000001f7da8 R08: ff1100035f8fffe8 R09: 0000000000027ffb [7679.293840] R10: 0000000000001f0a R11: ff1100035f840000 R12: ff11000109938000 [7679.306276] R13: 0000000000000002 R14: dead000000000122 R15: ffa00000001f7e18 [7679.318648] FS: 0000000000000000(0000) GS:ff11000361e00000(0000) knlGS:0000000000000000 [7679.332064] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [7679.342757] CR2: 00007ffff7fca168 CR3: 000000013b08a006 CR4: 0000000000471ef8 [7679.354984] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [7679.367207] DR3: 0000000000000000 DR6: 00000000fffe07f0 DR7: 0000000000000400 [7679.379370] PKRU: 55555554 [7679.386446] Call Trace: [7679.393152] <TASK> [7679.399363] ? __pfx_dev_watchdog+0x10/0x10 [7679.407870] call_timer_fn+0x31/0x110 [7679.415698] e ---truncated---	N/A	More Details
CVE-2023-54167	In the Linux kernel, the following vulnerability has been resolved: m68k: mm: Move initrd phys_to_virt handling after paging_init() When booting with an initial ramdisk on platforms where physical memory does not start at address zero (e.g. on Amiga): initrd: 0ef0602c - 0f800000 Zone ranges: DMA [mem 0x0000000008000000-0x000000f7ffffffff] Normal empty Movable zone start for each node Early memory node ranges node 0: [mem 0x0000000008000000-0x000000000f7fffff] Initmem setup node 0 [mem 0x0000000008000000-0x000000000f7fffff] Unable to handle kernel access at virtual address (ptrval) Oops: 00000000 Modules linked in: PC: [<00201d3c>] memcmp+0x28/0x56 As phys_to_virt() relies on m68k_memoffset and module_fixup(), it must not be called before paging_init(). Hence postpone the phys_to_virt handling for the initial ramdisk until after calling paging_init(). While at it, reduce #ifdef clutter by using IS_ENABLED() instead.	N/A	More Details
CVE-2023-54168	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx4: Prevent shift wrapping in set_user_sq_size() The ucmd->log_sq_bb_count variable is controlled by the user so this shift can wrap. Fix it by using check_shl_overflow() in the same way that it was done in commit 515f60004ed9 ("RDMA/hns: Prevent undefined behavior in hns_roce_set_user_sq_size()").	N/A	More Details
CVE-2023-54169	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: fix memory leak in mlx5e_ptp_open When kvzalloc_node or kvzalloc failed in mlx5e_ptp_open, the memory pointed by "c" or "cparams" is not freed, which can lead to a memory leak. Fix by freeing the array in the error path.	N/A	More Details
CVE-2023-54170	In the Linux kernel, the following vulnerability has been resolved: keys: Fix linking a duplicate key to a keyring's assoc_array When making a DNS query inside the kernel using dns_query(), the request code can in rare cases end up creating a duplicate index key in the assoc_array of the destination keyring. It is eventually found by a BUG_ON() check in the assoc_array implementation and results in a crash. Example report: [2158499.700025] kernel BUG at ../lib/assoc_array.c:652! [2158499.700039] invalid opcode: 0000 [#1] SMP PTI [2158499.700065] CPU: 3 PID: 31985 Comm: kworker/3:1 Kdump: loaded Not tainted 5.3.18-150300.59.90-default #1 SLE15-SP3 [2158499.700096] Hardware name: VMware, Inc. VMware Virtual Platform/440BX Desktop Reference Platform, BIOS 6.00 11/12/2020 [2158499.700351] Workqueue: cifsiod cifs_resolve_server [cifs] [2158499.700380] RIP: 0010:assoc_array_insert+0x85f/0xa40 [2158499.700401] Code: ff 74 2b 48 8b 3b 49 8b 45 18 4c 89 e6 48 83 e7 fe e8 95 ec 74 00 3b 45 88 7d db 85 c0 79 d4 0f 0b 0f 0b 0f 0b e8 41 f2 be ff <0f> 0b 0f 0b 81 7d 88 ff ff ff 7f 4c 89 eb 4c 8b ad 58 ff ff ff 0f [2158499.700448] RSP: 0018:ffffc0bd6187faf0 EFLAGS: 00010282 [2158499.700470] RAX: ffff9f1ea7da2fe8 RBX: ffff9f1ea7da2fc1 RCX: 0000000000000005 [2158499.700492] RDX: 0000000000000000 RSI: 0000000000000005 RDI: 0000000000000000 [2158499.700515] RBP: ffffc0bd6187fbb0 R08: ffff9f185faf1100 R09: 0000000000000000 [2158499.700538] R10: ffff9f1ea7da2cc0 R11: 000000005ed8cec8 R12: ffffc0bd6187fc28 [2158499.700561] R13: ffff9f15feb8d000 R14: ffff9f1ea7da2fc0 R15: ffff9f168dc0d740 [2158499.700585] FS: 0000000000000000(0000) GS:ffff9f185fac0000(0000) knlGS:0000000000000000 [2158499.700610] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [2158499.700630] CR2: 00007fdd94fca238 CR3: 0000000809d8c006 CR4: 00000000003706e0 [2158499.700702] Call Trace: [2158499.700741] ? key_alloc+0x447/0x4b0 [2158499.700768] ? __key_link_begin+0x43/0xa0 [2158499.700790] __key_link_begin+0x43/0xa0 [2158499.700814] request_key_and_link+0x2c7/0x730 [2158499.700847] ? dns_resolver_read+0x20/0x20 [dns_resolver] [2158499.700873] ? key_default_cmp+0x20/0x20 [2158499.700898] request_key_tag+0x43/0xa0 [2158499.700926] dns_query+0x114/0x2ca [dns_resolver] [2158499.701127] dns_resolve_server_name_to_ip+0x194/0x310 [cifs] [2158499.701164] ? scnprintf+0x49/0x90 [2158499.701190] ? __switch_to_asm+0x40/0x70 [2158499.701211] ? __switch_to_asm+0x34/0x70 [2158499.701405] reconn_set_ipaddr_from_hostname+0x81/0x2a0 [cifs] [2158499.701603] cifs_resolve_server+0x4b/0xd0 [cifs] [2158499.701632] process_one_work+0x1f8/0x3e0 [2158499.701658] worker_thread+0x2d/0x3f0 [2158499.701682] ?	N/A	More Details

	process_one_work+0x3e0/0x3e0 [2158499.701703] kthread+0x10d/0x130 [2158499.701723] ? kthread_park+0xb0/0xb0 [2158499.701746] ret_from_fork+0x1f/0x40 The situation occurs as follows: * Some kernel facility invokes dns_query() to resolve a hostname, for example, "abcdef". The function registers its global DNS resolver cache as current->cred.thread_keyring and passes the query to request_key_net() -> request_key_tag() -> request_key_and_link(). * Function request_key_and_link() creates a keyring_search_context object. Its match_data.cmp method gets set via a call to type->match_preparse() (resolves to dns_resolver_match_preparse()) to dns_resolver_cmp(). * Function request_key_and_link() continues and invokes search_process_keyrings_rcu() which returns that a given key was not found. The control is then passed to request_key_and_link() -> construct_alloc_key(). * Concurrently to that, a second task similarly makes a DNS query for "abcdef." and its result gets inserted into the DNS resolver cache. * Back on the first task, function construct_alloc_key() first runs __key_link_begin() to determine an assoc_array_edit operation to insert a new key. Index keys in the array are compared exactly as-is, using keyring_compare_object(). The operation ---truncated---		
CVE-2023-54171	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix memory leak of iter->temp when reading trace_pipe kmemleak reports: unreferenced object 0xffff88814d14e200 (size 256): comm "cat", pid 336, jiffies 4294871818 (age 779.490s) hex dump (first 32 bytes): 04 00 01 03 00 00 00 00 08 00 00 00 00 00 00 00 ..... 0c d8 c8 9b ff ff ff ff 04 5a ca 9b ff ff ff ff .....Z..... backtrace: [<ffffffff9bdff18f>] __kmalloc+0x4f/0x140 [<ffffffff9bc9238b>] trace_find_next_entry+0xbb/0x1d0 [<ffffffff9bc9caef>] trace_print_lat_context+0xaf/0x4e0 [<ffffffff9bc94490>] print_trace_line+0x3e0/0x950 [<ffffffff9bc95499>] tracing_read_pipe+0x2d9/0x5a0 [<ffffffff9bf03a43>] vfs_read+0x143/0x520 [<ffffffff9bf04c2d>] ksys_read+0xbd/0x160 [<ffffffff9d0f0edf>] do_syscall_64+0x3f/0x90 [<ffffffff9d2000aa>] entry_SYSCALL_64_after_hwframe+0x6e/0xd8 when reading file 'trace_pipe', 'iter->temp' is allocated or relocated in trace_find_next_entry() but not freed before 'trace_pipe' is closed. To fix it, free 'iter->temp' in tracing_release_pipe().	N/A	More Details
CVE-2023-54172	In the Linux kernel, the following vulnerability has been resolved: x86/hyperv: Disable IBT when hypercall page lacks ENDBR instruction On hardware that supports Indirect Branch Tracking (IBT), Hyper-V VMs with ConfigVersion 9.3 or later support IBT in the guest. However, current versions of Hyper-V have a bug in that there's not an ENDBR64 instruction at the beginning of the hypercall page. Since hypercalls are made with an indirect call to the hypercall page, all hypercall attempts fail with an exception and Linux panics. A Hyper-V fix is in progress to add ENDBR64. But guard against the Linux panic by clearing X86_FEATURE_IBT if the hypercall page doesn't start with ENDBR. The VM will boot and run without IBT. If future Linux 32-bit kernels were to support IBT, additional hypercall page hackery would be needed to make IBT work for such kernels in a Hyper-V VM.	N/A	More Details
CVE-2023-54173	In the Linux kernel, the following vulnerability has been resolved: bpf: Disable preemption in bpf_event_output We received report [1] of kernel crash, which is caused by using nesting protection without disabled preemption. The bpf_event_output can be called by programs executed by bpf_prog_run_array_cg function that disabled migration but keeps preemption enabled. This can cause task to be preempted by another one inside the nesting protection and lead eventually to two tasks using same perf_sample_data buffer and cause crashes like: BUG: kernel NULL pointer dereference, address: 0000000000000001 #PF: supervisor instruction fetch in kernel mode #PF: error_code(0x0010) - not-present page ... ? perf_output_sample+0x12a/0x9a0 ? finish_task_switch.isra.0+0x81/0x280 ? perf_event_output+0x66/0xa0 ? bpf_event_output+0x13a/0x190 ? bpf_event_output_data+0x22/0x40 ? bpf_prog_dfc84bbde731b257_cil_sock4_connect+0x40a/0xacb ? xa_load+0x87/0xe0 ? __cgroup_bpf_run_filter_sock_addr+0xc1/0x1a0 ? release_sock+0x3e/0x90 ? sk_setsockopt+0x1a1/0x12f0 ? udp_pre_connect+0x36/0x50 ? inet_dgram_connect+0x93/0xa0 ? __sys_connect+0xb4/0xe0 ? udp_setsockopt+0x27/0x40 ? __pfx_udp_push_pending_frames+0x10/0x10 ? __sys_setsockopt+0xdf/0x1a0 ? __x64_sys_connect+0xf/0x20 ? do_syscall_64+0x3a/0x90 ? entry_SYSCALL_64_after_hwframe+0x72/0xdc Fixing this by disabling preemption in bpf_event_output. [1] https://github.com/cilium/cilium/issues/26756	N/A	More Details
CVE-2023-54174	In the Linux kernel, the following vulnerability has been resolved: vfio: Fix NULL pointer dereference caused by uninitialized group->iommufd group->iommufd is not initialized for the iommufd_ctx_put() [20018.331541] BUG: kernel NULL pointer dereference, address: 0000000000000000 [20018.377508] RIP: 0010:iommufd_ctx_put+0x5/0x10 [iommufd] ... [20018.476483] Call Trace: [20018.479214] <TASK> [20018.481555] vfio_group_fops_unl_ioctl+0x506/0x690 [vfio] [20018.487586] __x64_sys_ioctl+0x6a/0xb0 [20018.491773] ? trace_hardirqs_on+0xc5/0xe0 [20018.496347] do_syscall_64+0x67/0x90 [20018.500340] entry_SYSCALL_64_after_hwframe+0x4b/0xb5	N/A	More Details
CVE-2023-54192	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix null pointer panic in tracepoint in __replace_atomic_write_block We got a kernel panic if old_addr is NULL. https://bugzilla.kernel.org/show_bug.cgi?id=217266 BUG: kernel NULL pointer dereference, address: 0000000000000000 Call Trace: <TASK> f2fs_commit_atomic_write+0x619/0x990 [f2fs a1b985b80f5babd6f3ea778384908880812bfa43] __f2fs_ioctl+0xd8e/0x4080 [f2fs a1b985b80f5babd6f3ea778384908880812bfa43] ? vfs_write+0x2ae/0x3f0 ? vfs_write+0x2ae/0x3f0 __x64_sys_ioctl+0x91/0xd0 do_syscall_64+0x5c/0x90 entry_SYSCALL_64_after_hwframe+0x72/0xdc RIP: 0033:0x7f69095fe53f	N/A	More Details
CVE-2023-54175	In the Linux kernel, the following vulnerability has been resolved: i2c: xiic: xiic_xfer(): Fix runtime PM leak on error path The xiic_xfer() function gets a runtime PM reference when the function is entered. This reference is released when the function is exited. There is currently one error path where the function exits directly, which leads to a leak of the runtime PM reference. Make sure that this error path also releases the runtime PM reference.	N/A	More Details
CVE-2023-54194	In the Linux kernel, the following vulnerability has been resolved: exfat: use kvmalloc_array/kvfree instead of kmalloc_array/kfree The call stack shown below is a scenario in the Linux 4.19 kernel. Allocating memory failed where exfat fs use kmalloc_array due to system memory fragmentation, while the u-disk was inserted without recognition. Devices such as u-disk using the exfat file system are pluggable and may be insert into the system at any time. However, long-term running systems cannot guarantee the continuity of physical memory. Therefore, it's necessary to address this issue. Binder:2632_6: page allocation failure: order:4, mode:0x6040c0(GFP_KERNEL|__GFP_COMP), nodemask=(null) Call trace: [242178.097582] dump_backtrace+0x0/0x4 [242178.097589] dump_stack+0xf4/0x134 [242178.097598] warn_alloc+0xd8/0x144 [242178.097603] __alloc_pages_nodemask+0x1364/0x1384 [242178.097608] kmalloc_order+0x2c/0x510 [242178.097612] kmalloc_order_trace+0x40/0x16c [242178.097618] __kmalloc+0x360/0x408 [242178.097624] load_alloc_bitmap+0x160/0x284 [242178.097628] exfat_fill_super+0xa3c/0xe7c [242178.097635] mount_bdev+0x2e8/0x3a0 [242178.097638] exfat_fs_mount+0x40/0x50 [242178.097643] mount_fs+0x138/0x2e8 [242178.097649] vfs_kern_mount+0x90/0x270 [242178.097655] do_mount+0x798/0x173c [242178.097659] ksys_mount+0x114/0x1ac [242178.097665] __arm64_sys_mount+0x24/0x34 [242178.097671] el0_svc_common+0xb8/0x1b8 [242178.097676] el0_svc_handler+0x74/0x90 [242178.097681] el0_svc+0x8/0x340 By analyzing the exfat code,we found that continuous physical memory is not required here,so kvmalloc_array is used can solve this problem.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net/sched: cls_api: remove block_cb from driver_list before freeing Error handler of tcf_block_bind() frees the whole bo->cb_list on error. However, by that time the flow_block_cb instances are already in the driver list because driver ndo_setup_tc() callback is called before that up the call chain in tcf_block_offload_cmd(). This leaves dangling pointers to freed objects in the list and causes use-after-free[0]. Fix it by also removing flow_block_cb instances from		

CVE-2023-54193	driver_list before deallocating them. [0]: [279.868433] == [279.869964] BUG: KASAN: slab-use-after-free in flow_block_cb_setup_simple+0x631/0x7c0 [279.871527] Read of size 8 at addr ffff888147e2bf20 by task tc/2963 [279.873151] CPU: 6 PID: 2963 Comm: tc Not tainted 6.3.0-rc6+ #4 [279.874273] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 [279.876295] Call Trace: [279.876882] <TASK> [279.877413] dump_stack_lvl+0x33/0x50 [279.878198] print_report+0xc2/0x610 [279.878987] ? flow_block_cb_setup_simple+0x631/0x7c0 [279.879994] kasan_report+0xae/0xe0 [279.880750] ? flow_block_cb_setup_simple+0x631/0x7c0 [279.881744] ? mlx5e_tc_reoffload_flows_work+0x240/0x240 [mlx5_core] [279.883047] flow_block_cb_setup_simple+0x631/0x7c0 [279.884027] tcf_block_offload_cmd.isra.0+0x189/0x2d0 [279.885037] ? tcf_block_setup+0x6b0/0x6b0 [279.885901] ? mutex_lock+0x7d/0xd0 [279.886669] ? __mutex_unlock_slowpath.constprop.0+0x2d0/0x2d0 [279.887844] ? ingress_init+0x1c0/0x1c0 [sch_ingress] [279.888846] tcf_block_get_ext+0x61c/0x1200 [279.889711] ingress_init+0x112/0x1c0 [sch_ingress] [279.890682] ? clsact_init+0x2b0/0x2b0 [sch_ingress] [279.891701] qdisc_create+0x401/0xea0 [279.892485] ? qdisc_tree_reduce_backlog+0x470/0x470 [279.893473] tc_modify_qdisc+0x6f7/0x16d0 [279.894344] ? tc_get_qdisc+0xac0/0xac0 [279.895213] ? mutex_lock+0x7d/0xd0 [279.896005] ? __mutex_lock_slowpath+0x10/0x10 [279.896910] rtnetlink_rcv_msg+0x5fe/0x9d0 [279.897770] ? rtnl_calcit.isra.0+0x2b0/0x2b0 [279.898672] ? __sys_sendmsg+0xb5/0x140 [279.899494] ? do_syscall_64+0x3d/0x90 [279.900302] ? entry_SYSCALL_64_after_hwframe+0x46/0xb0 [279.901337] ? kasan_save_stack+0x2e/0x40 [279.902177] ? kasan_save_stack+0x1e/0x40 [279.903058] ? kasan_set_track+0x21/0x30 [279.903913] ? kasan_save_free_info+0x2a/0x40 [279.904836] ? ____kasan_slab_free+0x11a/0x1b0 [279.905741] ? kmem_cache_free+0x179/0x400 [279.906599] netlink_rcv_skb+0x12c/0x360 [279.907450] ? rtnl_calcit.isra.0+0x2b0/0x2b0 [279.908360] ? netlink_ack+0x1550/0x1550 [279.909192] ? rhashtable_walk_peek+0x170/0x170 [279.910135] ? kmem_cache_alloc_node+0x1af/0x390 [279.911086] ? _copy_from_iter+0x3d6/0xc70 [279.912031] netlink_unicast+0x553/0x790 [279.912864] ? netlink_attachskb+0x6a0/0x6a0 [279.913763] ? netlink_recvmsg+0x416/0xb50 [279.914627] netlink_sendmsg+0x7a1/0xcb0 [279.915473] ? netlink_unicast+0x790/0x790 [279.916334] ? iovec_from_user.part.0+0x4d/0x220 [279.917293] ? netlink_unicast+0x790/0x790 [279.918159] sock_sendmsg+0xc5/0x190 [279.918938] ____sys_sendmsg+0x535/0x6b0 [279.919813] ? import_iovec+0x7/0x10 [279.920601] ? kernel_sendmsg+0x30/0x30 [279.921423] ? __copy_msghdr+0x3c0/0x3c0 [279.922254] ? import_iovec+0x7/0x10 [279.923041] ___sys_sendmsg+0xeb/0x170 [279.923854] ? copy_msghdr_from_user+0x110/0x110 [279.924797] ? ___sys_recvmsg+0xd9/0x130 [279.925630] ? __perf_event_task_sched_in+0x183/0x470 [279.926656] ? ___sys_sendmsg+0x170/0x170 [279.927529] ? ctx_sched_in+0x530/0x530 [279.928369] ? update_curr+0x283/0x4f0 [279.929185] ? perf_event_update_userpage+0x570/0x570 [279.930201] ? __fget_light+0x57/0x520 [279.931023] ? __switch_to+0x53d/0xe70 [27 ---truncated---	N/A	More Details
CVE-2023-54225	In the Linux kernel, the following vulnerability has been resolved: net: ipa: only reset hashed tables when supported Last year, the code that manages GSI channel transactions switched from using spinlock-protected linked lists to using indexes into the ring buffer used for a channel. Recently, Google reported seeing transaction reference count underflows occasionally during shutdown. Doug Anderson found a way to reproduce the issue reliably, and bisected the issue to the commit that eliminated the linked lists and the lock. The root cause was ultimately determined to be related to unused transactions being committed as part of the modem shutdown cleanup activity. Unused transactions are not normally expected (except in error cases). The modem uses some ranges of IPA-resident memory, and whenever it shuts down we zero those ranges. In ipa_filter_reset_table() a transaction is allocated to zero modem filter table entries. If hashing is not supported, hashed table memory should not be zeroed. But currently nothing prevents that, and the result is an unused transaction. Something similar occurs when we zero routing table entries for the modem. By preventing any attempt to clear hashed tables when hashing is not supported, the reference count underflow is avoided in this case. Note that there likely remains an issue with properly freeing unused transactions (if they occur due to errors). This patch addresses only the underflows that Google originally reported.	N/A	More Details
CVE-2023-54217	In the Linux kernel, the following vulnerability has been resolved: Revert "drm/msm: Add missing check and destroy for alloc_ordered_workqueue" This reverts commit 643b7d0869cc7f1f7a5ac7ca6bd25d88f54e31d0. A recent patch that tried to fix up the msm_drm_init() paths with respect to the workqueue but only ended up making things worse: First, the newly added calls to msm_drm_uninit() on early errors would trigger NULL-pointer dereferences, for example, as the kms pointer would not have been initialised. (Note that these paths were also modified by a second broken error handling patch which in effect cancelled out this part when merged.) Second, the newly added allocation sanity check would still leak the previously allocated drm device. Instead of trying to salvage what was badly broken (and clearly not tested), let's revert the bad commit so that clean and backportable fixes can be added in its place. Patchwork: https://patchwork.freedesktop.org/patch/525107/	N/A	More Details
CVE-2023-54218	In the Linux kernel, the following vulnerability has been resolved: net: Fix load-tearing on sk->sk_stamp in sock_recv_cmsgs(). KCSAN found a data race in sock_recv_cmsgs() where the read access to sk->sk_stamp needs READ_ONCE(). BUG: KCSAN: data-race in packet_recvmsg / packet_recvmsg write (marked) to 0xffff88803c81f258 of 8 bytes by task 19171 on cpu 0: sock_write_timestamp include/net/sock.h:2670 [inline] sock_recv_cmsgs include/net/sock.h:2722 [inline] packet_recvmsg+0xb97/0xd00 net/packet/af_packet.c:3489 sock_recvmsg_nosec net/socket.c:1019 [inline] sock_recvmsg+0x11a/0x130 net/socket.c:1040 sock_read_iter+0x176/0x220 net/socket.c:1118 call_read_iter include/linux/fs.h:1845 [inline] new_sync_read fs/read_write.c:389 [inline] vfs_read+0x5e0/0x630 fs/read_write.c:470 ksys_read+0x163/0x1a0 fs/read_write.c:613 __do_sys_read fs/read_write.c:623 [inline] __se_sys_read fs/read_write.c:621 [inline] __x64_sys_read+0x41/0x50 fs/read_write.c:621 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3b/0x90 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x72/0xdc read to 0xffff88803c81f258 of 8 bytes by task 19183 on cpu 1: sock_recv_cmsgs include/net/sock.h:2721 [inline] packet_recvmsg+0xb64/0xd00 net/packet/af_packet.c:3489 sock_recvmsg_nosec net/socket.c:1019 [inline] sock_recvmsg+0x11a/0x130 net/socket.c:1040 sock_read_iter+0x176/0x220 net/socket.c:1118 call_read_iter include/linux/fs.h:1845 [inline] new_sync_read fs/read_write.c:389 [inline] vfs_read+0x5e0/0x630 fs/read_write.c:470 ksys_read+0x163/0x1a0 fs/read_write.c:613 __do_sys_read fs/read_write.c:623 [inline] __se_sys_read fs/read_write.c:621 [inline] __x64_sys_read+0x41/0x50 fs/read_write.c:621 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3b/0x90 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x72/0xdc value changed: 0xffffffffc4653600 -> 0x0000000000000000 Reported by Kernel Concurrency Sanitizer on: CPU: 1 PID: 19183 Comm: syz-executor.5 Not tainted 6.3.0-rc7-02330-gca6270c12e20 #2 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: Revert "IB/isert: Fix incorrect release of isert connection" Commit: 699826f4e30a ("IB/isert: Fix incorrect release of isert connection") is causing problems on OPA when DEVICE_REMOVAL is happening. ------------[cut here]------------ WARNING: CPU: 52 PID: 2117247 at drivers/infiniband/core/cq.c:359 ib_cq_pool_cleanup+0xac/0xb0 [ib_core] Modules linked in: nfsd nfs_acl target_core_user uio tcm_fc libfc scsi_transport_fc tcm_loop target_core_pscsi target_core_iblock target_core_file rpcsec_gss_krb5 auth_rpcgss nfsv4 dns_resolver nfs lockd grace fscache netfs rfkill rpcrdma rdma_ucm ib_srpt sunrpc ib_isert iscsi_target_mod target_core_mod opa_vnic ib_iser libiscsi ib_umad scsi_transport_iscsi rdma_cm ib_ipoib iw_cm ib_cm hfi1(-) rdmavt ib_uverbs intel_rapl_msr intel_rapl_common sb_edac ib_core x86_pkg_temp_thermal intel_powerclamp coretemp i2c_i801 mxm_wmi rapl iTCO_wdt ipmi_si iTCO_vendor_support mei_me		

CVE-2023-54219	ipmi_devintf mei intel_cstate ioatdma intel_uncore i2c_smbus joydev pcspkr lpc_ich ipmi_msghandler acpi_power_meter acpi_pad xfs libcrc32c sr_mod sd_mod cdrom t10_pi sg crct10dif_pclmul crc32_pclmul crc32c_intel drm_kms_helper drm_shmem_helper ahci libahci ghash_clmulni_intel igb drm libata dca i2c_algo_bit wmi fuse CPU: 52 PID: 2117247 Comm: modprobe Not tainted 6.5.0-rc1+ #1 Hardware name: Intel Corporation S2600CWR/S2600CW, BIOS SE5C610.86B.01.01.0014.121820151719 12/18/2015 RIP: 0010:ib_cq_pool_cleanup+0xac/0xb0 [ib_core] Code: ff 48 8b 43 40 48 8d 7b 40 48 83 e8 40 4c 39 e7 75 b3 49 83 c4 10 4d 39 fc 75 94 5b 5d 41 5c 41 5d 41 5e 41 5f c3 cc cc cc cc <0f> 0b eb a1 90 90 90 90 90 90 90 90 90 90 90 90 90 90 90 90 0f 1f RSP: 0018:ffffc10bea13fc80 EFLAGS: 00010206 RAX: 000000000000010c RBX: ffff9bf5c7e66c00 RCX: 000000008020001d RDX: 000000008020001e RSI: fffff175221f9900 RDI: ffff9bf5c7e67640 RBP: ffff9bf5c7e67600 R08: ffff9bf5c7e64400 R09: 000000008020001d R10: 0000000040000000 R11: 0000000000000000 R12: ffff9bee4b1e8a18 R13: dead000000000122 R14: dead000000000100 R15: ffff9bee4b1e8a38 FS: 00007ff1e6d38740(0000) GS:ffff9bfd9fb00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00005652044ecc68 CR3: 0000000889b5c005 CR4: 00000000001706e0 Call Trace: <TASK> ? __warn+0x80/0x130 ? ib_cq_pool_cleanup+0xac/0xb0 [ib_core] ? report_bug+0x195/0x1a0 ? handle_bug+0x3c/0x70 ? exc_invalid_op+0x14/0x70 ? asm_exc_invalid_op+0x16/0x20 ? ib_cq_pool_cleanup+0xac/0xb0 [ib_core] disable_device+0x9d/0x160 [ib_core] __ib_unregister_device+0x42/0xb0 [ib_core] ib_unregister_device+0x22/0x30 [ib_core] rvt_unregister_device+0x20/0x90 [rdmavt] hfi1_unregister_ib_device+0x16/0xf0 [hfi1] remove_one+0x55/0x1a0 [hfi1] pci_device_remove+0x36/0xa0 device_release_driver_internal+0x193/0x200 driver_detach+0x44/0x90 bus_remove_driver+0x69/0xf0 pci_unregister_driver+0x2a/0xb0 hfi1_mod_cleanup+0xc/0x3c [hfi1] __do_sys_delete_module.constprop.0+0x17a/0x2f0 ? exit_to_user_mode_prepare+0xc4/0xd0 ? syscall_trace_enter.constprop.0+0x126/0x1a0 do_syscall_64+0x5c/0x90 ? syscall_exit_to_user_mode+0x12/0x30 ? do_syscall_64+0x69/0x90 ? syscall_exit_work+0x103/0x130 ? syscall_exit_to_user_mode+0x12/0x30 ? do_syscall_64+0x69/0x90 ? exc_page_fault+0x65/0x150 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 RIP: 0033:0x7ff1e643f5ab Code: 73 01 c3 48 8b 0d 75 a8 1b 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 1e fa b8 b0 00 00 00 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 8b 0d 45 a8 1b 00 f7 d8 64 89 01 48 RSP: 002b:00007ffec9103cc8 EFLAGS: 00000206 ORIG_RAX: 00000000000000b0 RAX: ffffffffffffffda RBX: 00005615267fdc50 RCX: 00007ff1e643f5ab RDX: 0000000000000000 RSI: 0000000000000800 RDI: 00005615267fdcb8 RBP: 00005615267fdc50 R08: 0000000000000000 R09: 0000000000000000 R10: 00007ff1e659eac0 R11: 0000000000000206 R12: 00005615267fdcb8 R13: 00000000000 ---truncated---	N/A	More Details
CVE-2023-54220	In the Linux kernel, the following vulnerability has been resolved: serial: 8250: Fix oops for port->pm on uart_change_pm() Unloading a hardware specific 8250 driver can produce error "Unable to handle kernel paging request at virtual address" about ten seconds after unloading the driver. This happens on uart_hangup() calling uart_change_pm(). Turns out commit 04e82793f068 ("serial: 8250: Reinit port->pm on port specific driver unbind") was only a partial fix. If the hardware specific driver has initialized port->pm function, we need to clear port->pm too. Just reinitializing port->ops does not do this. Otherwise serial8250_pm() will call port->pm() instead of serial8250_do_pm().	N/A	More Details
CVE-2023-54221	In the Linux kernel, the following vulnerability has been resolved: clk: imx93: fix memory leak and missing unwind goto in imx93_clocks_probe In function probe(), it returns directly without unregistered hws when error occurs. Fix this by adding 'goto unregister_hws;' on line 295 and line 310. Use devm_kzalloc() instead of kzalloc() to automatically free the memory using devm_kfree() when error occurs. Replace of_iomap() with devm_of_iomap() to automatically handle the unused ioremap region and delete 'iounmap(anatop_base);' in unregister_hws.	N/A	More Details
CVE-2023-54222	In the Linux kernel, the following vulnerability has been resolved: hte: tegra-194: Fix off by one in tegra_hte_map_to_line_id() The "map_sz" is the number of elements in the "m" array so the > comparison needs to be changed to >= to prevent an out of bounds read.	N/A	More Details
CVE-2023-54223	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: xsk: Fix invalid buffer access for legacy rq The below crash can be encountered when using xdpsock in rx mode for legacy rq: the buffer gets released in the XDP_REDIRECT path, and then once again in the driver. This fix sets the flag to avoid releasing on the driver side. XSK handling of buffers for legacy rq was relying on the caller to set the skip release flag. But the referenced fix started using fragment counts for pages instead of the skip flag. Crash log: general protection fault, probably for non-canonical address 0xffff8881217e3a: 0000 [#1] SMP CPU: 0 PID: 14 Comm: ksoftirqd/0 Not tainted 6.5.0-rc1+ #31 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 RIP: 0010:bpf_prog_03b13f331978c78c+0xf/0x28 Code: ... RSP: 0018:ffff88810082fc98 EFLAGS: 00010246 RAX: 0000000000000000 RBX: ffff888138404901 RCX: c0ffffc900027cbc RDX: fffffffa000b514 RSI: 00ffff8881217e32 RDI: ffff888138404901 RBP: ffff88810082fc98 R08: 0000000000091100 R09: 0000000000000006 R10: 0000000000000800 R11: 0000000000000800 R12: ffffc9000027a000 R13: ffff8881217e2dc0 R14: ffff8881217e2910 R15: ffff8881217e2f00 FS: 0000000000000000(0000) GS:ffff88852c800000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000564cb2e2cde0 CR3: 000000010e603004 CR4: 0000000000370eb0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> ? die_addr+0x32/0x80 ? exc_general_protection+0x192/0x390 ? asm_exc_general_protection+0x22/0x30 ? 0xffffffffa000b514 ? bpf_prog_03b13f331978c78c+0xf/0x28 mlx5e_xdp_handle+0x48/0x670 [mlx5_core] ? dev_gro_receive+0x3b5/0x6e0 mlx5e_xsk_skb_from_cqe_linear+0x6e/0x90 [mlx5_core] mlx5e_handle_rx_cqe+0x55/0x100 [mlx5_core] mlx5e_poll_rx_cq+0x87/0x6e0 [mlx5_core] mlx5e_napi_poll+0x45e/0x6b0 [mlx5_core] __napi_poll+0x25/0x1a0 net_rx_action+0x28a/0x300 __do_softirq+0xcd/0x279 ? sort_range+0x20/0x20 run_ksoftirqd+0x1a/0x20 smpboot_thread_fn+0xa2/0x130 kthread+0xc9/0xf0 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork+0x1f/0x30 </TASK> Modules linked in: mlx5_ib mlx5_core rpcrdma rdma_ucm ib_iser libiscsi scsi_transport_iscsi ib_umad rdma_cm ib_ipoib iw_cm ib_cm ib_uverbs ib_core xt_conntrack xt_MASQUERADE nf_conntrack_netlink nfnetlink xt_addrtype iptable_nat nf_nat br_netfilter overlay zram zsmalloc fuse [last unloaded: mlx5_core] ---[end trace 0000000000000000]---	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix lockdep splat and potential deadlock after failure running delayed items When running delayed items we are holding a delayed node's mutex and then we will attempt to modify a subvolume btree to insert/update/delete the delayed items. However if have an error during the insertions for example, btrfs_insert_delayed_items() may return with a path that has locked extent buffers (a leaf at the very least), and then we attempt to release the delayed node at __btrfs_run_delayed_items(), which requires taking the delayed node's mutex, causing an ABBA type of deadlock. This was reported by syzbot and the lockdep splat is the following: WARNING: possible circular locking dependency detected 6.5.0-rc7-syzkaller-00024-g93f5de5f648d #0 Not tainted -- syz-executor.2/13257 is trying to acquire lock: ffff88801835c0c0 (&delayed_node->mutex){+.+.}-{3:3}, at: __btrfs_release_delayed_node+0x9a/0xaa0 fs/btrfs/delayed-inode.c:256 but task is already holding lock: ffff88802a5ab8e8 (btrfs-tree-00){++++}-{3:3}, at: __btrfs_tree_lock+0x3c/0x2a0 fs/btrfs/locking.c:198 which lock already depends on the new lock. the existing dependency chain (in reverse order) is: -> #1 (btrfs-tree-00){++++}-{3:3}: __lock_release kernel/locking/lockdep.c:5475 [inline] lock_release+0x36f/0x9d0 kernel/locking/lockdep.c:5781 up_write+0x79/0x580 kernel/locking/rwsem.c:1625 btrfs_tree_unlock_rw fs/btrfs/locking.h:189 [inline] btrfs_unlock_up_safe+0x179/0x3b0 fs/btrfs/locking.c:239 search_leaf fs/btrfs/ctree.c:1986 [inline] btrfs_search_slot+0x2511/0x2f80 fs/btrfs/ctree.c:2230 btrfs_insert_empty_items+0x9c/0x180 fs/btrfs/ctree.c:4376 btrfs_insert_delayed_item fs/btrfs/delayed-inode.c:746 [inline] btrfs_insert_delayed_items fs/btrfs/delayed-inode.c:824 [inline]		More

2023-54224	__btrfs_commit_inode_delayed_items+0xd24/0x2410 fs/btrfs/delayed-inode.c:1111 __btrfs_run_delayed_items+0x1db/0x430 fs/btrfs/delayed-inode.c:1153 flush_space+0x269/0xe70 fs/btrfs/space-info.c:723 btrfs_async_reclaim_metadata_space+0x106/0x350 fs/btrfs/space-info.c:1078 process_one_work+0x92c/0x12c0 kernel/workqueue.c:2600 worker_thread+0xa63/0x1210 kernel/workqueue.c:2751 kthread+0x2b8/0x350 kernel/kthread.c:389 ret_from_fork+0x2e/0x60 arch/x86/kernel/process.c:145 ret_from_fork_asm+0x11/0x20 arch/x86/entry/entry_64.S:304 -> #0 (&delayed_node->mutex){+.+.}-{3:3}: check_prev_add kernel/locking/lockdep.c:3142 [inline] check_prevs_add kernel/locking/lockdep.c:3261 [inline] validate_chain kernel/locking/lockdep.c:3876 [inline] __lock_acquire+0x39ff/0x7f70 kernel/locking/lockdep.c:5144 lock_acquire+0x1e3/0x520 kernel/locking/lockdep.c:5761 __mutex_lock_common+0x1d8/0x2530 kernel/locking/mutex.c:603 __mutex_lock kernel/locking/mutex.c:747 [inline] mutex_lock_nested+0x1b/0x20 kernel/locking/mutex.c:799 __btrfs_release_delayed_node+0x9a/0xaa0 fs/btrfs/delayed-inode.c:256 btrfs_release_delayed_node fs/btrfs/delayed-inode.c:281 [inline] __btrfs_run_delayed_items+0x2b5/0x430 fs/btrfs/delayed-inode.c:1156 btrfs_commit_transaction+0x859/0x2ff0 fs/btrfs/transaction.c:2276 btrfs_sync_file+0xf56/0x1330 fs/btrfs/file.c:1988 vfs_fsync_range fs/sync.c:188 [inline] vfs_fsync fs/sync.c:202 [inline] do_fsync fs/sync.c:212 [inline] __do_sys_fsync fs/sync.c:220 [inline] __se_sys_fsync fs/sync.c:218 [inline] __x64_sys_fsync+0x196/0x1e0 fs/sync.c:218 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd other info that --- truncated---	N/A	Details
CVE-2023-54226	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix data races around sk->sk_shutdown. KCSAN found a data race around sk->sk_shutdown where unix_release_sock() and unix_shutdown() update it under unix_state_lock(), OTOH unix_poll() and unix_dgram_poll() read it locklessly. We need to annotate the writes and reads with WRITE_ONCE() and READ_ONCE(). BUG: KCSAN: data-race in unix_poll / unix_release_sock write to 0xffff88800d0f8aec of 1 bytes by task 264 on cpu 0: unix_release_sock+0x75c/0x910 net/unix/af_unix.c:631 unix_release+0x59/0x80 net/unix/af_unix.c:1042 __sock_release+0x7d/0x170 net/socket.c:653 sock_close+0x19/0x30 net/socket.c:1397 __fput+0x179/0x5e0 fs/file_table.c:321 ____fput+0x15/0x20 fs/file_table.c:349 task_work_run+0x116/0x1a0 kernel/task_work.c:179 resume_user_mode_work include/linux/resume_user_mode.h:49 [inline] exit_to_user_mode_loop kernel/entry/common.c:171 [inline] exit_to_user_mode_prepare+0x174/0x180 kernel/entry/common.c:204 __syscall_exit_to_user_mode_work kernel/entry/common.c:286 [inline] syscall_exit_to_user_mode+0x1a/0x30 kernel/entry/common.c:297 do_syscall_64+0x4b/0x90 arch/x86/entry/common.c:86 entry_SYSCALL_64_after_hwframe+0x72/0xdc read to 0xffff88800d0f8aec of 1 bytes by task 222 on cpu 1: unix_poll+0xa3/0x2a0 net/unix/af_unix.c:3170 sock_poll+0xcf/0x2b0 net/socket.c:1385 vfs_poll include/linux/poll.h:88 [inline] ep_item_poll.isra.0+0x78/0xc0 fs/eventpoll.c:855 ep_send_events fs/eventpoll.c:1694 [inline] ep_poll fs/eventpoll.c:1823 [inline] do_epoll_wait+0x6c4/0xea0 fs/eventpoll.c:2258 __do_sys_epoll_wait fs/eventpoll.c:2270 [inline] __se_sys_epoll_wait fs/eventpoll.c:2265 [inline] __x64_sys_epoll_wait+0xcc/0x190 fs/eventpoll.c:2265 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3b/0x90 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x72/0xdc value changed: 0x00 -> 0x03 Reported by Kernel Concurrency Sanitizer on: CPU: 1 PID: 222 Comm: dbus-broker Not tainted 6.3.0-rc7-02330-gca6270c12e20 #2 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014	N/A	More Details
CVE-2023-54215	In the Linux kernel, the following vulnerability has been resolved: virtio-vdpa: Fix cpumask memory leak in virtio_vdpa_find_vqs() Free the cpumask allocated by create_affinity_masks() before returning from the function.	N/A	More Details
CVE-2023-54227	In the Linux kernel, the following vulnerability has been resolved: blk-mq: fix tags leak when shrink nr_hw_queues Although we don't need to realloc set->tags[] when shrink nr_hw_queues, we need to free them. Or these tags will be leaked. How to reproduce: 1. mount -t configfs configfs /mnt 2. modprobe null_blk nr_devices=0 submit_queues=8 3. mkdir /mnt/nullb/nullb0 4. echo 1 > /mnt/nullb/nullb0/power 5. echo 4 > /mnt/nullb/nullb0/submit_queues 6. rmdir /mnt/nullb/nullb0 In step 4, will alloc 9 tags (8 submit queues and 1 poll queue), then in step 5, new_nr_hw_queues = 5 (4 submit queues and 1 poll queue). At last in step 6, only these 5 tags are freed, the other 4 tags leaked.	N/A	More Details
CVE-2023-54228	In the Linux kernel, the following vulnerability has been resolved: regulator: raa215300: Fix resource leak in case of error The clk_register_clkdev() allocates memory by calling vclkdev_alloc() and this memory is not freed in the error path. Similarly, resources allocated by clk_register_fixed_rate() are not freed in the error path. Fix these issues by using devm_clk_hw_register_fixed_rate() and devm_clk_hw_register_clkdev(). After this, the static variable clk is not needed. Replace it with local variable hw in probe() and drop calling clk_unregister_fixed_rate() from raa215300_rtc_unregister_device().	N/A	More Details
CVE-2023-54229	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: fix registration of 6Ghz-only phy without the full channel range Because of what seems to be a typo, a 6Ghz-only phy for which the BDF does not allow the 7115Mhz channel will fail to register: WARNING: CPU: 2 PID: 106 at net/wireless/core.c:907 wiphy_register+0x914/0x954 Modules linked in: ath11k_pci sbsa_gwdt CPU: 2 PID: 106 Comm: kworker/u8:5 Not tainted 6.3.0-rc7-next-20230418-00549-g1e096a17625a-dirty #9 Hardware name: Freebox V7R Board (DT) Workqueue: ath11k_qmi_driver_event ath11k_qmi_driver_event_work pstate: 60000005 (nZCv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : wiphy_register+0x914/0x954 lr : ieee80211_register_hw+0x67c/0xc10 sp : ffffff800b123aa0 x29: ffffff800b123aa0 x28: 0000000000000000 x27: 0000000000000000 x26: 0000000000000000 x25: 0000000000000006 x24: ffffffc008d51418 x23: ffffffc008cb0838 x22: ffffff80176c2460 x21: 0000000000000168 x20: ffffff80176c0000 x19: ffffff80176c03e0 x18: 0000000000000014 x17: 00000000cbef338c x16: 00000000d2a26f21 x15: 00000000ad6bb85f x14: 0000000000000020 x13: 0000000000000020 x12: 00000000ffffffbd x11: 0000000000000208 x10: 00000000fffffdf7 x9 : ffffffc009394718 x8 : ffffff80176c0528 x7 : 000000007fffffff x6 : 0000000000000006 x5 : 0000000000000005 x4 : ffffff800b304284 x3 : ffffff800b304284 x2 : ffffff800b304d98 x1 : 0000000000000000 x0 : 0000000000000000 Call trace: wiphy_register+0x914/0x954 ieee80211_register_hw+0x67c/0xc10 ath11k_mac_register+0x7c4/0xe10 ath11k_core_qmi_firmware_ready+0x1f4/0x570 ath11k_qmi_driver_event_work+0x198/0x590 process_one_work+0x1b8/0x328 worker_thread+0x6c/0x414 kthread+0x100/0x104 ret_from_fork+0x10/0x20 ---[end trace 0000000000000000]--- ath11k_pci 0002:01:00.0: ieee80211 registration failed: -22 ath11k_pci 0002:01:00.0: failed register the radio with mac80211: -22 ath11k_pci 0002:01:00.0: failed to create pdev core: -22	N/A	More Details
CVE-2023-54230	In the Linux kernel, the following vulnerability has been resolved: amba: bus: fix refcount leak commit 5de1540b7bc4 ("drivers/amba: create devices from device tree") increases the refcount of of_node, but not releases it in amba_device_release, so there is refcount leak. By using of_node_put to avoid refcount leak.	N/A	More Details
CVE-2023-54231	In the Linux kernel, the following vulnerability has been resolved: net: libwx: fix memory leak in wx_setup_rx_resources When wx_alloc_page_pool() failed in wx_setup_rx_resources(), it doesn't release DMA buffer. Add dma_free_coherent() in the error path to release the DMA buffer.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: m68k: Only force 030 bus error if PC not in exception table __get_kernel_nofault() does copy data in supervisor mode when forcing a task backtrace log through /proc/sysrq_trigger. This is		

CVE-2023-54232	expected cause a bus error exception on e.g. NULL pointer dereferencing when logging a kernel task has no workqueue associated. This bus error ought to be ignored. Our 030 bus error handler is ill equipped to deal with this: Whenever ssw indicates a kernel mode access on a data fault, we don't even attempt to handle the fault and instead always send a SEGV signal (or panic). As a result, the check for exception handling at the fault PC (buried in send_sig_fault() which gets called from do_page_fault() eventually) is never used. In contrast, both 040 and 060 access error handlers do not care whether a fault happened on supervisor mode access, and will call do_page_fault() on those, ultimately honoring the exception table. Add a check in bus_error030 to call do_page_fault() in case we do have an entry for the fault PC in our exception table. I had attempted a fix for this earlier in 2019 that did rely on testing pagefault_disabled() (see link below) to achieve the same thing, but this patch should be more generic. Tested on 030 Atari Falcon.	N/A	More Details
CVE-2023-54233	In the Linux kernel, the following vulnerability has been resolved: ASoC: SOF: avoid a NULL dereference with unsupported widgets If an IPC4 topology contains an unsupported widget, its .module_info field won't be set, then sof_ipc4_route_setup() will cause a kernel Oops trying to dereference it. Add a check for such cases.	N/A	More Details
CVE-2023-54234	In the Linux kernel, the following vulnerability has been resolved: scsi: mpi3mr: Fix missing mrioc->evtack_cmds initialization Commit c1af985d27da ("scsi: mpi3mr: Add Event acknowledgment logic") introduced an array mrioc->evtack_cmds but initialization of the array elements was missed. They are just zero cleared. The function mpi3mr_complete_evt_ack() refers host_tag field of the elements. Due to the zero value of the host_tag field, the function calls clear_bit() for mrico->evtack_cmds_bitmap with wrong bit index. This results in memory access to invalid address and "BUG: KASAN: use-after-free". This BUG was observed at eHBA-9600 firmware update to version 8.3.1.0. To fix it, add the missing initialization of mrioc->evtack_cmds.	N/A	More Details
CVE-2023-54216	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: TC, Fix using eswitch mapping in nic mode Cited patch is using the eswitch object mapping pool while in nic mode where it isn't initialized. This results in the trace below [0]. Fix that by using either nic or eswitch object mapping pool depending if eswitch is enabled or not. [0]: [826.446057] == [826.446729] BUG: KASAN: slab-use-after-free in mlx5_add_flow_rules+0x30/0x490 [mlx5_core] [826.447515] Read of size 8 at addr ffff888194485830 by task tc/6233 [826.448243] CPU: 16 PID: 6233 Comm: tc Tainted: G W 6.3.0-rc6+ #1 [826.448890] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 [826.449785] Call Trace: [826.450052] <TASK> [826.450302] dump_stack_lvl+0x33/0x50 [826.450650] print_report+0xc2/0x610 [826.450998] ? __virt_addr_valid+0xb1/0x130 [826.451385] ? mlx5_add_flow_rules+0x30/0x490 [mlx5_core] [826.451935] kasan_report+0xae/0xe0 [826.452276] ? mlx5_add_flow_rules+0x30/0x490 [mlx5_core] [826.452829] mlx5_add_flow_rules+0x30/0x490 [mlx5_core] [826.453368] ? __kmalloc_node+0x5a/0x120 [826.453733] esw_add_restore_rule+0x20f/0x270 [mlx5_core] [826.454288] ? mlx5_eswitch_add_send_to_vport_meta_rule+0x260/0x260 [mlx5_core] [826.455011] ? mutex_unlock+0x80/0xd0 [826.455361] ? __mutex_unlock_slowpath.constprop.0+0x210/0x210 [826.455862] ? mapping_add+0x2cb/0x440 [mlx5_core] [826.456425] mlx5e_tc_action_miss_mapping_get+0x139/0x180 [mlx5_core] [826.457058] ? mlx5e_tc_update_skb_nic+0xb0/0xb0 [mlx5_core] [826.457636] ? __kasan_kmalloc+0x77/0x90 [826.458000] ? __kmalloc+0x57/0x120 [826.458336] mlx5_tc_ct_flow_offload+0x325/0xe40 [mlx5_core] [826.458916] ? ct_kernel_enter.constprop.0+0x48/0xa0 [826.459360] ? mlx5_tc_ct_parse_action+0xf0/0xf0 [mlx5_core] [826.459933] ? mlx5e_mod_hdr_attach+0x491/0x520 [mlx5_core] [826.460507] ? mlx5e_mod_hdr_get+0x12/0x20 [mlx5_core] [826.461046] ? mlx5e_tc_attach_mod_hdr+0x154/0x170 [mlx5_core] [826.461635] mlx5e_configure_flower+0x969/0x2110 [mlx5_core] [826.462217] ? _raw_spin_lock_bh+0x85/0xe0 [826.462597] ? __mlx5e_add_fdb_flow+0x750/0x750 [mlx5_core] [826.463163] ? kasan_save_stack+0x2e/0x40 [826.463534] ? down_read+0x115/0x1b0 [826.463878] ? down_write_killable+0x110/0x110 [826.464288] ? tc_setup_action.part.0+0x9f/0x3b0 [826.464701] ? mlx5e_is_uplink_rep+0x4c/0x90 [mlx5_core] [826.465253] ? mlx5e_tc_reoffload_flows_work+0x130/0x130 [mlx5_core] [826.465878] tc_setup_cb_add+0x112/0x250 [826.466247] fl_hw_replace_filter+0x230/0x310 [cls_flower] [826.466724] ? fl_hw_destroy_filter+0x1a0/0x1a0 [cls_flower] [826.467212] fl_change+0x14e1/0x2030 [cls_flower] [826.467636] ? sock_def_readable+0x89/0x120 [826.468019] ? fl_tmplt_create+0x2d0/0x2d0 [cls_flower] [826.468509] ? kasan_unpoison+0x23/0x50 [826.468873] ? get_random_u16+0x180/0x180 [826.469244] ? __radix_tree_lookup+0x2b/0x130 [826.469640] ? fl_get+0x7b/0x140 [cls_flower] [826.470042] ? fl_mask_put+0x200/0x200 [cls_flower] [826.470478] ? __mutex_unlock_slowpath.constprop.0+0x210/0x210 [826.470973] ? fl_tmplt_create+0x2d0/0x2d0 [cls_flower] [826.471427] tc_new_tfilter+0x644/0x1050 [826.471795] ? tc_get_tfilter+0x860/0x860 [826.472170] ? __thaw_task+0x130/0x130 [826.472525] ? arch_stack_walk+0x98/0xf0 [826.472892] ? cap_capable+0x9f/0xd0 [826.473235] ? security_capable+0x47/0x60 [826.473608] rtnetlink_rcv_msg+0x1d5/0x550 [826.473985] ? rtnl_calcit.isra.0+0x1f0/0x1f0 [826.474383] ? __stack_depot_save+0x35/0x4c0 [826.474779] ? kasan_save_stack+0x2e/0x40 [826.475149] ? kasan_save_stack+0x1e/0x40 [826.475518] ? __kasan_record_aux_stack+0x9f/0xb0 [826.475939] ? task_work_add+0x77/0x1c0 [826.476305] netlink_rcv_skb+0xe0/0x210 ---truncated---	N/A	More Details
CVE-2023-54214	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: L2CAP: Fix potential user-after-free This fixes all instances of which requires to allocate a buffer calling alloc_skb which may release the chan lock and reacquire later which makes it possible that the chan is disconnected in the meantime.	N/A	More Details
CVE-2022-50883	In the Linux kernel, the following vulnerability has been resolved: bpf: Prevent decl_tag from being referenced in func_proto arg Syzkaller managed to hit another decl_tag issue: btf_func_proto_check kernel/bpf/btf.c:4506 [inline] btf_check_all_types kernel/bpf/btf.c:4734 [inline] btf_parse_type_sec+0x1175/0x1980 kernel/bpf/btf.c:4763 btf_parse kernel/bpf/btf.c:5042 [inline] btf_new_fd+0x65a/0xb00 kernel/bpf/btf.c:6709 bpf_btf_load+0x6f/0x90 kernel/bpf/syscall.c:4342 __sys_bpf+0x50a/0x6c0 kernel/bpf/syscall.c:5034 __do_sys_bpf kernel/bpf/syscall.c:5093 [inline] __se_sys_bpf kernel/bpf/syscall.c:5091 [inline] __x64_sys_bpf+0x7c/0x90 kernel/bpf/syscall.c:5091 do_syscall_64+0x54/0x70 arch/x86/entry/common.c:48 This seems similar to commit ea68376c8bed ("bpf: prevent decl_tag from being referenced in func_proto") but for the argument.	N/A	More Details
CVE-2023-54203	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix slab-out-of-bounds in init_smb2_rsp_hdr When smb1 mount fails, KASAN detect slab-out-of-bounds in init_smb2_rsp_hdr like the following one. For smb1 negotiate(56bytes) , init_smb2_rsp_hdr() for smb2 is called. The issue occurs while handling smb1 negotiate as smb2 server operations. Add smb server operations for smb1 (get_cmd_val, init_rsp_hdr, allocate_rsp_buf, check_user_session) to handle smb1 negotiate so that smb2 server operation does not handle it. [411.400423] CIFS: VFS: Use of the less secure dialect vers=1.0 is not recommended unless required for access to very old servers [411.400452] CIFS: Attempting to mount \\192.168.45.139\homes [411.479312] ksmbd: init_smb2_rsp_hdr : 492 [411.479323] == [411.479327] BUG: KASAN: slab-out-of-bounds in init_smb2_rsp_hdr+0x1e2/0x1f4 [ksmbd] [411.479369] Read of size 16 at addr ffff888488ed0734 by task kworker/14:1/199 [411.479379] CPU: 14 PID: 199 Comm: kworker/14:1 Tainted: G OE 6.1.21 #3 [411.479386] Hardware name: ASUSTeK COMPUTER INC. Z10PA-D8 Series/Z10PA-D8 Series, BIOS 3801 08/23/2019 [411.479390] Workqueue: ksmbd-io handle_ksmbd_work [ksmbd] [411.479425] Call Trace: [411.479428] <TASK> [411.479432] dump_stack_lvl+0x49/0x63 [411.479444] print_report+0x171/0x4a8 [411.479452] ? kasan_complete_mode_report_info+0x3c/0x200 [411.479463] ? init_smb2_rsp_hdr+0x1e2/0x1f4 [ksmbd] [411.479497] kasan_report+0xb4/0x130 [411.479503] ? init_smb2_rsp_hdr+0x1e2/0x1f4 [ksmbd] [411.479537] kasan_check_range+0x149/0x1e0 [411.479543] memcpy+0x24/0x70 [411.479550]	N/A	More Details

	init_smb2_rsp_hdr+0x1e2/0x1f4 [ksmbd] [411.479585] handle_ksmbd_work+0x109/0x760 [ksmbd] [411.479616] ? _raw_spin_unlock_irqrestore+0x50/0x50 [411.479624] ? smb3_encrypt_resp+0x340/0x340 [ksmbd] [411.479656] process_one_work+0x49c/0x790 [411.479667] worker_thread+0x2b1/0x6e0 [411.479674] ? process_one_work+0x790/0x790 [411.479680] kthread+0x177/0x1b0 [411.479686] ? kthread_complete_and_exit+0x30/0x30 [411.479692] ret_from_fork+0x22/0x30 [411.479702] </TASK>		
CVE-2023-54195	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix timeout of a call that hasn't yet been granted a channel afs_make_call() calls rxrpc_kernel_begin_call() to begin a call (which may get stalled in the background waiting for a connection to become available); it then calls rxrpc_kernel_set_max_life() to set the timeouts - but that starts the call timer so the call timer might then expire before we get a connection assigned - leading to the following oops if the call stalled: BUG: kernel NULL pointer dereference, address: 0000000000000000 ... CPU: 1 PID: 5111 Comm: krxrpcio/0 Not tainted 6.3.0-rc7-build3+ #701 RIP: 0010:rxrpc_alloc_txbuf+0xc0/0x157 ... Call Trace: <TASK> rxrpc_send_ACK+0x50/0x13b rxrpc_input_call_event+0x16a/0x67d rxrpc_io_thread+0x1b6/0x45f ? _raw_spin_unlock_irqrestore+0x1f/0x35 ? rxrpc_input_packet+0x519/0x519 kthread+0xe7/0xef ? kthread_complete_and_exit+0x1b/0x1b ret_from_fork+0x22/0x30 Fix this by noting the timeouts in struct rxrpc_call when the call is created. The timer will be started when the first packet is transmitted. It shouldn't be possible to trigger this directly from userspace through AF_RXRPC as sendmsg() will return EBUSY if the call is in the waiting-for-conn state if it dropped out of the wait due to a signal.	N/A	More Details
CVE-2023-54196	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix NULL pointer dereference in 'ni_write_inode' Syzbot found the following issue: Unable to handle kernel NULL pointer dereference at virtual address 0000000000000016 Mem abort info: ESR = 0x0000000096000006 EC = 0x25: DABT (current EL), IL = 32 bits SET = 0, FnV = 0 EA = 0, S1PTW = 0 FSC = 0x06: level 2 translation fault Data abort info: ISV = 0, ISS = 0x00000006 CM = 0, WnR = 0 user pgtable: 4k pages, 48-bit VAs, pgdp=000000010af56000 [0000000000000016] pgd=08000001090da003, p4d=08000001090da003, pud=08000001090ce003, pmd=0000000000000000 Internal error: Oops: 0000000096000006 [#1] PREEMPT SMP Modules linked in: CPU: 1 PID: 3036 Comm: syz-executor206 Not tainted 6.0.0-rc6-syzkaller-17739-g16c9f284e746 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 08/26/2022 pstate: 80400005 (Nzcv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : is_rec_inuse fs/ntfs3/ntfs.h:313 [inline] pc : ni_write_inode+0xac/0x798 fs/ntfs3/frecord.c:3232 lr : ni_write_inode+0xa0/0x798 fs/ntfs3/frecord.c:3226 sp : ffff8000126c3800 x29: ffff8000126c3860 x28: 0000000000000000 x27: ffff0000c8b02000 x26: ffff0000c7502320 x25: ffff0000c7502288 x24: 0000000000000000 x23: ffff80000cbec91c x22: ffff0000c8b03000 x21: ffff0000c8b02000 x20: 0000000000000001 x19: ffff0000c75024d8 x18: 00000000000000c0 x17: ffff80000dd1b198 x16: ffff80000db59158 x15: ffff0000c4b6b500 x14: 00000000000000b8 x13: 0000000000000000 x12: ffff0000c4b6b500 x11: ff80800008be1b60 x10: 0000000000000000 x9 : ffff0000c4b6b500 x8 : 0000000000000000 x7 : ffff800008be1b50 x6 : 0000000000000000 x5 : 0000000000000000 x4 : 0000000000000001 x3 : 0000000000000000 x2 : 0000000000000008 x1 : 0000000000000001 x0 : 0000000000000000 Call trace: is_rec_inuse fs/ntfs3/ntfs.h:313 [inline] ni_write_inode+0xac/0x798 fs/ntfs3/frecord.c:3232 ntfs_evict_inode+0x54/0x84 fs/ntfs3/inode.c:1744 evict+0xec/0x334 fs/inode.c:665 iput_final fs/inode.c:1748 [inline] iput+0x2c4/0x324 fs/inode.c:1774 ntfs_new_inode+0x7c/0xe0 fs/ntfs3/fsntfs.c:1660 ntfs_create_inode+0x20c/0xe78 fs/ntfs3/inode.c:1278 ntfs_create+0x54/0x74 fs/ntfs3/namei.c:100 lookup_open fs/namei.c:3413 [inline] open_last_lookups fs/namei.c:3481 [inline] path_openat+0x804/0x11c4 fs/namei.c:3688 do_filp_open+0xdc/0x1b8 fs/namei.c:3718 do_sys_openat2+0xb8/0x22c fs/open.c:1311 do_sys_open fs/open.c:1327 [inline] __do_sys_openat fs/open.c:1343 [inline] __se_sys_openat fs/open.c:1338 [inline] __arm64_sys_openat+0xb0/0xe0 fs/open.c:1338 __invoke_syscall arch/arm64/kernel/syscall.c:38 [inline] invoke_syscall arch/arm64/kernel/syscall.c:52 [inline] el0_svc_common+0x138/0x220 arch/arm64/kernel/syscall.c:142 do_el0_svc+0x48/0x164 arch/arm64/kernel/syscall.c:206 el0_svc+0x58/0x150 arch/arm64/kernel/entry-common.c:636 el0t_64_sync_handler+0x84/0xf0 arch/arm64/kernel/entry-common.c:654 el0t_64_sync+0x18c/0x190 Code: 97dafee4 340001b4 f9401328 2a1f03e0 (79402d14) ---[end trace 0000000000000000]--- Above issue may happens as follows: ntfs_new_inode mi_init mi->mrec = kmalloc(sbi->record_size, GFP_NOFS); -->failed to allocate memory if (!mi->mrec) return -ENOMEM; iput iput_final evict ntfs_evict_inode ni_write_inode is_rec_inuse(ni->mi.mrec)-> As 'ni->mi.mrec' is NULL trigger NULL-ptr-deref To solve above issue if new inode failed make inode bad before call 'iput()' in 'ntfs_new_inode()'.	N/A	More Details
CVE-2023-54197	In the Linux kernel, the following vulnerability has been resolved: Revert "Bluetooth: btsdio: fix use after free bug in btsdio_remove due to unfinished work" This reverts commit 1e9ac114c4428fdb7ff4635b45d4f46017e8916f. This patch introduces a possible null-ptr-def problem. Revert it. And the fixed bug by this patch have resolved by commit 73f7b171b7c0 ("Bluetooth: btsdio: fix use after free bug in btsdio_remove due to race condition").	N/A	More Details
CVE-2023-54198	In the Linux kernel, the following vulnerability has been resolved: tty: fix out-of-bounds access in tty_driver_lookup_tty() When specifying an invalid console= device like console=tty3270, tty_driver_lookup_tty() returns the tty struct without checking whether index is a valid number. To reproduce: qemu-system-x86_64 -enable-kvm -nographic -serial mon:stdio \ -kernel ../linux-build-x86/arch/x86/boot/bzImage \ -append "console=ttyS0 console=tty3270" This crashes with: [0.770599] BUG: kernel NULL pointer dereference, address: 00000000000000ef [0.771265] #PF: supervisor read access in kernel mode [0.771773] #PF: error_code(0x0000) - not-present page [0.772609] Oops: 0000 [#1] PREEMPT SMP PTI [0.774878] RIP: 0010:tty_open+0x268/0x6f0 [0.784013] chrdev_open+0xbd/0x230 [0.784444] ? cdev_device_add+0x80/0x80 [0.784920] do_dentry_open+0x1e0/0x410 [0.785389] path_openat+0xca9/0x1050 [0.785813] do_filp_open+0xaa/0x150 [0.786240] file_open_name+0x133/0x1b0 [0.786746] filp_open+0x27/0x50 [0.787244] console_on_rootfs+0x14/0x4d [0.787800] kernel_init_freeable+0x1e4/0x20d [0.788383] ? rest_init+0xc0/0xc0 [0.788881] kernel_init+0x11/0x120 [0.789356] ret_from_fork+0x22/0x30	N/A	More Details
CVE-2023-54199	In the Linux kernel, the following vulnerability has been resolved: drm/msm/adreno: Fix null ptr access in adreno_gpu_cleanup() Fix the below kernel panic due to null pointer access: [18.504431] Unable to handle kernel NULL pointer dereference at virtual address 0000000000000048 [18.513464] Mem abort info: [18.516346] ESR = 0x0000000096000005 [18.520204] EC = 0x25: DABT (current EL), IL = 32 bits [18.525706] SET = 0, FnV = 0 [18.528878] EA = 0, S1PTW = 0 [18.532117] FSC = 0x05: level 1 translation fault [18.537138] Data abort info: [18.540110] ISV = 0, ISS = 0x00000005 [18.544060] CM = 0, WnR = 0 [18.547109] user pgtable: 4k pages, 39-bit VAs, pgdp=0000000112826000 [18.553738] [0000000000000048] pgd=0000000000000000, p4d=0000000000000000, pud=0000000000000000 [18.562690] Internal error: Oops: 0000000096000005 [#1] PREEMPT SMP **Snip** [18.696758] Call trace: [18.699278] adreno_gpu_cleanup+0x30/0x88 [18.703396] a6xx_destroy+0xc0/0x130 [18.707066] a6xx_gpu_init+0x308/0x424 [18.710921] adreno_bind+0x178/0x288 [18.714590] component_bind_all+0xe0/0x214 [18.718797] msm_drm_bind+0x1d4/0x614 [18.722566] try_to_bring_up_aggregate_device+0x16c/0x1b8 [18.728105] __component_add+0xa0/0x158 [18.732048] component_add+0x20/0x2c [18.735719] adreno_probe+0x40/0xc0 [18.739300] platform_probe+0xb4/0xd4 [18.743068] really_probe+0xfc/0x284 [18.746738] __driver_probe_device+0xc0/0xec [18.751129] driver_probe_device+0x48/0x110 [18.755421] __device_attach_driver+0xa8/0xd0 [18.759900] bus_for_each_drv+0x90/0xdc [18.763843] __device_attach+0xfc/0x174 [18.767786] device_initial_probe+0x20/0x2c [18.772090] bus_probe_device+0x40/0xa0 [18.776032] deferred_probe_work_func+0x94/0xd0 [18.780686] process_one_work+0x190/0x3d0 [18.784805] worker_thread+0x280/0x3d4 [18.788659] kthread+0x104/0x1c0 [18.791981] ret_from_fork+0x10/0x20 [18.795654] Code: f9400408 aa0003f3 aa1f03f4 91142015 (f9402516) [18.801913] ---[end trace 0000000000000000]--- [18.809039] Kernel panic -	N/A	More Details

	not syncing: Oops: Fatal exception Patchwork: https://patchwork.freedesktop.org/patch/515605/		
CVE-2023-54200	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: always release netdev hooks from notifier This reverts "netfilter: nf_tables: skip netdev events generated on netns removal". The problem is that when a veth device is released, the veth release callback will also queue the peer netns device for removal. Its possible that the peer netns is also slated for removal. In this case, the device memory is already released before the pre_exit hook of the peer netns runs: BUG: KASAN: slab-use-after-free in nf_hook_entry_head+0x1b8/0x1d0 Read of size 8 at addr ffff88812c0124f0 by task kworker/u8:1/45 Workqueue: netns cleanup_net Call Trace: nf_hook_entry_head+0x1b8/0x1d0 __nf_unregister_net_hook+0x76/0x510 nft_netdev_unregister_hooks+0xa0/0x220 __nft_release_hook+0x184/0x490 nf_tables_pre_exit_net+0x12f/0x1b0 .. Order is: 1. First netns is released, veth_dellink() queues peer netns device for removal 2. peer netns is queued for removal 3. peer netns device is released, unreg event is triggered 4. unreg event is ignored because netns is going down 5. pre_exit hook calls nft_netdev_unregister_hooks but device memory might be free'd already.	N/A	More Details
CVE-2023-54201	In the Linux kernel, the following vulnerability has been resolved: RDMA/efa: Fix wrong resources deallocation order When trying to destroy QP or CQ, we first decrease the refcount and potentially free memory regions allocated for the object and then request the device to destroy the object. If the device fails, the object isn't fully destroyed so the user/IB core can try to destroy the object again which will lead to underflow when trying to decrease an already zeroed refcount. Deallocate resources in reverse order of allocating them to safely free them.	N/A	More Details
CVE-2023-54202	In the Linux kernel, the following vulnerability has been resolved: drm/i915: fix race condition UAF in i915_perf_add_config_ioctl Userspace can guess the id value and try to race oa_config object creation with config remove, resulting in a use-after-free if we dereference the object after unlocking the metrics_lock. For that reason, unlocking the metrics_lock must be done after we are done dereferencing the object. [tursulin: Manually added stable tag.] (cherry picked from commit 49f6f6483b652108bcb73accd0204a464b922395)	N/A	More Details
CVE-2023-54204	In the Linux kernel, the following vulnerability has been resolved: mmc: sunplus: fix return value check of mmc_add_host() mmc_add_host() may return error, if we ignore its return value, 1. the memory allocated in mmc_alloc_host() will be leaked 2. null-ptr-deref will happen when calling mmc_remove_host() in remove function spmmc_drv_remove() because deleting not added device. Fix this by checking the return value of mmc_add_host(). Moreover, I fixed the error handling path of spmmc_drv_probe() to clean up.	N/A	More Details
CVE-2023-54213	In the Linux kernel, the following vulnerability has been resolved: USB: sisusbvga: Add endpoint checks The syzbot fuzzer was able to provoke a WARNING from the sisusbvga driver: ------------[cut here]------------ usb 1-1: BOGUS urb xfer, pipe 3 != type 1 WARNING: CPU: 1 PID: 26 at drivers/usb/core/urb.c:504 usb_submit_urb+0xed6/0x1880 drivers/usb/core/urb.c:504 Modules linked in: CPU: 1 PID: 26 Comm: kworker/1:1 Not tainted 6.2.0-rc5-syzkaller-00199-g5af6ce704936 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/12/2023 Workqueue: usb_hub_wq hub_event RIP: 0010:usb_submit_urb+0xed6/0x1880 drivers/usb/core/urb.c:504 Code: 7c 24 18 e8 6c 50 80 fb 48 8b 7c 24 18 e8 62 1a 01 ff 41 89 d8 44 89 e1 4c 89 ea 48 89 c6 48 c7 c7 60 b1 fa 8a e8 84 b0 be 03 <0f> 0b e9 58 f8 ff ff e8 3e 50 80 fb 48 81 c5 c0 05 00 00 e9 84 f7 RSP: 0018:ffffc90000a1ed18 EFLAGS: 00010282 RAX: 0000000000000000 RBX: 0000000000000001 RCX: 0000000000000000 RDX: ffff888012783a80 RSI: ffffffff816680ec RDI: fffff52000143d95 RBP: ffff888079020000 R08: 0000000000000005 R09: 0000000000000000 R10: 0000000080000000 R11: 0000000000000000 R12: 0000000000000003 R13: ffff888017d33370 R14: 0000000000000003 R15: ffff888021213600 FS: 0000000000000000(0000) GS:ffff8880b9900000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00005592753a60b0 CR3: 0000000022899000 CR4: 00000000003506e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> sisusb_bulkout_msg drivers/usb/misc/sisusbvga/sisusbvga.c:224 [inline] sisusb_send_bulk_msg.constprop.0+0x904/0x1230 drivers/usb/misc/sisusbvga/sisusbvga.c:379 sisusb_send_bridge_packet drivers/usb/misc/sisusbvga/sisusbvga.c:567 [inline] sisusb_do_init_gfxdevice drivers/usb/misc/sisusbvga/sisusbvga.c:2077 [inline] sisusb_init_gfxdevice+0x87b/0x4000 drivers/usb/misc/sisusbvga/sisusbvga.c:2177 sisusb_probe+0x9cd/0xbe2 drivers/usb/misc/sisusbvga/sisusbvga.c:2869 ... The problem was caused by the fact that the driver does not check whether the endpoints it uses are actually present and have the appropriate types. This can be fixed by adding a simple check of the endpoints.	N/A	More Details
CVE-2023-54205	In the Linux kernel, the following vulnerability has been resolved: pinctrl: stm32: Fix refcount leak in stm32_pctrl_get_irq_domain of_irq_find_parent() returns a node pointer with refcount incremented, We should use of_node_put() on it when not needed anymore. Add missing of_node_put() to avoid refcount leak.	N/A	More Details
CVE-2023-54206	In the Linux kernel, the following vulnerability has been resolved: net/sched: flower: fix filter idr initialization The cited commit moved idr initialization too early in fl_change() which allows concurrent users to access the filter that is still being initialized and is in inconsistent state, which, in turn, can cause NULL pointer dereference [0]. Since there is no obvious way to fix the ordering without reverting the whole cited commit, alternative approach taken to first insert NULL pointer into idr in order to allocate the handle but still cause fl_get() to return NULL and prevent concurrent users from seeing the filter while providing miss-to-action infrastructure with valid handle id early in fl_change(). [152.434728] general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASAN [152.436163] KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] [152.437269] CPU: 4 PID: 3877 Comm: tc Not tainted 6.3.0-rc4+ #5 [152.438110] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 [152.439644] RIP: 0010:fl_dump_key+0x8b/0x1d10 [cls_flower] [152.440461] Code: 01 f2 02 f2 c7 40 08 04 f2 04 f2 c7 40 0c 04 f3 f3 f3 65 48 8b 04 25 28 00 00 00 48 89 84 24 00 01 00 00 48 89 c8 48 c1 e8 03 <0f> b6 04 10 84 c0 74 08 3c 03 0f 8e 98 19 00 00 8b 13 85 d2 74 57 [152.442885] RSP: 0018:ffff88817a28f158 EFLAGS: 00010246 [152.443851] RAX: 0000000000000000 RBX: 0000000000000000 RCX: 0000000000000000 [152.444826] RDX: dffffc0000000000 RSI: ffffffff8500ae80 RDI: ffff88810a987900 [152.445791] RBP: ffff888179d88240 R08: ffff888179d8845c R09: ffff888179d88240 [152.446780] R10: ffffed102f451e48 R11: 00000000ffffffff2 R12: ffff88810a987900 [152.447741] R13: ffffffff8500ae80 R14: ffff88810a987900 R15: ffff888149b3c738 [152.448756] FS: 00007f5eb2a34800(0000) GS:ffff88881ec00000(0000) knlGS:0000000000000000 [152.449888] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [152.450685] CR2: 000000000046ad19 CR3: 000000010b0bd006 CR4: 0000000000370ea0 [152.451641] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [152.452628] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [152.453588] Call Trace: [152.454032] <TASK> [152.454447] ? netlink_sendmsg+0x7a1/0xcb0 [152.455109] ? sock_sendmsg+0xc5/0x190 [152.455689] ? ____sys_sendmsg+0x535/0x6b0 [152.456320] ? ___sys_sendmsg+0xeb/0x170 [152.456916] ? do_syscall_64+0x3d/0x90 [152.457529] ? entry_SYSCALL_64_after_hwframe+0x46/0xb0 [152.458321] ? ___sys_sendmsg+0xeb/0x170 [152.458958] ? __sys_sendmsg+0xb5/0x140 [152.459564] ? do_syscall_64+0x3d/0x90 [152.460122] ? entry_SYSCALL_64_after_hwframe+0x46/0xb0 [152.460852] ? fl_dump_key_options.part.0+0xea0/0xea0 [cls_flower] [152.461710] ? _raw_spin_lock+0x7a/0xd0 [152.462299] ? _raw_read_lock_irq+0x30/0x30 [152.462924] ? nla_put+0x15e/0x1c0 [152.463480] fl_dump+0x228/0x650 [cls_flower] [152.464112] ? fl_tmplt_dump+0x210/0x210 [cls_flower] [152.464854] ? __kmem_cache_alloc_node+0x1a7/0x330 [152.465592] ? nla_put+0x15e/0x1c0 [152.466160] tcf_fill_node+0x515/0x9a0 [	N/A	More Details

	152.466766] ? tc_setup_offload_action+0xf0/0xf0 [152.467463] ? __alloc_skb+0x13c/0x2a0 [152.468067] ? __build_skb_around+0x330/0x330 [152.468814] ? fl_get+0x107/0x1a0 [cls_flower] [152.469503] tc_del_tfilter+0x718/0x1330 [152.470115] ? is_bpf_text_address+0xa/0x20 [152.470765] ? tc_ctl_chain+0xee0/0xee0 [152.471335] ? __kernel_text_address+0xe/0x30 [152.471948] ? unwind_get_return_address+0x56/0xa0 [152.472639] ? __thaw_task+0x150/0x150 [152.473218] ? arch_stack_walk+0x98/0xf0 [152.473839] ? __stack_depot_save+0x35/0x4c0 [152.474501] ? stack_trace_save+0x91/0xc0 [152.475119] ? security_capable+0x51/0x90 [152.475741] rtnetlink_rcv_msg+0x2c1/0x9d0 [152.476387] ? rtnl_calcit.isra.0+0x2b0/0x2b0 [152.477042] ---truncated---		
CVE-2023-54207	In the Linux kernel, the following vulnerability has been resolved: HID: uclogic: Correct devm device reference for hidinput input_dev name Reference the HID device rather than the input device for the devm allocation of the input_dev name. Referencing the input_dev would lead to a use-after-free when the input_dev was unregistered and subsequently fires a uevent that depends on the name. At the point of firing the uevent, the name would be freed by devres management. Use devm_kasprintf to simplify the logic for allocating memory and formatting the input_dev name string.	N/A	More Details
CVE-2023-54208	In the Linux kernel, the following vulnerability has been resolved: media: ov5675: Fix memleak in ov5675_init_controls() There is a kmemleak when testing the media/i2c/ov5675.c with bpf mock device: AssertionError: unreferenced object 0xffff888107362160 (size 16): comm "python3", pid 277, jiffies 4294832798 (age 20.722s) hex dump (first 16 bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<00000000abe7d67c>] __kmalloc_node+0x44/0x1b0 [<000000008a725aac>] kvmalloc_node+0x34/0x180 [<000000009a53cd11>] v4l2_ctrl_handler_init_class+0x11d/0x180 [videodev] [<0000000055b46db0>] ov5675_probe+0x38b/0x897 [ov5675] [<00000000153d886c>] i2c_device_probe+0x28d/0x680 [<000000004afb7e8f>] really_probe+0x17c/0x3f0 [<00000000ff2f18e4>] __driver_probe_device+0xe3/0x170 [<000000000a001029>] driver_probe_device+0x49/0x120 [<00000000e39743c7>] __device_attach_driver+0xf7/0x150 [<00000000d32fd070>] bus_for_each_drv+0x114/0x180 [<000000009083ac41>] __device_attach+0x1e5/0x2d0 [<0000000015b4a830>] bus_probe_device+0x126/0x140 [<000000007813deaf>] device_add+0x810/0x1130 [<000000007becb867>] i2c_new_client_device+0x386/0x540 [<000000007f9cf4b4>] of_i2c_register_device+0xf1/0x110 [<00000000ebfdd032>] of_i2c_notify+0xfc/0x1f0 ov5675_init_controls() won't clean all the allocated resources in fail path, which may causes the memleaks. Add v4l2_ctrl_handler_free() to prevent memleak.	N/A	More Details
CVE-2023-54209	In the Linux kernel, the following vulnerability has been resolved: block: fix blktrace debugfs entries leakage Commit 99d055b4fd4b ("block: remove per-disk debugfs files in blk_unregister_queue") moves blk_trace_shutdown() from blk_release_queue() to blk_unregister_queue(), this is safe if blktrace is created through sysfs, however, there is a regression in corner case. blktrace can still be enabled after del_gendisk() through ioctl if the disk is opened before del_gendisk(), and if blktrace is not shutdown through ioctl before closing the disk, debugfs entries will be leaked. Fix this problem by shutdown blktrace in disk_release(), this is safe because blk_trace_remove() is reentrant.	N/A	More Details
CVE-2023-54210	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_sync: Avoid use-after-free in dbg for hci_remove_adv_monitor() KASAN reports that there's a use-after-free in hci_remove_adv_monitor(). Trawling through the disassembly, you can see that the complaint is from the access in bt_dev_dbg() under the HCI_ADV_MONITOR_EXT_MSFT case. The problem case happens because msft_remove_monitor() can end up freeing the monitor structure. Specifically: hci_remove_adv_monitor() -> msft_remove_monitor() -> msft_remove_monitor_sync() -> msft_le_cancel_monitor_advertisement_cb() -> hci_free_adv_monitor() Let's fix the problem by just stashing the relevant data when it's still valid.	N/A	More Details
CVE-2023-54211	In the Linux kernel, the following vulnerability has been resolved: tracing: Fix warning in trace_buffered_event_disable() Warning happened in trace_buffered_event_disable() at WARN_ON_ONCE(!trace_buffered_event_ref) Call Trace: ? __warn+0xa5/0x1b0 ? trace_buffered_event_disable+0x189/0x1b0 __ftrace_event_enable_disable+0x19e/0x3e0 free_probe_data+0x3b/0xa0 unregister_ftrace_function_probe_func+0x6b8/0x800 event_enable_func+0x2f0/0x3d0 ftrace_process_regex.isra.0+0x12d/0x1b0 ftrace_filter_write+0xe6/0x140 vfs_write+0x1c9/0x6f0 [...] The cause of the warning is in __ftrace_event_enable_disable(), trace_buffered_event_enable() was called once while trace_buffered_event_disable() was called twice. Reproduction script show as below, for analysis, see the comments: ``` #!/bin/bash cd /sys/kernel/tracing/ # 1. Register a 'disable_event' command, then: # 1) SOFT_DISABLED_BIT was set; # 2) trace_buffered_event_enable() was called first time; echo 'cmdline_proc_show:disable_event:initcall:initcall_finish' > \ set_ftrace_filter # 2. Enable the event registered, then: # 1) SOFT_DISABLED_BIT was cleared; # 2) trace_buffered_event_disable() was called first time; echo 1 > events/initcall/initcall_finish/enable # 3. Try to call into cmdline_proc_show(), then SOFT_DISABLED_BIT was # set again!!! cat /proc/cmdline # 4. Unregister the 'disable_event' command, then: # 1) SOFT_DISABLED_BIT was cleared again; # 2) trace_buffered_event_disable() was called second time!!! echo '!cmdline_proc_show:disable_event:initcall:initcall_finish' > \ set_ftrace_filter ``` To fix it, IIUC, we can change to call trace_buffered_event_enable() at fist time soft-mode enabled, and call trace_buffered_event_disable() at last time soft-mode disabled.	N/A	More Details
CVE-2023-54212	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2022-50884	In the Linux kernel, the following vulnerability has been resolved: drm: Prevent drm_copy_field() to attempt copying a NULL pointer There are some struct drm_driver fields that are required by drivers since drm_copy_field() attempts to copy them to user-space via DRM_IOCTL_VERSION. But it can be possible that a driver has a bug and did not set some of the fields, which leads to drm_copy_field() attempting to copy a NULL pointer: [+10.395966] Unable to handle kernel access to user memory outside uaccess routines at virtual address 0000000000000000 [+0.010955] Mem abort info: [+0.002835] ESR = 0x0000000096000004 [+0.003872] EC = 0x25: DABT (current EL), IL = 32 bits [+0.005395] SET = 0, FnV = 0 [+0.003113] EA = 0, S1PTW = 0 [+0.003182] FSC = 0x04: level 0 translation fault [+0.004964] Data abort info: [+0.002919] ISV = 0, ISS = 0x00000004 [+0.003886] CM = 0, WnR = 0 [+0.003040] user pgtable: 4k pages, 48-bit VAs, pgdp=0000000115dad000 [+0.006536] [0000000000000000] pgd=0000000000000000, p4d=0000000000000000 [+0.006925] Internal error: Oops: 96000004 [#1] SMP ... [+0.011113] pstate: 80400005 (Nzcv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) [+0.007061] pc : __pi_strlen+0x14/0x150 [+0.003895] lr : drm_copy_field+0x30/0x1a4 [+0.004156] sp : ffff8000094b3a50 [+0.003355] x29: ffff8000094b3a50 x28: ffff8000094b3b70 x27: 0000000000000040 [+0.007242] x26: ffff443743c2ba00 x25: 0000000000000000 x24: 0000000000000040 [+0.007243] x23: ffff443743c2ba00 x22: ffff8000094b3b70 x21: 0000000000000000 [+0.007241] x20: 0000000000000000 x19: ffff8000094b3b90 x18: 0000000000000000 [+0.007241] x17: 0000000000000000 x16: 0000000000000000 x15: 0000aaab14b9af40 [+0.007241] x14: 0000000000000000 x13: 0000000000000000 x12: 0000000000000000 [+0.007239] x11: 0000000000000000 x10: 0000000000000000 x9 : ffffa524ad67d4d8 [+0.007242] x8 : 0101010101010101 x7 : 7f7f7f7f7f7f7f7f x6 : 6c6e6263606e7141 [+0.007239] x5 : 0000000000000000 x4 : 0000000000000000 x3 : 0000000000000000 [+0.007241] x2 : 0000000000000000 x1 : ffff8000094b3b90 x0 : 0000000000000000 [+0.007240] Call trace: [+0.002475] __pi_strlen+0x14/0x150 [+0.003537] drm_version+0x84/0xac [+0.003448] drm_ioctl_kernel+0xa8/0x16c [+0.003975] drm_ioctl+0x270/0x580 [+0.003448] __arm64_sys_ioctl+0xb8/0xfc [+0.003978] invoke_syscall+0x78/0x100 [+0.003799]	N/A	More Details

	el0_svc_common.constprop.0+0x4c/0xf4 [+0.004767] do_el0_svc+0x38/0x4c [+0.003357] el0_svc+0x34/0x100 [+0.003185] el0t_64_sync_handler+0x11c/0x150 [+0.004418] el0t_64_sync+0x190/0x194 [+0.003716] Code: 92402c04 b200c3e8 f13fc09f 5400088c (a9400c02) [+0.006180] ---[end trace 0000000000000000]---		
CVE-2022-50881	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: Fix use-after-free in ath9k_hif_usb_disconnect() This patch fixes a use-after-free in ath9k that occurs in ath9k_hif_usb_disconnect() when ath9k_destroy_wmi() is trying to access 'drv_priv' that has already been freed by ieee80211_free_hw(), called by ath9k_htc_hw_deinit(). The patch moves ath9k_destroy_wmi() before ieee80211_free_hw(). Note that urbs from the driver should be killed before freeing 'wmi' with ath9k_destroy_wmi() as their callbacks will access 'wmi'. Found by a modified version of syzkaller. == BUG: KASAN: use-after-free in ath9k_destroy_wmi+0x38/0x40 Read of size 8 at addr ffff8881069132a0 by task kworker/0:1/7 CPU: 0 PID: 7 Comm: kworker/0:1 Tainted: G O 5.14.0+ #131 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.12.1-0-ga5cab58e9a3f-prebuilt.qemu.org 04/01/2014 Workqueue: usb_hub_wq hub_event Call Trace: dump_stack_lvl+0x8e/0xd1 print_address_description.constprop.0.cold+0x93/0x334 ? ath9k_destroy_wmi+0x38/0x40 ? ath9k_destroy_wmi+0x38/0x40 kasan_report.cold+0x83/0xdf ? ath9k_destroy_wmi+0x38/0x40 ath9k_destroy_wmi+0x38/0x40 ath9k_hif_usb_disconnect+0x329/0x3f0 ? ath9k_hif_usb_suspend+0x120/0x120 ? usb_disable_interface+0xfc/0x180 usb_unbind_interface+0x19b/0x7e0 ? usb_autoresume_device+0x50/0x50 device_release_driver_internal+0x44d/0x520 bus_remove_device+0x2e5/0x5a0 device_del+0x5b2/0xe30 ? __device_link_del+0x370/0x370 ? usb_remove_ep_devs+0x43/0x80 ? remove_intf_ep_devs+0x112/0x1a0 usb_disable_device+0x1e3/0x5a0 usb_disconnect+0x267/0x870 hub_event+0x168d/0x3950 ? rcu_read_lock_sched_held+0xa1/0xd0 ? hub_port_debounce+0x2e0/0x2e0 ? check_irq_usage+0x860/0xf20 ? drain_workqueue+0x281/0x360 ? lock_release+0x640/0x640 ? rcu_read_lock_sched_held+0xa1/0xd0 ? rcu_read_lock_bh_held+0xb0/0xb0 ? lockdep_hardirqs_on_prepare+0x273/0x3e0 process_one_work+0x92b/0x1460 ? pwq_dec_nr_in_flight+0x330/0x330 ? rwlock_bug.part.0+0x90/0x90 worker_thread+0x95/0xe00 ? __kthread_parkme+0x115/0x1e0 ? process_one_work+0x1460/0x1460 kthread+0x3a1/0x480 ? set_kthread_struct+0x120/0x120 ret_from_fork+0x1f/0x30 The buggy address belongs to the page: page:ffffea00041a44c0 refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x106913 flags: 0x200000000000000(node=0|zone=2) raw: 0200000000000000 0000000000000000 dead000000000122 0000000000000000 raw: 0000000000000000 0000000000000000 00000000ffffffff 0000000000000000 page dumped because: kasan: bad access detected page_owner tracks the page as freed page last allocated via order 3, migratetype Unmovable, gfp_mask 0x40dc0(GFP_KERNEL|__GFP_COMP|__GFP_ZERO), pid 7, ts 38347963444, free_ts 41399957635 prep_new_page+0x1aa/0x240 get_page_from_freelist+0x159a/0x27c0 __alloc_pages+0x2da/0x6a0 alloc_pages+0xec/0x1e0 kmalloc_order+0x39/0xf0 kmalloc_order_trace+0x19/0x120 __kmalloc+0x308/0x390 wiphy_new_nm+0x6f5/0x1dd0 ieee80211_alloc_hw_nm+0x36d/0x2230 ath9k_htc_probe_device+0x9d/0x1e10 ath9k_htc_hw_init+0x34/0x50 ath9k_hif_usb_firmware_cb+0x25f/0x4e0 request_firmware_work_func+0x131/0x240 process_one_work+0x92b/0x1460 worker_thread+0x95/0xe00 kthread+0x3a1/0x480 page last free stack trace: free_pcp_prepare+0x3d3/0x7f0 free_unref_page+0x1e/0x3d0 device_release+0xa4/0x240 kobject_put+0x186/0x4c0 put_device+0x20/0x30 ath9k_htc_disconnect_device+0x1cf/0x2c0 ath9k_htc_hw_deinit+0x26/0x30 ath9k_hif_usb_disconnect+0x2d9/0x3f0 usb_unbind_interface+0x19b/0x7e0 device_release_driver_internal+0x44d/0x520 bus_remove_device+0x2e5/0x5a0 device_del+0x5b2/0xe30 usb_disable_device+0x1e3/0x5a0 usb_disconnect+0x267/0x870 hub_event+0x168d/0x3950 process_one_work+0x92b/0x1460 Memory state around the buggy address: ffff888106913180: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ffff888106913200: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff >ffff888 ---truncated---	N/A	More Details
CVE-2022-50882	In the Linux kernel, the following vulnerability has been resolved: media: uvcvideo: Fix memory leak in uvc_gpio_parse Previously the unit buffer was allocated before checking the IRQ for privacy GPIO. In case of error, the unit buffer was leaked. Allocate the unit buffer after the IRQ to avoid it. Addresses-Coverity-ID: 1474639 ("Resource leak")	N/A	More Details
CVE-2023-54236	In the Linux kernel, the following vulnerability has been resolved: net/net_failover: fix txq exceeding warning The failover txq is inited as 16 queues. when a packet is transmitted from the failover device firstly, the failover device will select the queue which is returned from the primary device if the primary device is UP and running. If the primary device txq is bigger than the default 16, it can lead to the following warning: eth0 selects TX queue 18, but real number of TX queues is 16 The warning backtrace is: [32.146376] CPU: 18 PID: 9134 Comm: chronyd Tainted: G E 6.2.8-1.el7.centos.x86_64 #1 [32.147175] Hardware name: Red Hat KVM, BIOS 1.10.2-3.el7_4.1 04/01/2014 [32.147730] Call Trace: [32.147971] <TASK> [32.148183] dump_stack_lvl+0x48/0x70 [32.148514] dump_stack+0x10/0x20 [32.148820] netdev_core_pick_tx+0xb1/0xe0 [32.149180] __dev_queue_xmit+0x529/0xcf0 [32.149533] ? __check_object_size.part.0+0x21c/0x2c0 [32.149967] ip_finish_output2+0x278/0x560 [32.150327] __ip_finish_output+0x1fe/0x2f0 [32.150690] ip_finish_output+0x2a/0xd0 [32.151032] ip_output+0x7a/0x110 [32.151337] ? __pfx_ip_finish_output+0x10/0x10 [32.151733] ip_local_out+0x5e/0x70 [32.152054] ip_send_skb+0x19/0x50 [32.152366] udp_send_skb.isra.0+0x163/0x3a0 [32.152736] udp_sendmsg+0xba8/0xec0 [32.153060] ? __folio_memcg_unlock+0x25/0x60 [32.153445] ? __pfx_ip_generic_getfrag+0x10/0x10 [32.153854] ? sock_has_perm+0x85/0xa0 [32.154190] inet_sendmsg+0x6d/0x80 [32.154508] ? inet_sendmsg+0x6d/0x80 [32.154838] sock_sendmsg+0x62/0x70 [32.155152] ____sys_sendmsg+0x134/0x290 [32.155499] ___sys_sendmsg+0x81/0xc0 [32.155828] ? _get_random_bytes.part.0+0x79/0x1a0 [32.156240] ? ip4_datagram_release_cb+0x5f/0x1e0 [32.156649] ? get_random_u16+0x69/0xf0 [32.156989] ? __fget_light+0xcf/0x110 [32.157326] __sys_sendmmsg+0xc4/0x210 [32.157657] ? __sys_connect+0xb7/0xe0 [32.157995] ? __audit_syscall_entry+0xce/0x140 [32.158388] ? syscall_trace_enter.isra.0+0x12c/0x1a0 [32.158820] __x64_sys_sendmmsg+0x24/0x30 [32.159171] do_syscall_64+0x38/0x90 [32.159493] entry_SYSCALL_64_after_hwframe+0x72/0xdc Fix that by reducing txq number as the non-existent primary-dev does.	N/A	More Details
CVE-2022-50818	In the Linux kernel, the following vulnerability has been resolved: scsi: pm8001: Fix running_req for internal abort commands Disabling the remote phy for a SATA disk causes a hang: root@(none)$ more /sys/class/sas_phy/phy-0:0:8/target_port_protocols sata root@(none)$ echo 0 > sys/class/sas_phy/phy-0:0:8/enable root@(none)$ [67.855950] sas: ex 500e004aaaaaaa1f phy08 change count has changed [67.920585] sd 0:0:2:0: [sdc] Synchronizing SCSI cache [67.925780] sd 0:0:2:0: [sdc] Synchronize Cache(10) failed: Result: hostbyte=0x04 driverbyte=DRIVER_OK [67.935094] sd 0:0:2:0: [sdc] Stopping disk [67.939305] sd 0:0:2:0: [sdc] Start/Stop Unit failed: Result: hostbyte=0x04 driverbyte=DRIVER_OK ... [123.998998] INFO: task kworker/u192:1:642 blocked for more than 30 seconds. [124.005960] Not tainted 6.0.0-rc1-205202-gf26f8f761e83 #218 [124.012049] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message. [124.019872] task:kworker/u192:1 state:D stack:0 pid: 642 ppid: 2 flags:0x00000008 [124.028223] Workqueue: 0000:04:00.0_event_q sas_port_event_worker [124.034319] Call trace: [124.036758] __switch_to+0x128/0x278 [124.040333] __schedule+0x434/0xa58 [124.043820] schedule+0x94/0x138 [124.047045] schedule_timeout+0x2fc/0x368 [124.051052] wait_for_completion+0xdc/0x200 [124.055234] __flush_workqueue+0x1a8/0x708 [124.059328] sas_porte_broadcast_rcvd+0xa8/0xc0 [124.063858] sas_port_event_worker+0x60/0x98 [124.068126] process_one_work+0x3f8/0x660 [124.072134] worker_thread+0x70/0x700 [124.075793] kthread+0x1a4/0x1b8 [124.079014] ret_from_fork+0x10/0x20 The issue is that the per-device running_req read in pm8001_dev_gone_notify() never goes to zero and we never make progress. This is caused by missing accounting for running_req for when an internal abort command completes. In commit 2cbbf489778e ("scsi: pm8001: Use libsas internal abort support") we started to send internal abort commands as a proper sas_task. In this when we deliver a sas_task to HW the per-device running_req is incremented in pm8001_queue_command().	N/A	More Details

	However it is never decremented for internal abort commnds, so decrement in pm8001_mpi_task_abort_resp().		
CVE-2022-50819	In the Linux kernel, the following vulnerability has been resolved: udmabuf: Set ubuf->sg = NULL if the creation of sg table fails When userspace tries to map the dmabuf and if for some reason (e.g. OOM) the creation of the sg table fails, ubuf->sg needs to be set to NULL. Otherwise, when the userspace subsequently closes the dmabuf fd, we'd try to erroneously free the invalid sg table from release_udmabuf resulting in the following crash reported by syzbot: general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] PREEMPT SMP KASAN KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] CPU: 0 PID: 3609 Comm: syz-executor487 Not tainted 5.19.0-syzkaller-13930-g7ebfc85e2cd7 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 07/22/2022 RIP: 0010:dma_unmap_sgtable include/linux/dma-mapping.h:378 [inline] RIP: 0010:put_sg_table drivers/dma-buf/udmabuf.c:89 [inline] RIP: 0010:release_udmabuf+0xcb/0x4f0 drivers/dma-buf/udmabuf.c:114 Code: 48 89 fa 48 c1 ea 03 80 3c 02 00 0f 85 2b 04 00 00 48 8d 7d 0c 4c 8b 63 30 48 b8 00 00 00 00 00 fc ff df 48 89 fa 48 c1 ea 03 <0f> b6 14 02 48 89 f8 83 e0 07 83 c0 03 38 d0 7c 08 84 d2 0f 85 e2 RSP: 0018:ffffc900037efd30 EFLAGS: 00010246 RAX: dffffc0000000000 RBX: ffffffff8cb67800 RCX: 0000000000000000 RDX: 0000000000000000 RSI: ffffffff84ad27e0 RDI: 0000000000000000 RBP: fffffffffffffff4 R08: 0000000000000005 R09: 0000000000000000 R10: 0000000000000000 R11: 000000000008c07c R12: ffff88801fa05000 R13: ffff888073db07e8 R14: ffff888025c25440 R15: 0000000000000000 FS: 0000555555fc4300(0000) GS:ffff8880b9a00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fc1c0ce06e4 CR3: 00000000715e6000 CR4: 0000000003506f0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> dma_buf_release+0x157/0x2d0 drivers/dma-buf/dma-buf.c:78 __dentry_kill+0x42b/0x640 fs/dcache.c:612 dentry_kill fs/dcache.c:733 [inline] dput+0x806/0xdb0 fs/dcache.c:913 __fput+0x39c/0x9d0 fs/file_table.c:333 task_work_run+0xdd/0x1a0 kernel/task_work.c:177 ptrace_notify+0x114/0x140 kernel/signal.c:2353 ptrace_report_syscall include/linux/ptrace.h:420 [inline] ptrace_report_syscall_exit include/linux/ptrace.h:482 [inline] syscall_exit_work kernel/entry/common.c:249 [inline] syscall_exit_to_user_mode_prepare+0x129/0x280 kernel/entry/common.c:276 __syscall_exit_to_user_mode_work kernel/entry/common.c:281 [inline] syscall_exit_to_user_mode+0x9/0x50 kernel/entry/common.c:294 do_syscall_64+0x42/0xb0 arch/x86/entry/common.c:86 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7fc1c0c35b6b Code: 0f 05 48 3d 00 f0 ff ff 77 45 c3 0f 1f 40 00 48 83 ec 18 89 7c 24 0c e8 63 fc ff ff 8b 7c 24 0c 41 89 c0 b8 03 00 00 00 0f 05 <48> 3d 00 f0 ff ff 77 35 44 89 c7 89 44 24 0c e8 a1 fc ff ff 8b 44 RSP: 002b:00007ffd78a06090 EFLAGS: 00000293 ORIG_RAX: 0000000000000003 RAX: 0000000000000000 RBX: 0000000000000007 RCX: 00007fc1c0c35b6b RDX: 0000000020000280 RSI: 0000000040086200 RDI: 0000000000000006 RBP: 0000000000000007 R08: 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000293 R12: 000000000000000c R13: 0000000000000003 R14: 00007fc1c0cfe4a0 R15: 00007ffd78a06140 </TASK> Modules linked in: ---[end trace 0000000000000000]--- RIP: 0010:dma_unmap_sgtable include/linux/dma-mapping.h:378 [inline] RIP: 0010:put_sg_table drivers/dma-buf/udmabuf.c:89 [inline] RIP: 0010:release_udmabuf+0xcb/0x4f0 drivers/dma-buf/udmabuf.c:114	N/A	More Details
CVE-2022-50820	In the Linux kernel, the following vulnerability has been resolved: perf/arm_dmc620: Fix hotplug callback leak in dmc620_pmu_init() dmc620_pmu_init() won't remove the callback added by cpuhp_setup_state_multi() when platform_driver_register() failed. Remove the callback by cpuhp_remove_multi_state() in fail path. Similar to the handling of arm_ccn_init() in commit 26242b330093 ("bus: arm-ccn: Prevent hotplug callback leak")	N/A	More Details
CVE-2022-50821	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Don't leak netobj memory when gss_read_proxy_verf() fails	N/A	More Details
CVE-2022-50822	In the Linux kernel, the following vulnerability has been resolved: RDMA/restrack: Release MR restrack when delete The MR restrack also needs to be released when delete it, otherwise it cause memory leak as the task struct won't be released.	N/A	More Details
CVE-2022-50823	In the Linux kernel, the following vulnerability has been resolved: clk: tegra: Fix refcount leak in tegra114_clock_init of_find_matching_node() returns a node pointer with refcount incremented, we should use of_node_put() on it when not need anymore. Add missing of_node_put() to avoid refcount leak.	N/A	More Details
CVE-2022-50824	In the Linux kernel, the following vulnerability has been resolved: tpm: tpm_tis: Add the missed acpi_put_table() to fix memory leak In check_acpi_tpm2(), we get the TPM2 table just to make sure the table is there, not used after the init, so the acpi_put_table() should be added to release the ACPI memory.	N/A	More Details
CVE-2022-50825	In the Linux kernel, the following vulnerability has been resolved: usb: typec: wusb3801: fix fwnode refcount leak in wusb3801_probe() I got the following report while doing fault injection test: OF: ERROR: memory leak, expected refcount 1 instead of 4, of_node_get()/of_node_put() unbalanced - destroy cset entry: attach overlay node /i2c/tcpc@60/connector If wusb3801_hw_init() fails, fwnode_handle_put() needs be called to avoid refcount leak.	N/A	More Details
CVE-2022-50826	In the Linux kernel, the following vulnerability has been resolved: ipu3-imgu: Fix NULL pointer dereference in imgu_subdev_set_selection() Calling v4l2_subdev_get_try_crop() and v4l2_subdev_get_try_compose() with a subdev state of NULL leads to a NULL pointer dereference. This can currently happen in imgu_subdev_set_selection() when the state passed in is NULL, as this method first gets pointers to both the "try" and "active" states and only then decides which to use. The same issue has been addressed for imgu_subdev_get_selection() with commit 30d03a0de650 ("ipu3-imgu: Fix NULL pointer dereference in active selection access"). However the issue still persists in imgu_subdev_set_selection(). Therefore, apply a similar fix as done in the aforementioned commit to imgu_subdev_set_selection(). To keep things a bit cleaner, introduce helper functions for "crop" and "compose" access and use them in both imgu_subdev_set_selection() and imgu_subdev_get_selection().	N/A	More Details
CVE-2022-50827	In the Linux kernel, the following vulnerability has been resolved: scsi: lpfc: Fix memory leak in lpfc_create_port() Commit 5e633302ace1 ("scsi: lpfc: vmid: Add support for VMID in mailbox command") introduced allocations for the VMID resources in lpfc_create_port() after the call to scsi_host_alloc(). Upon failure on the VMID allocations, the new code would branch to the 'out' label, which returns NULL without unwinding anything, thus skipping the call to scsi_host_put(). Fix the problem by creating a separate label 'out_free_vmid' to unwind the VMID resources and make the 'out_put_shost' label call only scsi_host_put(), as was done before the introduction of allocations for VMID.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: clk: zynqmp: Fix stack-out-of-bounds in strncpy` "BUG: KASAN: stack-out-of-bounds in strncpy+0x30/0x68" Linux-ATF interface is using 16 bytes of SMC payload. In case clock name is longer than 15 bytes, string terminated NULL character will not be received by Linux. Add explicit NULL character at last byte to fix issues when clock name is longer. This fixes below bug reported by KASAN: == BUG: KASAN: stack-out-of-bounds in strncpy+0x30/0x68 Read of size 1 at addr ffff0008c89a7410 by task swapper/0/1 CPU: 1 PID: 1 Comm: swapper/0 Not tainted 5.4.0-00396-g81ef9e7-dirty #3 Hardware name: Xilinx Versal vck190 Eval board revA (QSPI) (DT) Call trace:		

CVE-2022-50828	dump_backtrace+0x0/0x1e8 show_stack+0x14/0x20 dump_stack+0xd4/0x108 print_address_description.isra.0+0xbc/0x37c __kasan_report+0x144/0x198 kasan_report+0xc/0x18 __asan_load1+0x5c/0x68 strncpy+0x30/0x68 zynqmp_clock_probe+0x238/0x7b8 platform_drv_probe+0x6c/0xc8 really_probe+0x14c/0x418 driver_probe_device+0x74/0x130 __device_attach_driver+0xc4/0xe8 bus_for_each_drv+0xec/0x150 __device_attach+0x160/0x1d8 device_initial_probe+0x10/0x18 bus_probe_device+0xe0/0xf0 device_add+0x528/0x950 of_device_add+0x5c/0x80 of_platform_device_create_pdata+0x120/0x168 of_platform_bus_create+0x244/0x4e0 of_platform_populate+0x50/0xe8 zynqmp_firmware_probe+0x370/0x3a8 platform_drv_probe+0x6c/0xc8 really_probe+0x14c/0x418 driver_probe_device+0x74/0x130 device_driver_attach+0x94/0xa0 __driver_attach+0x70/0x108 bus_for_each_dev+0xe4/0x158 driver_attach+0x30/0x40 bus_add_driver+0x21c/0x2b8 driver_register+0xbc/0x1d0 __platform_driver_register+0x7c/0x88 zynqmp_firmware_driver_init+0x1c/0x24 do_one_initcall+0xa4/0x234 kernel_init_freeable+0x1b0/0x24c kernel_init+0x10/0x110 ret_from_fork+0x10/0x18 The buggy address belongs to the page: page:ffff0008f9be1c88 refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 raw: 0008d00000000000 ffff0008f9be1c90 ffff0008f9be1c90 0000000000000000 raw: 0000000000000000 0000000000000000 00000000ffffffff page dumped because: kasan: bad access detected addr ffff0008c89a7410 is located in stack of task swapper/0/1 at offset 112 in frame: zynqmp_clock_probe+0x0/0x7b8 this frame has 3 objects: [32, 44) 'response' [64, 80) 'ret_payload' [96, 112) 'name' Memory state around the buggy address: ffff0008c89a7300: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ffff0008c89a7380: 00 00 00 00 f1 f1 f1 f1 00 04 f2 f2 00 00 f2 f2 >ffff0008c89a7400: 00 00 f3 f3 00 00 00 00 00 00 00 00 00 00 00 00 ^ ffff0008c89a7480: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ffff0008c89a7500: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ==========	N/A	More Details
CVE-2022-50829	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: hif_usb: Fix use-after-free in ath9k_hif_usb_reg_in_cb() It is possible that skb is freed in ath9k_htc_rx_msg(), then usb_submit_urb() fails and we try to free skb again. It causes use-after-free bug. Moreover, if alloc_skb() fails, urb->context becomes NULL but rx_buf is not freed and there can be a memory leak. The patch removes unnecessary nskb and makes skb processing more clear: it is supposed that ath9k_htc_rx_msg() either frees old skb or passes its managing to another callback function. Found by Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2022-50830	In the Linux kernel, the following vulnerability has been resolved: auxdisplay: hd44780: Fix potential memory leak in hd44780_remove() hd44780_probe() allocates a memory chunk for hd with kzalloc() and makes "lcd->drvdata->hd44780" point to it. When we call hd44780_remove(), we should release all relevant memory and resource. But "lcd->drvdata ->hd44780" is not released, which will lead to a memory leak. We should release the "lcd->drvdata->hd44780" in hd44780_remove() to fix the memory leak bug.	N/A	More Details
CVE-2022-50831	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2022-50832	In the Linux kernel, the following vulnerability has been resolved: wifi: wilc1000: fix potential memory leak in wilc_mac_xmit() The wilc_mac_xmit() returns NETDEV_TX_OK without freeing skb, add dev_kfree_skb() to fix it. Compile tested only.	N/A	More Details
CVE-2022-50833	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: use hdev->workqueue when queuing hdev->{cmd,ncmd}_timer works syzbot is reporting attempt to schedule hdev->cmd_work work from system_wq WQ into hdev->workqueue WQ which is under draining operation [1], for commit c8efcc2589464ac7 ("workqueue: allow chained queueing during destruction") does not allow such operation. The check introduced by commit 877afadad2dce8aa ("Bluetooth: When HCI work queue is drained, only queue chained work") was incomplete. Use hdev->workqueue WQ when queuing hdev->{cmd,ncmd}_timer works because hci_{cmd,ncmd}_timeout() calls queue_work(hdev->workqueue). Also, protect the queuing operation with RCU read lock in order to avoid calling queue_delayed_work() after cancel_delayed_work() completed.	N/A	More Details
CVE-2022-50834	In the Linux kernel, the following vulnerability has been resolved: nfc: Fix potential resource leaks nfc_get_device() take reference for the device, add missing nfc_put_device() to release it when not need anymore. Also fix the style warnning by use error EOPNOTSUPP instead of ENOTSUPP.	N/A	More Details
CVE-2022-50835	In the Linux kernel, the following vulnerability has been resolved: jbd2: add miss release buffer head in fc_do_one_pass() In fc_do_one_pass() miss release buffer head after use which will lead to reference count leak.	N/A	More Details
CVE-2022-50836	In the Linux kernel, the following vulnerability has been resolved: remoteproc: sysmon: fix memory leak in qcom_add_sysmon_subdev() The kfree() should be called when of_irq_get_byname() fails or devm_request_threaded_irq() fails in qcom_add_sysmon_subdev(), otherwise there will be a memory leak, so add kfree() to fix it.	N/A	More Details
CVE-2022-50817	In the Linux kernel, the following vulnerability has been resolved: net: hsr: avoid possible NULL deref in skb_clone() syzbot got a crash [1] in skb_clone(), caused by a bug in hsr_get_untagged_frame(). When/if create_stripped_skb_hsr() returns NULL, we must not attempt to call skb_clone(). While we are at it, replace a WARN_ONCE() by netdev_warn_once(). [1] general protection fault, probably for non-canonical address 0xdffffc000000000f: 0000 [#1] PREEMPT SMP KASAN KASAN: null-ptr-deref in range [0x0000000000000078-0x000000000000007f] CPU: 1 PID: 754 Comm: syz-executor.0 Not tainted 6.0.0-syzkaller-02734-g0326074ff465 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 09/22/2022 RIP: 0010:skb_clone+0x108/0x3c0 net/core/skbuff.c:1641 Code: 93 02 00 00 49 83 7c 24 28 00 0f 85 e9 00 00 00 e8 5d 4a 29 fa 4c 8d 75 7e 48 b8 00 00 00 00 00 fc ff df 4c 89 f2 48 c1 ea 03 <0f> b6 04 02 4c 89 f2 83 e2 07 38 d0 7f 08 84 c0 0f 85 9e 01 00 00 RSP: 0018:ffffc90003ccf4e0 EFLAGS: 00010207 RAX: dffffc0000000000 RBX: ffffc90003ccf5f8 RCX: ffffc9000c24b000 RDX: 000000000000000f RSI: ffffffff8751cb13 RDI: 0000000000000000 RBP: 0000000000000000 R08: 00000000000000f0 R09: 0000000000000140 R10: fffffbfff181d972 R11: 0000000000000000 R12: ffff888161fc3640 R13: 0000000000000a20 R14: 000000000000007e R15: ffffffff8dc5f620 FS: 00007feb621e4700(0000) GS:ffff8880b9b00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007feb621e3ff8 CR3: 00000001643a9000 CR4: 00000000003506e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> hsr_get_untagged_frame+0x4e/0x610 net/hsr/hsr_forward.c:164 hsr_forward_do net/hsr/hsr_forward.c:461 [inline] hsr_forward_skb+0xcca/0x1d50 net/hsr/hsr_forward.c:623 hsr_handle_frame+0x588/0x7c0 net/hsr/hsr_slave.c:69 __netif_receive_skb_core+0x9fe/0x38f0 net/core/dev.c:5379 __netif_receive_skb_one_core+0xae/0x180 net/core/dev.c:5483 __netif_receive_skb+0x1f/0x1c0 net/core/dev.c:5599 netif_receive_skb_internal net/core/dev.c:5685 [inline] netif_receive_skb+0x12f/0x8d0 net/core/dev.c:5744 tun_rx_batched+0x4ab/0x7a0 drivers/net/tun.c:1544 tun_get_user+0x2686/0x3a00 drivers/net/tun.c:1995 tun_chr_write_iter+0xdb/0x200 drivers/net/tun.c:2025 call_write_iter include/linux/fs.h:2187 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_write+0x9e9/0xdd0 fs/read_write.c:584 ksys_write+0x127/0x250 fs/read_write.c:637 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd	N/A	More Details

CVE-2022-50816	In the Linux kernel, the following vulnerability has been resolved: ipv6: ensure sane device mtu in tunnels Another syzbot report [1] with no reproducer hints at a bug in ip6_gre tunnel (dev:ip6gretap0) Since ipv6 mcast code makes sure to read dev->mtu once and applies a sanity check on it (see commit b9b312a7a451 "ipv6: mcast: better catch silly mtu values"), a remaining possibility is that a layer is able to set dev->mtu to an underflowed value (high order bit set). This could happen indeed in ip6gre_tnl_link_config_route(), ip6_tnl_link_config() and ipip6_tunnel_bind_dev() Make sure to sanitize mtu value in a local variable before it is written once on dev->mtu, as lockless readers could catch wrong temporary value. [1] skbuff: skb_over_panic: text:ffff80000b7a2f38 len:40 put:40 head:ffff000149dcf200 data:ffff000149dcf2b0 tail:0xd8 end:0xc0 dev:ip6gretap0 ------------[cut here]------------ kernel BUG at net/core/skbuff.c:120 Internal error: Oops - BUG: 00000000f2000800 [#1] PREEMPT SMP Modules linked in: CPU: 1 PID: 10241 Comm: kworker/1:1 Not tainted 6.0.0-rc7-syzkaller-18095-gbbed346d5a96 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 09/30/2022 Workqueue: mld mld_ifc_work pstate: 60400005 (nZCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : skb_panic+0x4c/0x50 net/core/skbuff.c:116 lr : skb_panic+0x4c/0x50 net/core/skbuff.c:116 sp : ffff800020dd3b60 x29: ffff800020dd3b70 x28: 0000000000000000 x27: ffff00010df2a800 x26: 00000000000000c0 x25: 00000000000000b0 x24: ffff000149dcf200 x23: 00000000000000c0 x22: 00000000000000d8 x21: ffff80000b7a2f38 x20: ffff00014c2f7800 x19: 0000000000000028 x18: 00000000000001a9 x17: 0000000000000000 x16: ffff80000db49158 x15: ffff000113bf1a80 x14: 0000000000000000 x13: 00000000ffffffff x12: ffff000113bf1a80 x11: ff808000081c0d5c x10: 0000000000000000 x9 : 73f125dc5c63ba00 x8 : 73f125dc5c63ba00 x7 : ffff800008161d1c x6 : 0000000000000000 x5 : 0000000000000080 x4 : 0000000000000001 x3 : 0000000000000000 x2 : ffff0001fefddcd0 x1 : 0000000100000000 x0 : 0000000000000089 Call trace: skb_panic+0x4c/0x50 net/core/skbuff.c:116 skb_over_panic net/core/skbuff.c:125 [inline] skb_put+0xd4/0xdc net/core/skbuff.c:2049 ip6_mc_hdr net/ipv6/mcast.c:1714 [inline] mld_newpack+0x14c/0x270 net/ipv6/mcast.c:1765 add_grhead net/ipv6/mcast.c:1851 [inline] add_grec+0xa20/0xae0 net/ipv6/mcast.c:1989 mld_send_cr+0x438/0x5a8 net/ipv6/mcast.c:2115 mld_ifc_work+0x38/0x290 net/ipv6/mcast.c:2653 process_one_work+0x2d8/0x504 kernel/workqueue.c:2289 worker_thread+0x340/0x610 kernel/workqueue.c:2436 kthread+0x12c/0x158 kernel/kthread.c:376 ret_from_fork+0x10/0x20 arch/arm64/kernel/entry.S:860 Code: 91011400 aa0803e1 a90027ea 94373093 (d4210000)	N/A	More Details
CVE-2022-50815	In the Linux kernel, the following vulnerability has been resolved: ext2: Add sanity checks for group and filesystem size Add sanity check that filesystem size does not exceed the underlying device size and that group size is big enough so that metadata can fit into it. This avoid trying to mount some crafted filesystems with extremely large group counts.	N/A	More Details
CVE-2022-50609	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50601	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50602	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50603	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50604	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50605	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50606	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50607	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50608	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50610	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50814	In the Linux kernel, the following vulnerability has been resolved: crypto: hisilicon/zip - fix mismatch in get/set sgl_sge_nr KASAN reported this Bug: [17619.659757] BUG: KASAN: global-out-of-bounds in param_get_int+0x34/0x60 [17619.673193] Read of size 4 at addr fffff01332d7ed00 by task read_all/1507958 ... [17619.698934] The buggy address belongs to the variable: [17619.708371] sgl_sge_nr+0x0/0xfffffffffffffa300 [hisi_zip] There is a mismatch in hisi_zip when get/set the variable sgl_sge_nr. The type of sgl_sge_nr is u16, and get/set sgl_sge_nr by param_get/set_int. Replacing param_get/set_int to param_get/set_ushort can fix this bug.	N/A	More Details
CVE-2022-50611	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50612	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details

CVE-2022-50613	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50809	In the Linux kernel, the following vulnerability has been resolved: xhci: dbc: Fix memory leak in xhci_alloc_dbc() If DbC is already in use, then the allocated memory for the xhci_dbc struct doesn't get freed before returning NULL, which leads to a memleak.	N/A	More Details
CVE-2022-50810	In the Linux kernel, the following vulnerability has been resolved: rapidio: devices: fix missing put_device in mport_cdev_open When kfifo_alloc fails, the refcount of chdev->dev is left incremental. We should use put_device(&chdev->dev) to decrease the ref count of chdev->dev to avoid refcount leak.	N/A	More Details
CVE-2022-50811	In the Linux kernel, the following vulnerability has been resolved: erofs: fix missing unmap if z_erofs_get_extent_compressedlen() fails Otherwise, meta buffers could be leaked.	N/A	More Details
CVE-2022-50812	In the Linux kernel, the following vulnerability has been resolved: security: Restrict CONFIG_ZERO_CALL_USED_REGS to gcc or clang > 15.0.6 A bad bug in clang's implementation of -fzero-call-used-regs can result in NULL pointer dereferences (see the links above the check for more information). Restrict CONFIG_CC_HAS_ZERO_CALL_USED_REGS to either a supported GCC version or a clang newer than 15.0.6, which will catch both a theoretical 15.0.7 and the upcoming 16.0.0, which will both have the bug fixed.	N/A	More Details
CVE-2022-50813	In the Linux kernel, the following vulnerability has been resolved: drivers: mcb: fix resource leak in mcb_probe() When probe hook function failed in mcb_probe(), it doesn't put the device. Compiled test only.	N/A	More Details
CVE-2022-50837	In the Linux kernel, the following vulnerability has been resolved: net: dsa: tag_8021q: avoid leaking ctx on dsa_tag_8021q_register() error path If dsa_tag_8021q_setup() fails, for example due to the inability of the device to install a VLAN, the tag_8021q context of the switch will leak. Make sure it is freed on the error path.	N/A	More Details
CVE-2022-50838	In the Linux kernel, the following vulnerability has been resolved: net: stream: purge sk_error_queue in sk_stream_kill_queues() Changheon Lee reported TCP socket leaks, with a nice repro. It seems we leak TCP sockets with the following sequence: 1) SOF_TIMESTAMPING_TX_ACK is enabled on the socket. Each ACK will cook an skb put in error queue, from __skb_tstamp_tx(). __skb_tstamp_tx() is using skb_clone(), unless SOF_TIMESTAMPING_OPT_TSONLY was also requested. 2) If the application is also using MSG_ZEROCOPY, then we put in the error queue cloned skbs that had a struct ubuf_info attached to them. Whenever an struct ubuf_info is allocated, sock_zerocopy_alloc() does a sock_hold(). As long as the cloned skbs are still in sk_error_queue, socket refcount is kept elevated. 3) Application closes the socket, while error queue is not empty. Since tcp_close() no longer purges the socket error queue, we might end up with a TCP socket with at least one skb in error queue keeping the socket alive forever. This bug can be (ab)used to consume all kernel memory and freeze the host. We need to purge the error queue, with proper synchronization against concurrent writers.	N/A	More Details
CVE-2022-50839	In the Linux kernel, the following vulnerability has been resolved: jbd2: fix potential buffer head reference count leak As in 'jbd2_fc_wait_bufs' if buffer isn't uptodate, will return -EIO without update 'journal->j_fc_off'. But 'jbd2_fc_release_bufs' will release buffer head from ‘j_fc_off - 1’ if 'bh' is NULL will terminal release which will lead to buffer head buffer head reference count leak. To solve above issue, update 'journal->j_fc_off' before return -EIO.	N/A	More Details
CVE-2022-50871	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: Fix qmi_msg_handler data structure initialization qmi_msg_handler is required to be null terminated by QMI module. There might be a case where a handler for a msg id is not present in the handlers array which can lead to infinite loop while searching the handler and therefore out of bound access in qmi_invoke_handler(). Hence update the initialization in qmi_msg_handler data structure. Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-01100-QCAHKSWPL_SILICONZ-1	N/A	More Details
CVE-2022-50863	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw89: free unused skb to prevent memory leak This avoid potential memory leak under power saving mode.	N/A	More Details
CVE-2022-50864	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix shift-out-of-bounds due to too large exponent of block size If field s_log_block_size of superblock data is corrupted and too large, init_nilfs() and load_nilfs() still can trigger a shift-out-of-bounds warning followed by a kernel panic (if panic_on_warn is set): shift exponent 38973 is too large for 32-bit type 'int' Call Trace: <TASK> dump_stack_lvl+0xcd/0x134 ubsan_epilogue+0xb/0x50 __ubsan_handle_shift_out_of_bounds.cold.12+0x17b/0x1f5 init_nilfs.cold.11+0x18/0x1d [nilfs2] nilfs_mount+0x9b5/0x12b0 [nilfs2] ... This fixes the issue by adding and using a new helper function for getting block size with sanity check.	N/A	More Details
CVE-2022-50865	In the Linux kernel, the following vulnerability has been resolved: tcp: fix a signed-integer-overflow bug in tcp_add_backlog() The type of sk_rcvbuf and sk_sndbuf in struct sock is int, and in tcp_add_backlog(), the variable limit is caculated by adding sk_rcvbuf, sk_sndbuf and 64 * 1024, it may exceed the max value of int and overflow. This patch reduces the limit budget by halving the sndbuf to solve this issue since ACK packets are much smaller than the payload.	N/A	More Details
CVE-2022-50866	In the Linux kernel, the following vulnerability has been resolved: ASoC: pxa: fix null-pointer dereference in filter() kasprintf() would return NULL pointer when kmalloc() fail to allocate. Need to check the return pointer before calling strcmp().	N/A	More Details
CVE-2022-50867	In the Linux kernel, the following vulnerability has been resolved: drm/msm/a6xx: Fix kvzalloc vs state_kcalloc usage adreno_show_object() is a trap! It will re-allocate the pointer it is passed on first call, when the data is ascii85 encoded, using kvmalloc/ kvfree(). Which means the data *passed* to it must be kvmalloc'd, ie. we cannot use the state_kcalloc() helper. This partially reverts commit ec8f1813bf8d ("drm/msm/a6xx: Replace kcalloc() with kvzalloc()"), but adds the missing kvfree() to fix the memory leak that was present previously. And adds a warning comment. Patchwork: https://patchwork.freedesktop.org/patch/507014/	N/A	More Details
CVE-2022-50868	In the Linux kernel, the following vulnerability has been resolved: hwrng: amd - Fix PCI device refcount leak for_each_pci_dev() is implemented by pci_get_device(). The comment of pci_get_device() says that it will increase the reference count for the returned pci_dev and also decrease the reference count for the input pci_dev @from if it is not NULL. If we break for_each_pci_dev() loop with pdev not NULL, we need to call pci_dev_put() to decrease the reference count. Add the missing pci_dev_put() for the normal and error path.	N/A	More Details

CVE-2022-50869	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix slab-out-of-bounds in r_page When PAGE_SIZE is 64K, if read_log_page is called by log_read_rst for the first time, the size of *buffer would be equal to DefaultLogPageSize(4K).But for *buffer operations like memcpy, if the memory area size(n) which being assigned to buffer is larger than 4K (log->page_size(64K) or bytes(64K-page_off)), it will cause an out of boundary error. Call trace: [...] kasan_report+0x44/0x130 check_memory_region+0xf8/0x1a0 memcpy+0xc8/0x100 ntfs_read_run_nb+0x20c/0x460 read_log_page+0xd0/0x1f4 log_read_rst+0x110/0x75c log_replay+0x1e8/0x4aa0 ntfs_loadlog_and_replay+0x290/0x2d0 ntfs_fill_super+0x508/0xec0 get_tree_bdev+0x1fc/0x34c [...] Fix this by setting variable r_page to NULL in log_read_rst.	N/A	More Details
CVE-2022-50870	In the Linux kernel, the following vulnerability has been resolved: powerpc/rtas: avoid device tree lookups in rtas_os_term() rtas_os_term() is called during panic. Its behavior depends on a couple of conditions in the /rtas node of the device tree, the traversal of which entails locking and local IRQ state changes. If the kernel panics while devtree_lock is held, rtas_os_term() as currently written could hang. Instead of discovering the relevant characteristics at panic time, cache them in file-static variables at boot. Note the lookup for "ibm,extended-os-term" is converted to of_property_read_bool() since it is a boolean property, not an RTAS function token. [mpe: Incorporate suggested change from Nick]	N/A	More Details
CVE-2022-50872	In the Linux kernel, the following vulnerability has been resolved: ARM: OMAP2+: Fix memory leak in realtime_counter_init() The "sys_clk" resource is malloced by clk_get(), it is not released when the function return.	N/A	More Details
CVE-2022-50861	In the Linux kernel, the following vulnerability has been resolved: NFSD: Finish converting the NFSv2 GETACL result encoder The xdr_stream conversion inadvertently left some code that set the page_len of the send buffer. The XDR stream encoders should handle this automatically now. This oversight adds garbage past the end of the Reply message. Clients typically ignore the garbage, but NFSD does not need to send it, as it leaks stale memory contents onto the wire.	N/A	More Details
CVE-2022-50873	In the Linux kernel, the following vulnerability has been resolved: vdpa/vp_vdpa: fix kfree a wrong pointer in vp_vdpa_remove In vp_vdpa_remove(), the code kfree(&vp_vdpa_mgtdev->mgtdev.id_table) uses a reference of pointer as the argument of kfree, which is the wrong pointer and then may hit crash like this: Unable to handle kernel paging request at virtual address 00ffff003363e30c Internal error: Oops: 96000004 [#1] SMP Call trace: rb_next+0x20/0x5c ext4_readdir+0x494/0x5c4 [ext4] iterate_dir+0x168/0x1b4 __se_sys_getdents64+0x68/0x170 __arm64_sys_getdents64+0x24/0x30 el0_svc_common.constprop.0+0x7c/0x1bc do_el0_svc+0x2c/0x94 el0_svc+0x20/0x30 el0_sync_handler+0xb0/0xb4 el0_sync+0x160/0x180 Code: 54000220 f9400441 b4000161 aa0103e0 (f9400821) SMP: stopping secondary CPUs Starting crashdump kernel...	N/A	More Details
CVE-2022-50874	In the Linux kernel, the following vulnerability has been resolved: RDMA/erdma: Fix refcount leak in erdma_mmap rdma_user_mmap_entry_get() take reference, we should release it when not need anymore, add the missing rdma_user_mmap_entry_put() in the error path to fix it.	N/A	More Details
CVE-2022-50875	In the Linux kernel, the following vulnerability has been resolved: of: overlay: fix null pointer dereferencing in find_dup_cset_node_entry() and find_dup_cset_prop() When kmalloc() fail to allocate memory in kasprintf(), fn_1 or fn_2 will be NULL, and strcmp() will cause null pointer dereference.	N/A	More Details
CVE-2022-50876	In the Linux kernel, the following vulnerability has been resolved: usb: musb: Fix musb_gadget.c rxstate overflow bug The usb function device call musb_gadget_queue() adds the passed request to musb_ep::req_list,If the (request->length > musb_ep->packet_sz) and (is_buffer_mapped(req) return false),the rxstate() will copy all data in fifo to request->buf which may cause request->buf out of bounds. Fix it by add the length check : fifocnt = min_t(unsigned, request->length - request->actual, fifocnt);	N/A	More Details
CVE-2022-50877	In the Linux kernel, the following vulnerability has been resolved: net: broadcom: bcm4908_enet: update TX stats after actual transmission Queueing packets doesn't guarantee their transmission. Update TX stats after hardware confirms consuming submitted data. This also fixes a possible race and NULL dereference. bcm4908_enet_start_xmit() could try to access skb after freeing it in the bcm4908_enet_poll_tx().	N/A	More Details
CVE-2022-50878	In the Linux kernel, the following vulnerability has been resolved: gpu: lontium-lt9611: Fix NULL pointer dereference in lt9611_connector_init() A NULL check for bridge->encoder shows that it may be NULL, but it already been dereferenced on all paths leading to the check. 812 if (!bridge->encoder) { Dereference the pointer bridge->encoder. 810 drm_connector_attach_encoder(&lt9611->connector, bridge->encoder);	N/A	More Details
CVE-2022-50879	In the Linux kernel, the following vulnerability has been resolved: objtool: Fix SEGFAULT find_insn() will return NULL in case of failure. Check insn in order to avoid a kernel Oops for NULL pointer dereference.	N/A	More Details
CVE-2022-50880	In the Linux kernel, the following vulnerability has been resolved: wifi: ath10k: add peer map clean up for peer delete in ath10k_sta_state() When peer delete failed in a disconnect operation, use-after-free detected by KFENCE in below log. It is because for each vdev_id and address, it has only one struct ath10k_peer, it is allocated in ath10k_peer_map_event(). When connected to an AP, it has more than one HTT_T2H_MSG_TYPE_PEER_MAP reported from firmware, then the array peer_map of struct ath10k will be set muti-elements to the same ath10k_peer in ath10k_peer_map_event(). When peer delete failed in ath10k_sta_state(), the ath10k_peer will be free for the 1st peer id in array peer_map of struct ath10k, and then use-after-free happened for the 2nd peer id because they map to the same ath10k_peer. And clean up all peers in array peer_map for the ath10k_peer, then user-after-free disappeared peer map event log: [306.911021] wlan0: authenticate with b0:2a:43:e6:75:0e [306.957187] ath10k_pci 0000:01:00.0: mac vdev 0 peer create b0:2a:43:e6:75:0e (new sta) sta 1 / 32 peer 1 / 33 [306.957395] ath10k_pci 0000:01:00.0: htt peer map vdev 0 peer b0:2a:43:e6:75:0e id 246 [306.957404] ath10k_pci 0000:01:00.0: htt peer map vdev 0 peer b0:2a:43:e6:75:0e id 198 [306.986924] ath10k_pci 0000:01:00.0: htt peer map vdev 0 peer b0:2a:43:e6:75:0e id 166 peer unmap event log: [435.715691] wlan0: deauthenticating from b0:2a:43:e6:75:0e by local choice (Reason: 3=DEAUTH_LEAVING) [435.716802] ath10k_pci 0000:01:00.0: mac vdev 0 peer delete b0:2a:43:e6:75:0e sta ffff990e0e9c2b50 (sta gone) [435.717177] ath10k_pci 0000:01:00.0: htt peer unmap vdev 0 peer b0:2a:43:e6:75:0e id 246 [435.717186] ath10k_pci 0000:01:00.0: htt peer unmap vdev 0 peer b0:2a:43:e6:75:0e id 198 [435.717193] ath10k_pci 0000:01:00.0: htt peer unmap vdev 0 peer b0:2a:43:e6:75:0e id 166 use-after-free log: [21705.888627] wlan0: deauthenticating from d0:76:8f:82:be:75 by local choice (Reason: 3=DEAUTH_LEAVING) [21713.799910] ath10k_pci 0000:01:00.0: failed to delete peer d0:76:8f:82:be:75 for vdev 0: -110 [21713.799925] ath10k_pci 0000:01:00.0: found sta peer d0:76:8f:82:be:75 (ptr 0000000000000000 id 102) entry on vdev 0 after it was supposedly removed [21713.799968] == [21713.799991] BUG: KFENCE: use-after-free read in ath10k_sta_state+0x265/0xb8a [ath10k_core] [21713.799991] [21713.799997] Use-after-free read at 0x00000000abe1c75e (in kfence-#69): [21713.800010] ath10k_sta_state+0x265/0xb8a [ath10k_core] [21713.800041] drv_sta_state+0x115/0x677 [mac80211] [21713.800059] __sta_info_destroy_part2+0xb1/0x133 [mac80211] [21713.800076] __sta_info_flush+0x11d/0x162 [mac80211] [21713.800093] ieee80211_set_disassoc+0x12d/0x2f4 [mac80211] [21713.800110] ieee80211_mgd_deauth+0x26c/0x29b [mac80211] [21713.800137] cfg80211_mlme_deauth+0x13f/0x1bb [cfg80211] [21713.800153] nl80211_deauthenticate+0xf8/0x121 [cfg80211] [21713.800161] genl_rcv_msg+0x38e/0x3be	N/A	More Details

	[21713.800166] netlink_rcv_skb+0x89/0xf7 [21713.800171] genl_rcv+0x28/0x36 [21713.800176] netlink_unicast+0x179/0x24b [21713.800181] netlink_sendmsg+0x3a0/0x40e [21713.800187] sock_sendmsg+0x72/0x76 [21713.800192] ____sys_sendmsg+0x16d/0x1e3 [21713.800196] ___sys_sendmsg+0x95/0xd1 [21713.800200] __sys_sendmsg+0x85/0xbf [21713.800205] do_syscall_64+0x43/0x55 [21713.800210] entry_SYSCALL_64_after_hwframe+0x44/0xa9 [21713.800213] [21713.800219] kfence-#69: 0x000000009149b0d5-0x000000004c0697fb, size=1064, cache=kmalloc-2k [21713.800219] [21713.800224] allocated by task 13 on cpu 0 at 21705.501373s: [21713.800241] ath10k_peer_map_event+0x7e/0x154 [ath10k_core] [21713.800254] ath10k_htt_t2h_msg_handler+0x586/0x1039 [ath10k_core] [21713.800265] ath10k_htt_htc_t2h_msg_handler+0x12/0x28 [ath10k_core] [21713.800277] ath10k_htc_rx_completion_handler+0x14c/0x1b5 [ath10k_core] [21713.800283] ath10k_pci_process_rx_cb+0x195/0x1d ---truncated---		
CVE-2022-50862	In the Linux kernel, the following vulnerability has been resolved: bpf: prevent decl_tag from being referenced in func_proto Syzkaller was able to hit the following issue: ------------[cut here]------------ WARNING: CPU: 0 PID: 3609 at kernel/bpf/btf.c:1946 btf_type_id_size+0x2d5/0x9d0 kernel/bpf/btf.c:1946 Modules linked in: CPU: 0 PID: 3609 Comm: syz-executor361 Not tainted 6.0.0-syzkaller-02734-g0326074ff465 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 09/22/2022 RIP: 0010:btf_type_id_size+0x2d5/0x9d0 kernel/bpf/btf.c:1946 Code: ef e8 7f 8e e4 ff 41 83 ff 0b 77 28 f6 44 24 10 18 75 3f e8 6d 91 e4 ff 44 89 fe bf 0e 00 00 00 e8 20 8e e4 ff e8 5b 91 e4 ff <0f> 0b 45 31 f6 e9 98 02 00 00 41 83 ff 12 74 18 e8 46 91 e4 ff 44 RSP: 0018:ffffc90003cefb40 EFLAGS: 00010293 RAX: 0000000000000000 RBX: 0000000000000002 RCX: 0000000000000000 RDX: ffff8880259c0000 RSI: ffffffff81968415 RDI: 0000000000000005 RBP: ffff88801270ca00 R08: 0000000000000005 R09: 000000000000000e R10: 0000000000000011 R11: 0000000000000000 R12: 0000000000000000 R13: 0000000000000011 R14: ffff888026ee6424 R15: 0000000000000011 FS: 000055555641b300(0000) GS:ffff8880b9a00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000f2e258 CR3: 000000007110e000 CR4: 00000000003506f0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> btf_func_proto_check kernel/bpf/btf.c:4447 [inline] btf_check_all_types kernel/bpf/btf.c:4723 [inline] btf_parse_type_sec kernel/bpf/btf.c:4752 [inline] btf_parse kernel/bpf/btf.c:5026 [inline] btf_new_fd+0x1926/0x1e70 kernel/bpf/btf.c:6892 bpf_btf_load kernel/bpf/syscall.c:4324 [inline] __sys_bpf+0xb7d/0x4cf0 kernel/bpf/syscall.c:5010 __do_sys_bpf kernel/bpf/syscall.c:5069 [inline] __se_sys_bpf kernel/bpf/syscall.c:5067 [inline] __x64_sys_bpf+0x75/0xb0 kernel/bpf/syscall.c:5067 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7f0fbae41c69 Code: 28 c3 e8 2a 14 00 00 66 2e 0f 1f 84 00 00 00 00 00 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d8 64 89 01 48 RSP: 002b:00007ffc8aeb6228 EFLAGS: 00000246 ORIG_RAX: 0000000000000141 RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 00007f0fbae41c69 RDX: 0000000000000020 RSI: 0000000020000140 RDI: 0000000000000012 RBP: 00007f0fbae05e10 R08: 0000000000000000 R09: 0000000000000000 R10: 00000000ffffffff R11: 0000000000000246 R12: 00007f0fbae05ea0 R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000000 </TASK> Looks like it tries to create a func_proto which return type is decl_tag. For the details, see Martin's spot on analysis in [0]. 0: https://lore.kernel.org/bpf/CAKH8qBuQDLva_hHxxBuZzyAcYNO4ejhovz6TQeVSk8HY-2SO6g@mail.gmail.com/T/#mea6524b3fcd6298347432226e81b1e6155efc62c	N/A	More Details
CVE-2022-50860	In the Linux kernel, the following vulnerability has been resolved: apparmor: Fix memleak in alloc_ns() After changes in commit a1bd627b46d1 ("apparmor: share profile name on replacement"), the hname member of struct aa_policy is not valid slab object, but a subset of that, it can not be freed by kfree_sensitive(), use aa_policy_destroy() to fix it.	N/A	More Details
CVE-2022-50840	In the Linux kernel, the following vulnerability has been resolved: scsi: snic: Fix possible UAF in snic_tgt_create() Smatch reports a warning as follows: drivers/scsi/snic/snic_disc.c:307 snic_tgt_create() warn: '&tgt->list' not removed from list If device_add() fails in snic_tgt_create(), tgt will be freed, but tgt->list will not be removed from snic->disc.tgt_list, then list traversal may cause UAF. Remove from snic->disc.tgt_list before free().	N/A	More Details
CVE-2022-50849	In the Linux kernel, the following vulnerability has been resolved: pstore: Avoid kcore oops by vmap()ing with VM_IOREMAP An oops can be induced by running 'cat /proc/kcore > /dev/null' on devices using pstore with the ram backend because kmap_atomic() assumes lowmem pages are accessible with __va(). Unable to handle kernel paging request at virtual address ffffff807ff2b000 Mem abort info: ESR = 0x96000006 EC = 0x25: DABT (current EL), IL = 32 bits SET = 0, FnV = 0 EA = 0, S1PTW = 0 FSC = 0x06: level 2 translation fault Data abort info: ISV = 0, ISS = 0x00000006 CM = 0, WnR = 0 swapper pgtable: 4k pages, 39-bit VAs, pgdp=0000000081d87000 [ffffff807ff2b000] pgd=180000017fe18003, p4d=180000017fe18003, pud=180000017fe18003, pmd=0000000000000000 Internal error: Oops: 96000006 [#1] PREEMPT SMP Modules linked in: dm_integrity CPU: 7 PID: 21179 Comm: perf Not tainted 5.15.67-10882-ge4eb2eb988cd #1 baa443fb8e8477896a370b31a821eb2009f9bfba Hardware name: Google Lazor (rev3 - 8) (DT) pstate: a0400009 (NzCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : __memcpy+0x110/0x260 lr : vread+0x194/0x294 sp : ffffffc013ee39d0 x29: ffffffc013ee39f0 x28: 0000000000001000 x27: ffffff807ff2b000 x26: 0000000000001000 x25: ffffffc0085a2000 x24: ffffff802d4b3000 x23: ffffff80f8a60000 x22: ffffff802d4b3000 x21: ffffffc0085a2000 x20: ffffff8080b7bc68 x19: 0000000000001000 x18: 0000000000000000 x17: 0000000000000000 x16: 0000000000000000 x15: ffffffd3073f2e60 x14: ffffffffad588000 x13: 0000000000000000 x12: 0000000000000001 x11: 00000000000001a2 x10: 00680000fff2bf0b x9 : 03ffffff807ff2b x8 : 0000000000000001 x7 : 0000000000000000 x6 : 0000000000000000 x5 : ffffff802d4b4000 x4 : ffffff807ff2c000 x3 : ffffffc013ee3a78 x2 : 0000000000001000 x1 : ffffff807ff2b000 x0 : ffffff802d4b3000 Call trace: __memcpy+0x110/0x260 read_kcore+0x584/0x778 proc_reg_read+0xb4/0xe4 During early boot, memblock reserves the pages for the ramoops reserved memory node in DT that would otherwise be part of the direct lowmem mapping. Pstore's ram backend reuses those reserved pages to change the memory type (writeback or non-cached) by passing the pages to vmap() (see pfn_to_page() usage in persistent_ram_vmap() for more details) with specific flags. When read_kcore() starts iterating over the vmalloc region, it runs over the virtual address that vmap() returned for ramoops. In aligned_vread() the virtual address is passed to vmalloc_to_page() which returns the page struct for the reserved lowmem area. That lowmem page is passed to kmap_atomic(), which effectively calls page_to_virt() that assumes a lowmem page struct must be directly accessible with __va() and friends. These pages are mapped via vmap() though, and the lowmem mapping was never made, so accessing them via the lowmem virtual address oopses like above. Let's side-step this problem by passing VM_IOREMAP to vmap(). This will tell vread() to not include the ramoops region in the kcore. Instead the area will look like a bunch of zeros. The alternative is to teach kmap() about vmalloc areas that intersect with lowmem. Presumably such a change isn't a one-liner, and there isn't much interest in inspecting the ramoops region in kcore files anyway, so the most expedient route is taken for now.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Add overflow check for attribute size The offset addition could overflow and pass the used size check given an attribute with very large size (e.g., 0xffffff7f) while parsing MFT attributes. This could lead to out-of-bound memory R/W if we try to access the next attribute derived by Add2Ptr(attr, asize) [32.963847] BUG: unable to handle page fault for address: ffff956a83c76067 [32.964301] #PF: supervisor read access in kernel mode [32.964526] #PF: error_code(0x0000) - not-present page [32.964893] PGD 4dc01067 P4D 4dc01067 PUD 0 [32.965316] Oops: 0000 [#1] PREEMPT SMP NOPTI [32.965727] CPU: 0 PID: 243 Comm: mount Not tainted 5.19.0+ #6 [32.966050] Hardware name: QEMU		

CVE-2022-50841	Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org 04/01/2014 [32.966628] RIP: 0010:mi_enum_attr+0x44/0x110 [32.967239] Code: 89 f0 48 29 c8 48 89 c1 39 c7 0f 86 94 00 00 00 8b 56 04 83 fa 17 0f 86 88 00 00 00 89 d0 01 ca 48 01 f0 8d 4a 08 39 f9a [32.968101] RSP: 0018:ffffba15c06a7c38 EFLAGS: 00000283 [32.968364] RAX: ffff956a83c76067 RBX: ffff956983c76050 RCX: 000000000000006f [32.968651] RDX: 0000000000000067 RSI: ffff956983c760e8 RDI: 00000000000001c8 [32.968963] RBP: ffffba15c06a7c38 R08: 0000000000000064 R09: 00000000ffffff7f [32.969249] R10: 0000000000000007 R11: ffff956983c760e8 R12: ffff95698225e000 [32.969870] R13: 0000000000000000 R14: ffffba15c06a7cd8 R15: ffff95698225e170 [32.970655] FS: 00007fdab8189e40(0000) GS:ffff9569fdc00000(0000) knlGS:0000000000000000 [32.971098] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [32.971378] CR2: ffff956a83c76067 CR3: 0000000002c58000 CR4: 00000000000006f0 [32.972098] Call Trace: [32.972842] <TASK> [32.973341] ni_enum_attr_ex+0xda/0xf0 [32.974087] ntfs_iget5+0x1db/0xde0 [32.974386] ? slab_post_alloc_hook+0x53/0x270 [32.974778] ? ntfs_fill_super+0x4c7/0x12a0 [32.975115] ntfs_fill_super+0x5d6/0x12a0 [32.975336] get_tree_bdev+0x175/0x270 [32.975709] ? put_ntfs+0x150/0x150 [32.975956] ntfs_fs_get_tree+0x15/0x20 [32.976191] vfs_get_tree+0x2a/0xc0 [32.976374] ? capable+0x19/0x20 [32.976572] path_mount+0x484/0xaa0 [32.977025] ? putname+0x57/0x70 [32.977380] do_mount+0x80/0xa0 [32.977555] __x64_sys_mount+0x8b/0xe0 [32.978105] do_syscall_64+0x3b/0x90 [32.978830] entry_SYSCALL_64_after_hwframe+0x63/0xcd [32.979311] RIP: 0033:0x7fdab72e948a [32.980015] Code: 48 8b 0d 11 fa 2a 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 49 89 ca b8 a5 00 00 008 [32.981251] RSP: 002b:00007ffd15b87588 EFLAGS: 00000206 ORIG_RAX: 00000000000000a5 [32.981832] RAX: ffffffffffffffda RBX: 0000557de0aaf060 RCX: 00007fdab72e948a [32.982234] RDX: 0000557de0aaf260 RSI: 0000557de0aaf2e0 RDI: 0000557de0ab7ce0 [32.982714] RBP: 0000000000000000 R08: 0000557de0aaf280 R09: 0000000000000020 [32.983046] R10: 00000000c0ed0000 R11: 0000000000000206 R12: 0000557de0ab7ce0 [32.983494] R13: 0000557de0aaf260 R14: 0000000000000000 R15: 00000000ffffffff [32.984094] </TASK> [32.984352] Modules linked in: [32.984753] CR2: ffff956a83c76067 [32.985911] ---[end trace 0000000000000000]--- [32.986555] RIP: 0010:mi_enum_attr+0x44/0x110 [32.987217] Code: 89 f0 48 29 c8 48 89 c1 39 c7 0f 86 94 00 00 00 8b 56 04 83 fa 17 0f 86 88 00 00 00 89 d0 01 ca 48 01 f0 8d 4a 08 39 f9a [32.988232] RSP: 0018:ffffba15c06a7c38 EFLAGS: 00000283 [32.988532] RAX: ffff956a83c76067 RBX: ffff956983c76050 RCX: 000000000000006f [32.988916] RDX: 0000000000000067 RSI: ffff956983c760e8 RDI: 00000000000001c8 [32.989356] RBP: ffffba15c06a7c38 R08: 0000000000000064 R09: 00000000ffffff7f [32.989994] R10: 0000000000000007 R11: ffff956983c760e8 R12: ffff95698225e000 [32.990415] R13: 0000000000000000 R14: ffffba15c06a7cd8 R15: ffff95698225e170 [32.991011] FS: ---truncated---	N/A	More Details
CVE-2022-50842	In the Linux kernel, the following vulnerability has been resolved: drm/virtio: Check whether transferred 2D BO is shmem Transferred 2D BO always must be a shmem BO. Add check for that to prevent NULL dereference if userspace passes a VRAM BO.	N/A	More Details
CVE-2022-50843	In the Linux kernel, the following vulnerability has been resolved: dm clone: Fix UAF in clone_dtr() Dm_clone also has the same UAF problem when dm_resume() and dm_destroy() are concurrent. Therefore, cancelling timer again in clone_dtr().	N/A	More Details
CVE-2022-50844	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix type of second parameter in odn_edit_dpm_table() callback With clang's kernel control flow integrity (kCFI, CONFIG_CFI_CLANG), indirect call targets are validated against the expected function pointer prototype to make sure the call target is valid to help mitigate ROP attacks. If they are not identical, there is a failure at run time, which manifests as either a kernel panic or thread getting killed. A proposed warning in clang aims to catch these at compile time, which reveals: drivers/gpu/drm/amd/amdgpu/../pm/swsmu/amdgpu_smu.c:3008:29: error: incompatible function pointer types initializing 'int (*)(void *, uint32_t, long *, uint32_t)' (aka 'int (*)(void *, unsigned int, long *, unsigned int)') with an expression of type 'int (void *, enum PP_OD_DPM_TABLE_COMMAND, long *, uint32_t)' (aka 'int (void *, enum PP_OD_DPM_TABLE_COMMAND, long *, unsigned int)') [-Werror,-Wincompatible-function-pointer-types-strict] .odn_edit_dpm_table = smu_od_edit_dpm_table, ^~~~~~~~~~~~~~~~~~~~~~ 1 error generated. There are only two implementations of ->odn_edit_dpm_table() in 'struct amd_pm_funcs': smu_od_edit_dpm_table() and pp_odn_edit_dpm_table(). One has a second parameter type of 'enum PP_OD_DPM_TABLE_COMMAND' and the other uses 'u32'. Ultimately, smu_od_edit_dpm_table() calls ->od_edit_dpm_table() from 'struct pptable_funcs' and pp_odn_edit_dpm_table() calls ->odn_edit_dpm_table() from 'struct pp_hwmgr_func', which both have a second parameter type of 'enum PP_OD_DPM_TABLE_COMMAND'. Update the type parameter in both the prototype in 'struct amd_pm_funcs' and pp_odn_edit_dpm_table() to 'enum PP_OD_DPM_TABLE_COMMAND', which cleans up the warning.	N/A	More Details
CVE-2022-50845	In the Linux kernel, the following vulnerability has been resolved: ext4: fix inode leak in ext4_xattr_inode_create() on an error path There is issue as follows when do setxattr with inject fault: [localhost]# fsck.ext4 -fn /dev/sda e2fsck 1.46.6-rc1 (12-Sep-2022) Pass 1: Checking inodes, blocks, and sizes Pass 2: Checking directory structure Pass 3: Checking directory connectivity Pass 4: Checking reference counts Unattached zero-length inode 15. Clear? no Unattached inode 15 Connect to /lost+found? no Pass 5: Checking group summary information /dev/sda: ********** WARNING: Filesystem still has errors ********** /dev/sda: 15/655360 files (0.0% non-contiguous), 66755/2621440 blocks This occurs in 'ext4_xattr_inode_create()'. If 'ext4_mark_inode_dirty()' fails, dropping i_nlink of the inode is needed. Or will lead to inode leak.	N/A	More Details
CVE-2022-50846	In the Linux kernel, the following vulnerability has been resolved: mmc: via-sdmmc: fix return value check of mmc_add_host() mmc_add_host() may return error, if we ignore its return value, it will lead two issues: 1. The memory that allocated in mmc_alloc_host() is leaked. 2. In the remove() path, mmc_remove_host() will be called to delete device, but it's not added yet, it will lead a kernel crash because of null-ptr-deref in device_del(). Fix this by checking the return value and goto error path which will call mmc_free_host().	N/A	More Details
CVE-2022-50847	In the Linux kernel, the following vulnerability has been resolved: drm/bridge: it6505: Initialize AUX channel in it6505_i2c_probe During device boot, the HPD interrupt could be triggered before the DRM subsystem registers it6505 as a DRM bridge. In such cases, the driver tries to access AUX channel and causes NULL pointer dereference. Initializing the AUX channel earlier to prevent such error.	N/A	More Details
CVE-2022-50848	In the Linux kernel, the following vulnerability has been resolved: drivers: dio: fix possible memory leak in dio_init() If device_register() returns error, the 'dev' and name needs be freed. Add a release function, and then call put_device() in the error path, so the name is freed in kobject_cleanup() and to the 'dev' is freed in release function.	N/A	More Details
CVE-2022-50850	In the Linux kernel, the following vulnerability has been resolved: scsi: ipr: Fix WARNING in ipr_init() ipr_init() will not call unregister_reboot_notifier() when pci_register_driver() fails, which causes a WARNING. Call unregister_reboot_notifier() when pci_register_driver() fails. notifier callback ipr_halt [ipr] already registered WARNING: CPU: 3 PID: 299 at kernel/notifier.c:29 notifier_chain_register+0x16d/0x230 Modules linked in: ipr(+) xhci_pci_renesas xhci_hcd ehci_hcd usbcore led_class gpu_sched drm_buddy video wmi drm_ttm_helper ttm drm_display_helper drm_kms_helper drm drm_panel_orientation_quirks agpgart cfbft CPU: 3 PID: 299 Comm: modprobe Tainted: G W 6.1.0-rc1-00190-g39508d23b672-dirty #332 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.15.0-0-g2dd4b9b3f840-prebuilt.qemu.org 04/01/2014 RIP:	N/A	More Details

	0010:notifier_chain_register+0x16d/0x230 Call Trace: <TASK> __blocking_notifier_chain_register+0x73/0xb0 ipr_init+0x30/0x1000 [ipr] do_one_initcall+0xdb/0x480 do_init_module+0x1cf/0x680 load_module+0x6a50/0x70a0 __do_sys_finit_module+0x12f/0x1c0 do_syscall_64+0x3f/0x90 entry_SYSCALL_64_after_hwframe+0x63/0xcd		
CVE-2022-50859	In the Linux kernel, the following vulnerability has been resolved: cifs: Fix the error length of VALIDATE_NEGOTIATE_INFO message Commit d5c7076b772a ("smb3: add smb3.1.1 to default dialect list") extend the dialects from 3 to 4, but forget to decrease the extended length when specific the dialect, then the message length is larger than expected. This maybe leak some info through network because not initialize the message body. After apply this patch, the VALIDATE_NEGOTIATE_INFO message length is reduced from 28 bytes to 26 bytes.	N/A	More Details
CVE-2022-50851	In the Linux kernel, the following vulnerability has been resolved: vhost_vdpa: fix the crash in unmap a large memory While testing in vIOMMU, sometimes Guest will unmap very large memory, which will cause the crash. To fix this, add a new function vhost_vdpa_general_unmap(). This function will only unmap the memory that saved in iotlb. Call Trace: [647.820144] ------------[cut here]------------ [647.820848] kernel BUG at drivers/iommu/intel/iommu.c:1174! [647.821486] invalid opcode: 0000 [#1] PREEMPT SMP PTI [647.822082] CPU: 10 PID: 1181 Comm: qemu-system-x86 Not tainted 6.0.0-rc1home_lulu_2452_lulu7_vhost+ #62 [647.823139] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.15.0-29-g6a62e0cb0dfe-prebuilt.qem4 [647.824365] RIP: 0010:domain_unmap+0x48/0x110 [647.825424] Code: 48 89 fb 8d 4c f6 1e 39 c1 0f 4f c8 83 e9 0c 83 f9 3f 7f 18 48 89 e8 48 d3 e8 48 85 c0 75 59 [647.828064] RSP: 0018:ffffae5340c0bbf0 EFLAGS: 00010202 [647.828973] RAX: 0000000000000001 RBX: ffff921793d10540 RCX: 000000000000001b [647.830083] RDX: 00000000080000ff RSI: 0000000000000001 RDI: ffff921793d10540 [647.831214] RBP: 0000000007fc0100 R08: ffffae5340c0bcd0 R09: 0000000000000003 [647.832388] R10: 0000007fc0100000 R11: 0000000000100000 R12: 00000000080000ff [647.833668] R13: ffffae5340c0bcd0 R14: ffff921793d10590 R15: 0000008000100000 [647.834782] FS: 00007f772ec90640(0000) GS:ffff921ce7a80000(0000) knlGS:0000000000000000 [647.836004] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [647.836990] CR2: 00007f02c27a3a20 CR3: 0000000101b0c006 CR4: 0000000000372ee0 [647.838107] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [647.839283] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [647.840666] Call Trace: [647.841437] <TASK> [647.842107] intel_iommu_unmap_pages+0x93/0x140 [647.843112] __iommu_unmap+0x91/0x1b0 [647.844003] iommu_unmap+0x6a/0x95 [647.844885] vhost_vdpa_unmap+0x1de/0x1f0 [vhost_vdpa] [647.845985] vhost_vdpa_process_iotlb_msg+0xf0/0x90b [vhost_vdpa] [647.847235] ? _raw_spin_unlock+0x15/0x30 [647.848181] ? _copy_from_iter+0x8c/0x580 [647.849137] vhost_chr_write_iter+0xb3/0x430 [vhost] [647.850126] vfs_write+0x1e4/0x3a0 [647.850897] ksys_write+0x53/0xd0 [647.851688] do_syscall_64+0x3a/0x90 [647.852508] entry_SYSCALL_64_after_hwframe+0x63/0xcd [647.853457] RIP: 0033:0x7f7734ef9f4f [647.854408] Code: 89 54 24 18 48 89 74 24 10 89 7c 24 08 e8 29 76 f8 ff 48 8b 54 24 18 48 8b 74 24 10 41 89 c8 [647.857217] RSP: 002b:00007f772ec8f040 EFLAGS: 00000293 ORIG_RAX: 0000000000000001 [647.858486] RAX: ffffffffffffffda RBX: 00000000fef00000 RCX: 00007f7734ef9f4f [647.859713] RDX: 0000000000000048 RSI: 00007f772ec8f090 RDI: 0000000000000010 [647.860942] RBP: 00007f772ec8f1a0 R08: 0000000000000000 R09: 0000000000000000 [647.862206] R10: 0000000000000001 R11: 0000000000000293 R12: 0000000000000010 [647.863446] R13: 0000000000000002 R14: 0000000000000000 R15: ffffffff01100000 [647.864692] </TASK> [647.865458] Modules linked in: rpcsec_gss_krb5 auth_rpcgss nfsv4 dns_resolver nfs lockd grace fscache netfs v] [647.874688] ---[end trace 0000000000000000]---	N/A	More Details
CVE-2022-50852	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7921: fix use after free in mt7921_acpi_read() Don't dereference "sar_root" after it has been freed.	N/A	More Details
CVE-2022-50853	In the Linux kernel, the following vulnerability has been resolved: NFSv4: Fix a credential leak in _nfs4_discover_trunking()	N/A	More Details
CVE-2022-50854	In the Linux kernel, the following vulnerability has been resolved: nfc: virtual_ncidev: Fix memory leak in virtual_nci_send() skb should be free in virtual_nci_send(), otherwise kmemleak will report memleak. Steps for reproduction (simulated in qemu): cd tools/testing/selftests/nci make ./nci_dev BUG: memory leak unreferenced object 0xffff888107588000 (size 208): comm "nci_dev", pid 206, jiffies 4294945376 (age 368.248s) hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<000000008d94c8fd>] __alloc_skb+0x1da/0x290 [<00000000278bc7f8>] nci_send_cmd+0xa3/0x350 [<0000000081256a22>] nci_reset_req+0x6b/0xa0 [<000000009e721112>] __nci_request+0x90/0x250 [<000000005d556e59>] nci_dev_up+0x217/0x5b0 [<00000000e618ce62>] nfc_dev_up+0x114/0x220 [<00000000981e226b>] nfc_genl_dev_up+0x94/0xe0 [<000000009bb03517>] genl_family_rcv_msg_doit.isra.14+0x228/0x2d0 [<00000000b7f8c101>] genl_rcv_msg+0x35c/0x640 [<00000000c94075ff>] netlink_rcv_skb+0x11e/0x350 [<00000000440cfb1e>] genl_rcv+0x24/0x40 [<0000000062593b40>] netlink_unicast+0x43f/0x640 [<000000001d0b13cc>] netlink_sendmsg+0x73a/0xbf0 [<000000003272487f>] __sys_sendto+0x324/0x370 [<00000000ef9f1747>] __x64_sys_sendto+0xdd/0x1b0 [<000000001e437841>] do_syscall_64+0x3f/0x90	N/A	More Details
CVE-2022-50855	In the Linux kernel, the following vulnerability has been resolved: bpf: prevent leak of lsm program after failed attach In [0], we added the ability to bpf_prog_attach LSM programs to cgroups, but in our validation to make sure the prog is meant to be attached to BPF_LSM_CGROUP, we return too early if the check fails. This results in lack of decrementing prog's refcnt (through bpf_prog_put) leaving the LSM program alive past the point of the expected lifecycle. This fix allows for the decrement to take place. [0] https://lore.kernel.org/all/20220628174314.1216643-4-sdf@google.com/	N/A	More Details
CVE-2022-50856	In the Linux kernel, the following vulnerability has been resolved: cifs: Fix xid leak in cifs_ses_add_channel() Before return, should free the xid, otherwise, the xid will be leaked.	N/A	More Details
CVE-2022-50857	In the Linux kernel, the following vulnerability has been resolved: rapidio: rio: fix possible name leak in rio_register_mport() If device_register() returns error, the name allocated by dev_set_name() need be freed. It should use put_device() to give up the reference in the error path, so that the name can be freed in kobject_cleanup(), and list_del() is called to delete the port from rio_mports.	N/A	More Details
CVE-2022-50858	In the Linux kernel, the following vulnerability has been resolved: mmc: alcor: fix return value check of mmc_add_host() mmc_add_host() may return error, if we ignore its return value, the memory that allocated in mmc_alloc_host() will be leaked and it will lead a kernel crash because of deleting not added device in the remove path. So fix this by checking the return value and calling mmc_free_host() in the error path.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: PCI/DOE: Fix destroy_work_on_stack() race The following debug object splat was observed in testing: ODEBUG: free active (active state 0) object: 0000000097d23782 object type: work_struct hint: doe_statemachine_work+0x0/0x510 WARNING: CPU: 1 PID: 71 at lib/debugobjects.c:514 debug_print_object+0x7d/0xb0 ...		

CVE-2023-54235	Workqueue: pci 0000:36:00.0 DOE [1 doe_statemachine_work RIP: 0010:debug_print_object+0x7d/0xb0 ... Call Trace: ? debug_print_object+0x7d/0xb0 ? __pfx_doe_statemachine_work+0x10/0x10 debug_object_free.part.0+0x11b/0x150 doe_statemachine_work+0x45e/0x510 process_one_work+0x1d4/0x3c0 This occurs because destroy_work_on_stack() was called after signaling the completion in the calling thread. This creates a race between destroy_work_on_stack() and the task->work struct going out of scope in pci_doe(). Signal the work complete after destroying the work struct. This is safe because signal_task_complete() is the final thing the work item does and the workqueue code is careful not to access the work struct after.	N/A	More Details
CVE-2023-54239	In the Linux kernel, the following vulnerability has been resolved: iommufd: Check for uptr overflow syzkaller found that setting up a map with a user VA that wraps past zero can trigger WARN_ONs, particularly from pin_user_pages weirdly returning 0 due to invalid arguments. Prevent creating a pages with a uptr and size that would math overflow. WARNING: CPU: 0 PID: 518 at drivers/iommu/iommufd/pages.c:793 pfn_reader_user_pin+0x2e6/0x390 Modules linked in: CPU: 0 PID: 518 Comm: repro Not tainted 6.3.0-rc2-eeac8ede1755+ #1 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 RIP: 0010:pfn_reader_user_pin+0x2e6/0x390 Code: b1 11 e9 25 fe ff ff e8 28 e4 0f ff 31 ff 48 89 de e8 2e e6 0f ff 48 85 db 74 0a e8 14 e4 0f ff e9 4d ff ff ff e8 0a e4 0f ff <0f> 0b bb f2 ff ff ff e9 3c ff ff ff e8 f9 e3 0f ff ba 01 00 00 00 RSP: 0018:ffffc90000f9fa30 EFLAGS: 00010246 RAX: 0000000000000000 RBX: 0000000000000000 RCX: ffffffff821e2b72 RDX: 0000000000000000 RSI: ffff888014184680 RDI: 0000000000000002 RBP: ffffc90000f9fa78 R08: 00000000000000ff R09: 0000000079de6f4e R10: ffffc90000f9f790 R11: ffff888014185418 R12: ffffc90000f9fc60 R13: 0000000000000002 R14: ffff888007879800 R15: 0000000000000000 FS: 00007f4227555740(0000) GS:ffff88807dc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000020000043 CR3: 000000000e748005 CR4: 0000000000770ef0 PKRU: 55555554 Call Trace: <TASK> pfn_reader_next+0x14a/0x7b0 ? interval_tree_double_span_iter_update+0x11a/0x140 pfn_reader_first+0x140/0x1b0 iopt_pages_rw_slow+0x71/0x280 ? __this_cpu_preempt_check+0x20/0x30 iopt_pages_rw_access+0x2b2/0x5b0 iommufd_access_rw+0x19f/0x2f0 iommufd_test+0xd11/0x16f0 ? write_comp_data+0x2f/0x90 iommufd_fops_ioctl+0x206/0x330 __x64_sys_ioctl+0x10e/0x160 ? __pfx_iommufd_fops_ioctl+0x10/0x10 do_syscall_64+0x3b/0x90 entry_SYSCALL_64_after_hwframe+0x72/0xdc	N/A	More Details
CVE-2023-54237	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix potential panic dues to unprotected smc_llc_srv_add_link() There is a certain chance to trigger the following panic: PID: 5900 TASK: ffff88c1c8af4100 CPU: 1 COMMAND: "kworker/1:48" #0 [ffff9456c1cc79a0] machine_kexec at ffffffff870665b7 #1 [ffff9456c1cc79f0] __crash_kexec at ffffffff871b4c7a #2 [ffff9456c1cc7ab0] crash_kexec at ffffffff871b5b60 #3 [ffff9456c1cc7ac0] oops_end at ffffffff87026ce7 #4 [ffff9456c1cc7ae0] page_fault_oops at ffffffff87075715 #5 [ffff9456c1cc7b58] exc_page_fault at ffffffff87ad0654 #6 [ffff9456c1cc7b80] asm_exc_page_fault at ffffffff87c00b62 [exception RIP: ib_alloc_mr+19] RIP: ffffffffc0c9cce3 RSP: ffff9456c1cc7c38 RFLAGS: 00010202 RAX: 0000000000000000 RBX: 0000000000000002 RCX: 0000000000000004 RDX: 0000000000000010 RSI: 0000000000000000 RDI: 0000000000000000 RBP: ffff88c1ea281d00 R8: 000000020a34ffff R9: ffff88c1350bbb20 R10: 0000000000000000 R11: 0000000000000001 R12: 0000000000000000 R13: 0000000000000010 R14: ffff88c1ab040a50 R15: ffff88c1ea281d00 ORIG_RAX: ffffffffffffffff CS: 0010 SS: 0018 #7 [ffff9456c1cc7c60] smc_ib_get_memory_region at ffffffffc0aff6df [smc] #8 [ffff9456c1cc7c88] smcr_buf_map_link at ffffffffc0b0278c [smc] #9 [ffff9456c1cc7ce0] __smc_buf_create at ffffffffc0b03586 [smc] The reason here is that when the server tries to create a second link, smc_llc_srv_add_link() has no protection and may add a new link to link group. This breaks the security environment protected by llc_conf_mutex.	N/A	More Details
CVE-2025-66848	JD Cloud NAS routers AX1800 (4.3.1.r4308 and earlier), AX3000 (4.3.1.r4318 and earlier), AX6600 (4.5.1.r4533 and earlier), BE6500 (4.4.1.r4308 and earlier), ER1 (4.5.1.r4518 and earlier), and ER2 (4.5.1.r4518 and earlier) contain an unauthorized remote command execution vulnerability.	N/A	More Details
CVE-2023-53992	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: ocb: don't leave if not joined If there's no OCB state, don't ask the driver/mac80211 to leave, since that's just confusing. Since set/clear the chandef state, that's a simple check.	N/A	More Details
CVE-2023-53991	In the Linux kernel, the following vulnerability has been resolved: drm/msm/dpu: Disallow unallocated resources to be returned In the event that the topology requests resources that have not been created by the system (because they are typically not represented in dpu_mdss_cfg ^1), the resource(s) in global_state (in this case DSC blocks, until their allocation/assignment is being sanity-checked in "drm/msm/dpu: Reject topologies for which no DSC blocks are available") remain NULL but will still be returned out of dpu_rm_get_assigned_resources, where the caller expects to get an array containing num_blks valid pointers (but instead gets these NULLs). To prevent this from happening, where null-pointer dereferences typically result in a hard-to-debug platform lockup, num_blks shouldn't increase past NULL blocks and will print an error and break instead. After all, max_blks represents the static size of the maximum number of blocks whereas the actual amount varies per platform. ^1: which can happen after a git rebase ended up moving additions to _dpu_cfg to a different struct which has the same patch context. Patchwork: https://patchwork.freedesktop.org/patch/517636/	N/A	More Details
CVE-2023-53990	In the Linux kernel, the following vulnerability has been resolved: SMB3: Add missing locks to protect deferred close file list cifs_del_deferred_close function has a critical section which modifies the deferred close file list. We must acquire deferred_lock before calling cifs_del_deferred_close function.	N/A	More Details
CVE-2023-54134	In the Linux kernel, the following vulnerability has been resolved: autofs: fix memory leak of waitqueues in autofs_catatonic_mode Syzkaller reports a memory leak: BUG: memory leak unreferenced object 0xffff88810b279e00 (size 96): comm "syz-executor399", pid 3631, jiffies 4294964921 (age 23.870s) hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 08 9e 27 0b 81 88 ff ff ....'..... 08 9e 27 0b 81 88 ff ff 00 00 00 00 00 00 00 00 ..'..... backtrace: [<ffffffff814cfc90>] kmalloc_trace+0x20/0x90 mm/slab_common.c:1046 [<ffffffff81bb75ca>] kmalloc include/linux/slab.h:576 [inline] [<ffffffff81bb75ca>] autofs_wait+0x3fa/0x9a0 fs/autofs/waitq.c:378 [<ffffffff81bb88a7>] autofs_do_expire_multi+0xa7/0x3e0 fs/autofs/expire.c:593 [<ffffffff81bb8c33>] autofs_expire_multi+0x53/0x80 fs/autofs/expire.c:619 [<ffffffff81bb6972>] autofs_root_ioctl_unlocked+0x322/0x3b0 fs/autofs/root.c:897 [<ffffffff81bb6a95>] autofs_root_ioctl+0x25/0x30 fs/autofs/root.c:910 [<ffffffff81602a9c>] vfs_ioctl fs/ioctl.c:51 [inline] [<ffffffff81602a9c>] __do_sys_ioctl fs/ioctl.c:870 [inline] [<ffffffff81602a9c>] __se_sys_ioctl fs/ioctl.c:856 [inline] [<ffffffff81602a9c>] __x64_sys_ioctl+0xfc/0x140 fs/ioctl.c:856 [<ffffffff84608225>] do_syscall_x64 arch/x86/entry/common.c:50 [inline] [<ffffffff84608225>] do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80 [<ffffffff84800087>] entry_SYSCALL_64_after_hwframe+0x63/0xcd autofs_wait_queue structs should be freed if their wait_ctr becomes zero. Otherwise they will be lost. In this case an AUTOFS_IOC_EXPIRE_MULTI ioctl is done, then a new waitqueue struct is allocated in autofs_wait(), its initial wait_ctr equals 2. After that wait_event_killable() is interrupted (it returns -ERESTARTSYS), so that 'wq->name.name == NULL' condition may be not satisfied. Actually, this condition can be satisfied when autofs_wait_release() or autofs_catatonic_mode() is called and, what is also important, wait_ctr is decremented in those places. Upon the exit of autofs_wait(), wait_ctr is decremented to 1. Then the unmounting process begins: kill_sb calls autofs_catatonic_mode(), which should have freed the waitqueues, but it only decrements its usage counter to zero which is not a correct behaviour. edit:imk This description is of course not correct. The umount performed as a result of an expire is a umount of a mount that has been automounted, it's not the autofs mount itself. They happen independently, usually after everything mounted within the autofs file system has been expired away. If everything hasn't been expired away the automount daemon can still exit leaving mounts in place.	N/A	More Details

	But expires done in both cases will result in a notification that calls autofs_wait_release() with a result status. The problem case is the summary execution of of the automount daemon. In this case any waiting processes won't be woken up until either they are terminated or the mount is umounted. end edit: imk So in catatonic mode we should free waitqueues which counter becomes zero. edit: imk Initially I was concerned that the calling of autofs_wait_release() and autofs_catatonic_mode() was not mutually exclusive but that can't be the case (obviously) because the queue entry (or entries) is removed from the list when either of these two functions are called. Consequently the wait entry will be freed by only one of these functions or by the woken process in autofs_wait() depending on the order of the calls. end edit: imk		
CVE-2023-54133	In the Linux kernel, the following vulnerability has been resolved: nfp: clean mc addresses in application firmware when closing port When moving devices from one namespace to another, mc addresses are cleaned in software while not removed from application firmware. Thus the mc addresses are remained and will cause resource leak. Now use `__dev_mc_unsync` to clean mc addresses when closing port.	N/A	More Details
CVE-2025-65925	An issue was discovered in Zeroheight (SaaS) prior to 2025-06-13. A legacy user creation API pathway allowed accounts to be created without completing the intended email verification step. While unverified accounts could not access product functionality, the behavior bypassed intended verification controls and allowed unintended account creation. This could have enabled spam/fake account creation or resource usage impact. No data exposure or unauthorized access to existing accounts was reported.	N/A	More Details
CVE-2023-53989	In the Linux kernel, the following vulnerability has been resolved: arm64: mm: fix VA-range sanity check Both create_mapping_noalloc() and update_mapping_prot() sanity-check their 'virt' parameter, but the check itself doesn't make much sense. The condition used today appears to be a historical accident. The sanity-check condition: if ((virt >= PAGE_END) && (virt < VMALLOC_START)) { [... warning here ...] return; } ... can only be true for the KASAN shadow region or the module region, and there's no reason to exclude these specifically for creating and updateing mappings. When arm64 support was first upstreamed in commit: c1cc1552616d0f35 ("arm64: MMU initialisation") ... the condition was: if (virt < VMALLOC_START) { [... warning here ...] return; } At the time, VMALLOC_START was the lowest kernel address, and this was checking whether 'virt' would be translated via TTBR1. Subsequently in commit: 14c127c957c1c607 ("arm64: mm: Flip kernel VA space") ... the condition was changed to: if ((virt >= VA_START) && (virt < VMALLOC_START)) { [... warning here ...] return; } This appear to have been a thinko. The commit moved the linear map to the bottom of the kernel address space, with VMALLOC_START being at the halfway point. The old condition would warn for changes to the linear map below this, and at the time VA_START was the end of the linear map. Subsequently we cleaned up the naming of VA_START in commit: 77ad4ce69321abbe ("arm64: memory: rename VA_START to PAGE_END") ... keeping the erroneous condition as: if ((virt >= PAGE_END) && (virt < VMALLOC_START)) { [... warning here ...] return; } Correct the condition to check against the start of the TTBR1 address space, which is currently PAGE_OFFSET. This simplifies the logic, and more clearly matches the "outside kernel range" message in the warning.	N/A	More Details
CVE-2023-53988	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix slab-out-of-bounds read in hdr_delete_de() Here is a BUG report from syzbot: BUG: KASAN: slab-out-of-bounds in hdr_delete_de+0xe0/0x150 fs/ntfs3/index.c:806 Read of size 16842960 at addr ffff888079cc0600 by task syz-executor934/3631 Call Trace: memmove+0x25/0x60 mm/kasan/shadow.c:54 hdr_delete_de+0xe0/0x150 fs/ntfs3/index.c:806 indx_delete_entry+0x74f/0x3670 fs/ntfs3/index.c:2193 ni_remove_name+0x27a/0x980 fs/ntfs3/frecord.c:2910 ntfs_unlink_inode+0x3d4/0x720 fs/ntfs3/inode.c:1712 ntfs_rename+0x41a/0xcb0 fs/ntfs3/namei.c:276 Before using the meta-data in struct INDEX_HDR, we need to check index header valid or not. Otherwise, the corruptedi (or malicious) fs image can cause out-of-bounds access which could make kernel panic.	N/A	More Details
CVE-2023-54132	In the Linux kernel, the following vulnerability has been resolved: erofs: stop parsing non-compact HEAD index if clusterofs is invalid Syzbot generated a crafted image [1] with a non-compact HEAD index of clusterofs 33024 while valid numbers should be 0 ~ lclustersize-1, which causes the following unexpected behavior as below: BUG: unable to handle page fault for address: fffff52101a3fff9 #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 23ffed067 P4D 23ffed067 PUD 0 Oops: 0000 [#1] PREEMPT SMP KASAN CPU: 1 PID: 4398 Comm: kworker/u5:1 Not tainted 6.3.0-rc6-syzkaller-g09a9639e56c0 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 03/30/2023 Workqueue: erofs_worker z_erofs_decompressqueue_work RIP: 0010:z_erofs_decompress_queue+0xb7e/0x2b40 ... Call Trace: <TASK> z_erofs_decompressqueue_work+0x99/0xe0 process_one_work+0x8f6/0x1170 worker_thread+0xa63/0x1210 kthread+0x270/0x300 ret_from_fork+0x1f/0x30 Note that normal images or images using compact indexes are not impacted. Let's fix this now. [1] https://lore.kernel.org/r/000000000000ec75b005ee97fbaa@google.com	N/A	More Details
CVE-2023-54135	In the Linux kernel, the following vulnerability has been resolved: maple_tree: fix potential out-of-bounds access in mas_wr_end_piv() Check the write offset end bounds before using it as the offset into the pivot array. This avoids a possible out-of-bounds access on the pivot array if the write extends to the last slot in the node, in which case the node maximum should be used as the end pivot. akpm: this doesn't affect any current callers, but new users of mapletree may encounter this problem if backported into earlier kernels, so let's fix it in -stable kernels in case of this.	N/A	More Details
CVE-2023-54131	In the Linux kernel, the following vulnerability has been resolved: wifi: rt2x00: Fix memory leak when handling surveys When removing a rt2x00 device, its associated channel surveys are not freed, causing a memory leak observable with kmemleak: unreferenced object 0xffff9620f0881a00 (size 512): comm "systemd-udevd", pid 2290, jiffies 4294906974 (age 33.768s) hex dump (first 32 bytes): 70 44 12 00 00 00 00 00 92 8a 00 00 00 00 00 00 pD.............. 00 00 00 00 00 00 00 00 ab 87 01 00 00 00 00 00 ..... backtrace: [<ffffffffb0ed858b>] __kmalloc+0x4b/0x130 [<ffffffffc1b0f29b>] rt2800_probe_hw+0xc2b/0x1380 [rt2800lib] [<ffffffffc1a9496e>] rt2800usb_probe_hw+0xe/0x60 [rt2800usb] [<ffffffffc1ae491a>] rt2x00lib_probe_dev+0x21a/0x7d0 [rt2x00lib] [<ffffffffc1b3b83e>] rt2x00usb_probe+0x1be/0x980 [rt2x00usb] [<ffffffffc05981e2>] usb_probe_interface+0xe2/0x310 [usbcore] [<ffffffffb13be2d5>] really_probe+0x1a5/0x410 [<ffffffffb13be5c8>] __driver_probe_device+0x78/0x180 [<ffffffffb13be6fe>] driver_probe_device+0x1e/0x90 [<ffffffffb13be972>] __driver_attach+0xd2/0x1c0 [<ffffffffb13bbc57>] bus_for_each_dev+0x77/0xd0 [<ffffffffb13bd2a2>] bus_add_driver+0x112/0x210 [<ffffffffb13bfc6c>] driver_register+0x5c/0x120 [<ffffffffc0596ae8>] usb_register_driver+0x88/0x150 [usbcore] [<ffffffffb0c011c4>] do_one_initcall+0x44/0x220 [<ffffffffb0d6134c>] do_init_module+0x4c/0x220 Fix this by freeing the channel surveys on device removal. Tested with a RT3070 based USB wireless adapter.	N/A	More Details
CVE-2023-54130	In the Linux kernel, the following vulnerability has been resolved: hfs/hfsplus: avoid WARN_ON() for sanity check, use proper error handling Commit 55d1cbbbb29e ("hfs/hfsplus: use WARN_ON for sanity check") fixed a build warning by turning a comment into a WARN_ON(), but it turns out that syzbot then complains because it can trigger said warning with a corrupted hfs image. The warning actually does warn about a bad situation, but we are much better off just handling it as the error it is. So rather than warn about us doing bad things, stop doing the bad things and return -EIO. While at it, also fix a memory leak that was introduced by an earlier fix for a similar syzbot warning situation, and add a check for one case that historically wasn't handled at all (ie neither comment nor subsequent WARN_ON).	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: Add validation for lmac type Upon physical link		

CVE-2023-54129	change, firmware reports to the kernel about the change along with the details like speed, lmac_type_id, etc. Kernel derives lmac_type based on lmac_type_id received from firmware. In a few scenarios, firmware returns an invalid lmac_type_id, which is resulting in below kernel panic. This patch adds the missing validation of the lmac_type_id field. Internal error: Oops: 96000005 [#1] PREEMPT SMP [35.321595] Modules linked in: [35.328982] CPU: 0 PID: 31 Comm: kworker/0:1 Not tainted 5.4.210-g2e3169d8e1bc-dirty #17 [35.337014] Hardware name: Marvell CN103XX board (DT) [35.344297] Workqueue: events work_for_cpu_fn [35.352730] pstate: 40400089 (nZcv daIf +PAN -UAO) [35.360267] pc : strncpy+0x10/0x30 [35.366595] lr : cgx_link_change_handler+0x90/0x180	N/A	More Details
CVE-2023-54128	In the Linux kernel, the following vulnerability has been resolved: fs: drop peer group ids under namespace lock When cleaning up peer group ids in the failure path we need to make sure to hold on to the namespace lock. Otherwise another thread might just turn the mount from a shared into a non-shared mount concurrently.	N/A	More Details
CVE-2025-56332	Authentication Bypass in fosrl/pangolin v1.6.2 and before allows attackers to access Pangolin resource via Insecure Default Configuration	N/A	More Details
CVE-2025-65409	A divide-by-zero in the encryption/decryption routines of GNU Recutils v1.9 allows attackers to cause a Denial of Service (DoS) via inputting an empty value as a password.	N/A	More Details
CVE-2025-65411	A NULL pointer dereference in the src/path.c component of GNU Unrtf v0.21.10 allows attackers to cause a Denial of Service (DoS) via injecting a crafted payload into the search_path parameter.	N/A	More Details
CVE-2023-53987	In the Linux kernel, the following vulnerability has been resolved: ping: Fix potentail NULL deref for /proc/net/icmp. After commit dbca1596bbb0 ("ping: convert to RCU lookups, get rid of rwlock"), we use RCU for ping sockets, but we should use spinlock for /proc/net/icmp to avoid a potential NULL deref mentioned in the previous patch. Let's go back to using spinlock there. Note we can convert ping sockets to use hlist instead of hlist_nulls because we do not use SLAB_TYPESAFE_BY_RCU for ping sockets.	N/A	More Details
CVE-2025-67746	Composer is a dependency manager for PHP. In versions on the 2.x branch prior to 2.2.26 and 2.9.3, attackers controlling remote sources that Composer downloads from might in some way inject ANSI control characters in the terminal output of various Composer commands, causing mangled output and potentially leading to confusion or DoS of the terminal application. There is no proven exploit and this has thus a low severity but we still publish a CVE as it has potential for abuse, and we want to be on the safe side informing users that they should upgrade. Versions 2.2.26 and 2.9.3 contain a patch for the issue.	N/A	More Details
CVE-2025-64528	Discourse is an open source discussion platform. Prior to versions 3.5.3, 2025.11.1, and 2025.12.0, an attacker who knows part of a username can find the user and their full name via UI or API, even when `enable_names` is disabled. Versions 3.5.3, 2025.11.1, and 2025.12.0 contain a fix.	N/A	More Details
CVE-2023-54324	In the Linux kernel, the following vulnerability has been resolved: dm: fix a race condition in retrieve_deps There's a race condition in the multipath target when retrieve_deps races with multipath_message calling dm_get_device and dm_put_device. retrieve_deps walks the list of open devices without holding any lock but multipath may add or remove devices to the list while it is running. The end result may be memory corruption or use-after-free memory access. See this description of a UAF with multipath_message(): https://listman.redhat.com/archives/dm-devel/2022-October/052373.html Fix this bug by introducing a new rw semaphore "devices_lock". We grab devices_lock for read in retrieve_deps and we grab it for write in dm_get_device and dm_put_device.	N/A	More Details
CVE-2023-54139	In the Linux kernel, the following vulnerability has been resolved: tracing/user_events: Ensure write index cannot be negative The write index indicates which event the data is for and accesses a per-file array. The index is passed by user processes during write() calls as the first 4 bytes. Ensure that it cannot be negative by returning -EINVAL to prevent out of bounds accesses. Update ftrace self-test to ensure this occurs properly.	N/A	More Details
CVE-2023-54326	In the Linux kernel, the following vulnerability has been resolved: misc: pci_endpoint_test: Free IRQs before removing the device In pci_endpoint_test_remove(), freeing the IRQs after removing the device creates a small race window for IRQs to be received with the test device memory already released, causing the IRQ handler to access invalid memory, resulting in an oops. Free the device IRQs before removing the device to avoid this issue.	N/A	More Details
CVE-2024-58242	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2024-58243	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2024-58244	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2024-58245	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2024-58246	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2024-58247	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2024. Notes: none.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix WARNING in mark_buffer_dirty due to discarded buffer reuse A syzbot stress test using a corrupted disk image reported that mark_buffer_dirty() called from __nilfs_mark_inode_dirty() or nilfs_palloc_commit_alloc_entry() may output a kernel warning, and can panic if the kernel is booted with panic_on_warn. This is because nilfs2 keeps buffer pointers in local structures for some metadata and reuses them, but such buffers may be forcibly discarded by nilfs_clear_dirty_page() in some critical situations. This issue is reported to appear after commit 28a65b49eb53 ("nilfs2:	N/A	More Details

54140	do not write dirty data after degenerating to read-only"), but the issue has potentially existed before. Fix this issue by checking the uptodate flag when attempting to reuse an internally held buffer, and reloading the metadata instead of reusing the buffer if the flag was lost.		
CVE-2023-54138	In the Linux kernel, the following vulnerability has been resolved: drm/msm: fix NULL-deref on irq uninstall In case of early initialisation errors and on platforms that do not use the DPU controller, the deinitilisation code can be called with the kms pointer set to NULL. Patchwork: https://patchwork.freedesktop.org/patch/525104/	N/A	More Details
CVE-2023-53993	In the Linux kernel, the following vulnerability has been resolved: PCI/DOE: Fix memory leak with CONFIG_DEBUG_OBJECTS=y After a pci_doe_task completes, its work_struct needs to be destroyed to avoid a memory leak with CONFIG_DEBUG_OBJECTS=y.	N/A	More Details
CVE-2023-54137	In the Linux kernel, the following vulnerability has been resolved: vfio/type1: fix cap_migration information leak Fix an information leak where an uninitialized hole in struct vfio_iommu_type1_info_cap_migration on the stack is exposed to userspace. The definition of struct vfio_iommu_type1_info_cap_migration contains a hole as shown in this pahole(1) output: struct vfio_iommu_type1_info_cap_migration { struct vfio_info_cap_header header; /* 0 8 */ __u32 flags; /* 8 4 */ /* XXX 4 bytes hole, try to pack */ __u64 pgsize_bitmap; /* 16 8 */ __u64 max_dirty_bitmap_size; /* 24 8 */ /* size: 32, cachelines: 1, members: 4 */ /* sum members: 28, holes: 1, sum holes: 4 */ /* last cacheline: 32 bytes */ }; The cap_mig variable is filled in without initializing the hole: static int vfio_iommu_migration_build_caps(struct vfio_iommu *iommu, struct vfio_info_cap *caps) { struct vfio_iommu_type1_info_cap_migration cap_mig; cap_mig.header.id = VFIO_IOMMU_TYPE1_INFO_CAP_MIGRATION; cap_mig.header.version = 1; cap_mig.flags = 0; /* support minimum pgsize */ cap_mig.pgsize_bitmap = (size_t)1 << __ffs(iommu->pgsize_bitmap); cap_mig.max_dirty_bitmap_size = DIRTY_BITMAP_SIZE_MAX; return vfio_info_add_capability(caps, &cap_mig.header, sizeof(cap_mig)); } The structure is then copied to a temporary location on the heap. At this point it's already too late and ioctl(VFIO_IOMMU_GET_INFO) copies it to userspace later: int vfio_info_add_capability(struct vfio_info_cap *caps, struct vfio_info_cap_header *cap, size_t size) { struct vfio_info_cap_header *header; header = vfio_info_cap_add(caps, size, cap->id, cap->version); if (IS_ERR(header)) return PTR_ERR(header); memcpy(header + 1, cap + 1, size - sizeof(*header)); return 0; } This issue was found by code inspection.	N/A	More Details
CVE-2023-54136	In the Linux kernel, the following vulnerability has been resolved: serial: sprd: Fix DMA buffer leak issue Release DMA buffer when _probe() returns failure to avoid memory leak.	N/A	More Details
CVE-2023-53998	In the Linux kernel, the following vulnerability has been resolved: hwrng: virtio - Fix race on data_avail and actual data The virtio rng device kicks off a new entropy request whenever the data available reaches zero. When a new request occurs at the end of a read operation, that is, when the result of that request is only needed by the next reader, then there is a race between the writing of the new data and the next reader. This is because there is no synchronisation whatsoever between the writer and the reader. Fix this by writing data_avail with smp_store_release and reading it with smp_load_acquire when we first enter read. The subsequent reads are safe because they're either protected by the first load acquire, or by the completion mechanism. Also remove the redundant zeroing of data_idx in random_recv_done (data_idx must already be zero at this point) and data_avail in request_entropy (ditto).	N/A	More Details
CVE-2023-53997	In the Linux kernel, the following vulnerability has been resolved: thermal: of: fix double-free on unregistration Since commit 3d439b1a2ad3 ("thermal/core: Alloc-copy-free the thermal zone parameters structure"), thermal_zone_device_register() allocates a copy of the tzp argument and frees it when unregistering, so thermal_of_zone_register() now ends up leaking its original tzp and double-freeing the tzp copy. Fix this by locating tzp on stack instead.	N/A	More Details
CVE-2023-53996	In the Linux kernel, the following vulnerability has been resolved: x86/sev: Make enc_dec_hypercall() accept a size instead of npages enc_dec_hypercall() accepted a page count instead of a size, which forced its callers to round up. As a result, non-page aligned vaddrs caused pages to be spuriously marked as decrypted via the encryption status hypercall, which in turn caused consistent corruption of pages during live migration. Live migration requires accurate encryption status information to avoid migrating pages from the wrong perspective.	N/A	More Details
CVE-2025-61557	nixseparatedebuginfod before v0.4.1 is vulnerable to Directory Traversal.	N/A	More Details
CVE-2023-53995	In the Linux kernel, the following vulnerability has been resolved: net: ipv4: fix one memleak in __inet_del_ifa() I got the below warning when do fuzzing test: unregister_netdevice: waiting for bond0 to become free. Usage count = 2 It can be repoduced via: ip link add bond0 type bond sysctl -w net.ipv4.conf.bond0.promote_secondaries=1 ip addr add 4.117.174.103/0 scope 0x40 dev bond0 ip addr add 192.168.100.111/255.255.255.254 scope 0 dev bond0 ip addr add 0.0.0.4/0 scope 0x40 secondary dev bond0 ip addr del 4.117.174.103/0 scope 0x40 dev bond0 ip link delete bond0 type bond In this reproduction test case, an incorrect 'last_prim' is found in __inet_del_ifa(), as a result, the secondary address(0.0.0.4/0 scope 0x40) is lost. The memory of the secondary address is leaked and the reference of in_device and net_device is leaked. Fix this problem: Look for 'last_prim' starting at location of the deleted IP and inserting the promoted IP into the location of 'last_prim'.	N/A	More Details
CVE-2023-53994	In the Linux kernel, the following vulnerability has been resolved: ionic: remove WARN_ON to prevent panic_on_warn Remove unnecessary early code development check and the WARN_ON that it uses. The irq alloc and free paths have long been cleaned up and this check shouldn't have stuck around so long.	N/A	More Details
CVE-2023-53986	In the Linux kernel, the following vulnerability has been resolved: mips: bmips: BCM6358: disable RAC flush for TP1 RAC flush causes kernel panics on BCM6358 with EHCI/OHCI when booting from TP1: [3.881739] usb 1-1: new high-speed USB device number 2 using ehci-platform [3.895011] Reserved instruction in kernel code[#1]: [3.900113] CPU: 0 PID: 1 Comm: init Not tainted 5.10.16 #0 [3.905829] $ 0 : 00000000 10008700 00000000 77d94060 [3.911238] $ 4 : 7fd1f088 00000000 81431cac 81431ca0 [3.916641] $ 8 : 00000000 ffffefff 8075cd34 00000000 [3.922043] $12 : 806f8d40 f3e812b7 00000000 000d9aaa [3.927446] $16 : 7fd1f068 7fd1f080 7ff559b8 81428470 [3.932848] $20 : 00000000 00000000 55590000 77d70000 [3.938251] $24 : 00000018 00000010 [3.943655] $28 : 81430000 81431e60 81431f28 800157fc [3.949058] Hi : 00000000 [3.952013] Lo : 00000000 [3.955019] epc : 80015808 setup_sigcontext+0x54/0x24c [3.960464] ra : 800157fc setup_sigcontext+0x48/0x24c [3.965913] Status: 10008703 KERNEL EXL IE [3.970216] Cause : 00800028 (ExcCode 0a) [3.974340] PrId : 0002a010 (Broadcom BMIPS4350) [3.979170] Modules linked in: ohci_platform ohci_hcd fsl_mph_dr_of ehci_platform ehci_fsl ehci_hcd gpio_button_hotplug usbcore nls_base usb_common [3.992907] Process init (pid: 1, threadinfo=(ptrval), task=(ptrval), tls=77e22ec8) [4.000776] Stack : 81431ef4 7fd1f080 81431f28 81428470 7fd1f068 81431edc 7ff559b8 81428470 [4.009467] 81431f28 7fd1f080 55590000 77d70000 77d5498c 80015c70 806f0000 8063ae74 [4.018149] 08100002 81431f28 0000000a 08100002 81431f28 0000000a 77d6b418 00000003 [4.026831] ffffffff 80016414 80080734 81431ecc 81431ecc 00000001 00000000 04000000 [4.035512] 77d54874 00000000 00000000 00000000 00000000 00000012 00000002 00000000 [4.044196] ... [4.046706] Call Trace: [4.049238]	N/A	More Details

	[<80015808>] setup_sigcontext+0x54/0x24c [4.054356] [<80015c70>] setup_frame+0xdc/0x124 [4.059015] [<80016414>] do_notify_resume+0x1dc/0x288 [4.064207] [<80011b50>] work_notifysig+0x10/0x18 [4.069036] [4.070538] Code: 8fc300b4 00001025 26240008 <ac820000> ac830004 3c048063 0c0228aa 24846a00 26240010 [4.080686] [4.082517] ---[end trace 22a8edb41f5f983b]--- [4.087374] Kernel panic - not syncing: Fatal exception [4.092753] Rebooting in 1 seconds.. Because the bootloader (CFE) is not initializing the Read-ahead cache properly on the second thread (TP1). Since the RAC was not initialized properly, we should avoid flushing it at the risk of corrupting the instruction stream as seen in the trace above.		
CVE-2025-66824	A Stored Cross-Site Scripting (XSS) vulnerability exists in the Meeting location field of the Create/Edit Conference functionality in TrueConf Server v5.5.2.10813. The injected payload is stored via the meeting_room parameter and executed when users visit the Conference Info page, allowing attackers to achieve full Account Takeover (ATO). This issue is caused by improper sanitization of user-supplied input in the meeting_room field.	N/A	More Details
CVE-2025-66834	A CSV Formula Injection vulnerability in TrueConf Server v5.5.2.10813 allows a normal user to inject malicious spreadsheet formulas into exported chat logs via crafted Display Name.	N/A	More Details
CVE-2023-54122	In the Linux kernel, the following vulnerability has been resolved: drm/msm/dpu: Add check for cstate As kzalloc may fail and return NULL pointer, it should be better to check cstate in order to avoid the NULL pointer dereference in __drm_atomic_helper_crtc_reset. Patchwork: https://patchwork.freedesktop.org/patch/514163/	N/A	More Details
CVE-2022-50700	In the Linux kernel, the following vulnerability has been resolved: wifi: ath10k: Delay the unmapping of the buffer On WCN3990, we are seeing a rare scenario where copy engine hardware is sending a copy complete interrupt to the host driver while still processing the buffer that the driver has sent, this is leading into an SMMU fault triggering kernel panic. This is happening on copy engine channel 3 (CE3) where the driver normally enqueues WMI commands to the firmware. Upon receiving a copy complete interrupt, host driver will immediately unmap and frees the buffer presuming that hardware has processed the buffer. In the issue case, upon receiving copy complete interrupt, host driver will unmap and free the buffer but since hardware is still accessing the buffer (which in this case got unmapped in parallel), SMMU hardware will trigger an SMMU fault resulting in a kernel panic. In order to avoid this, as a work around, add a delay before unmapping the copy engine source DMA buffer. This is conditionally done for WCN3990 and only for the CE3 channel where issue is seen. Below is the crash signature: wifi smmu error: kernel: [10.120965] arm-smmu 15000000.iommu: Unhandled context fault: fsr=0x402, iova=0x7fdfd8ac0, fsynr=0x500003,cbfrsynra=0xc1, cb=6 arm-smmu 15000000.iommu: Unhandled context fault:fsr=0x402, iova=0x7fe06fdc0, fsynr=0x710003, cbfrsynra=0xc1, cb=6 qcom-q6v5-mss 4080000.remoteproc: fatal error received: err_qdi.c:1040:EF:wlan_process:0x1:WLAN RT:0x2091: cmnos_thread.c:3998:Asserted in copy_engine.c:AXI_ERROR_DETECTED:2149 remoteproc remoteproc0: crash detected in 4080000.remoteproc: type fatal error <3> remoteproc remoteproc0: handling crash #1 in 4080000.remoteproc pc : __arm_lpae_unmap+0x500/0x514 lr : __arm_lpae_unmap+0x4bc/0x514 sp : ffffffc011ffb530 x29: ffffffc011ffb590 x28: 0000000000000000 x27: 0000000000000000 x26: 0000000000000004 x25: 0000000000000003 x24: ffffffc011ffb890 x23: ffffffa762ef9be0 x22: ffffffa77244ef00 x21: 0000000000000009 x20: 00000007fff7c000 x19: 0000000000000003 x18: 0000000000000000 x17: 0000000000000004 x16: ffffffd7a357d9f0 x15: 0000000000000000 x14: 00fd5d4fa7ffffff x13: 000000000000000e x12: 0000000000000000 x11: 00000000ffffffff x10: 00000000fffffe00 x9 : 000000000000017c x8 : 000000000000000c x7 : 0000000000000000 x6 : ffffffa762ef9000 x5 : 0000000000000003 x4 : 0000000000000004 x3 : 0000000000001000 x2 : 00000007fff7c000 x1 : ffffffc011ffb890 x0 : 0000000000000000 Call trace: __arm_lpae_unmap+0x500/0x514 __arm_lpae_unmap+0x4bc/0x514 __arm_lpae_unmap+0x4bc/0x514 arm_lpae_unmap_pages+0x78/0xa4 arm_smmu_unmap_pages+0x78/0x104 __iommu_unmap+0xc8/0x1e4 iommu_unmap_fast+0x38/0x48 __iommu_dma_unmap+0x84/0x104 iommu_dma_free+0x34/0x50 dma_free_attrs+0xa4/0xd0 ath10k_htt_rx_free+0xc4/0xf4 [ath10k_core] ath10k_core_stop+0x64/0x7c [ath10k_core] ath10k_halt+0x11c/0x180 [ath10k_core] ath10k_stop+0x54/0x94 [ath10k_core] drv_stop+0x48/0x1c8 [mac80211] ieee80211_do_open+0x638/0x77c [mac80211] ieee80211_open+0x48/0x5c [mac80211] __dev_open+0xb4/0x174 __dev_change_flags+0xc4/0x1dc dev_change_flags+0x3c/0x7c devinet_ioctl+0x2b4/0x580 inet_ioctl+0xb0/0x1b4 sock_do_ioctl+0x4c/0x16c compat_ifreq_ioctl+0x1cc/0x35c compat_sock_ioctl+0x110/0x2ac __arm64_compat_sys_ioctl+0xf4/0x3e0 el0_svc_common+0xb4/0x17c el0_svc_compat_handler+0x2c/0x58 el0_svc_compat+0x8/0x2c Tested-on: WCN3990 hw1.0 SNOC WLAN.HL.2.0-01387-QCAHLSWMTPLZ-1	N/A	More Details
CVE-2022-50699	In the Linux kernel, the following vulnerability has been resolved: selinux: enable use of both GFP_KERNEL and GFP_ATOMIC in convert_context() The following warning was triggered on a hardware environment: SELinux: Converting 162 SID table entries... BUG: sleeping function called from invalid context at __might_sleep+0x60/0x74 0x0 in_atomic(): 1, irqs_disabled(): 128, non_block: 0, pid: 5943, name: tar CPU: 7 PID: 5943 Comm: tar Tainted: P O 5.10.0 #1 Call trace: dump_backtrace+0x0/0x1c8 show_stack+0x18/0x28 dump_stack+0xe8/0x15c ___might_sleep+0x168/0x17c __might_sleep+0x60/0x74 __kmalloc_track_caller+0xa0/0x7dc kstrdup+0x54/0xac convert_context+0x48/0x2e4 sidtab_context_to_sid+0x1c4/0x36c security_context_to_sid_core+0x168/0x238 security_context_to_sid_default+0x14/0x24 inode_doinit_use_xattr+0x164/0x1e4 inode_doinit_with_dentry+0x1c0/0x488 selinux_d_instantiate+0x20/0x34 security_d_instantiate+0x70/0xbc d_splice_alias+0x4c/0x3c0 ext4_lookup+0x1d8/0x200 [ext4] __lookup_slow+0x12c/0x1e4 walk_component+0x100/0x200 path_lookupat+0x88/0x118 filename_lookup+0x98/0x130 user_path_at_empty+0x48/0x60 vfs_statx+0x84/0x140 vfs_fstatat+0x20/0x30 __se_sys_newfstatat+0x30/0x74 __arm64_sys_newfstatat+0x1c/0x2c el0_svc_common.constprop.0+0x100/0x184 do_el0_svc+0x1c/0x2c el0_svc+0x20/0x34 el0_sync_handler+0x80/0x17c el0_sync+0x13c/0x140 SELinux: Context system_u:object_r:pssp_rsyslog_log_t:s0:c0 is not valid (left unmapped). It was found that within a critical section of spin_lock_irqsave in sidtab_context_to_sid(), convert_context() (hooked by sidtab_convert_params.func) might cause the process to sleep via allocating memory with GFP_KERNEL, which is problematic. As Ondrej pointed out [1], convert_context()/sidtab_convert_params.func has another caller sidtab_convert_tree(), which is okay with GFP_KERNEL. Therefore, fix this problem by adding a gfp_t argument for convert_context()/sidtab_convert_params.func and pass GFP_KERNEL/_ATOMIC properly in individual callers. [PM: wrap long BUG() output lines, tweak subject line]	N/A	More Details
CVE-2022-50698	In the Linux kernel, the following vulnerability has been resolved: ASoC: da7219: Fix an error handling path in da7219_register_dai_clks() If clk_hw_register() fails, the corresponding clk should not be unregistered. To handle errors from loops, clean up partial iterations before doing the goto. So add a clk_hw_unregister(). Then use a while (--i >= 0) loop in the unwind section.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: mrp: introduce active flags to prevent UAF when applicant uninit The caller of del_timer_sync must prevent restarting of the timer, If we have no this synchronization, there is a small probability that the cancellation will not be successful. And syzbot report the fellowing crash: == BUG: KASAN: use-after-free in hlist_add_head include/linux/list.h:929 [inline] BUG: KASAN: use-after-free in enqueue_timer+0x18/0xa4 kernel/time/timer.c:605 Write at addr f9ff000024df6058 by task syz-fuzzer/2256 Pointer tag: [f9], memory tag: [fe] CPU: 1 PID: 2256 Comm: syz-fuzzer Not tainted 6.1.0-rc5-syzkaller-00008- ge01d50cbd6ee #0 Hardware name: linux,dummy-virt (DT) Call trace:		

CVE-2022-50697	dump_backtrace.part.0+0xe0/0xf0 arch/arm64/kernel/stacktrace.c:156 dump_backtrace arch/arm64/kernel/stacktrace.c:162 [inline] show_stack+0x18/0x40 arch/arm64/kernel/stacktrace.c:163 __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x68/0x84 lib/dump_stack.c:106 print_address_description mm/kasan/report.c:284 [inline] print_report+0x1a8/0x4a0 mm/kasan/report.c:395 kasan_report+0x94/0xb4 mm/kasan/report.c:495 __do_kernel_fault+0x164/0x1e0 arch/arm64/mm/fault.c:320 do_bad_area arch/arm64/mm/fault.c:473 [inline] do_tag_check_fault+0x78/0x8c arch/arm64/mm/fault.c:749 do_mem_abort+0x44/0x94 arch/arm64/mm/fault.c:825 el1_abort+0x40/0x60 arch/arm64/kernel/entry-common.c:367 el1h_64_sync_handler+0xd8/0xe4 arch/arm64/kernel/entry-common.c:427 el1h_64_sync+0x64/0x68 arch/arm64/kernel/entry.S:576 hlist_add_head include/linux/list.h:929 [inline] enqueue_timer+0x18/0xa4 kernel/time/timer.c:605 mod_timer+0x14/0x20 kernel/time/timer.c:1161 mrp_periodic_timer_arm net/802/mrp.c:614 [inline] mrp_periodic_timer+0xa0/0xc0 net/802/mrp.c:627 call_timer_fn.constprop.0+0x24/0x80 kernel/time/timer.c:1474 expire_timers+0x98/0xc4 kernel/time/timer.c:1519 To fix it, we can introduce a new active flags to make sure the timer will not restart.	N/A	More Details
CVE-2023-54126	In the Linux kernel, the following vulnerability has been resolved: crypto: safexcel - Cleanup ring IRQ workqueues on load failure A failure loading the safexcel driver results in the following warning on boot, because the IRQ affinity has not been correctly cleaned up. Ensure we clean up the affinity and workqueues on a failure to load the driver. crypto-safexcel: probe of f2800000.crypto failed with error -2 ------------[cut here]------------ WARNING: CPU: 1 PID: 232 at kernel/irq/manage.c:1913 free_irq+0x300/0x340 Modules linked in: hwmon mdio_i2c crypto_safexcel(+) md5 sha256_generic libsha256 authenc libdes omap_rng rng_core nft_masq nft_nat nft_chain_nat nf_nat nft_ct nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 nf_tables libcrc32c nfnetlink fuse autofs4 CPU: 1 PID: 232 Comm: systemd-udevd Tainted: G W 6.1.6-00002-g9d4898824677 #3 Hardware name: MikroTik RB5009 (DT) pstate: 600000c5 (nZCv daIF -PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : free_irq+0x300/0x340 lr : free_irq+0x2e0/0x340 sp : ffff800008fa3890 x29: ffff800008fa3890 x28: 0000000000000000 x27: 0000000000000000 x26: ffff8000008e6dc0 x25: ffff000009034cac x24: ffff000009034d50 x23: 0000000000000000 x22: 000000000000004a x21: ffff0000093e0d80 x20: ffff000009034c00 x19: ffff00000615fc00 x18: 0000000000000000 x17: 0000000000000000 x16: 0000000000000000 x15: 000075f5c1584c5e x14: 0000000000000017 x13: 0000000000000000 x12: 0000000000000040 x11: ffff000000579b60 x10: ffff000000579b62 x9 : ffff800008bbe370 x8 : ffff000000579dd0 x7 : 0000000000000000 x6 : ffff000000579e18 x5 : ffff000000579da8 x4 : ffff800008ca0000 x3 : ffff800008ca0188 x2 : 0000000013033204 x1 : ffff000009034c00 x0 : ffff8000087eadf0 Call trace: free_irq+0x300/0x340 devm_irq_release+0x14/0x20 devres_release_all+0xa0/0x100 device_unbind_cleanup+0x14/0x60 really_probe+0x198/0x2d4 __driver_probe_device+0x74/0xdc driver_probe_device+0x3c/0x110 __driver_attach+0x8c/0x190 bus_for_each_dev+0x6c/0xc0 driver_attach+0x20/0x30 bus_add_driver+0x148/0x1fc driver_register+0x74/0x120 __platform_driver_register+0x24/0x30 safexcel_init+0x48/0x1000 [crypto_safexcel] do_one_initcall+0x4c/0x1b0 do_init_module+0x44/0x1cc load_module+0x1724/0x1be4 __do_sys_finit_module+0xbc/0x110 __arm64_sys_finit_module+0x1c/0x24 invoke_syscall+0x44/0x110 el0_svc_common.constprop.0+0xc0/0xe0 do_el0_svc+0x20/0x80 el0_svc+0x14/0x4c el0t_64_sync_handler+0xb0/0xb4 el0t_64_sync+0x148/0x14c ---[end trace 0000000000000000]---	N/A	More Details
CVE-2023-54125	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Return error for inconsistent extended attributes ntfs_read_ea is called when we want to read extended attributes. There are some sanity checks for the validity of the EAs. However, it fails to return a proper error code for the inconsistent attributes, which might lead to unpredicted memory accesses after return. [138.916927] BUG: KASAN: use-after-free in ntfs_set_ea+0x453/0xbf0 [138.923876] Write of size 4 at addr ffff88800205cfac by task poc/199 [138.931132] [138.933016] CPU: 0 PID: 199 Comm: poc Not tainted 6.2.0-rc1+ #4 [138.938070] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 [138.947327] Call Trace: [138.949557] <TASK> [138.951539] dump_stack_lvl+0x4d/0x67 [138.956834] print_report+0x16f/0x4a6 [138.960798] ? ntfs_set_ea+0x453/0xbf0 [138.964437] ? kasan_complete_mode_report_info+0x7d/0x200 [138.969793] ? ntfs_set_ea+0x453/0xbf0 [138.973523] kasan_report+0xb8/0x140 [138.976740] ? ntfs_set_ea+0x453/0xbf0 [138.980578] __asan_store4+0x76/0xa0 [138.984669] ntfs_set_ea+0x453/0xbf0 [138.988115] ? __pfx_ntfs_set_ea+0x10/0x10 [138.993390] ? kernel_text_address+0xd3/0xe0 [138.998270] ? __kernel_text_address+0x16/0x50 [139.002121] ? unwind_get_return_address+0x3e/0x60 [139.005659] ? __pfx_stack_trace_consume_entry+0x10/0x10 [139.010177] ? arch_stack_walk+0xa2/0x100 [139.013657] ? filter_irq_stacks+0x27/0x80 [139.017018] ntfs_setxattr+0x405/0x440 [139.022151] ? __pfx_ntfs_setxattr+0x10/0x10 [139.026569] ? kvmalloc_node+0x2d/0x120 [139.030329] ? kasan_save_stack+0x41/0x60 [139.033883] ? kasan_save_stack+0x2a/0x60 [139.037338] ? kasan_set_track+0x29/0x40 [139.040163] ? kasan_save_alloc_info+0x1f/0x30 [139.043588] ? __kasan_kmalloc+0x8b/0xa0 [139.047255] ? __kmalloc_node+0x68/0x150 [139.051264] ? kvmalloc_node+0x2d/0x120 [139.055301] ? vmemdup_user+0x2b/0xa0 [139.058584] __vfs_setxattr+0x121/0x170 [139.062617] ? __pfx___vfs_setxattr+0x10/0x10 [139.066282] __vfs_setxattr_noperm+0x97/0x300 [139.070061] __vfs_setxattr_locked+0x145/0x170 [139.073580] vfs_setxattr+0x137/0x2a0 [139.076641] ? __pfx_vfs_setxattr+0x10/0x10 [139.080223] ? __kasan_check_write+0x18/0x20 [139.084234] do_setxattr+0xce/0x150 [139.087768] setxattr+0x126/0x140 [139.091250] ? __pfx_setxattr+0x10/0x10 [139.094948] ? __virt_addr_valid+0xcb/0x140 [139.097838] ? __call_rcu_common.constprop.0+0x1c7/0x330 [139.102688] ? debug_smp_processor_id+0x1b/0x30 [139.105985] ? kasan_quarantine_put+0x5b/0x190 [139.109980] ? putname+0x84/0xa0 [139.113886] ? __kasan_slab_free+0x11e/0x1b0 [139.117961] ? putname+0x84/0xa0 [139.121316] ? preempt_count_sub+0x1c/0xd0 [139.124427] ? __mnt_want_write+0xae/0x100 [139.127836] ? mnt_want_write+0x8f/0x150 [139.130954] path_setxattr+0x164/0x180 [139.133998] ? __pfx_path_setxattr+0x10/0x10 [139.137853] ? __pfx_ksys_pwrite64+0x10/0x10 [139.141299] ? debug_smp_processor_id+0x1b/0x30 [139.145714] ? fpregs_assert_state_consistent+0x6b/0x80 [139.150796] __x64_sys_setxattr+0x71/0x90 [139.155407] do_syscall_64+0x3f/0x90 [139.159035] entry_SYSCALL_64_after_hwframe+0x72/0xdc [139.163843] RIP: 0033:0x7f108cae4469 [139.166481] Code: 00 f3 c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 40 00 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 088 [139.183764] RSP: 002b:00007fff87588388 EFLAGS: 00000286 ORIG_RAX: 00000000000000bc [139.190657] RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 00007f108cae4469 [139.196586] RDX: 00007fff875883b0 RSI: 00007fff875883d1 RDI: 00007fff875883b6 [139.201716] RBP: 00007fff8758c530 R08: 0000000000000001 R09: 00007fff8758c618 [139.207940] R10: 0000000000000006 R11: 0000000000000286 R12: 00000000004004c0 [139.214007] R13: 00007fff8758c610 R14: 0000000000000000 R15 --- truncated---	N/A	More Details
CVE-2023-54124	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to drop all dirty pages during umount() if cp_error is set xfstest generic/361 reports a bug as below: f2fs_bug_on(sbi, sbi->fsync_node_num); kernel BUG at fs/f2fs/super.c:1627! RIP: 0010:f2fs_put_super+0x3a8/0x3b0 Call Trace: generic_shutdown_super+0x8c/0x1b0 kill_block_super+0x2b/0x60 kill_f2fs_super+0x87/0x110 deactivate_locked_super+0x39/0x80 deactivate_super+0x46/0x50 cleanup_mnt+0x109/0x170 __cleanup_mnt+0x16/0x20 task_work_run+0x65/0xa0 exit_to_user_mode_prepare+0x175/0x190 syscall_exit_to_user_mode+0x25/0x50 do_syscall_64+0x4c/0x90 entry_SYSCALL_64_after_hwframe+0x72/0xdc During umount(), if cp_error is set, f2fs_wait_on_all_pages() should not stop waiting all F2FS_WB_CP_DATA pages to be writebacked, otherwise, fsync_node_num can be non-zero after f2fs_wait_on_all_pages() causing this bug. In this case, to avoid deadloop in f2fs_wait_on_all_pages(), it needs to drop all dirty pages rather than redirtying them.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: md/raid10: fix memleak for 'conf->bio_split' In the error path of raid10_run(), 'conf' need be freed, however, 'conf->bio_split' is missed and memory will be leaked. Since there are 3 places to free	N/A	More

54123	'conf', factor out a helper to fix the problem.		Details
CVE-2025-68695	Rejected reason: Not used	N/A	More Details
CVE-2025-66835	TrueConf Client 8.5.2 is vulnerable to DLL hijacking via crafted wfapi.dll allowing local attackers to execute arbitrary code within the user's context.	N/A	More Details
CVE-2025-68694	Rejected reason: Not used	N/A	More Details
CVE-2025-68693	Rejected reason: Not used	N/A	More Details
CVE-2025-68692	Rejected reason: Not used	N/A	More Details
CVE-2025-68691	Rejected reason: Not used	N/A	More Details
CVE-2025-68690	Rejected reason: Not used	N/A	More Details
CVE-2025-68689	Rejected reason: Not used	N/A	More Details
CVE-2025-68688	Rejected reason: Not used	N/A	More Details
CVE-2025-68687	Rejected reason: Not used	N/A	More Details
CVE-2022-50701	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7921s: fix slab-out-of-bounds access in sdio host SDIO may need addtional 511 bytes to align bus operation. If the tailroom of this skb is not big enough, we would access invalid memory region. For low level operation, increase skb size to keep valid memory access in SDIO host. Error message: [69.951] BUG: KASAN: slab-out-of-bounds in sg_copy_buffer+0xe9/0x1a0 [69.951] Read of size 64 at addr ffff88811c9cf000 by task kworker/u16:7/451 [69.951] CPU: 4 PID: 451 Comm: kworker/u16:7 Tainted: G W OE 6.1.0-rc5 #1 [69.951] Workqueue: kvub300c vub300_cmndwork_thread [vub300] [69.951] Call Trace: [69.951] <TASK> [69.952] dump_stack_lvl+0x49/0x63 [69.952] print_report+0x171/0x4a8 [69.952] kasan_report+0xb4/0x130 [69.952] kasan_check_range+0x149/0x1e0 [69.952] memcpy+0x24/0x70 [69.952] sg_copy_buffer+0xe9/0x1a0 [69.952] sg_copy_to_buffer+0x12/0x20 [69.952] __command_write_data.isra.0+0x23c/0xbf0 [vub300] [69.952] vub300_cmndwork_thread+0x17f3/0x58b0 [vub300] [69.952] process_one_work+0x7ee/0x1320 [69.952] worker_thread+0x53c/0x1240 [69.952] kthread+0x2b8/0x370 [69.952] ret_from_fork+0x1f/0x30 [69.952] </TASK> [69.952] Allocated by task 854: [69.952] kasan_save_stack+0x26/0x50 [69.952] kasan_set_track+0x25/0x30 [69.952] kasan_save_alloc_info+0x1b/0x30 [69.952] __kasan_kmalloc+0x87/0xa0 [69.952] __kmalloc_node_track_caller+0x63/0x150 [69.952] kmalloc_reserve+0x31/0xd0 [69.952] __alloc_skb+0xfc/0x2b0 [69.952] __mt76_mcu_msg_alloc+0xbf/0x230 [mt76] [69.952] mt76_mcu_send_and_get_msg+0xab/0x110 [mt76] [69.952] __mt76_mcu_send_firmware.cold+0x94/0x15d [mt76] [69.952] mt76_connac_mcu_send_ram_firmware+0x415/0x54d [mt76_connac_lib] [69.952] mt76_connac2_load_ram.cold+0x118/0x4bc [mt76_connac_lib] [69.952] mt7921_run_firmware.cold+0x2e9/0x405 [mt7921_common] [69.952] mt7921s_mcu_init+0x45/0x80 [mt7921s] [69.953] mt7921_init_work+0xe1/0x2a0 [mt7921_common] [69.953] process_one_work+0x7ee/0x1320 [69.953] worker_thread+0x53c/0x1240 [69.953] kthread+0x2b8/0x370 [69.953] ret_from_fork+0x1f/0x30 [69.953] The buggy address belongs to the object at ffff88811c9ce800 which belongs to the cache kmalloc-2k of size 2048 [69.953] The buggy address is located 0 bytes to the right of 2048-byte region [ffff88811c9ce800, ffff88811c9cf000) [69.953] Memory state around the buggy address: [69.953] ffff88811c9cef00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 [69.953] ffff88811c9cef80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 [69.953] >ffff88811c9cf000: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc [69.953] ^ [69.953] ffff88811c9cf080: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc [69.953] ffff88811c9cf100: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc	N/A	More Details
CVE-2022-50702	In the Linux kernel, the following vulnerability has been resolved: vdpa_sim: fix possible memory leak in vdpasim_net_init() and vdpasim_blk_init() Inject fault while probing module, if device_register() fails in vdpasim_net_init() or vdpasim_blk_init(), but the refcount of kobject is not decreased to 0, the name allocated in dev_set_name() is leaked. Fix this by calling put_device(), so that name can be freed in callback function kobject_cleanup(). (vdpa_sim_net) unreferenced object 0xffff88807eebc370 (size 16): comm "modprobe", pid 3848, jiffies 4362982860 (age 18.153s) hex dump (first 16 bytes): 76 64 70 61 73 69 6d 5f 6e 65 74 00 6b 6b 6b a5 vdpasim_net.kkk. backtrace: [<ffffffff8174f19e>] __kmalloc_node_track_caller+0x4e/0x150 [<ffffffff81731d53>] kstrdup+0x33/0x60 [<ffffffff83a5d421>] kobject_set_name_vargs+0x41/0x110 [<ffffffff82d87aab>] dev_set_name+0xab/0xe0 [<ffffffff82d91a23>] device_add+0xe3/0x1a80 [<ffffffffa0270013>] 0xffffffffa0270013 [<ffffffff81001c27>] do_one_initcall+0x87/0x2e0 [<ffffffff813739cb>] do_init_module+0x1ab/0x640 [<ffffffff81379d20>] load_module+0x5d00/0x77f0 [<ffffffff8137bc40>] __do_sys_finit_module+0x110/0x1b0 [<ffffffff83c4d505>] do_syscall_64+0x35/0x80 [<ffffffff83e0006a>] entry_SYSCALL_64_after_hwframe+0x46/0xb0 (vdpa_sim_blk) unreferenced object 0xffff8881070c1250 (size 16): comm "modprobe", pid 6844, jiffies 4364069319 (age 17.572s) hex dump (first 16 bytes): 76 64 70 61 73 69 6d 5f 62 6c 6b 00 6b 6b 6b a5 vdpasim_blk.kkk. backtrace: [<ffffffff8174f19e>] __kmalloc_node_track_caller+0x4e/0x150 [<ffffffff81731d53>] kstrdup+0x33/0x60 [<ffffffff83a5d421>] kobject_set_name_vargs+0x41/0x110 [<ffffffff82d87aab>] dev_set_name+0xab/0xe0 [<ffffffff82d91a23>] device_add+0xe3/0x1a80 [<ffffffffa0220013>] 0xffffffffa0220013 [<ffffffff81001c27>] do_one_initcall+0x87/0x2e0 [<ffffffff813739cb>] do_init_module+0x1ab/0x640 [<ffffffff81379d20>] load_module+0x5d00/0x77f0 [<ffffffff8137bc40>]	N/A	More Details

	__do_sys_finit_module+0x110/0x1b0 [<ffffffff83c4d505>] do_syscall_64+0x35/0x80 [<ffffffff83e0006a>] entry_SYSCALL_64_after_hwframe+0x46/0xb0		
CVE-2022-50703	In the Linux kernel, the following vulnerability has been resolved: soc: qcom: smsm: Fix refcount leak bugs in qcom_smsm_probe() There are two refcount leak bugs in qcom_smsm_probe(): (1) The 'local_node' is escaped out from for_each_child_of_node() as the break of iteration, we should call of_node_put() for it in error path or when it is not used anymore. (2) The 'node' is escaped out from for_each_available_child_of_node() as the 'goto', we should call of_node_put() for it in goto target.	N/A	More Details
CVE-2022-50704	In the Linux kernel, the following vulnerability has been resolved: USB: gadget: Fix use-after-free during usb config switch In the process of switching USB config from rndis to other config, if the hardware does not support the ->pullup callback, or the hardware encounters a low probability fault, both of them may cause the ->pullup callback to fail, which will then cause a system panic (use after free). The gadget drivers sometimes need to be unloaded regardless of the hardware's behavior. Analysis as follows: == (1) write /config/usb_gadget/g1/UDC "none" gether_disconnect+0x2c/0x1f8 rndis_disable+0x4c/0x74 composite_disconnect+0x74/0xb0 configfs_composite_disconnect+0x60/0x7c usb_gadget_disconnect+0x70/0x124 usb_gadget_unregister_driver+0xc8/0x1d8 gadget_dev_desc_UDC_store+0xec/0x1e4 (2) rm /config/usb_gadget/g1/configs/b.1/f1 rndis_deregister+0x28/0x54 rndis_free+0x44/0x7c usb_put_function+0x14/0x1c config_usb_cfg_unlink+0xc4/0xe0 configfs_unlink+0x124/0x1c8 vfs_unlink+0x114/0x1dc (3) rmdir /config/usb_gadget/g1/functions/rndis.gs4 panic+0x1fc/0x3d0 do_page_fault+0xa8/0x46c do_mem_abort+0x3c/0xac el1_sync_handler+0x40/0x78 0xffffff801138f880 rndis_close+0x28/0x34 eth_stop+0x74/0x110 dev_close_many+0x48/0x194 rollback_registered_many+0x118/0x814 unregister_netdev+0x20/0x30 gether_cleanup+0x1c/0x38 rndis_attr_release+0xc/0x14 kref_put+0x74/0xb8 configfs_rmdir+0x314/0x374 If gadget->ops->pullup() return an error, function rndis_close() will be called, then it will causes a use-after-free problem. ==	N/A	More Details
CVE-2023-53867	In the Linux kernel, the following vulnerability has been resolved: ceph: fix potential use-after-free bug when trimming caps When trimming the caps and just after the 'session->s_cap_lock' is released in ceph_iterate_session_caps() the cap maybe removed by another thread, and when using the stale cap memory in the callbacks it will trigger use-after-free crash. We need to check the existence of the cap just after the 'ci->i_ceph_lock' being acquired. And do nothing if it's already removed.	N/A	More Details
CVE-2022-50711	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: mtk_eth_soc: fix possible memory leak in mtk_probe() If mtk_wed_add_hw() has been called, mtk_wed_exit() needs be called in error path or removing module to free the memory allocated in mtk_wed_add_hw().	N/A	More Details
CVE-2022-50710	In the Linux kernel, the following vulnerability has been resolved: ice: set tx_tstamps when creating new Tx rings via ethtool When the user changes the number of queues via ethtool, the driver allocates new rings. This allocation did not initialize tx_tstamps. This results in the tx_tstamps field being zero (due to kcalloc allocation), and would result in a NULL pointer dereference when attempting a transmit timestamp on the new ring.	N/A	More Details
CVE-2025-50343	An issue was discovered in matio 1.5.28. A heap-based memory corruption can occur in Mat_VarCreateStruct() when the nfields value does not match the actual number of strings in the fields array. This leads to out-of-bounds reads and invalid memory frees during cleanup, potentially causing a segmentation fault or heap corruption.	N/A	More Details
CVE-2025-66823	An HTML Injection vulnerability in TrueConf server 5.5.2.10813 in the conference description field allows an attacker to inject arbitrary HTML in the Create/Edit conference functionality. The payload will be triggered when the victim opens the Conference Info page ([conference url]/info).	N/A	More Details
CVE-2025-69210	FacturaScripts is open-source enterprise resource planning and accounting software. Prior to version 2025.7, a stored cross-site scripting (XSS) vulnerability exists in the product file upload functionality. Authenticated users can upload crafted XML files containing executable JavaScript. These files are later rendered by the application without sufficient sanitization or content-type enforcement, allowing arbitrary JavaScript execution when the file is accessed. Because product files uploaded by regular users are visible to administrative users, this vulnerability can be leveraged to execute malicious JavaScript in an administrator’s browser session. Version 2025.7 fixes the issue.	N/A	More Details
CVE-2022-50709	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: avoid uninit memory read in ath9k_htc_rx_msg() syzbot is reporting uninit value at ath9k_htc_rx_msg() [1], for ioctl(USB_RAW_IOCTL_EP_WRITE) can call ath9k_hif_usb_rx_stream() with pkt_len = 0 but ath9k_hif_usb_rx_stream() uses __dev_alloc_skb(pkt_len + 32, GFP_ATOMIC) based on an assumption that pkt_len is valid. As a result, ath9k_hif_usb_rx_stream() allocates skb with uninitialized memory and ath9k_htc_rx_msg() is reading from uninitialized memory. Since bytes accessed by ath9k_htc_rx_msg() is not known until ath9k_htc_rx_msg() is called, it would be difficult to check minimal valid pkt_len at "if (pkt_len > 2 * MAX_RX_BUF_SIZE) {" line in ath9k_hif_usb_rx_stream(). We have two choices. One is to workaround by adding __GFP_ZERO so that ath9k_htc_rx_msg() sees 0 if pkt_len is invalid. The other is to let ath9k_htc_rx_msg() validate pkt_len before accessing. This patch chose the latter. Note that I'm not sure threshold condition is correct, for I can't find details on possible packet length used by this protocol.	N/A	More Details
CVE-2025-69261	WasmEdge is a WebAssembly runtime. Prior to version 0.16.0-alpha.3, a multiplication in `WasmEdge/include/runtime/instance/memory.h` can wrap, causing `checkAccessBound()` to incorrectly allow the access. This leads to a segmentation fault. Version 0.16.0-alpha.3 contains a patch for the issue.	N/A	More Details
CVE-2025-14986	When frontend.enableExecuteMultiOperation is enabled, the server can apply namespace-scoped validation and feature gates for the embedded StartWorkflowExecutionRequest using its Namespace field rather than the outer, authorized ExecuteMultiOperationRequest.Namespace. This allows a caller authorized for one namespace to bypass that namespace's limits/policies by setting the embedded start request's namespace to a different namespace. The workflow is still created in the outer (authorized) namespace; only validation/gating is performed under the wrong namespace context. This issue affects Temporal: from 1.24.0 through 1.29.1. Fixed in 1.27.4, 1.28.2, 1.29.2.	N/A	More Details
CVE-2025-14987	When system.enableCrossNamespaceCommands is enabled (on by default), the Temporal server permits certain workflow task commands (e.g. StartChildWorkflowExecution, SignalExternalWorkflowExecution, RequestCancelExternalWorkflowExecution) to target a different namespace than the namespace authorized at the gRPC boundary. The frontend authorizes RespondWorkflowTaskCompleted based on the outer request namespace, but the history service later resolves and executes the command using the namespace embedded in command attributes without authorizing the caller for that target namespace. This can allow a worker authorized for one namespace to create, signal, or cancel workflows in another namespace. This issue affects Temporal: through 1.29.1. Fixed in 1.27.4, 1.28.2, 1.29.2.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: HSI: ssi_protocol: fix potential resource leak in ssip_pn_open()		

CVE-2022-50708	ssip_pn_open() claims the HSI client's port with hsi_claim_port(). When hsi_register_port_event() gets some error and returns a negetive value, the HSI client's port should be released with hsi_release_port(). Fix it by calling hsi_release_port() when hsi_register_port_event() fails.	N/A	More Details
CVE-2022-50707	In the Linux kernel, the following vulnerability has been resolved: virtio-crypto: fix memory leak in virtio_crypto_alg_skcipher_close_session() 'vc_ctrl_req' is alloced in virtio_crypto_alg_skcipher_close_session(), and should be freed in the invalid ctrl_status->status error handling case. Otherwise there is a memory leak.	N/A	More Details
CVE-2025-61594	URI is a module providing classes to handle Uniform Resource Identifiers. In versions prior to 0.12.5, 0.13.3, and 1.0.4, a bypass exists for the fix to CVE-2025-27221 that can expose user credentials. When using the `+` operator to combine URIs, sensitive information like passwords from the original URI can be leaked, violating RFC3986 and making applications vulnerable to credential exposure. Versions 0.12.5, 0.13.3, and 1.0.4 fix the issue.	N/A	More Details
CVE-2025-66723	inMusic Brands Engine DJ 4.3.0 suffers from Insecure Permissions due to exposed HTTP service in the Remote Library, which allows attackers to access all files and network paths.	N/A	More Details
CVE-2023-54127	In the Linux kernel, the following vulnerability has been resolved: fs/jfs: prevent double-free in dbUnmount() after failed jfs_remount() Syzkaller reported the following issue: ========== BUG: KASAN: double-free in slab_free mm/slub.c:3787 [inline] BUG: KASAN: double-free in __kmem_cache_free+0x71/0x110 mm/slub.c:3800 Free of addr ffff888086408000 by task syz-executor.4/12750 [...] Call Trace: <TASK> [...] kasan_report_invalid_free+0xac/0xd0 mm/kasan/report.c:482 ____kasan_slab_free+0xfb/0x120 kasan_slab_free include/linux/kasan.h:177 [inline] slab_free_hook mm/slub.c:1781 [inline] slab_free_freelist_hook+0x12e/0x1a0 mm/slub.c:1807 slab_free mm/slub.c:3787 [inline] __kmem_cache_free+0x71/0x110 mm/slub.c:3800 dbUnmount+0xf4/0x110 fs/jfs/jfs_dmap.c:264 jfs_umount+0x248/0x3b0 fs/jfs/jfs_umount.c:87 jfs_put_super+0x86/0x190 fs/jfs/super.c:194 generic_shutdown_super+0x130/0x310 fs/super.c:492 kill_block_super+0x79/0xd0 fs/super.c:1386 deactivate_locked_super+0xa7/0xf0 fs/super.c:332 cleanup_mnt+0x494/0x520 fs/namespace.c:1291 task_work_run+0x243/0x300 kernel/task_work.c:179 resume_user_mode_work include/linux/resume_user_mode.h:49 [inline] exit_to_user_mode_loop+0x124/0x150 kernel/entry/common.c:171 exit_to_user_mode_prepare+0xb2/0x140 kernel/entry/common.c:203 __syscall_exit_to_user_mode_work kernel/entry/common.c:285 [inline] syscall_exit_to_user_mode+0x26/0x60 kernel/entry/common.c:296 do_syscall_64+0x49/0xb0 arch/x86/entry/common.c:86 entry_SYSCALL_64_after_hwframe+0x63/0xcd [...] </TASK> Allocated by task 13352: kasan_save_stack mm/kasan/common.c:45 [inline] kasan_set_track+0x3d/0x60 mm/kasan/common.c:52 ____kasan_kmalloc mm/kasan/common.c:371 [inline] __kasan_kmalloc+0x97/0xb0 mm/kasan/common.c:380 kmalloc include/linux/slab.h:580 [inline] dbMount+0x54/0x980 fs/jfs/jfs_dmap.c:164 jfs_mount+0x1dd/0x830 fs/jfs/jfs_mount.c:121 jfs_fill_super+0x590/0xc50 fs/jfs/super.c:556 mount_bdev+0x26c/0x3a0 fs/super.c:1359 legacy_get_tree+0xea/0x180 fs/fs_context.c:610 vfs_get_tree+0x88/0x270 fs/super.c:1489 do_new_mount+0x289/0xad0 fs/namespace.c:3145 do_mount fs/namespace.c:3488 [inline] __do_sys_mount fs/namespace.c:3697 [inline] __se_sys_mount+0x2d3/0x3c0 fs/namespace.c:3674 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3d/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd Freed by task 13352: kasan_save_stack mm/kasan/common.c:45 [inline] kasan_set_track+0x3d/0x60 mm/kasan/common.c:52 kasan_save_free_info+0x27/0x40 mm/kasan/generic.c:518 ____kasan_slab_free+0xd6/0x120 mm/kasan/common.c:236 kasan_slab_free include/linux/kasan.h:177 [inline] slab_free_hook mm/slub.c:1781 [inline] slab_free_freelist_hook+0x12e/0x1a0 mm/slub.c:1807 slab_free mm/slub.c:3787 [inline] __kmem_cache_free+0x71/0x110 mm/slub.c:3800 dbUnmount+0xf4/0x110 fs/jfs/jfs_dmap.c:264 jfs_mount_rw+0x545/0x740 fs/jfs/jfs_mount.c:247 jfs_remount+0x3db/0x710 fs/jfs/super.c:454 reconfigure_super+0x3bc/0x7b0 fs/super.c:935 vfs_fsconfig_locked fs/fsopen.c:254 [inline] __do_sys_fsconfig fs/fsopen.c:439 [inline] __se_sys_fsconfig+0xad5/0x1060 fs/fsopen.c:314 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3d/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd [...] JFS_SBI(ipbmap->i_sb)->bmap wasn't set to NULL after kfree() in dbUnmount(). Syzkaller uses faultinject to reproduce this KASAN double-free warning. The issue is triggered if either diMount() or dbMount() fail in jfs_remount(), since diUnmount() or dbUnmount() already happened in such a case - they will do double-free on next execution: jfs_umount or jfs_remount. Tested on both upstream and jfs-next by syzkaller.	N/A	More Details
CVE-2022-50706	In the Linux kernel, the following vulnerability has been resolved: net/ieee802154: don't warn zero-sized raw_sendmsg() syzbot is hitting skb_assert_len() warning at __dev_queue_xmit() [1], for PF_IEEE802154 socket's zero-sized raw_sendmsg() request is hitting __dev_queue_xmit() with skb->len == 0. Since PF_IEEE802154 socket's zero-sized raw_sendmsg() request was able to return 0, don't call __dev_queue_xmit() if packet length is 0. ---------- #include <sys/socket.h> #include <netinet/in.h> int main(int argc, char *argv[]) { struct sockaddr_in addr = { .sin_family = AF_INET, .sin_addr.s_addr = htonl(INADDR_LOOPBACK) }; struct iovec iov = { }; struct msghdr hdr = { .msg_name = &addr, .msg_namelen = sizeof(addr), .msg_iov = &iov, .msg_iovlen = 1 }; sendmsg(socket(PF_IEEE802154, SOCK_RAW, 0), &hdr, 0); return 0; } ---------- Note that this might be a sign that commit fd1894224407c484 ("bpf: Don't redirect packets with invalid pkt_len") should be reverted, for skb->len == 0 was acceptable for at least PF_IEEE802154 socket.	N/A	More Details
CVE-2022-50705	In the Linux kernel, the following vulnerability has been resolved: io_uring/rw: defer fsnotify calls to task context We can't call these off the kiocb completion as that might be off soft/hard irq context. Defer the calls to when we process the task_work for this request. That avoids valid complaints like: stack backtrace: CPU: 1 PID: 0 Comm: swapper/1 Not tainted 6.0.0-rc6-syzkaller-00321-g105a36f3694e #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 08/26/2022 Call Trace: <IRQ> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106 print_usage_bug kernel/locking/lockdep.c:3961 [inline] valid_state kernel/locking/lockdep.c:3973 [inline] mark_lock_irq kernel/locking/lockdep.c:4176 [inline] mark_lock.part.0.cold+0x18/0xd8 kernel/locking/lockdep.c:4632 mark_lock kernel/locking/lockdep.c:4596 [inline] mark_usage kernel/locking/lockdep.c:4527 [inline] __lock_acquire+0x11d9/0x56d0 kernel/locking/lockdep.c:5007 lock_acquire kernel/locking/lockdep.c:5666 [inline] lock_acquire+0x1ab/0x570 kernel/locking/lockdep.c:5631 __fs_reclaim_acquire mm/page_alloc.c:4674 [inline] fs_reclaim_acquire+0x115/0x160 mm/page_alloc.c:4688 might_alloc include/linux/sched/mm.h:271 [inline] slab_pre_alloc_hook mm/slab.h:700 [inline] slab_alloc mm/slab.c:3278 [inline] __kmem_cache_alloc_lru mm/slab.c:3471 [inline] kmem_cache_alloc+0x39/0x520 mm/slab.c:3491 fanotify_alloc_fid_event fs/notify/fanotify/fanotify.c:580 [inline] fanotify_alloc_event fs/notify/fanotify/fanotify.c:813 [inline] fanotify_handle_event+0x1130/0x3f40 fs/notify/fanotify/fanotify.c:948 send_to_group fs/notify/fsnotify.c:360 [inline] fsnotify+0xafb/0x1680 fs/notify/fsnotify.c:570 __fsnotify_parent+0x62f/0xa60 fs/notify/fsnotify.c:230 fsnotify_parent include/linux/fsnotify.h:77 [inline] fsnotify_file include/linux/fsnotify.h:99 [inline] fsnotify_access include/linux/fsnotify.h:309 [inline] __io_complete_rw_common+0x485/0x720 io_uring/rw.c:195 io_complete_rw+0x1a/0x1f0 io_uring/rw.c:228 iomap_dio_complete_work fs/iomap/direct-io.c:144 [inline] iomap_dio_bio_end_io+0x438/0x5e0 fs/iomap/direct-io.c:178 bio_endio+0x5f9/0x780 block/bio.c:1564 req_bio_endio block/blk-mq.c:695 [inline] blk_update_request+0x3fc/0x1300 block/blk-mq.c:825 scsi_end_request+0x7a/0x9a0 drivers/scsi/scsi_lib.c:541 scsi_io_completion+0x173/0x1f70 drivers/scsi/scsi_lib.c:971 scsi_complete+0x122/0x3b0 drivers/scsi/scsi_lib.c:1438	N/A	More Details

	blk_complete_reqs+0xad/0xe0 block/blk-mq.c:1022 __do_softirq+0x1d3/0x9c6 kernel/softirq.c:571 invoke_softirq kernel/softirq.c:445 [inline] __irq_exit_rcu+0x123/0x180 kernel/softirq.c:650 irq_exit_rcu+0x5/0x20 kernel/softirq.c:662 common_interrupt+0xa9/0xc0 arch/x86/kernel/irq.c:240		
CVE-2023-54325	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - fix out-of-bounds read When preparing an AER-CTR request, the driver copies the key provided by the user into a data structure that is accessible by the firmware. If the target device is QAT GEN4, the key size is rounded up by 16 since a rounded up size is expected by the device. If the key size is rounded up before the copy, the size used for copying the key might be bigger than the size of the region containing the key, causing an out-of-bounds read. Fix by doing the copy first and then update the keylen. This is to fix the following warning reported by KASAN: [138.150574] BUG: KASAN: global-out-of-bounds in qat_alg_skcipher_init_com.isra.0+0x197/0x250 [intel_qat] [138.150641] Read of size 32 at addr ffffffff88c402c0 by task cryptomgr_test/2340 [138.150651] CPU: 15 PID: 2340 Comm: cryptomgr_test Not tainted 6.2.0-rc1+ #45 [138.150659] Hardware name: Intel Corporation ArcherCity/ArcherCity, BIOS EGSDCRB1.86B.0087.D13.2208261706 08/26/2022 [138.150663] Call Trace: [138.150668] <TASK> [138.150922] kasan_check_range+0x13a/0x1c0 [138.150931] memcpy+0x1f/0x60 [138.150940] qat_alg_skcipher_init_com.isra.0+0x197/0x250 [intel_qat] [138.151006] qat_alg_skcipher_init_sessions+0xc1/0x240 [intel_qat] [138.151073] crypto_skcipher_setkey+0x82/0x160 [138.151085] ? prepare_keybuf+0xa2/0xd0 [138.151095] test_skcipher_vec_cfg+0x2b8/0x800	N/A	More Details
CVE-2023-54323	In the Linux kernel, the following vulnerability has been resolved: cxl/pmem: Fix nvdimm registration races A loop of the form: while true; do modprobe cxl_pci; modprobe -r cxl_pci; done ...fails with the following crash signature: BUG: kernel NULL pointer dereference, address: 0000000000000040 [..] RIP: 0010:cxl_internal_send_cmd+0x5/0xb0 [cxl_core] [..] Call Trace: <TASK> cxl_pmem_ctl+0x121/0x240 [cxl_pmem] nvdimm_get_config_data+0xd6/0x1a0 [libnvdimm] nd_label_data_init+0x135/0x7e0 [libnvdimm] nvdimm_probe+0xd6/0x1c0 [libnvdimm] nvdimm_bus_probe+0x7a/0x1e0 [libnvdimm] really_probe+0xde/0x380 __driver_probe_device+0x78/0x170 driver_probe_device+0x1f/0x90 __device_attach_driver+0x85/0x110 bus_for_each_drv+0x7d/0xc0 __device_attach+0xb4/0x1e0 bus_probe_device+0x9f/0xc0 device_add+0x445/0x9c0 nd_async_device_register+0xe/0x40 [libnvdimm] async_run_entry_fn+0x30/0x130 ...namely that the bottom half of async nvdimm device registration runs after the CXL has already torn down the context that cxl_pmem_ctl() needs. Unlike the ACPI NFIT case that benefits from launching multiple nvdimm device registrations in parallel from those listed in the table, CXL is already marked PROBE_PREFER_ASYNCHRONOUS. So provide for a synchronous registration path to preclude this scenario.	N/A	More Details
CVE-2023-54238	In the Linux kernel, the following vulnerability has been resolved: mlx5: fix skb leak while fifo resync and push During ptp resync operation SKBs were poped from the fifo but were never freed neither by napi_consume nor by dev_kfree_skb_any. Add call to napi_consume_skb to properly free SKBs. Another leak was happening because mlx5e_skb_fifo_has_room() had an error in the check. Comparing free running counters works well unless C promotes the types to something wider than the counter. In this case counters are u16 but the result of the substraction is promouted to int and it causes wrong result (negative value) of the check when producer have already overlapped but consumer haven't yet. Explicit cast to u16 fixes the issue.	N/A	More Details
CVE-2023-54269	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: double free xprt_ctxt while still in use When an RPC request is deferred, the rq_xprt_ctxt pointer is moved out of the svc_rqst into the svc_deferred_req. When the deferred request is revisited, the pointer is copied into the new svc_rqst - and also remains in the svc_deferred_req. In the (rare?) case that the request is deferred a second time, the old svc_deferred_req is reused - it still has all the correct content. However in that case the rq_xprt_ctxt pointer is NOT cleared so that when xpo_release_xprt is called, the ctxt is freed (UDP) or possible added to a free list (RDMA). When the deferred request is revisited for a second time, it will reference this ctxt which may be invalid, and the free the object a second time which is likely to oops. So change svc_defer() to *always* clear rq_xprt_ctxt, and assert that the value is now stored in the svc_deferred_req.	N/A	More Details
CVE-2023-54261	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Add missing gfx11 MQD manager callbacks mqd_stride function was introduced in commit 2f77b9a242a2 ("drm/amdkfd: Update MQD management on multi XCC setup") but not assigned for gfx11. Fixes a NULL dereference in debugfs.	N/A	More Details
CVE-2023-54262	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Don't clone flow post action attributes second time The code already clones post action attributes in mlx5e_clone_flow_attr_for_post_act(). Creating another copy in mlx5e_tc_post_act_add() is a erroneous leftover from original implementation. Instead, assign handle->attribute to post_attr provided by the caller. Note that cloning the attribute second time is not just wasteful but also causes issues like second copy not being properly updated in neigh update code which leads to following use-after-free: Feb 21 09:02:00 c-237-177-40-045 kernel: BUG: KASAN: use-after-free in mlx5_cmd_set_fte+0x200d/0x24c0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_report+0xbb/0x1a0 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_save_stack+0x1e/0x40 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_set_track+0x21/0x30 Feb 21 09:02:00 c-237-177-40-045 kernel: __kasan_kmalloc+0x7a/0x90 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_save_stack+0x1e/0x40 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_set_track+0x21/0x30 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_save_free_info+0x2a/0x40 Feb 21 09:02:00 c-237-177-40-045 kernel: ____kasan_slab_free+0x11a/0x1b0 Feb 21 09:02:00 c-237-177-40-045 kernel: page dumped because: kasan: bad access detected Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_core 0000:08:00.0: mlx5_cmd_out_err:803:(pid 8833): SET_FLOW_TABLE_ENTRY(0x936) op_mod(0x0) failed, status bad resource state(0x9), syndrome (0xf2ff71), err(-22) Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_core 0000:08:00.0 enp8s0f0: Failed to add post action rule Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_core 0000:08:00.0: mlx5e_tc_encap_flows_add:190:(pid 8833): Failed to update flow post acts, -22 Feb 21 09:02:00 c-237-177-40-045 kernel: Call Trace: Feb 21 09:02:00 c-237-177-40-045 kernel: <TASK> Feb 21 09:02:00 c-237-177-40-045 kernel: dump_stack_lvl+0x57/0x7d Feb 21 09:02:00 c-237-177-40-045 kernel: print_report+0x170/0x471 Feb 21 09:02:00 c-237-177-40-045 kernel: ? mlx5_cmd_set_fte+0x200d/0x24c0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_report+0xbb/0x1a0 Feb 21 09:02:00 c-237-177-40-045 kernel: ? mlx5_cmd_set_fte+0x200d/0x24c0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_cmd_set_fte+0x200d/0x24c0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: ? __module_address.part.0+0x62/0x200 Feb 21 09:02:00 c-237-177-40-045 kernel: ? mlx5_cmd_stub_create_flow_table+0xd0/0xd0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: ? __raw_spin_lock_init+0x3b/0x110 Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_cmd_create_fte+0x80/0xb0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: add_rule_fg+0xe80/0x19c0 [mlx5_core] -- Feb 21 09:02:00 c-237-177-40-045 kernel: Allocated by task 13476: Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_save_stack+0x1e/0x40 Feb 21 09:02:00 c-237-177-40-045 kernel: kasan_set_track+0x21/0x30 Feb 21 09:02:00 c-237-177-40-045 kernel: __kasan_kmalloc+0x7a/0x90 Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5_packet_reformat_alloc+0x7b/0x230 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_tc_tun_create_header_ipv4+0x977/0xf10 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_attach_encap+0x15b4/0x1e10 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: post_process_attr+0x305/0xa30 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_tc_add_fdb_flow+0x4c0/0xcf0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: __mlx5e_add_fdb_flow+0x7cf/0xe90 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_configure_flower+0xcaa/0x4b90 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_rep_setup_tc_cls_flower+0x99/0x1b0 [mlx5_core] Feb 21 09:02:00 c-237-177-40-045 kernel: mlx5e_rep_setup_tc_cb+0x133/0x1e0 [mlx5_core] -- Feb 21 09:02:00 c-237-177-40-045 kernel: Freed by task 8833: Feb 21 09:02:00	N/A	More Details

	c-237-177-40-045 kernel: kasan_save_s ---truncated---		
CVE-2023-54263	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/kms/nv50-: init hpd_irq_lock for PIOR DP Fixes OOPS on boards with ANX9805 DP encoders.	N/A	More Details
CVE-2023-54264	In the Linux kernel, the following vulnerability has been resolved: fs/sysv: Null check to prevent null-ptr-deref bug sb_getblk(inode->i_sb, parent) return a null ptr and taking lock on that leads to the null-ptr-deref bug.	N/A	More Details
CVE-2023-54265	In the Linux kernel, the following vulnerability has been resolved: ipv6: Fix an uninit variable access bug in __ip6_make_skb() Syzbot reported a bug as following: === BUG: KMSAN: uninit-value in arch_atomic64_inc arch/x86/include/asm/atomic64_64.h:88 [inline] BUG: KMSAN: uninit-value in arch_atomic_long_inc include/linux/atomic/atomic-long.h:161 [inline] BUG: KMSAN: uninit-value in atomic_long_inc include/linux/atomic/atomic-instrumented.h:1429 [inline] BUG: KMSAN: uninit-value in __ip6_make_skb+0x2f37/0x30f0 net/ipv6/ip6_output.c:1956 arch_atomic64_inc arch/x86/include/asm/atomic64_64.h:88 [inline] arch_atomic_long_inc include/linux/atomic/atomic-long.h:161 [inline] atomic_long_inc include/linux/atomic/atomic-instrumented.h:1429 [inline] __ip6_make_skb+0x2f37/0x30f0 net/ipv6/ip6_output.c:1956 ip6_finish_skb include/net/ipv6.h:1122 [inline] ip6_push_pending_frames+0x10e/0x550 net/ipv6/ip6_output.c:1987 rawv6_push_pending_frames+0xb12/0xb90 net/ipv6/raw.c:579 rawv6_sendmsg+0x297e/0x2e60 net/ipv6/raw.c:922 inet_sendmsg+0x101/0x180 net/ipv4/af_inet.c:827 sock_sendmsg_nosec net/socket.c:714 [inline] sock_sendmsg net/socket.c:734 [inline] ____sys_sendmsg+0xa8e/0xe70 net/socket.c:2476 ___sys_sendmsg+0x2a1/0x3f0 net/socket.c:2530 __sys_sendmsg net/socket.c:2559 [inline] __do_sys_sendmsg net/socket.c:2568 [inline] __se_sys_sendmsg net/socket.c:2566 [inline] __x64_sys_sendmsg+0x367/0x540 net/socket.c:2566 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3d/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd Uninit was created at: slab_post_alloc_hook mm/slab.h:766 [inline] slab_alloc_node mm/slub.c:3452 [inline] __kmem_cache_alloc_node+0x71f/0xce0 mm/slub.c:3491 __do_kmalloc_node mm/slab_common.c:967 [inline] __kmalloc_node_track_caller+0x114/0x3b0 mm/slab_common.c:988 kmalloc_reserve net/core/skbuff.c:492 [inline] __alloc_skb+0x3af/0x8f0 net/core/skbuff.c:565 alloc_skb include/linux/skbuff.h:1270 [inline] __ip6_append_data+0x51c1/0x6bb0 net/ipv6/ip6_output.c:1684 ip6_append_data+0x411/0x580 net/ipv6/ip6_output.c:1854 rawv6_sendmsg+0x2882/0x2e60 net/ipv6/raw.c:915 inet_sendmsg+0x101/0x180 net/ipv4/af_inet.c:827 sock_sendmsg_nosec net/socket.c:714 [inline] sock_sendmsg net/socket.c:734 [inline] ____sys_sendmsg+0xa8e/0xe70 net/socket.c:2476 ___sys_sendmsg+0x2a1/0x3f0 net/socket.c:2530 __sys_sendmsg net/socket.c:2559 [inline] __do_sys_sendmsg net/socket.c:2568 [inline] __se_sys_sendmsg net/socket.c:2566 [inline] __x64_sys_sendmsg+0x367/0x540 net/socket.c:2566 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3d/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd It is because icmp6hdr does not in skb linear region under the scenario of SOCK_RAW socket. Access icmp6_hdr(skb)->icmp6_type directly will trigger the uninit variable access bug. Use a local variable icmp6_type to carry the correct value in different scenarios.	N/A	More Details
CVE-2023-54266	In the Linux kernel, the following vulnerability has been resolved: media: dvb-usb: m920x: Fix a potential memory leak in m920x_i2c_xfer() 'read' is freed when it is known to be NULL, but not when a read error occurs. Revert the logic to avoid a small leak, should a m920x_read() call fail.	N/A	More Details
CVE-2023-54267	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: Rework lppaca_shared_proc() to avoid DEBUG_PREEMPT lppaca_shared_proc() takes a pointer to the lppaca which is typically accessed through get_lppaca(). With DEBUG_PREEMPT enabled, this leads to checking if preemption is enabled, for example: BUG: using smp_processor_id() in preemptible [00000000] code: grep/10693 caller is lparcfg_data+0x408/0x19a0 CPU: 4 PID: 10693 Comm: grep Not tainted 6.5.0-rc3 #2 Call Trace: dump_stack_lvl+0x154/0x200 (unreliable) check_preemption_disabled+0x214/0x220 lparcfg_data+0x408/0x19a0 ... This isn't actually a problem however, as it does not matter which lppaca is accessed, the shared proc state will be the same. vcpudispatch_stats_procfs_init() already works around this by disabling preemption, but the lparcfg code does not, erroring any time /proc/powerpc/lparcfg is accessed with DEBUG_PREEMPT enabled. Instead of disabling preemption on the caller side, rework lppaca_shared_proc() to not take a pointer and instead directly access the lppaca, bypassing any potential preemption checks. [mpe: Rework to avoid needing a definition in paca.h and lppaca.h]	N/A	More Details
CVE-2023-54268	In the Linux kernel, the following vulnerability has been resolved: debugobjects: Don't wake up kswapd from fill_pool() syzbot is reporting a lockdep warning in fill_pool() because the allocation from debugobjects is using GFP_ATOMIC, which is (__GFP_HIGH | __GFP_KSWAPD_RECLAIM) and therefore tries to wake up kswapd, which acquires kswapd_wait::lock. Since fill_pool() might be called with arbitrary locks held, fill_pool() should not assume that acquiring kswapd_wait::lock is safe. Use __GFP_HIGH instead and remove __GFP_NORETRY as it is pointless for !__GFP_DIRECT_RECLAIM allocation.	N/A	More Details
CVE-2023-54270	In the Linux kernel, the following vulnerability has been resolved: media: usb: siano: Fix use after free bugs caused by do_submit_urb There are UAF bugs caused by do_submit_urb(). One of the KASan reports is shown below: [36.403605] BUG: KASAN: use-after-free in worker_thread+0x4a2/0x890 [36.406105] Read of size 8 at addr ffff8880059600e8 by task kworker/0:2/49 [36.408316] [36.408867] CPU: 0 PID: 49 Comm: kworker/0:2 Not tainted 6.2.0-rc3-15798-g5a41237ad1d4-dir8 [36.411696] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g15584 [36.416157] Workqueue: 0x0 (events) [36.417654] Call Trace: [36.418546] <TASK> [36.419320] dump_stack_lvl+0x96/0xd0 [36.420522] print_address_description+0x75/0x350 [36.421992] print_report+0x11b/0x250 [36.423174] ? _raw_spin_lock_irqsave+0x87/0xd0 [36.424806] ? __virt_addr_valid+0xcf/0x170 [36.426069] ? worker_thread+0x4a2/0x890 [36.427355] kasan_report+0x131/0x160 [36.428556] ? worker_thread+0x4a2/0x890 [36.430053] worker_thread+0x4a2/0x890 [36.431297] ? worker_clr_flags+0x90/0x90 [36.432479] kthread+0x166/0x190 [36.433493] ? kthread_blkcg+0x50/0x50 [36.434669] ret_from_fork+0x22/0x30 [36.435923] </TASK> [36.436684] [36.437215] Allocated by task 24: [36.438289] kasan_set_track+0x50/0x80 [36.439436] __kasan_kmalloc+0x89/0xa0 [36.440566] smsusb_probe+0x374/0xc90 [36.441920] usb_probe_interface+0x2d1/0x4c0 [36.443253] really_probe+0x1d5/0x580 [36.444539] __driver_probe_device+0xe3/0x130 [36.446085] driver_probe_device+0x49/0x220 [36.447423] __device_attach_driver+0x19e/0x1b0 [36.448931] bus_for_each_drv+0xcb/0x110 [36.450217] __device_attach+0x132/0x1f0 [36.451470] bus_probe_device+0x59/0xf0 [36.452563] device_add+0x4ec/0x7b0 [36.453830] usb_set_configuration+0xc63/0xe10 [36.455230] usb_generic_driver_probe+0x3b/0x80 [36.456166] printk: console [ttyGS0] disabled [36.456569] usb_probe_device+0x90/0x110 [36.459523] really_probe+0x1d5/0x580 [36.461027] __driver_probe_device+0xe3/0x130 [36.462465] driver_probe_device+0x49/0x220 [36.463847] __device_attach_driver+0x19e/0x1b0 [36.465229] bus_for_each_drv+0xcb/0x110 [36.466466] __device_attach+0x132/0x1f0 [36.467799] bus_probe_device+0x59/0xf0 [36.469010] device_add+0x4ec/0x7b0 [36.470125] usb_new_device+0x863/0xa00 [36.471374] hub_event+0x18c7/0x2220 [36.472746] process_one_work+0x34c/0x5b0 [36.474041] worker_thread+0x4b7/0x890 [36.475216] kthread+0x166/0x190 [36.476267] ret_from_fork+0x22/0x30 [36.477447] [36.478160] Freed by task 24: [36.479239] kasan_set_track+0x50/0x80 [36.480512] kasan_save_free_info+0x2b/0x40 [36.481808] ____kasan_slab_free+0x122/0x1a0 [36.483173] __kmem_cache_free+0xc4/0x200 [36.484563] smsusb_term_device+0xcd/0xf0 [36.485896] smsusb_probe+0xc85/0xc90 [36.486976] usb_probe_interface+0x2d1/0x4c0 [36.488303] really_probe+0x1d5/0x580 [36.489498]	N/A	More Details

	__driver_probe_device+0xe3/0x130 [36.491140] driver_probe_device+0x49/0x220 [36.492475] __device_attach_driver+0x19e/0x1b0 [36.493988] bus_for_each_drv+0xcb/0x110 [36.495171] __device_attach+0x132/0x1f0 [36.496617] bus_probe_device+0x59/0xf0 [36.497875] device_add+0x4ec/0x7b0 [36.498972] usb_set_configuration+0xc63/0xe10 [36.500264] usb_generic_driver_probe+0x3b/0x80 [36.501740] usb_probe_device+0x90/0x110 [36.503084] really_probe+0x1d5/0x580 [36.504241] __driver_probe_device+0xe3/0x130 [36.505548] driver_probe_device+0x49/0x220 [36.506766] __device_attach_driver+0x19e/0x1b0 [36.508368] bus_for_each_drv+0xcb/0x110 [36.509646] __device_attach+0x132/0x1f0 [36.510911] bus_probe_device+0x59/0xf0 [36.512103] device_add+0x4ec/0x7b0 [36.513215] usb_new_device+0x863/0xa00 [36.514736] hub_event+0x18c7/0x2220 [36.516130] process_one_work+ ---truncated---		
CVE-2023-54259	In the Linux kernel, the following vulnerability has been resolved: soundwire: bus: Fix unbalanced pm_runtime_put() causing usage count underflow This reverts commit 443a98e649b4 ("soundwire: bus: use pm_runtime_resume_and_get()") Change calls to pm_runtime_resume_and_get() back to pm_runtime_get_sync(). This fixes a usage count underrun caused by doing a pm_runtime_put() even though pm_runtime_resume_and_get() returned an error. The three affected functions ignore -EACCES error from trying to get pm_runtime, and carry on, including a put at the end of the function. But pm_runtime_resume_and_get() does not increment the usage count if it returns an error. So in the -EACCES case you must not call pm_runtime_put(). The documentation for pm_runtime_get_sync() says: "Consider using pm_runtime_resume_and_get() ... as this is likely to result in cleaner code." In this case I don't think it results in cleaner code because the pm_runtime_put() at the end of the function would have to be conditional on the return value from pm_runtime_resume_and_get() at the top of the function. pm_runtime_get_sync() doesn't have this problem because it always increments the count, so always needs a put. The code can just flow through and do the pm_runtime_put() unconditionally.	N/A	More Details
CVE-2023-54271	In the Linux kernel, the following vulnerability has been resolved: blk-cgroup: Fix NULL deref caused by blkg_policy_data being installed before init blk-iocost sometimes causes the following crash: BUG: kernel NULL pointer dereference, address: 00000000000000e0 ... RIP: 0010:_raw_spin_lock+0x17/0x30 Code: be 01 02 00 00 e8 79 38 39 ff 31 d2 89 d0 5d c3 0f 1f 00 0f 1f 44 00 00 55 48 89 e5 65 ff 05 48 d0 34 7e b9 01 00 00 00 31 c0 <f0> 0f b1 0f 75 02 5d c3 89 c6 e8 ea 04 00 00 5d c3 0f 1f 84 00 00 RSP: 0018:ffffc900023b3d40 EFLAGS: 00010046 RAX: 0000000000000000 RBX: 00000000000000e0 RCX: 0000000000000001 RDX: ffffc900023b3d20 RSI: ffffc900023b3cf0 RDI: 00000000000000e0 RBP: ffffc900023b3d40 R08: ffffc900023b3c10 R09: 0000000000000003 R10: 0000000000000064 R11: 000000000000000a R12: ffff888102337000 R13: fffffffffffffff2 R14: ffff88810af408c8 R15: ffff8881070c3600 FS: 00007faaaf364fc0(0000) GS:ffff88842fdc0000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00000000000000e0 CR3: 00000001097b1000 CR4: 0000000000350ea0 Call Trace: <TASK> ioc_weight_write+0x13d/0x410 cgroup_file_write+0x7a/0x130 kernfs_fop_write_iter+0xf5/0x170 vfs_write+0x298/0x370 ksys_write+0x5f/0xb0 __x64_sys_write+0x1b/0x20 do_syscall_64+0x3d/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0 This happens because iocg->ioc is NULL. The field is initialized by ioc_pd_init() and never cleared. The NULL deref is caused by blkcg_activate_policy() installing blkg_policy_data before initializing it. blkcg_activate_policy() was doing the following: 1. Allocate pd's for all existing blkg's and install them in blkg->pd[]. 2. Initialize all pd's. 3. Online all pd's. blkcg_activate_policy() only grabs the queue_lock and may release and re-acquire the lock as allocation may need to sleep. ioc_weight_write() grabs blkcg->lock and iterates all its blkg's. The two can race and if ioc_weight_write() runs during #1 or between #1 and #2, it can encounter a pd which is not initialized yet, leading to crash. The crash can be reproduced with the following script: #!/bin/bash echo +io > /sys/fs/cgroup/cgroup.subtree_control systemd-run --unit touch-sda --scope dd if=/dev/sda of=/dev/null bs=1M count=1 iflag=direct echo 100 > /sys/fs/cgroup/system.slice/io.weight bash -c "echo '8:0 enable=1' > /sys/fs/cgroup/io.cost.qos" & sleep .2 echo 100 > /sys/fs/cgroup/system.slice/io.weight with the following patch applied: > diff --git a/block/blk-cgroup.c b/block/blk-cgroup.c > index fc49be622e05..38d671d5e10c 100644 > --- a/block/blk-cgroup.c > +++ b/block/blk-cgroup.c > @@ -1553,6 +1553,12 @@ int blkcg_activate_policy(struct gendisk *disk, const struct blkcg_policy *pol) > pd->online = false; > } > > + if (system_state == SYSTEM_RUNNING) { > + spin_unlock_irq(&q->queue_lock); > + ssleep(1); > + spin_lock_irq(&q->queue_lock); > + } > + > /* all allocated, init in the same order */ > if (pol->pd_init_fn) > list_for_each_entry_reverse(blkg, &q->blkg_list, q_node) I don't see a reason why all pd's should be allocated, initialized and onlined together. The only ordering requirement is that parent blkgs to be initialized and onlined before children, which is guaranteed from the walking order. Let's fix the bug by allocating, initializing and onlining pd for each blkg and holding blkcg->lock over initialization and onlining. This ensures that an installed blkg is always fully initialized and onlined removing the the race window.	N/A	More Details
CVE-2023-54272	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix a possible null-pointer dereference in ni_clear() In a previous commit c1006bd13146, ni->mi.mrec in ni_write_inode() could be NULL, and thus a NULL check is added for this variable. However, in the same call stack, ni->mi.mrec can be also dereferenced in ni_clear(): ntfs_evict_inode(inode) ni_write_inode(inode, ...) ni = ntfs_i(inode); is_rec_inuse(ni->mi.mrec) -> Add a NULL check by previous commit ni_clear(ntfs_i(inode)) is_rec_inuse(ni->mi.mrec) -> No check Thus, a possible null-pointer dereference may exist in ni_clear(). To fix it, a NULL check is added in this function.	N/A	More Details
CVE-2023-54273	In the Linux kernel, the following vulnerability has been resolved: xfrm: Fix leak of dev tracker At the stage of direction checks, the netdev reference tracker is already initialized, but released with wrong *_put() call.	N/A	More Details
CVE-2023-54274	In the Linux kernel, the following vulnerability has been resolved: RDMA/srpt: Add a check for valid 'mad_agent' pointer When unregistering MAD agent, srpt module has a non-null check for 'mad_agent' pointer before invoking ib_unregister_mad_agent(). This check can pass if 'mad_agent' variable holds an error value. The 'mad_agent' can have an error value for a short window when srpt_add_one() and srpt_remove_one() is executed simultaneously. In srpt module, added a valid pointer check for 'sport->mad_agent' before unregistering MAD agent. This issue can hit when RoCE driver unregisters ib_device Stack Trace: ---------- BUG: kernel NULL pointer dereference, address: 000000000000004d PGD 145003067 P4D 145003067 PUD 2324fe067 PMD 0 Oops: 0002 [#1] PREEMPT SMP NOPTI CPU: 10 PID: 4459 Comm: kworker/u80:0 Kdump: loaded Tainted: P Hardware name: Dell Inc. PowerEdge R640/06NR82, BIOS 2.5.4 01/13/2020 Workqueue: bnxt_re bnxt_re_task [bnxt_re] RIP: 0010:_raw_spin_lock_irqsave+0x19/0x40 Call Trace: ib_unregister_mad_agent+0x46/0x2f0 [ib_core] IPv6: ADDRCONF(NETDEV_CHANGE): bond0: link becomes ready ? __schedule+0x20b/0x560 srpt_unregister_mad_agent+0x93/0xd0 [ib_srpt] srpt_remove_one+0x20/0x150 [ib_srpt] remove_client_context+0x88/0xd0 [ib_core] bond0: (slave p2p1): link status definitely up, 100000 Mbps full duplex disable_device+0x8a/0x160 [ib_core] bond0: active interface up! ? kernfs_name_hash+0x12/0x80 (NULL device *): Bonding Info Received: rdev: 000000006c0b8247 __ib_unregister_device+0x42/0xb0 [ib_core] (NULL device *): Master: mode: 4 num_slaves:2 ib_unregister_device+0x22/0x30 [ib_core] (NULL device *): Slave: id: 105069936 name:p2p1 link:0 state:0 bnxt_re_stopqps_and_ib_uninit+0x83/0x90 [bnxt_re] bnxt_re_alloc_lag+0x12e/0x4e0 [bnxt_re]	N/A	More Details
CVE-2023-54275	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: Fix memory leak in ath11k_peer_rx_frag_setup crypto_alloc_shash() allocates resources, which should be released by crypto_free_shash(). When ath11k_peer_find() fails, there has memory leak. Add missing crypto_free_shash() to fix this.	N/A	More Details

CVE-2023-54276	In the Linux kernel, the following vulnerability has been resolved: nfsd: move init of percpu reply_cache_stats counters back to nfsd_init_net Commit f5f9d4a314da ("nfsd: move reply cache initialization into nfsd startup") moved the initialization of the reply cache into nfsd startup, but didn't account for the stats counters, which can be accessed before nfsd is ever started. The result can be a NULL pointer dereference when someone accesses /proc/fs/nfsd/reply_cache_stats while nfsd is still shut down. This is a regression and a user-triggerable oops in the right situation: - non-x86_64 arch - /proc/fs/nfsd is mounted in the namespace - nfsd is not started in the namespace - unprivileged user calls "cat /proc/fs/nfsd/reply_cache_stats" Although this is easy to trigger on some arches (like aarch64), on x86_64, calling this_cpu_ptr(NULL) evidently returns a pointer to the fixed_percpu_data. That struct looks just enough like a newly initialized percpu var to allow nfsd_reply_cache_stats_show to access it without Oopsing. Move the initialization of the per-net+per-cpu reply-cache counters back into nfsd_init_net, while leaving the rest of the reply cache allocations to be done at nfsd startup time. Kudos to Eirik who did most of the legwork to track this down.	N/A	More Details
CVE-2023-54277	In the Linux kernel, the following vulnerability has been resolved: fbdev: udlfb: Fix endpoint check The syzbot fuzzer detected a problem in the udlfb driver, caused by an endpoint not having the expected type: usb 1-1: Read EDID byte 0 failed: -71 usb 1-1: Unable to get valid EDID from device/display ------------[cut here]------------ usb 1-1: BOGUS urb xfer, pipe 3 != type 1 WARNING: CPU: 0 PID: 9 at drivers/usb/core/urb.c:504 usb_submit_urb+0xed6/0x1880 drivers/usb/core/urb.c:504 Modules linked in: CPU: 0 PID: 9 Comm: kworker/0:1 Not tainted 6.4.0-rc1-syzkaller-00016-ga4422ff22142 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 04/28/2023 Workqueue: usb_hub_wq hub_event RIP: 0010:usb_submit_urb+0xed6/0x1880 drivers/usb/core/urb.c:504 ... Call Trace: <TASK> dlfb_submit_urb+0x92/0x180 drivers/video/fbdev/udlfb.c:1980 dlfb_set_video_mode+0x21f0/0x2950 drivers/video/fbdev/udlfb.c:315 dlfb_ops_set_par+0x2a7/0x8d0 drivers/video/fbdev/udlfb.c:1111 dlfb_usb_probe+0x149a/0x2710 drivers/video/fbdev/udlfb.c:1743 The current approach for this issue failed to catch the problem because it only checks for the existence of a bulk-OUT endpoint; it doesn't check whether this endpoint is the one that the driver will actually use. We can fix the problem by instead checking that the endpoint used by the driver does exist and is bulk-OUT.	N/A	More Details
CVE-2023-54278	In the Linux kernel, the following vulnerability has been resolved: s390/vmem: split pages when debug pagealloc is enabled Since commit bb1520d581a3 ("s390/mm: start kernel with DAT enabled") the kernel crashes early during boot when debug pagealloc is enabled: mem auto-init: stack:off, heap alloc:off, heap free:off addressing exception: 0005 ilc:2 [#1] SMP DEBUG_PAGEALLOC Modules linked in: CPU: 0 PID: 0 Comm: swapper Not tainted 6.5.0-rc3-09759-gc5666c912155 #630 [..] Krnl Code: 0000000001325f6: ec5600248064 cgrj %r5,%r6,8,000000000013263e 00000000001325fc: eb880002000c srlg %r8,%r8,2 #0000000000132602: b2210051 ipte %r5,%r1,%r0,0 >0000000000132606: b90400d1 lgr %r13,%r1 000000000013260a: 41605008 la %r6,8(%r5) 000000000013260e: a7db1000 aghi %r13,4096 0000000000132612: b221006d ipte %r6,%r13,%r0,0 0000000000132616: e3d0d0000171 lay %r13,4096(%r13) Call Trace: __kernel_map_pages+0x14e/0x320 __free_pages_ok+0x23a/0x5a8) free_low_memory_core_early+0x214/0x2c8 memblock_free_all+0x28/0x58 mem_init+0xb6/0x228 mm_core_init+0xb6/0x3b0 start_kernel+0x1d2/0x5a8 startup_continue+0x36/0x40 Kernel panic - not syncing: Fatal exception: panic_on_oops This is caused by using large mappings on machines with EDAT1/EDAT2. Add the code to split the mappings into 4k pages if debug pagealloc is enabled by CONFIG_DEBUG_PAGEALLOC_ENABLE_DEFAULT or the debug_pagealloc kernel command line option.	N/A	More Details
CVE-2023-54260	In the Linux kernel, the following vulnerability has been resolved: cifs: Fix lost destroy smbd connection when MR allocate failed If the MR allocate failed, the smb direct connection info is NULL, then smbd_destroy() will directly return, then the connection info will be leaked. Let's set the smb direct connection info to the server before call smbd_destroy().	N/A	More Details
CVE-2023-54258	In the Linux kernel, the following vulnerability has been resolved: cifs: fix potential oops in cifs_oplock_break With deferred close we can have closes that race with lease breaks, and so with the current checks for whether to send the lease response, oplock_response(), this can mean that an unmount (kill_sb) can occur just before we were checking if the tcon->ses is valid. See below: [Fri Aug 4 04:12:50 2023] RIP: 0010:cifs_oplock_break+0x1f7/0x5b0 [cifs] [Fri Aug 4 04:12:50 2023] Code: 7d a8 48 8b 7d c0 c0 e9 02 48 89 45 b8 41 89 cf e8 3e f5 ff ff 4c 89 f7 41 83 e7 01 e8 82 b3 03 f2 49 8b 45 50 48 85 c0 74 5e <48> 83 78 60 00 74 57 45 84 ff 75 52 48 8b 43 98 48 83 eb 68 48 39 [Fri Aug 4 04:12:50 2023] RSP: 0018:ffffb30607ddbdf8 EFLAGS: 00010206 [Fri Aug 4 04:12:50 2023] RAX: 632d223d32612022 RBX: ffff97136944b1e0 RCX: 0000000080100009 [Fri Aug 4 04:12:50 2023] RDX: 0000000000000001 RSI: 0000000080100009 RDI: ffff97136944b188 [Fri Aug 4 04:12:50 2023] RBP: ffffb30607ddbe58 R08: 0000000000000001 R09: ffffffffc08e0900 [Fri Aug 4 04:12:50 2023] R10: 0000000000000001 R11: 000000000000000f R12: ffff97136944b138 [Fri Aug 4 04:12:50 2023] R13: ffff97149147c000 R14: ffff97136944b188 R15: 0000000000000000 [Fri Aug 4 04:12:50 2023] FS: 0000000000000000(0000) GS:ffff9714f7c00000(0000) knlGS:0000000000000000 [Fri Aug 4 04:12:50 2023] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [Fri Aug 4 04:12:50 2023] CR2: 00007fd8de9c7590 CR3: 000000011228e000 CR4: 0000000000350ef0 [Fri Aug 4 04:12:50 2023] Call Trace: [Fri Aug 4 04:12:50 2023] <TASK> [Fri Aug 4 04:12:50 2023] process_one_work+0x225/0x3d0 [Fri Aug 4 04:12:50 2023] worker_thread+0x4d/0x3e0 [Fri Aug 4 04:12:50 2023] ? process_one_work+0x3d0/0x3d0 [Fri Aug 4 04:12:50 2023] kthread+0x12a/0x150 [Fri Aug 4 04:12:50 2023] ? set_kthread_struct+0x50/0x50 [Fri Aug 4 04:12:50 2023] ret_from_fork+0x22/0x30 [Fri Aug 4 04:12:50 2023] </TASK> To fix this change the ordering of the checks before sending the oplock_response to first check if the openFileList is empty.	N/A	More Details
CVE-2023-54322	In the Linux kernel, the following vulnerability has been resolved: arm64: set __exception_irq_entry with __irq_entry as a default filter_irq_stacks() is supposed to cut entries which are related irq entries from its call stack. And in_irqentry_text() which is called by filter_irq_stacks() uses __irqentry_text_start/end symbol to find irq entries in callstack. But it doesn't work correctly as without "CONFIG_FUNCTION_GRAPH_TRACER", arm64 kernel doesn't include gic_handle_irq which is entry point of arm64 irq between __irqentry_text_start and __irqentry_text_end as we discussed in below link. https://lore.kernel.org/all/CACT4Y+aReMGLYua2rCLHgFpS9io5cZC04Q8GLs-uNmrn1ezxYQ@mail.gmail.com/#t This problem can makes unintentional deep call stack entries especially in KASAN enabled situation as below. [2479.383395]I[0:launcher-loader: 1719] Stack depot reached limit capacity [2479.383538]I[0:launcher-loader: 1719] WARNING: CPU: 0 PID: 1719 at lib/stackdepot.c:129 __stack_depot_save+0x464/0x46c [2479.385693]I[0:launcher-loader: 1719] pstate: 624000c5 (nZCv daIF +PAN -UAO +TCO -DIT -SSBS BTYPE=--) [2479.385724]I[0:launcher-loader: 1719] pc : __stack_depot_save+0x464/0x46c [2479.385751]I[0:launcher-loader: 1719] lr : __stack_depot_save+0x460/0x46c [2479.385774]I[0:launcher-loader: 1719] sp : ffffffc0080073c0 [2479.385793]I[0:launcher-loader: 1719] x29: ffffffc0080073e0 x28: ffffffd00b78a000 x27: 0000000000000000 [2479.385839]I[0:launcher-loader: 1719] x26: 000000000004d1dd x25: ffffff891474f000 x24: 00000000ca64d1dd [2479.385882]I[0:launcher-loader: 1719] x23: 0000000000000200 x22: 0000000000000220 x21: 0000000000000040 [2479.385925]I[0:launcher-loader: 1719] x20: ffffffc008007440 x19: 0000000000000000 x18: 0000000000000000 [2479.385969]I[0:launcher-loader: 1719] x17: 2065726568207475 x16: 000000000000005e x15: 2d2d2d2d2d2d2d20 [2479.386013]I[0:launcher-loader: 1719] x14: 5d393313731203a72 x13: 00000000002f6b30 x12: 00000000002f6af8 [2479.386057]I[0:launcher-loader: 1719] x11: 00000000ffffffff x10: ffffffb90aacf000 x9 : e8a74a6c16008800 [2479.386101]I[0:launcher-loader: 1719] x8 : e8a74a6c16008800 x7 : 00000000002f6b30 x6 : 00000000002f6af8 [2479.386145]I[0:launcher-loader: 1719] x5 : ffffffc0080070c8 x4 : ffffffd00b192380 x3 : ffffffd0092b313c [2479.386189]I[0:launcher-loader: 1719] x2 : 0000000000000001 x1 : 0000000000000004 x0 : 0000000000000022 [	N/A	More Details

	2479.386231]I[0:launcher-loader: 1719] Call trace: [2479.386248]I[0:launcher-loader: 1719] __stack_depot_save+0x464/0x46c [2479.386273]I[0:launcher-loader: 1719] kasan_save_stack+0x58/0x70 [2479.386303]I[0:launcher-loader: 1719] save_stack_info+0x34/0x138 [2479.386331]I[0:launcher-loader: 1719] kasan_save_free_info+0x18/0x24 [2479.386358]I[0:launcher-loader: 1719] ____kasan_slab_free+0x16c/0x170 [2479.386385]I[0:launcher-loader: 1719] __kasan_slab_free+0x10/0x20 [2479.386410]I[0:launcher-loader: 1719] kmem_cache_free+0x238/0x53c [2479.386435]I[0:launcher-loader: 1719] mempool_free_slab+0x1c/0x28 [2479.386460]I[0:launcher-loader: 1719] mempool_free+0x7c/0x1a0 [2479.386484]I[0:launcher-loader: 1719] bvec_free+0x34/0x80 [2479.386514]I[0:launcher-loader: 1719] bio_free+0x60/0x98 [2479.386540]I[0:launcher-loader: 1719] bio_put+0x50/0x21c [2479.386567]I[0:launcher-loader: 1719] f2fs_write_end_io+0x4ac/0x4d0 [2479.386594]I[0:launcher-loader: 1719] bio_endio+0x2dc/0x300 [2479.386622]I[0:launcher-loader: 1719] __dm_io_complete+0x324/0x37c [2479.386650]I[0:launcher-loader: 1719] dm_io_dec_pending+0x60/0xa4 [2479.386676]I[0:launcher-loader: 1719] clone_endio+0xf8/0x2f0 [2479.386700]I[0:launcher-loader: 1719] bio_endio+0x2dc/0x300 [2479.386727]I[0:launcher-loader: 1719] blk_update_request+0x258/0x63c [2479.386754]I[0:launcher-loader: 1719] scsi_end_request+0x50/0x304 [2479.386782]I[0:launcher-loader: 1719] scsi_io_completion+0x88/0x160 [2479.386808]I[0:launcher-loader: 1719] scsi_finish_command+0x17c/0x194 [2479.386833]I ---truncated---		
CVE-2023-54247	In the Linux kernel, the following vulnerability has been resolved: bpf: Silence a warning in btf_type_id_size() syzbot reported a warning in [1] with the following stacktrace: WARNING: CPU: 0 PID: 5005 at kernel/bpf/btf.c:1988 btf_type_id_size+0x2d9/0x9d0 kernel/bpf/btf.c:1988 ... RIP: 0010:btf_type_id_size+0x2d9/0x9d0 kernel/bpf/btf.c:1988 ... Call Trace: <TASK> map_check_btf kernel/bpf/syscall.c:1024 [inline] map_create+0x1157/0x1860 kernel/bpf/syscall.c:1198 __sys_bpf+0x127f/0x5420 kernel/bpf/syscall.c:5040 __do_sys_bpf kernel/bpf/syscall.c:5162 [inline] __se_sys_bpf kernel/bpf/syscall.c:5160 [inline] __x64_sys_bpf+0x79/0xc0 kernel/bpf/syscall.c:5160 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x39/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd With the following btf [1] DECL_TAG 'a' type_id=4 component_idx=-1 [2] PTR '(anon)' type_id=0 [3] TYPE_TAG 'a' type_id=2 [4] VAR 'a' type_id=3, linkage=static and when the bpf_attr.btf_key_type_id = 1 (DECL_TAG), the following WARN_ON_ONCE in btf_type_id_size() is triggered: if (WARN_ON_ONCE(!btf_type_is_modifier(size_type) && !btf_type_is_var(size_type))) return NULL; Note that 'return NULL' is the correct behavior as we don't want a DECL_TAG type to be used as a btf_{key,value}_type_id even for the case like 'DECL_TAG -> STRUCT'. So there is no correctness issue here, we just want to silence warning. To silence the warning, I added DECL_TAG as one of kinds in btf_type_nosize() which will cause btf_type_id_size() returning NULL earlier without the warning. [1] https://lore.kernel.org/bpf/000000000000e0df8d05fc75ba86@google.com/	N/A	More Details
CVE-2022-50599	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2023-54240	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: mtk_eth_soc: fix possible NULL pointer dereference in mtk_hwlro_get_fdir_all() rule_locs is allocated in ethtool_get_rxnfc and the size is determined by rule_cnt from user space. So rule_cnt needs to be check before using rule_locs to avoid NULL pointer dereference.	N/A	More Details
CVE-2023-54241	In the Linux kernel, the following vulnerability has been resolved: MIPS: KVM: Fix NULL pointer dereference After commit 45c7e8af4a5e3f0bea4ac209 ("MIPS: Remove KVM_TE support") we get a NULL pointer dereference when creating a KVM guest: [146.243409] Starting KVM with MIPS VZ extensions [149.849151] CPU 3 Unable to handle kernel paging request at virtual address 0000000000000300, epc == ffffffffc06356ec, ra == ffffffffc063568c [149.849177] Oops[#1]: [149.849182] CPU: 3 PID: 2265 Comm: qemu-system-mip Not tainted 6.4.0-rc3+ #1671 [149.849188] Hardware name: THTF CX TL630 Series/THTF-LS3A4000-7A1000-ML4A, BIOS KL4.1F.TF.D.166.201225.R 12/25/2020 [149.849192] $ 0 : 0000000000000000 000000007400cce0 0000000000400004 ffffffff8119c740 [149.849209] $ 4 : 000000007400cce1 000000007400cce1 0000000000000000 0000000000000000 [149.849221] $ 8 : 000000240058bb36 ffffffff81421ac0 0000000000000000 0000000000400dc0 [149.849233] $12 : 9800000102a07cc8 ffffffff80e40e38 0000000000000001 0000000000400dc0 [149.849245] $16 : 0000000000000000 9800000106cd0000 9800000106cd0000 9800000100cce000 [149.849257] $20 : ffffffffc0632b28 ffffffffc05b31b0 9800000100ccca00 0000000000400000 [149.849269] $24 : 9800000106cd09ce ffffffff802f69d0 [149.849281] $28 : 9800000102a04000 9800000102a07cd0 98000001106a8000 ffffffffc063568c [149.849293] Hi : 00000335b2111e66 [149.849295] Lo : 6668d90061ae0ae9 [149.849298] epc : ffffffffc06356ec kvm_vz_vcpu_setup+0xc4/0x328 [kvm] [149.849324] ra : ffffffffc063568c kvm_vz_vcpu_setup+0x64/0x328 [kvm] [149.849336] Status: 7400cce3 KX SX UX KERNEL EXL IE [149.849351] Cause : 1000000c (ExcCode 03) [149.849354] BadVA : 0000000000000300 [149.849357] PrId : 0014c004 (ICT Loongson-3) [149.849360] Modules linked in: kvm nfnetlink_queue nfnetlink_log nfnetlink fuse sha256_generic libsha256 cfg80211 rfkill binfmt_misc vfat fat snd_hda_codec_hdmi input_leds led_class snd_hda_intel snd_intel_dspcfg snd_hda_codec snd_hda_core snd_pcm snd_timer snd serio_raw xhci_pci radeon drm_suballoc_helper drm_display_helper xhci_hcd ip_tables x_tables [149.849432] Process qemu-system-mip (pid: 2265, threadinfo=00000000ae2982d2, task=0000000038e09ad4, tls=000000ffeba16030) [149.849439] Stack : 9800000000000003 9800000100ccca00 9800000100ccc000 ffffffffc062cef4 [149.849453] 9800000102a07d18 c89b63a7ab338e00 0000000000000000 ffffffff811a0000 [149.849465] 0000000000000000 9800000106cd0000 ffffffff80e59938 98000001106a8920 [149.849476] ffffffff80e57f30 ffffffffc062854c ffffffff811a0000 9800000102bf4240 [149.849488] ffffffffc05b0000 ffffffff80e3a798 000000ff78000000 000000ff78000010 [149.849500] 0000000000000255 98000001021f7de0 98000001023f0078 ffffffff81434000 [149.849511] 0000000000000000 0000000000000000 9800000102ae0000 980000025e92ae28 [149.849523] 0000000000000000 c89b63a7ab338e00 0000000000000001 ffffffff8119dce0 [149.849535] 000000ff78000010 ffffffff804f3d3c 9800000102a07eb0 0000000000000255 [149.849546] 0000000000000000 ffffffff8049460c 000000ff78000010 0000000000000255 [149.849558] ... [149.849565] Call Trace: [149.849567] [<ffffffffc06356ec>] kvm_vz_vcpu_setup+0xc4/0x328 [kvm] [149.849586] [<ffffffffc062cef4>] kvm_arch_vcpu_create+0x184/0x228 [kvm] [149.849605] [<ffffffffc062854c>] kvm_vm_ioctl+0x64c/0xf28 [kvm] [149.849623] [<ffffffff805209c0>] sys_ioctl+0xc8/0x118 [149.849631] [<ffffffff80219eb0>] syscall_common+0x34/0x58 The root cause is the deletion of kvm_mips_commpage_init() leaves vcpu ->arch.cop0 NULL. So fix it by making cop0 from a pointer to an embedded object.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: block, bfq: Fix division by zero error on zero wsum When the weighted sum is zero the calculation of limit causes a division by zero error. Fix this by continuing to the next level. This was discovered by running as root: stress-ng --ioprio 0 Fixes divison by error oops: [521.450556] divide error: 0000 [#1] SMP NOPTI [521.450766] CPU: 2 PID: 2684464 Comm: stress-ng-iopri Not tainted 6.2.1-1280.native #1 [521.451117] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.1-0-g3208b098f51a-prebuilt.qemu.org 04/01/2014 [521.451627] RIP: 0010:bfqq_request_over_limit+0x207/0x400 [521.451875] Code: 01 48 8d 0c c8 74 0b 48 8b 82 98 00 00 00 48 8d 0c c8 8b 85 34 ff ff ff 48 89 ca 41 0f af 41 50 48 d1 ea 48 98 48 01 d0 31 d2 <48> f7 f1 41 39 41 48 89 85 34 ff ff ff 0f 8c 7b 01 00 00 49 8b 44 [521.452699] RSP: 0018:ffffb1af84eb3948 EFLAGS: 00010046 [521.452938] RAX: 000000000000003c RBX: 0000000000000000 RCX: 0000000000000000 [521.453262] RDX: 0000000000000000 RSI: 0000000000000000 RDI: ffffb1af84eb3978 [521.453584] RBP: ffffb1af84eb3a30 R08: 0000000000000001 R09: ffff8f88ab8a4ba0 [521.453905] R10: 0000000000000000 R11: 0000000000000001 R12: ffff8f88ab8a4b18 [521.454224] R13: ffff8f8699093000 R14: 0000000000000001 R15: ffffb1af84eb3970 [521.454549] FS: 00005640b6b0b580(0000) GS:ffff8f88b3880000(0000) knlGS:0000000000000000 [521.454912] CS: 0010 DS:	N/A	More

54242	0000 ES: 0000 CR0: 0000000080050033 [521.455170] CR2: 00007ffcbcae4e38 CR3: 00000002e46de001 CR4: 0000000000770ee0 [521.455491] PKRU: 55555554 [521.455619] Call Trace: [521.455736] <TASK> [521.455837] ? bfq_request_merge+0x3a/0xc0 [521.456027] ? elv_merge+0x115/0x140 [521.456191] bfq_limit_depth+0xc8/0x240 [521.456366] __blk_mq_alloc_requests+0x21a/0x2c0 [521.456577] blk_mq_submit_bio+0x23c/0x6c0 [521.456766] __submit_bio+0xb8/0x140 [521.457236] submit_bio_noacct_nocheck+0x212/0x300 [521.457748] submit_bio_noacct+0x1a6/0x580 [521.458220] submit_bio+0x43/0x80 [521.458660] ext4_io_submit+0x23/0x80 [521.459116] ext4_do_writepages+0x40a/0xd00 [521.459596] ext4_writepages+0x65/0x100 [521.460050] do_writepages+0xb7/0x1c0 [521.460492] __filemap_fdatawrite_range+0xa6/0x100 [521.460979] file_write_and_wait_range+0xbf/0x140 [521.461452] ext4_sync_file+0x105/0x340 [521.461882] __x64_sys_fsync+0x67/0x100 [521.462305] ? syscall_exit_to_user_mode+0x2c/0x1c0 [521.462768] do_syscall_64+0x3b/0xc0 [521.463165] entry_SYSCALL_64_after_hwframe+0x5a/0xc4 [521.463621] RIP: 0033:0x5640b6c56590 [521.464006] Code: 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 80 3d 71 70 0e 00 00 74 17 b8 4a 00 00 00 0f 05 <48> 3d 00 f0 ff ff 77 48 c3 0f 1f 80 00 00 00 00 48 83 ec 18 89 7c		Details
CVE-2023-54243	In the Linux kernel, the following vulnerability has been resolved: netfilter: ebtables: fix table blob use-after-free We are not allowed to return an error at this point. Looking at the code it looks like ret is always 0 at this point, but its not. t = find_table_lock(net, repl->name, &ret, &ebt_mutex); ... this can return a valid table, with ret != 0. This bug causes update of table->private with the new blob, but then frees the blob right away in the caller. Syzbot report: BUG: KASAN: vmalloc-out-of-bounds in __ebt_unregister_table+0xc00/0xcd0 net/bridge/netfilter/ebtables.c:1168 Read of size 4 at addr ffffc90005425000 by task kworker/u4:4/74 Workqueue: netns cleanup_net Call Trace: kasan_report+0xbf/0x1f0 mm/kasan/report.c:517 __ebt_unregister_table+0xc00/0xcd0 net/bridge/netfilter/ebtables.c:1168 ebt_unregister_table+0x35/0x40 net/bridge/netfilter/ebtables.c:1372 ops_exit_list+0xb0/0x170 net/core/net_namespace.c:169 cleanup_net+0x4ee/0xb10 net/core/net_namespace.c:613 ... ip(6)tables appears to be ok (ret should be 0 at this point) but make this more obvious.	N/A	More Details
CVE-2023-54244	In the Linux kernel, the following vulnerability has been resolved: ACPI: EC: Fix oops when removing custom query handlers When removing custom query handlers, the handler might still be used inside the EC query workqueue, causing a kernel oops if the module holding the callback function was already unloaded. Fix this by flushing the EC query workqueue when removing custom query handlers. Tested on a Acer Travelmate 4002WLMi	N/A	More Details
CVE-2023-54245	In the Linux kernel, the following vulnerability has been resolved: ASoC: codecs: tx-macro: Fix for KASAN: slab-out-of-bounds When we run syzkaller we get below Out of Bound. "KASAN: slab-out-of-bounds Read in regcache_flat_read" Below is the backtrace of the issue: dump_backtrace+0x0/0x4c8 show_stack+0x34/0x44 dump_stack_lvl+0xd8/0x118 print_address_description+0x30/0x2d8 kasan_report+0x158/0x198 __asan_report_load4_noabort+0x44/0x50 regcache_flat_read+0x10c/0x110 regcache_read+0xf4/0x180 _regmap_read+0xc4/0x278 _regmap_update_bits+0x130/0x290 regmap_update_bits_base+0xc0/0x15c snd_soc_component_update_bits+0xa8/0x22c snd_soc_component_write_field+0x68/0xd4 tx_macro_digital_mute+0xec/0x140 Actually There is no need to have decimator with 32 bits. By limiting the variable with short type u8 issue is resolved.	N/A	More Details
CVE-2023-54246	In the Linux kernel, the following vulnerability has been resolved: rcuscale: Move rcu_scale_writer() schedule_timeout_uninterruptible() to _idle() The rcuscale.holdoff module parameter can be used to delay the start of rcu_scale_writer() kthread. However, the hung-task timeout will trigger when the timeout specified by rcuscale.holdoff is greater than hung_task_timeout_secs: runqemu kvm nographic slirp qemuparams="-smp 4 -m 2048M" bootparams="rcuscale.shutdown=0 rcuscale.holdoff=300" [247.071753] INFO: task rcu_scale_write:59 blocked for more than 122 seconds. [247.072529] Not tainted 6.4.0-rc1-00134-gb9ed6de8d4ff #7 [247.073400] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message. [247.074331] task:rcu_scale_write state:D stack:30144 pid:59 ppid:2 flags:0x00004000 [247.075346] Call Trace: [247.075660] <TASK> [247.075965] __schedule+0x635/0x1280 [247.076448] ? __pfx___schedule+0x10/0x10 [247.076967] ? schedule_timeout+0x2dc/0x4d0 [247.077471] ? __pfx_lock_release+0x10/0x10 [247.078018] ? enqueue_timer+0xe2/0x220 [247.078522] schedule+0x84/0x120 [247.078957] schedule_timeout+0x2e1/0x4d0 [247.079447] ? __pfx_schedule_timeout+0x10/0x10 [247.080032] ? __pfx_rcu_scale_writer+0x10/0x10 [247.080591] ? __pfx_process_timeout+0x10/0x10 [247.081163] ? __pfx_sched_set_fifo_low+0x10/0x10 [247.081760] ? __pfx_rcu_scale_writer+0x10/0x10 [247.082287] rcu_scale_writer+0x6b1/0x7f0 [247.082773] ? mark_held_locks+0x29/0xa0 [247.083252] ? __pfx_rcu_scale_writer+0x10/0x10 [247.083865] ? __pfx_rcu_scale_writer+0x10/0x10 [247.084412] kthread+0x179/0x1c0 [247.084759] ? __pfx_kthread+0x10/0x10 [247.085098] ret_from_fork+0x2c/0x50 [247.085433] </TASK> This commit therefore replaces schedule_timeout_uninterruptible() with schedule_timeout_idle().	N/A	More Details
CVE-2023-54248	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Add check for kmemdup Since the kmemdup may return NULL pointer, it should be better to add check for the return value in order to avoid NULL pointer dereference.	N/A	More Details
CVE-2023-54257	In the Linux kernel, the following vulnerability has been resolved: net: macb: fix a memory corruption in extended buffer descriptor mode For quite some time we were chasing a bug which looked like a sudden permanent failure of networking and mmc on some of our devices. The bug was very sensitive to any software changes and even more to any kernel debug options. Finally we got a setup where the problem was reproducible with CONFIG_DMA_API_DEBUG=y and it revealed the issue with the rx dma: [16.992082] ----------[cut here]------------ [16.996779] DMA-API: macb ff0b0000.ethernet: device driver tries to free DMA memory it has not allocated [device address=0x0000000875e3e244] [size=1536 bytes] [17.011049] WARNING: CPU: 0 PID: 85 at kernel/dma/debug.c:1011 check_unmap+0x6a0/0x900 [17.018977] Modules linked in: xxxxx [17.038823] CPU: 0 PID: 85 Comm: irq/55-8000f000 Not tainted 5.4.0 #28 [17.045345] Hardware name: xxxxx [17.049528] pstate: 60000005 (nZCv daif -PAN -UAO) [17.054322] pc : check_unmap+0x6a0/0x900 [17.058243] lr : check_unmap+0x6a0/0x900 [17.062163] sp : ffffffc010003c40 [17.065470] x29: ffffffc010003c40 x28: 000000004000c03c [17.070783] x27: ffffffc010da7048 x26: ffffff8878e38800 [17.076095] x25: ffffff8879d22810 x24: ffffffc010003cc8 [17.081407] x23: 0000000000000000 x22: ffffffc010a08750 [17.086719] x21: ffffff8878e3c7c0 x20: ffffffc010acb000 [17.092032] x19: 0000000875e3e244 x18: 0000000000000010 [17.097343] x17: 0000000000000000 x16: 0000000000000000 [17.102647] x15: ffffff8879e4a988 x14: 0720072007200720 [17.107959] x13: 0720072007200720 x12: 0720072007200720 [17.113261] x11: 0720072007200720 x10: 0720072007200720 [17.118565] x9 : 0720072007200720 x8 : 000000000000022d [17.123869] x7 : 0000000000000015 x6 : 0000000000000098 [17.129173] x5 : 0000000000000000 x4 : 0000000000000000 [17.134475] x3 : 00000000ffffffff x2 : ffffffc010a1d370 [17.139778] x1 : b420c9d75d27bb00 x0 : 0000000000000000 [17.145082] Call trace: [17.147524] check_unmap+0x6a0/0x900 [17.151091] debug_dma_unmap_page+0x88/0x90 [17.155266] gem_rx+0x114/0x2f0 [17.158396] macb_poll+0x58/0x100 [17.161705] net_rx_action+0x118/0x400 [17.165445] __do_softirq+0x138/0x36c [17.169100] irq_exit+0x98/0xc0 [17.172234] __handle_domain_irq+0x64/0xc0 [17.176320] gic_handle_irq+0x5c/0xc0 [17.179974] el1_irq+0xb8/0x140 [17.183109] xiic_process+0x5c/0xe30 [17.186677] irq_thread_fn+0x28/0x90 [17.190244] irq_thread+0x208/0x2a0 [17.193724] kthread+0x130/0x140 [17.196945] ret_from_fork+0x10/0x20 [17.200510] ---[end trace 7240980785f81d6f]--- [237.021490] ----------[cut here]------------ [237.026129] DMA-API: exceeded 7 overlapping mappings of cacheline 0x0000000021d79e7b [237.033886] WARNING: CPU: 0 PID: 0 at kernel/dma/debug.c:499 add_dma_entry+0x214/0x240 [237.041802] Modules linked in: xxxxx [237.061637] CPU: 0 PID: 0 Comm: swapper/0 Tainted: G W 5.4.0 #28 [237.068941] Hardware name: xxxxx [237.073116]	N/A	More Details

	pstate: 80000085 (Nzcv dalf -PAN -UAO) [237.077900] pc : add_dma_entry+0x214/0x240 [237.081986] lr : add_dma_entry+0x214/0x240 [237.086072] sp : ffffffc010003c30 [237.089379] x29: ffffffc010003c30 x28: ffffff8878a0be00 [237.094683] x27: 0000000000000180 x26: ffffff8878e387c0 [237.099987] x25: 0000000000000002 x24: 0000000000000000 [237.105290] x23: 000000000000003b x22: ffffffc010a0fa00 [237.110594] x21: 0000000021d79e7b x20: ffffffc010abe600 [237.115897] x19: 00000000ffffffef x18: 0000000000000010 [237.121201] x17: 0000000000000000 x16: 0000000000000000 [237.126504] x15: ffffffc010a0fdc8 x14: 0720072007200720 [237.131807] x13: 0720072007200720 x12: 0720072007200720 [237.137111] x11: 0720072007200720 x10: 0720072007200720 [237.142415] x9 : 0720072007200720 x8 : 0000000000000259 [237.147718] x7 : 0000000000000001 x6 : 0000000000000000 [237.15302 ---truncated---		
CVE-2023-54249	In the Linux kernel, the following vulnerability has been resolved: bus: mhi: ep: Only send -ENOTCONN status if client driver is available For the STOP and RESET commands, only send the channel disconnect status -ENOTCONN if client driver is available. Otherwise, it will result in null pointer dereference.	N/A	More Details
CVE-2023-54250	In the Linux kernel, the following vulnerability has been resolved: ksmbd: avoid out of bounds access in decode_preauth_ctxt() Confirm that the accessed pneg_ctxt->HashAlgorithms address sits within the SMB request boundary; deassemble_neg_contexts() only checks that the eight byte smb2_neg_context header + (client controlled) DataLength are within the packet boundary, which is insufficient. Checking for sizeof(struct smb2_preauth_neg_context) is overkill given that the type currently assumes SMB311_SALT_SIZE bytes of trailing Salt.	N/A	More Details
CVE-2023-54251	In the Linux kernel, the following vulnerability has been resolved: net/sched: taprio: Limit TCA_TAPRIO_ATTR_SCHED_CYCLE_TIME to INT_MAX. syzkaller found zero division error [0] in div_s64_rem() called from get_cycle_time_elapsed(), where sched->cycle_time is the divisor. We have tests in parse_taprio_schedule() so that cycle_time will never be 0, and actually cycle_time is not 0 in get_cycle_time_elapsed(). The problem is that the types of divisor are different; cycle_time is s64, but the argument of div_s64_rem() is s32. syzkaller fed this input and 0x100000000 is cast to s32 to be 0. @TCA_TAPRIO_ATTR_SCHED_CYCLE_TIME={0xc, 0x8, 0x100000000} We use s64 for cycle_time to cast it to ktime_t, so let's keep it and set max for cycle_time. While at it, we prevent overflow in setup_txtime() and add another test in parse_taprio_schedule() to check if cycle_time overflows. Also, we add a new tdc test case for this issue. [0]: divide error: 0000 [#1] PREEMPT SMP KASAN NOPTI CPU: 1 PID: 103 Comm: kworker/1:3 Not tainted 6.5.0-rc1-00330-g60cc1f7d0605 #3 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 Workqueue: ipv6_addrconf addrconf_dad_work RIP: 0010:div_s64_rem include/linux/math64.h:42 [inline] RIP: 0010:get_cycle_time_elapsed net/sched/sch_taprio.c:223 [inline] RIP: 0010:find_entry_to_transmit+0x252/0x7e0 net/sched/sch_taprio.c:344 Code: 3c 02 00 0f 85 5e 05 00 00 48 8b 4c 24 08 4d 8b bd 40 01 00 00 48 8b 7c 24 48 48 89 c8 4c 29 f8 48 63 f7 48 99 48 89 74 24 70 <48> f7 fe 48 29 d1 48 8d 04 0f 49 89 cc 48 89 44 24 20 49 8d 85 10 RSP: 0018:ffffc90000acf260 EFLAGS: 00010206 RAX: 177450e0347560cf RBX: 0000000000000000 RCX: 177450e0347560cf RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000100000000 RBP: 0000000000000056 R08: 0000000000000000 R09: ffffed10020a0934 R10: ffff8880105049a7 R11: ffff88806cf3a520 R12: ffff888010504800 R13: ffff88800c00d800 R14: ffff8880105049a0 R15: 0000000000000000 FS: 0000000000000000(0000) GS:ffff88806cf00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007f0edf84f0e8 CR3: 000000000d73c002 CR4: 0000000000770ee0 PKRU: 55555554 Call Trace: <TASK> get_packet_txtime net/sched/sch_taprio.c:508 [inline] taprio_enqueue_one+0x900/0xff0 net/sched/sch_taprio.c:577 taprio_enqueue+0x378/0xae0 net/sched/sch_taprio.c:658 dev_qdisc_enqueue+0x46/0x170 net/core/dev.c:3732 __dev_xmit_skb net/core/dev.c:3821 [inline] __dev_queue_xmit+0x1b2f/0x3000 net/core/dev.c:4169 dev_queue_xmit include/linux/netdevice.h:3088 [inline] neigh_resolve_output net/core/neighbour.c:1552 [inline] neigh_resolve_output+0x4a7/0x780 net/core/neighbour.c:1532 neigh_output include/net/neighbour.h:544 [inline] ip6_finish_output2+0x924/0x17d0 net/ipv6/ip6_output.c:135 __ip6_finish_output+0x620/0xaa0 net/ipv6/ip6_output.c:196 ip6_finish_output net/ipv6/ip6_output.c:207 [inline] NF_HOOK_COND include/linux/netfilter.h:292 [inline] ip6_output+0x206/0x410 net/ipv6/ip6_output.c:228 dst_output include/net/dst.h:458 [inline] NF_HOOK.constprop.0+0xea/0x260 include/linux/netfilter.h:303 ndisc_send_skb+0x872/0xe80 net/ipv6/ndisc.c:508 ndisc_send_ns+0xb5/0x130 net/ipv6/ndisc.c:666 addrconf_dad_work+0xc14/0x13f0 net/ipv6/addrconf.c:4175 process_one_work+0x92c/0x13a0 kernel/workqueue.c:2597 worker_thread+0x60f/0x1240 kernel/workqueue.c:2748 kthread+0x2fe/0x3f0 kernel/kthread.c:389 ret_from_fork+0x2c/0x50 arch/x86/entry/entry_64.S:308 </TASK> Modules linked in:	N/A	More Details
CVE-2023-54252	In the Linux kernel, the following vulnerability has been resolved: platform/x86: think-lmi: Fix memory leaks when parsing ThinkStation WMI strings My previous commit introduced a memory leak where the item allocated from tlmi_setting was not freed. This commit also renames it to avoid confusion with the similarly name variable in the same function.	N/A	More Details
CVE-2023-54253	In the Linux kernel, the following vulnerability has been resolved: btrfs: set page extent mapped after read_folio in relocate_one_page One of the CI runs triggered the following panic assertion failed: PagePrivate(page) && page->private, in fs/btrfs/subpage.c:229 ------------[cut here]------------ kernel BUG at fs/btrfs/subpage.c:229! Internal error: Oops - BUG: 00000000f2000800 [#1] SMP CPU: 0 PID: 923660 Comm: btrfs Not tainted 6.5.0-rc3+ #1 pstate: 61400005 (nZCv daif +PAN -UAO -TCO +DIT -SSBS BTYPE=--) pc : btrfs_subpage_assert+0xbc/0xf0 lr : btrfs_subpage_assert+0xbc/0xf0 sp : ffff800093213720 x29: ffff800093213720 x28: ffff8000932138b4 x27: 000000000c280000 x26: 00000001b5d00000 x25: 000000000c281000 x24: 000000000c281fff x23: 0000000000001000 x22: 0000000000000000 x21: ffffff42b95bf880 x20: ffff42b9528e0000 x19: 0000000000001000 x18: ffffffffffffffff x17: 667274622f736620 x16: 6e69202c65746176 x15: 0000000000000028 x14: 0000000000000003 x13: 00000000002672d7 x12: 0000000000000000 x11: ffffcd3f0ccd9204 x10: ffffcd3f0554ae50 x9 : ffffcd3f0379528c x8 : ffff800093213428 x7 : 0000000000000000 x6 : ffffcd3f091771e8 x5 : ffff42b97f333948 x4 : 0000000000000000 x3 : 0000000000000000 x2 : 0000000000000000 x1 : ffff42b9556cde80 x0 : 000000000000004f Call trace: btrfs_subpage_assert+0xbc/0xf0 btrfs_subpage_set_dirty+0x38/0xa0 btrfs_page_set_dirty+0x58/0x88 relocate_one_page+0x204/0x5f0 relocate_file_extent_cluster+0x11c/0x180 relocate_data_extent+0xd0/0xf8 relocate_block_group+0x3d0/0x4e8 btrfs_relocate_block_group+0x2d8/0x490 btrfs_relocate_chunk+0x54/0x1a8 btrfs_balance+0x7f4/0x1150 btrfs_ioctl+0x10f0/0x20b8 __arm64_sys_ioctl+0x120/0x11d8 invoke_syscall.constprop.0+0x80/0xd8 do_el0_svc+0x6c/0x158 el0_svc+0x50/0x1b0 el0t_64_sync_handler+0x120/0x130 el0t_64_sync+0x194/0x198 Code: 91098021 b0007fa0 91346000 97e9c6d2 (d4210000) This is the same problem outlined in 17b17fcd6d44 ("btrfs: set_page_extent_mapped after read_folio in btrfs_cont_expand") , and the fix is the same. I originally looked for the same pattern elsewhere in our code, but mistakenly skipped over this code because I saw the page cache readahead before we set_page_extent_mapped, not realizing that this was only in the !page case, that we can still end up with a !uptodate page and then do the btrfs_read_folio further down. The fix here is the same as the above mentioned patch, move the set_page_extent_mapped call to after the btrfs_read_folio() block to make sure that we have the subpage blocksize stuff setup properly before using the page.	N/A	More Details
CVE-2023-54254	In the Linux kernel, the following vulnerability has been resolved: drm/ttm: Don't leak a resource on eviction error On eviction errors other than -EMULTIHOP we were leaking a resource. Fix. v2: - Avoid yet another goto (Andi Shyti)	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: sh: dma: Fix DMA channel offset calculation Various SoCs of the SH3, SH4 and SH4A family, which use this driver, feature a differing number of DMA channels, which can be distributed between up		

CVE-2023-54255	to two DMAC modules. The existing implementation fails to correctly accommodate for all those variations, resulting in wrong channel offset calculations and leading to kernel panics. Rewrite dma_base_addr() in order to properly calculate channel offsets in a DMAC module. Fix dmaor_read_reg() and dmaor_write_reg(), so that the correct DMAC module base is selected for the DMAOR register.	N/A	More Details
CVE-2023-54256	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2023-54279	In the Linux kernel, the following vulnerability has been resolved: MIPS: fw: Allow firmware to pass a empty env fw_getenv will use env entry to determine style of env, however it is legal for firmware to just pass a empty list. Check if first entry exist before running strchr to avoid null pointer dereference.	N/A	More Details
CVE-2023-54280	In the Linux kernel, the following vulnerability has been resolved: cifs: fix potential race when tree connecting ipc Protect access of TCP_Server_Info::hostname when building the ipc tree name as it might get freed in cifsd thread and thus causing an use-after-free bug in __tree_connect_dfs_target(). Also, while at it, update status of IPC tcon on success and then avoid any extra tree connects.	N/A	More Details
CVE-2023-54281	In the Linux kernel, the following vulnerability has been resolved: btrfs: release path before inode lookup during the ino lookup ioctl During the ino lookup ioctl we can end up calling btrfs_iget() to get an inode reference while we are holding on a root's btree. If btrfs_iget() needs to lookup the inode from the root's btree, because it's not currently loaded in memory, then it will need to lock another or the same path in the same root btree. This may result in a deadlock and trigger the following lockdep splat: WARNING: possible circular locking dependency detected 6.5.0-rc7-syzkaller-00004-gf7757129e3de #0 Not tainted -- syz-executor277/5012 is trying to acquire lock: ffff88802df41710 (btrfs-tree-01){++++}-{3:3}, at: __btrfs_tree_read_lock+0x2f/0x220 fs/btrfs/locking.c:136 but task is already holding lock: ffff88802df418e8 (btrfs-tree-00){++++}-{3:3}, at: __btrfs_tree_read_lock+0x2f/0x220 fs/btrfs/locking.c:136 which lock already depends on the new lock. the existing dependency chain (in reverse order) is: -> #1 (btrfs-tree-00){++++}-{3:3}: down_read_nested+0x49/0x2f0 kernel/locking/rwsem.c:1645 __btrfs_tree_read_lock+0x2f/0x220 fs/btrfs/locking.c:136 btrfs_search_slot+0x13a4/0x2f80 fs/btrfs/ctree.c:2302 btrfs_init_root_free_objectid+0x148/0x320 fs/btrfs/disk-io.c:4955 btrfs_init_fs_root fs/btrfs/disk-io.c:1128 [inline] btrfs_get_root_ref+0x5ae/0xae0 fs/btrfs/disk-io.c:1338 btrfs_get_fs_root fs/btrfs/disk-io.c:1390 [inline] open_ctree+0x29c8/0x3030 fs/btrfs/disk-io.c:3494 btrfs_fill_super+0x1c7/0x2f0 fs/btrfs/super.c:1154 btrfs_mount_root+0x7e0/0x910 fs/btrfs/super.c:1519 legacy_get_tree+0xef/0x190 fs/fs_context.c:611 vfs_get_tree+0x8c/0x270 fs/super.c:1519 fc_mount fs/namespace.c:1112 [inline] vfs_kern_mount+0xbc/0x150 fs/namespace.c:1142 btrfs_mount+0x39f/0xb50 fs/btrfs/super.c:1579 legacy_get_tree+0xef/0x190 fs/fs_context.c:611 vfs_get_tree+0x8c/0x270 fs/super.c:1519 do_new_mount+0x28f/0xae0 fs/namespace.c:3335 do_mount fs/namespace.c:3675 [inline] __do_sys_mount fs/namespace.c:3884 [inline] __se_sys_mount+0x2d9/0x3c0 fs/namespace.c:3861 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd -> #0 (btrfs-tree-01){++++}-{3:3}: check_prev_add kernel/locking/lockdep.c:3142 [inline] check_prevs_add kernel/locking/lockdep.c:3261 [inline] validate_chain kernel/locking/lockdep.c:3876 [inline] __lock_acquire+0x39ff/0x7f70 kernel/locking/lockdep.c:5144 lock_acquire+0x1e3/0x520 kernel/locking/lockdep.c:5761 down_read_nested+0x49/0x2f0 kernel/locking/rwsem.c:1645 __btrfs_tree_read_lock+0x2f/0x220 fs/btrfs/locking.c:136 btrfs_tree_read_lock fs/btrfs/locking.c:142 [inline] btrfs_read_lock_root_node+0x292/0x3c0 fs/btrfs/locking.c:281 btrfs_search_slot_get_root fs/btrfs/ctree.c:1832 [inline] btrfs_search_slot+0x4ff/0x2f80 fs/btrfs/ctree.c:2154 btrfs_lookup_inode+0xdc/0x480 fs/btrfs/inode-item.c:412 btrfs_read_locked_inode fs/btrfs/inode.c:3892 [inline] btrfs_iget_path+0x2d9/0x1520 fs/btrfs/inode.c:5716 btrfs_search_path_in_tree_user fs/btrfs/ioctl.c:1961 [inline] btrfs_ioctl_ino_lookup_user+0x77a/0xf50 fs/btrfs/ioctl.c:2105 btrfs_ioctl+0xb0b/0xd40 fs/btrfs/ioctl.c:4683 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:870 [inline] __se_sys_ioctl+0xf8/0x170 fs/ioctl.c:856 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd other info ---truncated---	N/A	More Details
CVE-2023-54312	In the Linux kernel, the following vulnerability has been resolved: samples/bpf: Fix buffer overflow in tcp_basertt Using sizeof(nv) or strlen(nv)+1 is correct.	N/A	More Details
CVE-2023-54304	In the Linux kernel, the following vulnerability has been resolved: firmware: meson_sm: fix to avoid potential NULL pointer dereference of_match_device() may fail and returns a NULL pointer. Fix this by checking the return value of of_match_device.	N/A	More Details
CVE-2023-54305	In the Linux kernel, the following vulnerability has been resolved: ext4: refuse to create ea block when umounted The ea block expansion need to access s_root while it is already set as NULL when umount is triggered. Refuse this request to avoid panic.	N/A	More Details
CVE-2023-54306	In the Linux kernel, the following vulnerability has been resolved: net: tls: avoid hanging tasks on the tx_lock syzbot sent a hung task report and Eric explains that adversarial receiver may keep RWIN at 0 for a long time, so we are not guaranteed to make forward progress. Thread which took tx_lock and went to sleep may not release tx_lock for hours. Use interruptible sleep where possible and reschedule the work if it can't take the lock. Testing: existing selftest passes	N/A	More Details
CVE-2023-54307	In the Linux kernel, the following vulnerability has been resolved: ptp_qoriq: fix memory leak in probe() Smatch complains that: drivers/ptp/ptp_qoriq.c ptp_qoriq_probe() warn: 'base' from ioremap() not released. Fix this by revising the parameter from 'ptp_qoriq->base' to 'base'. This is only a bug if ptp_qoriq_init() returns on the first -ENODEV error path. For other error paths ptp_qoriq->base and base are the same. And this change makes the code more readable.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ALSA: ymfpci: Create card with device-managed snd_devm_card_new() snd_card_ymfpci_remove() was removed in commit c6e6bb5eab74 ("ALSA: ymfpci: Allocate resources with device-managed APIs"), but the call to snd_card_new() was not replaced with snd_devm_card_new(). Since there was no longer a call to snd_card_free, unloading the module would eventually result in Oops: [697561.532887] BUG: unable to handle page fault for address: ffffffffc0924480 [697561.532893] #PF: supervisor read access in kernel mode [697561.532896] #PF: error_code(0x0000) - not-present page [697561.532899] PGD ae1e15067 P4D ae1e15067 PUD ae1e17067 PMD 11a8f5067 PTE 0 [697561.532905] Oops: 0000 [#1] PREEMPT SMP NOPTI [697561.532909] CPU: 21 PID: 5080 Comm: wireplumber Tainted: G W OE 6.2.7 #1 [697561.532914] Hardware name: System manufacturer System Product Name/TUF GAMING X570-PLUS, BIOS 4408 10/28/2022 [697561.532916] RIP: 0010:try_module_get.part.0+0x1a/0xe0 [697561.532924] Code: 90 90 90 90 90 90 90 90 90 90 90 90 90 90 90 90 0f 1f 44 00 00 55 48 89 e5 41 55 41 54 49 89 fc bf 01 00 00 00 e8 56 3c f8 ff <41> 83 3c 24 02 0f 84 96 00 00 00 41 8b 84 24 30 03 00 00 85 c0 0f [697561.532927] RSP: 0018:ffffbe9b858c3bd8 EFLAGS: 00010246 [697561.532930] RAX: ffff9815d14f1900 RBX: ffff9815c14e6000 RCX: 0000000000000000 [697561.532933] RDX: 0000000000000000 RSI: ffffffffc055092c RDI:		

CVE-2023-54308	ffffffffb3778c1a [697561.532935] RBP: ffffbe9b858c3be8 R08: 0000000000000040 R09: ffff981a1a741380 [697561.532937] R10: ffffbe9b858c3c80 R11: 00000009d56533a6 R12: ffffffffc0924480 [697561.532939] R13: ffff9823439d8500 R14: 0000000000000025 R15: ffff9815cd109f80 [697561.532942] FS: 00007f13084f1f80(0000) GS:ffff9824aef40000(0000) knlGS:0000000000000000 [697561.532945] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [697561.532947] CR2: ffffffffc0924480 CR3: 0000000145344000 CR4: 0000000000350ee0 [697561.532949] Call Trace: [697561.532951] <TASK> [697561.532955] try_module_get+0x13/0x30 [697561.532960] snd_ctl_open+0x61/0x1c0 [snd] [697561.532976] snd_open+0xb4/0x1e0 [snd] [697561.532989] chrdev_open+0xc7/0x240 [697561.532995] ? fsnotify_perm.part.0+0x6e/0x160 [697561.533000] ? __pfx_chrdev_open+0x10/0x10 [697561.533005] do_dentry_open+0x169/0x440 [697561.533009] vfs_open+0x2d/0x40 [697561.533012] path_openat+0xa9d/0x10d0 [697561.533017] ? debug_smp_processor_id+0x17/0x20 [697561.533022] ? trigger_load_balance+0x65/0x370 [697561.533026] do_filp_open+0xb2/0x160 [697561.533032] ? _raw_spin_unlock+0x19/0x40 [697561.533036] ? alloc_fd+0xa9/0x190 [697561.533040] do_sys_openat2+0x9f/0x160 [697561.533044] __x64_sys_openat+0x55/0x90 [697561.533048] do_syscall_64+0x3b/0x90 [697561.533052] entry_SYSCALL_64_after_hwframe+0x72/0xdc [697561.533056] RIP: 0033:0x7f1308a40db4 [697561.533059] Code: 24 20 eb 8f 66 90 44 89 54 24 0c e8 46 68 f8 ff 44 8b 54 24 0c 44 89 e2 48 89 ee 41 89 c0 bf 9c ff ff ff b8 01 01 00 00 0f 05 <48> 3d 00 f0 ff ff 77 32 44 89 c7 89 44 24 0c e8 78 68 f8 ff 8b 44 [697561.533062] RSP: 002b:00007ffcce664450 EFLAGS: 00000293 ORIG_RAX: 0000000000000101 [697561.533066] RAX: ffffffffffffffda RBX: 0000000000000003 RCX: 00007f1308a40db4 [697561.533068] RDX: 0000000000080000 RSI: 00007ffcce664690 RDI: 00000000ffffff9c [697561.533070] RBP: 00007ffcce664690 R08: 0000000000000000 R09: 0000000000000012 [697561.533072] R10: 0000000000000000 R11: 0000000000000293 R12: 0000000000080000 [697561.533074] R13: 00007f13054b069b R14: 0000565209f83200 R15: 0000000000000000 [697561.533078] </TASK>	N/A	More Details
CVE-2023-54309	In the Linux kernel, the following vulnerability has been resolved: tpm: tpm_vtpm_proxy: fix a race condition in /dev/vtpmx creation /dev/vtpmx is made visible before 'workqueue' is initialized, which can lead to a memory corruption in the worst case scenario. Address this by initializing 'workqueue' as the very first step of the driver initialization.	N/A	More Details
CVE-2023-54310	In the Linux kernel, the following vulnerability has been resolved: scsi: message: mptlan: Fix use after free bug in mptlan_remove() due to race condition mptlan_probe() calls mpt_register_lan_device() which initializes the &priv->post_buckets_task workqueue. A call to mpt_lan_wake_post_buckets_task() will subsequently start the work. During driver unload in mptlan_remove() the following race may occur: CPU0 CPU1 |mpt_lan_post_receive_buckets_work() mptlan_remove() | free_netdev() | kfree(dev); | | dev->mtu | //use Fix this by finishing the work prior to cleaning up in mptlan_remove(). [mkp: we really should remove mptlan instead of attempting to fix it]	N/A	More Details
CVE-2023-54311	In the Linux kernel, the following vulnerability has been resolved: ext4: fix deadlock when converting an inline directory in nojournal mode In no journal mode, ext4_finish_convert_inline_dir() can self-deadlock by calling ext4_handle_dirty_dirblock() when it already has taken the directory lock. There is a similar self-deadlock in ext4_incvert_inline_data_nolock() for data files which we'll fix at the same time. A simple reproducer demonstrating the problem: mke2fs -Fq -t ext2 -O inline_data -b 4k /dev/vdc 64 mount -t ext4 -o dirsync /dev/vdc /vdc cd /vdc mkdir file0 cd file0 touch file0 touch file1 attr -s BurnSpaceInEA -V abcde . touch supercalifragilisticexpialidocious	N/A	More Details
CVE-2023-54313	In the Linux kernel, the following vulnerability has been resolved: ovl: fix null pointer dereference in ovl_get_acl_rcu() Following process: P1 P2 path_openat link_path_walk may_lookup inode_permission(rcu) ovl_permission acl_permission_check check_acl get_cached_acl_rcu ovl_get_inode_acl realinode = ovl_inode_real(ovl_inode) drop_cache __dentry_kill(ovl_dentry) iput(ovl_inode) ovl_destroy_inode(ovl_inode) dput(oi->__upperdentry) dentry_kill(upperdentry) dentry_unlink_inode upperdentry->d_inode = NULL ovl_inode_upper upperdentry = ovl_i_dentry_upper(ovl_inode) d_inode(upperdentry) // returns NULL IS_POSIXACL(realinode) // NULL pointer dereference , will trigger an null pointer dereference at realinode: [205.472797] BUG: kernel NULL pointer dereference, address: 0000000000000028 [205.476701] CPU: 2 PID: 2713 Comm: ls Not tainted 6.3.0-12064-g2edfa098e750-dirty #1216 [205.478754] RIP: 0010:do_ovl_get_acl+0x5d/0x300 [205.489584] Call Trace: [205.489812] <TASK> [205.490014] ovl_get_inode_acl+0x26/0x30 [205.490466] get_cached_acl_rcu+0x61/0xa0 [205.490908] generic_permission+0x1bf/0x4e0 [205.491447] ovl_permission+0x79/0x1b0 [205.491917] inode_permission+0x15e/0x2c0 [205.492425] link_path_walk+0x115/0x550 [205.493311] path_lookupat.isra.0+0xb2/0x200 [205.493803] filename_lookup+0xda/0x240 [205.495747] vfs_fstatat+0x7b/0xb0 Fetch a reproducer in [Link]. Use the helper ovl_i_path_realinode() to get realinode and then do non-nullptr checking.	N/A	More Details
CVE-2023-54282	In the Linux kernel, the following vulnerability has been resolved: media: tuners: qt1010: replace BUG_ON with a regular error BUG_ON is unnecessary here, and in addition it confuses smatch. Replacing this with an error return help resolve this smatch warning: drivers/media/tuners/qt1010.c:350 qt1010_init() error: buffer overflow 'i2c_data' 34 <= 34	N/A	More Details
CVE-2023-54314	In the Linux kernel, the following vulnerability has been resolved: media: af9005: Fix null-ptr-deref in af9005_i2c_xfer In af9005_i2c_xfer, msg is controlled by user. When msg[i].buf is null and msg[i].len is zero, former checks on msg[i].buf would be passed. Malicious data finally reach af9005_i2c_xfer. If accessing msg[i].buf[0] without sanity check, null ptr deref would happen. We add check on msg[i].len to prevent crash. Similar commit: commit 0ed554fd769a ("media: dvb-usb: az6027: fix null-ptr-deref in az6027_i2c_xfer()")	N/A	More Details
CVE-2023-54315	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv/sriov: perform null check on iov before dereferencing iov Currently pointer iov is being dereferenced before the null check of iov which can lead to null pointer dereference errors. Fix this by moving the iov null check before the dereferencing. Detected using cppcheck static analysis: linux/arch/powerpc/platforms/powernv/pci-sriov.c:597:12: warning: Either the condition '!iov' is redundant or there is possible null pointer dereference: iov. [nullPointerRedundantCheck] num_vfs = iov->num_vfs; ^	N/A	More Details
CVE-2023-54316	In the Linux kernel, the following vulnerability has been resolved: refscale: Fix uninitalized use of wait_queue_head_t Running the refscale test occasionally crashes the kernel with the following error: [8569.952896] BUG: unable to handle page fault for address: fffffffffffffe8 [8569.952900] #PF: supervisor read access in kernel mode [8569.952902] #PF: error_code(0x0000) - not-present page [8569.952904] PGD c4b048067 P4D c4b049067 PUD c4b04b067 PMD 0 [8569.952910] Oops: 0000 [#1] PREEMPT_RT SMP NOPTI [8569.952916] Hardware name: Dell Inc. PowerEdge R750/0WMWCR, BIOS 1.2.4 05/28/2021 [8569.952917] RIP: 0010:prepare_to_wait_event+0x101/0x190 : [8569.952940] Call Trace: [8569.952941] <TASK> [8569.952944] ref_scale_reader+0x380/0x4a0 [refscale] [8569.952959] kthread+0x10e/0x130 [8569.952966] ret_from_fork+0x1f/0x30 [8569.952973] </TASK> The likely cause is that init_waitqueue_head() is called after the call to the torture_create_kthread() function that creates the ref_scale_reader kthread. Although this init_waitqueue_head() call will very likely complete before this kthread is created and starts running, it is possible that the calling kthread will be delayed between the calls to torture_create_kthread() and init_waitqueue_head(). In this case, the new kthread will use the waitqueue head before it is properly initialized, which is not good for the kernel's health and well-being. The above crash happened here: static inline void __add_wait_queue(...) { : if (!(wq->flags & WQ_FLAG_PRIORITY)) <=== Crash here The offset of flags from list_head entry in wait_queue_entry is -0x18. If reader_tasks[i].wq.head.next is NULL as allocated reader_task structure is zero initialized, the instruction will try to access address	N/A	More Details

	0xfffffffffffffe8, which is exactly the fault address listed above. This commit therefore invokes init_waitqueue_head() before creating the kthread.		
CVE-2023-54317	In the Linux kernel, the following vulnerability has been resolved: dm flakey: don't corrupt the zero page When we need to zero some range on a block device, the function __blkdev_issue_zero_pages submits a write bio with the bio vector pointing to the zero page. If we use dm-flakey with corrupt bio writes option, it will corrupt the content of the zero page which results in crashes of various userspace programs. Glibc assumes that memory returned by mmap is zeroed and it uses it for calloc implementation; if the newly mapped memory is not zeroed, calloc will return non-zeroed memory. Fix this bug by testing if the page is equal to ZERO_PAGE(0) and avoiding the corruption in this case.	N/A	More Details
CVE-2023-54318	In the Linux kernel, the following vulnerability has been resolved: net/smc: use smc_lgr_list.lock to protect smc_lgr_list.list iterate in smcr_port_add While doing smcr_port_add, there maybe linkgroup add into or delete from smc_lgr_list.list at the same time, which may result kernel crash. So, use smc_lgr_list.lock to protect smc_lgr_list.list iterate in smcr_port_add. The crash calltrace show below: BUG: kernel NULL pointer dereference, address: 0000000000000000 PGD 0 P4D 0 Oops: 0000 [#1] SMP NOPTI CPU: 0 PID: 559726 Comm: kworker/0:92 Kdump: loaded Tainted: G Hardware name: Alibaba Cloud Alibaba Cloud ECS, BIOS 449e491 04/01/2014 Workqueue: events smc_ib_port_event_work [smc] RIP: 0010:smcr_port_add+0xa6/0xf0 [smc] RSP: 0000:ffffa5a2c8f67de0 EFLAGS: 00010297 RAX: 0000000000000001 RBX: ffff9935e0650000 RCX: 0000000000000000 RDX: 0000000000000010 RSI: ffff9935e0654290 RDI: ffff9935c8560000 RBP: 0000000000000000 R08: 0000000000000000 R09: ffff9934c0401918 R10: 0000000000000000 R11: ffffffffb4a5c278 R12: ffff99364029aae4 R13: ffff99364029aa00 R14: 00000000ffffffed R15: ffff99364029ab08 FS: 0000000000000000(0000) GS:ffff994380600000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000000 CR3: 0000000f06a10003 CR4: 0000000002770ef0 PKRU: 55555554 Call Trace: smc_ib_port_event_work+0x18f/0x380 [smc] process_one_work+0x19b/0x340 worker_thread+0x30/0x370 ? process_one_work+0x340/0x340 kthread+0x114/0x130 ? __kthread_cancel_work+0x50/0x50 ret_from_fork+0x1f/0x30	N/A	More Details
CVE-2023-54319	In the Linux kernel, the following vulnerability has been resolved: pinctrl: at91-pio4: check return value of devm_kasprintf() devm_kasprintf() returns a pointer to dynamically allocated memory. Pointer could be NULL in case allocation fails. Check pointer validity. Identified with coccinelle (kmerr.cocci script). Depends-on: 1c4e5c470a56 ("pinctrl: at91: use devm_kasprintf() to avoid potential leaks") Depends-on: 5a8f9cf269e8 ("pinctrl: at91-pio4: use proper format specifier for unsigned int")	N/A	More Details
CVE-2023-54320	In the Linux kernel, the following vulnerability has been resolved: platform/x86/amd: pmc: Fix memory leak in amd_pmc_stb_debugfs_open_v2() Function amd_pmc_stb_debugfs_open_v2() may be called when the STB debug mechanism enabled. When amd_pmc_send_cmd() fails, the 'buf' needs to be released.	N/A	More Details
CVE-2023-54321	In the Linux kernel, the following vulnerability has been resolved: driver core: fix potential null-ptr-deref in device_add() I got the following null-ptr-deref report while doing fault injection test: BUG: kernel NULL pointer dereference, address: 0000000000000058 CPU: 2 PID: 278 Comm: 37-i2c-ds2482 Tainted: G B W N 6.1.0-rc3+ RIP: 0010:klist_put+0x2d/0xd0 Call Trace: <TASK> klist_remove+0xf1/0x1c0 device_release_driver_internal+0x196/0x210 bus_remove_device+0x1bd/0x240 device_add+0xd3d/0x1100 w1_add_master_device+0x476/0x490 [wire] ds2482_probe+0x303/0x3e0 [ds2482] This is how it happened: w1_alloc_dev() // The dev->driver is set to w1_master_driver. memcpy(&dev->dev, device, sizeof(struct device)); device_add() bus_add_device() dpm_sysfs_add() // It fails, calls bus_remove_device. // error path bus_remove_device() // The dev->driver is not null, but driver is not bound. __device_release_driver() klist_remove(&dev->p->knode_driver) <-- It causes null-ptr-deref. // normal path bus_probe_device() // It's not called yet. device_bind_driver() If dev->driver is set, in the error path after calling bus_add_device() in device_add(), bus_remove_device() is called, then the device will be detached from driver. But device_bind_driver() is not called yet, so it causes null-ptr-deref while access the 'knode_driver'. To fix this, set dev->driver to null in the error path before calling bus_remove_device().	N/A	More Details
CVE-2023-54303	In the Linux kernel, the following vulnerability has been resolved: bpf: Disable preemption in bpf_perf_event_output The nesting protection in bpf_perf_event_output relies on disabled preemption, which is guaranteed for kprobes and tracepoints. However bpf_perf_event_output can be also called from uprobes context through bpf_prog_run_array_sleepable function which disables migration, but keeps preemption enabled. This can cause task to be preempted by another one inside the nesting protection and lead eventually to two tasks using same perf_sample_data buffer and cause crashes like: kernel tried to execute NX-protected page - exploit attempt? (uid: 0) BUG: unable to handle page fault for address: ffffffff82be3eea ... Call Trace: ? __die+0x1f/0x70 ? page_fault_oops+0x176/0x4d0 ? exc_page_fault+0x132/0x230 ? asm_exc_page_fault+0x22/0x30 ? perf_output_sample+0x12b/0x910 ? perf_event_output+0xd0/0x1d0 ? bpf_perf_event_output+0x162/0x1d0 ? bpf_prog_c6271286d9a4c938_krava1+0x76/0x87 ? __uprobe_perf_func+0x12b/0x540 ? uprobe_dispatcher+0x2c4/0x430 ? uprobe_notify_resume+0x2da/0xce0 ? atomic_notifier_call_chain+0x7b/0x110 ? exit_to_user_mode_prepare+0x13e/0x290 ? irqentry_exit_to_user_mode+0x5/0x30 ? asm_exc_int3+0x35/0x40 Fixing this by disabling preemption in bpf_perf_event_output.	N/A	More Details
CVE-2023-54302	In the Linux kernel, the following vulnerability has been resolved: RDMA/irdma: Fix data race on CQP completion stats CQP completion statistics is read lockesly in irdma_wait_event and irdma_check_cqp_progress while it can be updated in the completion thread irdma_sc_ccq_get_cqe_info on another CPU as KCSAN reports. Make completion statistics an atomic variable to reflect coherent updates to it. This will also avoid load/store tearing logic bug potentially possible by compiler optimizations. [77346.170861] BUG: KCSAN: data-race in irdma_handle_cqp_op [irdma] / irdma_sc_ccq_get_cqe_info [irdma] [77346.171383] write to 0xffff8a3250b108e0 of 8 bytes by task 9544 on cpu 4: [77346.171483] irdma_sc_ccq_get_cqe_info+0x27a/0x370 [irdma] [77346.171658] irdma_cqp_ce_handler+0x164/0x270 [irdma] [77346.171835] cqp_compl_worker+0x1b/0x20 [irdma] [77346.172009] process_one_work+0x4d1/0xa40 [77346.172024] worker_thread+0x319/0x700 [77346.172037] kthread+0x180/0x1b0 [77346.172054] ret_from_fork+0x22/0x30 [77346.172136] read to 0xffff8a3250b108e0 of 8 bytes by task 9838 on cpu 2: [77346.172234] irdma_handle_cqp_op+0xf4/0x4b0 [irdma] [77346.172413] irdma_cqp_aeq_cmd+0x75/0xa0 [irdma] [77346.172592] irdma_create_aeq+0x390/0x45a [irdma] [77346.172769] irdma_rt_init_hw.cold+0x212/0x85d [irdma] [77346.172944] irdma_probe+0x54f/0x620 [irdma] [77346.173122] auxiliary_bus_probe+0x66/0xa0 [77346.173137] really_probe+0x140/0x540 [77346.173154] __driver_probe_device+0xc7/0x220 [77346.173173] driver_probe_device+0x5f/0x140 [77346.173190] __driver_attach+0xf0/0x2c0 [77346.173208] bus_for_each_dev+0xa8/0xf0 [77346.173225] driver_attach+0x29/0x30 [77346.173240] bus_add_driver+0x29c/0x2f0 [77346.173255] driver_register+0x10f/0x1a0 [77346.173272] __auxiliary_driver_register+0xbc/0x140 [77346.173287] irdma_init_module+0x55/0x1000 [irdma] [77346.173460] do_one_initcall+0x7d/0x410 [77346.173475] do_init_module+0x81/0x2c0 [77346.173491] load_module+0x1232/0x12c0 [77346.173506] __do_sys_finit_module+0x101/0x180 [77346.173522] __x64_sys_finit_module+0x3c/0x50 [77346.173538] do_syscall_64+0x39/0x90 [77346.173553] entry_SYSCALL_64_after_hwframe+0x63/0xcd [77346.173634] value changed: 0x0000000000000094 -> 0x0000000000000095	N/A	More Details
CVE-2023-54301	In the Linux kernel, the following vulnerability has been resolved: serial: 8250_bcm7271: fix leak in `brcmuart_probe` Smatch reports: drivers/tty/serial/8250/8250_bcm7271.c:1120 brcmuart_probe() warn: 'baud_mux_clk' from clk_prepare_enable() not released on lines: 1032. The issue is fixed by using a managed clock.	N/A	More Details

CVE-2023-54300	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: avoid referencing uninit memory in ath9k_wmi_ctrl_rx For the reasons also described in commit b383e8abed41 ("wifi: ath9k: avoid uninit memory read in ath9k_htc_rx_msg()"), ath9k_htc_rx_msg() should validate pkt_len before accessing the SKB. For example, the obtained SKB may have been badly constructed with pkt_len = 8. In this case, the SKB can only contain a valid htc_frame_hdr but after being processed in ath9k_htc_rx_msg() and passed to ath9k_wmi_ctrl_rx() endpoint RX handler, it is expected to have a WMI command header which should be located inside its data payload. Implement sanity checking inside ath9k_wmi_ctrl_rx(). Otherwise, uninit memory can be referenced. Tested on Qualcomm Atheros Communications AR9271 802.11n . Found by Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2023-54283	In the Linux kernel, the following vulnerability has been resolved: bpf: Address KCSAN report on bpf_lru_list KCSAN reported a data-race when accessing node->ref. Although node->ref does not have to be accurate, take this chance to use a more common READ_ONCE() and WRITE_ONCE() pattern instead of data_race(). There is an existing bpf_lru_node_is_ref() and bpf_lru_node_set_ref(). This patch also adds bpf_lru_node_clear_ref() to do the WRITE_ONCE(node->ref, 0) also. ========== BUG: KCSAN: data-race in __bpf_lru_list_rotate / __htab_lru_percpu_map_update_elem write to 0xffff888137038deb of 1 bytes by task 11240 on cpu 1: __bpf_lru_node_move kernel/bpf/bpf_lru_list.c:113 [inline] __bpf_lru_list_rotate_active kernel/bpf/bpf_lru_list.c:149 [inline] __bpf_lru_list_rotate+0x1bf/0x750 kernel/bpf/bpf_lru_list.c:240 bpf_lru_list_pop_free_to_local kernel/bpf/bpf_lru_list.c:329 [inline] bpf_common_lru_pop_free kernel/bpf/bpf_lru_list.c:447 [inline] bpf_lru_pop_free+0x638/0xe20 kernel/bpf/bpf_lru_list.c:499 prealloc_lru_pop kernel/bpf/hashtab.c:290 [inline] __htab_lru_percpu_map_update_elem+0xe7/0x820 kernel/bpf/hashtab.c:1316 bpf_percpu_hash_update+0x5e/0x90 kernel/bpf/hashtab.c:2313 bpf_map_update_value+0x2a9/0x370 kernel/bpf/syscall.c:200 generic_map_update_batch+0x3ae/0x4f0 kernel/bpf/syscall.c:1687 bpf_map_do_batch+0x2d9/0x3d0 kernel/bpf/syscall.c:4534 __sys_bpf+0x338/0x810 __do_sys_bpf kernel/bpf/syscall.c:5096 [inline] __se_sys_bpf kernel/bpf/syscall.c:5094 [inline] __x64_sys_bpf+0x43/0x50 kernel/bpf/syscall.c:5094 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd read to 0xffff888137038deb of 1 bytes by task 11241 on cpu 0: bpf_lru_node_set_ref kernel/bpf/bpf_lru_list.h:70 [inline] __htab_lru_percpu_map_update_elem+0x2f1/0x820 kernel/bpf/hashtab.c:1332 bpf_percpu_hash_update+0x5e/0x90 kernel/bpf/hashtab.c:2313 bpf_map_update_value+0x2a9/0x370 kernel/bpf/syscall.c:200 generic_map_update_batch+0x3ae/0x4f0 kernel/bpf/syscall.c:1687 bpf_map_do_batch+0x2d9/0x3d0 kernel/bpf/syscall.c:4534 __sys_bpf+0x338/0x810 __do_sys_bpf kernel/bpf/syscall.c:5096 [inline] __se_sys_bpf kernel/bpf/syscall.c:5094 [inline] __x64_sys_bpf+0x43/0x50 kernel/bpf/syscall.c:5094 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd value changed: 0x01 -> 0x00 Reported by Kernel Concurrency Sanitizer on: CPU: 0 PID: 11241 Comm: syz-executor.3 Not tainted 6.3.0-rc7-syzkaller-00136-g6a66fdd29ea1 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 03/30/2023 ==========	N/A	More Details
CVE-2023-54284	In the Linux kernel, the following vulnerability has been resolved: media: av7110: prevent underflow in write_ts_to_decoder() The buf[4] value comes from the user via ts_play(). It is a value in the u8 range. The final length we pass to av7110_ipack_instant_repack() is "len - (buf[4] + 1) - 4" so add a check to ensure that the length is not negative. It's not clear that passing a negative len value does anything bad necessarily, but it's not best practice. With the new bounds checking the "if (!len)" condition is no longer possible or required so remove that.	N/A	More Details
CVE-2023-54285	In the Linux kernel, the following vulnerability has been resolved: iomap: Fix possible overflow condition in iomap_write_delalloc_scan folio_next_index() returns an unsigned long value which left shifted by PAGE_SHIFT could possibly cause an overflow on 32-bit system. Instead use folio_pos(folio) + folio_size(folio), which does this correctly.	N/A	More Details
CVE-2023-54286	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: dvm: Fix memcpy: detected field-spanning write backtrace A received TKIP key may be up to 32 bytes because it may contain MIC rx/tx keys too. These are not used by iwl and copying these over overflows the iwl_keyinfo.key field. Add a check to not copy more data to iwl_keyinfo.key then will fit. This fixes backtraces like this one: memcpy: detected field-spanning write (size 32) of single field "sta_cmd.key.key" at drivers/net/wireless/intel/iwlwifi/dvm/sta.c:1103 (size 16) WARNING: CPU: 1 PID: 946 at drivers/net/wireless/intel/iwlwifi/dvm/sta.c:1103 iwlagn_send_sta_key+0x375/0x390 [iwldvm] <snip> Hardware name: Dell Inc. Latitude E6430/0H3MT5, BIOS A21 05/08/2017 RIP: 0010:iwlagn_send_sta_key+0x375/0x390 [iwldvm] <snip> Call Trace: <TASK> iwl_set_dynamic_key+0x1f0/0x220 [iwldvm] iwlagn_mac_set_key+0x1e4/0x280 [iwldvm] drv_set_key+0xa4/0x1b0 [mac80211] ieee80211_key_enable_hw_accel+0xa8/0x2d0 [mac80211] ieee80211_key_replace+0x22d/0x8e0 [mac80211] <snip>	N/A	More Details
CVE-2023-54287	In the Linux kernel, the following vulnerability has been resolved: tty: serial: imx: disable Ageing Timer interrupt request irq There maybe pending USR interrupt before requesting irq, however uart_add_one_port has not executed, so there will be kernel panic: [0.795668] Unable to handle kernel NULL pointer dereference at virtual addre ss 0000000000000080 [0.802701] Mem abort info: [0.805367] ESR = 0x0000000096000004 [0.808950] EC = 0x25: DABT (current EL), IL = 32 bits [0.814033] SET = 0, FnV = 0 [0.816950] EA = 0, S1PTW = 0 [0.819950] FSC = 0x04: level 0 translation fault [0.824617] Data abort info: [0.827367] ISV = 0, ISS = 0x00000004 [0.831033] CM = 0, WnR = 0 [0.833866] [0000000000000080] user address but active_mm is swapper [0.839951] Internal error: Oops: 0000000096000004 [#1] PREEMPT SMP [0.845953] Modules linked in: [0.848869] CPU: 0 PID: 1 Comm: swapper/0 Not tainted 6.1.1+g56321e101aca #1 [0.855617] Hardware name: Freescale i.MX8MP EVK (DT) [0.860452] pstate: 000000c5 (nzcv daIF -PAN -UAO -TCO -DIT -SSBS BTYPE=--) [0.867117] pc : __imx_uart_rxint.constprop.0+0x11c/0x2c0 [0.872283] lr : imx_uart_int+0xf8/0x1ec The issue only happends in the inmate linux when Jailhouse hypervisor enabled. The test procedure is: while true; do jailhouse enable imx8mp.cell jailhouse cell linux xxxx sleep 10 jailhouse cell destroy 1 jailhouse disable sleep 5 done And during the upper test, press keys to the 2nd linux console. When `jailhouse cell destroy 1`, the 2nd linux has no chance to put the uart to a quiese state, so USR1/2 may has pending interrupts. Then when `jailhosue cell linux xx` to start 2nd linux again, the issue trigger. In order to disable irqs before requesting them, both UCR1 and UCR2 irqs should be disabled, so here fix that, disable the Ageing Timer interrupt in UCR2 as UCR1 does.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fortify the spinlock against deadlock by interrupt In the function ieee80211_tx_dequeue() there is a particular locking sequence: begin: spin_lock(&local->queue_stop_reason_lock); q_stopped = local->queue_stop_reasons[q]; spin_unlock(&local->queue_stop_reason_lock); However small the chance (increased by ftracetest), an asynchronous interrupt can occur in between of spin_lock() and spin_unlock(), and the interrupt routine will attempt to lock the same &local->queue_stop_reason_lock again. This will cause a costly reset of the CPU and the wifi device or an altogether hang in the single CPU and single core scenario. The only remaining spin_lock(&local->queue_stop_reason_lock) that did not disable interrupts was patched, which should prevent any deadlocks on the same CPU/core and the same wifi device. This is the probable trace of the deadlock: kernel: ========== kernel: WARNING: inconsistent lock state kernel: 6.3.0-rc6-mt-20230401-00001-gf86822a1170f #4 Tainted: G W kernel: ---------- kernel: inconsistent {IN-SOFTIRQ-W} -> {SOFTIRQ-ON-W} usage. kernel: kworker/5:0/25656 [HC0[0]:SC0[0]:HE1:SE1] takes: kernel: ffff9d6190779478 (&local->queue_stop_reason_lock){+.?.}-{2:2}, at: return_to_handler+0x0/0x40 kernel: {IN-SOFTIRQ-W} state was registered at: kernel: lock_acquire+0xc7/0x2d0 kernel: _raw_spin_lock+0x36/0x50 kernel: ieee80211_tx_dequeue+0xb4/0x1330 [mac80211] kernel:		

CVE-2023-54288	iwl_mvm_mac_itxq_xmit+0xae/0x210 [iwlmvm] kernel: iwl_mvm_mac_wake_tx_queue+0x2d/0xd0 [iwlmvm] kernel: ieee80211_queue_skb+0x450/0x730 [mac80211] kernel: __ieee80211_xmit_fast.constprop.66+0x834/0xa50 [mac80211] kernel: __ieee80211_subif_start_xmit+0x217/0x530 [mac80211] kernel: ieee80211_subif_start_xmit+0x60/0x580 [mac80211] kernel: dev_hard_start_xmit+0xb5/0x260 kernel: __dev_queue_xmit+0xdbe/0x1200 kernel: neigh_resolve_output+0x166/0x260 kernel: ip_finish_output2+0x216/0xb80 kernel: __ip_finish_output+0x2a4/0x4d0 kernel: ip_finish_output+0x2d/0xd0 kernel: ip_output+0x82/0x2b0 kernel: ip_local_out+0xec/0x110 kernel: igmpv3_sendpack+0x5c/0x90 kernel: igmp_ifc_timer_expire+0x26e/0x4e0 kernel: call_timer_fn+0xa5/0x230 kernel: run_timer_softirq+0x27f/0x550 kernel: __do_softirq+0xb4/0x3a4 kernel: irq_exit_rcu+0x9b/0xc0 kernel: sysvec_apic_timer_interrupt+0x80/0xa0 kernel: asm_sysvec_apic_timer_interrupt+0x1f/0x30 kernel: _raw_spin_unlock_irqrestore+0x3f/0x70 kernel: free_to_partial_list+0x3d6/0x590 kernel: __slab_free+0x1b7/0x310 kernel: kmem_cache_free+0x52d/0x550 kernel: putname+0x5d/0x70 kernel: do_sys_openat2+0x1d7/0x310 kernel: do_sys_open+0x51/0x80 kernel: __x64_sys_openat+0x24/0x30 kernel: do_syscall_64+0x5c/0x90 kernel: entry_SYSCALL_64_after_hwframe+0x72/0xdc kernel: irq event stamp: 5120729 kernel: hardirqs last enabled at (5120729): [<ffffffff9d149936>] trace_graph_return+0xd6/0x120 kernel: hardirqs last disabled at (5120728): [<ffffffff9d149950>] trace_graph_return+0xf0/0x120 kernel: softirqs last enabled at (5069900): [<ffffffff9cf65b60>] return_to_handler+0x0/0x40 kernel: softirqs last disabled at (5067555): [<ffffffff9cf65b60>] return_to_handler+0x0/0x40 kernel: other info that might help us debug this: kernel: Possible unsafe locking scenario: kernel: CPU0 kernel: ---- kernel: lock(&local->queue_stop_reason_lock); kernel: <Interrupt> kernel: lock(&local->queue_stop_reason_lock); kernel: *** DEADLOCK *** kernel: 8 locks held by kworker/5:0/25656: kernel: #0: ffff9d618009d138 ((wq_completion)events_freezable){+.+.}-{0:0}, at: process_one_work+0x1ca/0x530 kernel: #1: ffffb1ef4637fe68 ((work_completion)(&local->restart_work)){+.+.}-{0:0}, at: process_one_work+0x1ce/0x530 kernel: #2: ffffffff9f166548 (rtnl_mutex){+.+.}-{3:3}, at: return_to_handler+0x0/0x40 kernel: #3: ffff9d619 ---truncated---	N/A	More Details
CVE-2023-54289	In the Linux kernel, the following vulnerability has been resolved: scsi: qedf: Fix NULL dereference in error handling Smatch reported: drivers/scsi/qedf/qedf_main.c:3056 qedf_alloc_global_queues() warn: missing unwind goto? At this point in the function, nothing has been allocated so we can return directly. In particular the "qedf->global_queues" have not been allocated so calling qedf_free_global_queues() will lead to a NULL dereference when we check if (!gl[i]) and "gl" is NULL.	N/A	More Details
CVE-2023-54290	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2023-54291	In the Linux kernel, the following vulnerability has been resolved: vduse: fix NULL pointer dereference vduse_vdpa_set_vq_affinity callback can be called with NULL value as cpu_mask when deleting the vduse device. This patch resets virtqueue's IRQ affinity mask value to set all CPUs instead of dereferencing NULL cpu_mask. [4760.952149] BUG: kernel NULL pointer dereference, address: 0000000000000000 [4760.959110] #PF: supervisor read access in kernel mode [4760.964247] #PF: error_code(0x0000) - not-present page [4760.969385] PGD 0 P4D 0 [4760.971927] Oops: 0000 [#1] PREEMPT SMP PTI [4760.976112] CPU: 13 PID: 2346 Comm: vdpa Not tainted 6.4.0-rc6+ #4 [4760.982291] Hardware name: Dell Inc. PowerEdge R640/0W23H8, BIOS 2.8.1 06/26/2020 [4760.989769] RIP: 0010:memcpy_orig+0xc5/0x130 [4760.994049] Code: 16 f8 4c 89 07 4c 89 4f 08 4c 89 54 17 f0 4c 89 5c 17 f8 c3 cc cc cc cc 66 66 2e 0f 1f 84 00 00 00 00 00 66 90 83 fa 08 72 1b <4c> 8b 06 4c 8b 4c 16 f8 4c 89 07 4c 89 4c 17 f8 c3 cc cc cc cc 66 [4761.012793] RSP: 0018:ffffb1d565abb830 EFLAGS: 00010246 [4761.018020] RAX: ffff9f4bf6b27898 RBX: ffff9f4be23969c0 RCX: ffff9f4bcadf6400 [4761.025152] RDX: 0000000000000008 RSI: 0000000000000000 RDI: ffff9f4bf6b27898 [4761.032286] RBP: 0000000000000000 R08: 0000000000000008 R09: 0000000000000000 [4761.039416] R10: 0000000000000000 R11: 0000000000000600 R12: 0000000000000000 [4761.046549] R13: 0000000000000000 R14: 0000000000000080 R15: ffffb1d565abbb10 [4761.053680] FS: 00007f64c2ec2740(0000) GS:ffff9f635f980000(0000) knlGS:0000000000000000 [4761.061765] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [4761.067513] CR2: 0000000000000000 CR3: 0000001875270006 CR4: 00000000007706e0 [4761.074645] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [4761.081775] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [4761.088909] PKRU: 55555554 [4761.091620] Call Trace: [4761.094074] <TASK> [4761.096180] ? __die+0x1f/0x70 [4761.099238] ? page_fault_oops+0x171/0x4f0 [4761.103340] ? exc_page_fault+0x7b/0x180 [4761.107265] ? asm_exc_page_fault+0x22/0x30 [4761.111460] ? memcpy_orig+0xc5/0x130 [4761.115126] vduse_vdpa_set_vq_affinity+0x3e/0x50 [vduse] [4761.120533] virtnet_clean_affinity.part.0+0x3d/0x90 [virtio_net] [4761.126635] remove_vq_common+0x1a4/0x250 [virtio_net] [4761.131781] virtnet_remove+0x5d/0x70 [virtio_net] [4761.136580] virtio_dev_remove+0x3a/0x90 [4761.140509] device_release_driver_internal+0x19b/0x200 [4761.145742] bus_remove_device+0xc2/0x130 [4761.149755] device_del+0x158/0x3e0 [4761.153245] ? kernfs_find_ns+0x35/0xc0 [4761.157086] device_unregister+0x13/0x60 [4761.161010] unregister_virtio_device+0x11/0x20 [4761.165543] device_release_driver_internal+0x19b/0x200 [4761.170770] bus_remove_device+0xc2/0x130 [4761.174782] device_del+0x158/0x3e0 [4761.178276] ? __pfx_vdpa_name_match+0x10/0x10 [vdpa] [4761.183336] device_unregister+0x13/0x60 [4761.187260] vdpa_nl_cmd_dev_del_set_doit+0x63/0xe0 [vdpa]	N/A	More Details
CVE-2023-54292	In the Linux kernel, the following vulnerability has been resolved: RDMA/irdma: Fix data race on CQP request done KCSAN detects a data race on cqp_request->request_done memory location which is accessed locklessly in irdma_handle_cqp_op while being updated in irdma_cqp_ce_handler. Annotate lockless intent with READ_ONCE/WRITE_ONCE to avoid any compiler optimizations like load fusing and/or KCSAN warning. [222808.417128] BUG: KCSAN: data-race in irdma_cqp_ce_handler [irdma] / irdma_wait_event [irdma] [222808.417532] write to 0xffff8e44107019dc of 1 bytes by task 29658 on cpu 5: [222808.417610] irdma_cqp_ce_handler+0x21e/0x270 [irdma] [222808.417725] cqp_compl_worker+0x1b/0x20 [irdma] [222808.417827] process_one_work+0x4d1/0xa40 [222808.417835] worker_thread+0x319/0x700 [222808.417842] kthread+0x180/0x1b0 [222808.417852] ret_from_fork+0x22/0x30 [222808.417918] read to 0xffff8e44107019dc of 1 bytes by task 29688 on cpu 1: [222808.417995] irdma_wait_event+0x1e2/0x2c0 [irdma] [222808.418099] irdma_handle_cqp_op+0xae/0x170 [irdma] [222808.418202] irdma_cqp_cq_destroy_cmd+0x70/0x90 [irdma] [222808.418308] irdma_puda_dele_rsrc+0x46d/0x4d0 [irdma] [222808.418411] irdma_rt_deinit_hw+0x179/0x1d0 [irdma] [222808.418514] irdma_ib_dealloc_device+0x11/0x40 [irdma] [222808.418618] ib_dealloc_device+0x2a/0x120 [ib_core] [222808.418823] __ib_unregister_device+0xde/0x100 [ib_core] [222808.418981] ib_unregister_device+0x22/0x40 [ib_core] [222808.419142] irdma_ib_unregister_device+0x70/0x90 [irdma] [222808.419248] i40iw_close+0x6f/0xc0 [irdma] [222808.419352] i40e_client_device_unregister+0x14a/0x180 [i40e] [222808.419450] i40iw_remove+0x21/0x30 [irdma] [222808.419554] auxiliary_bus_remove+0x31/0x50 [222808.419563] device_remove+0x69/0xb0 [222808.419572] device_release_driver_internal+0x293/0x360 [222808.419582] driver_detach+0x7c/0xf0 [222808.419592] bus_remove_driver+0x8c/0x150 [222808.419600] driver_unregister+0x45/0x70 [222808.419610] auxiliary_driver_unregister+0x16/0x30 [222808.419618] irdma_exit_module+0x18/0x1e [irdma] [222808.419733] __do_sys_delete_module.constprop.0+0x1e2/0x310 [222808.419745] __x64_sys_delete_module+0x1b/0x30 [222808.419755] do_syscall_64+0x39/0x90 [222808.419763] entry_SYSCALL_64_after_hwframe+0x63/0xcd [222808.419829] value changed: 0x01 -> 0x03	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: bcache: fixup btree_cache_wait list damage We get a kernel crash about "list_add corruption. next->prev should be prev (ffff9c801bc01210), but was ffff9c77b688237c. (next=ffffae586d8afe68)."		

CVE-2023-54293	crash> struct list_head 0xffff9c801bc01210 struct list_head { next = 0xffffae586d8afe68, prev = 0xffffae586d8afe68 } crash> struct list_head 0xffff9c77b688237c struct list_head { next = 0x0, prev = 0x0 } crash> struct list_head 0xffffae586d8afe68 struct list_head struct: invalid kernel virtual address: ffffae586d8afe68 type: "gdb_readmem_callback" Cannot access memory at address 0xffffae586d8afe68 [230469.019492] Call Trace: [230469.032041] prepare_to_wait+0x8a/0xb0 [230469.044363] ? bch_btree_keys_free+0x6c/0xc0 [escache] [230469.056533] mca_cannibalize_lock+0x72/0x90 [escache] [230469.068788] mca_alloc+0x2ae/0x450 [escache] [230469.080790] bch_btree_node_get+0x136/0x2d0 [escache] [230469.092681] bch_btree_check_thread+0x1e1/0x260 [escache] [230469.104382] ? finish_wait+0x80/0x80 [230469.115884] ? bch_btree_check_recurse+0x1a0/0x1a0 [escache] [230469.127259] kthread+0x112/0x130 [230469.138448] ? kthread_flush_work_fn+0x10/0x10 [230469.149477] ret_from_fork+0x35/0x40 bch_btree_check_thread() and bch_dirty_init_thread() may call mca_cannibalize() to cannibalize other cached btree nodes. Only one thread can do it at a time, so the op of other threads will be added to the btree_cache_wait list. We must call finish_wait() to remove op from btree_cache_wait before free it's memory address. Otherwise, the list will be damaged. Also should call bch_cannibalize_unlock() to release the btree_cache_alloc_lock and wake_up other waiters.	N/A	More Details
CVE-2023-54294	In the Linux kernel, the following vulnerability has been resolved: md/raid10: fix memleak of md thread In raid10_run(), if setup_conf() succeed and raid10_run() failed before setting 'mddev->thread', then in the error path 'conf->thread' is not freed. Fix the problem by setting 'mddev->thread' right after setup_conf().	N/A	More Details
CVE-2023-54295	In the Linux kernel, the following vulnerability has been resolved: mtd: spi-nor: Fix shift-out-of-bounds in spi_nor_set_erase_type spi_nor_set_erase_type() was used either to set or to mask out an erase type. When we used it to mask out an erase type a shift-out-of-bounds was hit: UBSAN: shift-out-of-bounds in drivers/mtd/spi-nor/core.c:2237:24 shift exponent 4294967295 is too large for 32-bit type 'int' The setting of the size_{shift, mask} and of the opcode are unnecessary when the erase size is zero, as throughout the code just the erase size is considered to determine whether an erase type is supported or not. Setting the opcode to 0xFF was wrong too as nobody guarantees that 0xFF is an unused opcode. Thus when masking out an erase type, just set the erase size to zero. This will fix the shift-out-of-bounds. [ta: refine changes, new commit message, fix compilation error]	N/A	More Details
CVE-2023-54296	In the Linux kernel, the following vulnerability has been resolved: KVM: SVM: Get source vCPUs from source VM for SEV-ES intrahost migration Fix a goof where KVM tries to grab source vCPUs from the destination VM when doing intrahost migration. Grabbing the wrong vCPU not only hoses the guest, it also crashes the host due to the VMSA pointer being left NULL. BUG: unable to handle page fault for address: ffffe38687000000 #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: 0000 [#1] SMP NOPTI CPU: 39 PID: 17143 Comm: sev_migrate_tes Tainted: GO 6.5.0-smp--fff2e47e6c3b-next #151 Hardware name: Google, Inc. Arcadia_IT_80/Arcadia_IT_80, BIOS 34.28.0 07/10/2023 RIP: 0010:__free_pages+0x15/0xd0 RSP: 0018:ffff923fcf6e3c78 EFLAGS: 00010246 RAX: 0000000000000000 RBX: ffffe38687000000 RCX: 0000000000000100 RDX: 0000000000000100 RSI: 0000000000000000 RDI: ffffe38687000000 RBP: ffff923fcf6e3c88 R08: ffff923fcafb0000 R09: 0000000000000000 R10: 0000000000000000 R11: ffffffff83619b90 R12: ffff923fa9540000 R13: 0000000000080007 R14: ffff923f6d35d000 R15: 0000000000000000 FS: 0000000000000000(0000) GS:ffff929d0d7c0000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: ffffe38687000000 CR3: 0000005224c34005 CR4: 0000000000770ee0 PKRU: 55555554 Call Trace: <TASK> sev_free_vcpu+0xcb/0x110 [kvm_amd] svm_vcpu_free+0x75/0xf0 [kvm_amd] kvm_arch_vcpu_destroy+0x36/0x140 [kvm] kvm_destroy_vcpus+0x67/0x100 [kvm] kvm_arch_destroy_vm+0x161/0x1d0 [kvm] kvm_put_kvm+0x276/0x560 [kvm] kvm_vm_release+0x25/0x30 [kvm] __fput+0x106/0x280 ____fput+0x12/0x20 task_work_run+0x86/0xb0 do_exit+0x2e3/0x9c0 do_group_exit+0xb1/0xc0 __x64_sys_exit_group+0x1b/0x20 do_syscall_64+0x41/0x90 entry_SYSCALL_64_after_hwframe+0x63/0xcd </TASK> CR2: ffffe38687000000	N/A	More Details
CVE-2023-54297	In the Linux kernel, the following vulnerability has been resolved: btrfs: zoned: fix memory leak after finding block group with super blocks At exclude_super_stripes(), if we happen to find a block group that has super blocks mapped to it and we are on a zoned filesystem, we error out as this is not supposed to happen, indicating either a bug or maybe some memory corruption for example. However we are exiting the function without freeing the memory allocated for the logical address of the super blocks. Fix this by freeing the logical address.	N/A	More Details
CVE-2023-54298	In the Linux kernel, the following vulnerability has been resolved: thermal: intel: quark_dts: fix error pointer dereference If alloc_soc_dts() fails, then we can just return. Trying to free "soc_dts" will lead to an Oops.	N/A	More Details
CVE-2023-54299	In the Linux kernel, the following vulnerability has been resolved: usb: typec: bus: verify partner exists in typec_altmode_attention Some usb hubs will negotiate DisplayPort Alt mode with the device but will then negotiate a data role swap after entering the alt mode. The data role swap causes the device to unregister all alt modes, however the usb hub will still send Attention messages even after failing to reregister the Alt Mode. type_altmode_attention currently does not verify whether or not a device's altmode partner exists, which results in a NULL pointer error when dereferencing the typec_altmode and typec_altmode_ops belonging to the altmode partner. Verify the presence of a device's altmode partner before sending the Attention message to the Alt Mode driver.	N/A	More Details
CVE-2022-50600	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2023-54162	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix possible memory leak in smb2_lock() argv needs to be free when setup_async_work fails or when the current process is woken up.	N/A	More Details
CVE-2022-50598	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2022-50731	In the Linux kernel, the following vulnerability has been resolved: crypto: akcipher - default implementation for setting a private key Changes from v1: * removed the default implementation from set_pub_key: it is assumed that an implementation must always have this callback defined as there are no use case for an algorithm, which doesn't need a public key Many akcipher implementations (like ECDSA) support only signature verifications, so they don't have all callbacks defined. Commit 78a0324f4a53 ("crypto: akcipher - default implementations for request callbacks") introduced default callbacks for sign/verify operations, which just return an error code. However, these are not enough, because before calling sign the caller would likely call set_priv_key first on the instantiated transform (as the in-kernel testmgr does). This function does not have a default stub, so the kernel crashes, when trying to set a private key on an akcipher, which doesn't support signature generation. I've noticed this, when trying to add a KAT vector for ECDSA signature to the testmgr. With this patch the testmgr returns an error in dmesg (as it should) instead of crashing the kernel NULL ptr dereference.	N/A	More Details

CVE-2025-68744	In the Linux kernel, the following vulnerability has been resolved: bpf: Free special fields when update [lru_,]percpu_hash maps As [lru_,]percpu_hash maps support BPF_KPTR_{REF,PERCPU}, missing calls to 'bpf_obj_free_fields()' in 'pcpu_copy_value()' could cause the memory referenced by BPF_KPTR_{REF,PERCPU} fields to be held until the map gets freed. Fix this by calling 'bpf_obj_free_fields()' after 'copy_map_value[,_long]()' in 'pcpu_copy_value()'.	N/A	More Details
CVE-2022-50736	In the Linux kernel, the following vulnerability has been resolved: RDMA/siw: Fix immediate work request flush to completion queue Correctly set send queue element opcode during immediate work request flushing in post sendqueue operation, if the QP is in ERROR state. An undefined ocode value results in out-of-bounds access to an array for mapping the opcode between siw internal and RDMA core representation in work completion generation. It resulted in a KASAN BUG report of type 'global-out-of-bounds' during NFSoRDMA testing. This patch further fixes a potential case of a malicious user which may write undefined values for completion queue elements status or opcode, if the CQ is memory mapped to user land. It avoids the same out-of-bounds access to arrays for status and opcode mapping as described above.	N/A	More Details
CVE-2025-68743	In the Linux kernel, the following vulnerability has been resolved: mshv: Fix create memory region overlap check The current check is incorrect; it only checks if the beginning or end of a region is within an existing region. This doesn't account for userspace specifying a region that begins before and ends after an existing region. Change the logic to a range intersection check against gfns and uaddrs for each region. Remove mshv_partition_region_by_uaddr() as it is no longer used.	N/A	More Details
CVE-2022-50735	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: do not run mt76u_status_worker if the device is not running Fix the following NULL pointer dereference avoiding to run mt76u_status_worker thread if the device is not running yet. KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] CPU: 0 PID: 98 Comm: kworker/u2:2 Not tainted 5.14.0+ #78 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.12.1-0-ga5cab58e9a3f-prebuilt.qemu.org 04/01/2014 Workqueue: mt76 mt76u_tx_status_data RIP: 0010:mt76x02_mac_fill_tx_status.isra.0+0x82c/0x9e0 Code: c5 48 b8 00 00 00 00 00 fc ff df 80 3c 02 00 0f 85 94 01 00 00 48 b8 00 00 00 00 00 fc ff df 4d 8b 34 24 4c 89 f2 48 c1 ea 03 <0f> b6 04 02 84 c0 74 08 3c 03 0f 8e 89 01 00 00 41 8b 16 41 0f b7 RSP: 0018:ffffc900005af988 EFLAGS: 00010246 RAX: dfffc0000000000 RBX: ffffc900005afae8 RCX: 0000000000000000 RDX: 0000000000000000 RSI: ffffffff832fc661 RDI: ffffc900005afc2a RBP: ffffc900005afae0 R08: 0000000000000001 R09: fffff520000b5f3c R10: 0000000000000003 R11: fffff520000b5f3b R12: ffff88810b6132d8 R13: 000000000000ffff R14: 0000000000000000 R15: ffffc900005afc28 FS: 0000000000000000(0000) GS:ffff88811aa00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fa0eda6a000 CR3: 0000000118f17000 CR4: 0000000000750ef0 PKRU: 55555554 Call Trace: mt76x02_send_tx_status+0x1d2/0xeb0 mt76x02_tx_status_data+0x8e/0xd0 mt76u_tx_status_data+0xe1/0x240 process_one_work+0x92b/0x1460 worker_thread+0x95/0xe00 kthread+0x3a1/0x480 ret_from_fork+0x1f/0x30 Modules linked in: --[end trace 8df5d20fc5040f65]-- RIP: 0010:mt76x02_mac_fill_tx_status.isra.0+0x82c/0x9e0 Code: c5 48 b8 00 00 00 00 00 fc ff df 80 3c 02 00 0f 85 94 01 00 00 48 b8 00 00 00 00 00 fc ff df 4d 8b 34 24 4c 89 f2 48 c1 ea 03 <0f> b6 04 02 84 c0 74 08 3c 03 0f 8e 89 01 00 00 41 8b 16 41 0f b7 RSP: 0018:ffffc900005af988 EFLAGS: 00010246 RAX: dfffc0000000000 RBX: ffffc900005afae8 RCX: 0000000000000000 RDX: 0000000000000000 RSI: ffffffff832fc661 RDI: ffffc900005afc2a RBP: ffffc900005afae0 R08: 0000000000000001 R09: fffff520000b5f3c R10: 0000000000000003 R11: fffff520000b5f3b R12: ffff88810b6132d8 R13: 000000000000ffff R14: 0000000000000000 R15: ffffc900005afc28 FS: 0000000000000000(0000) GS:ffff88811aa00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fa0eda6a000 CR3: 0000000118f17000 CR4: 0000000000750ef0 PKRU: 55555554 Moreover move stat_work schedule out of the for loop.	N/A	More Details
CVE-2022-50734	In the Linux kernel, the following vulnerability has been resolved: nvmem: core: Fix memleak in nvmem_register() dev_set_name will alloc memory for nvmem->dev.kobj.name in nvmem_register, when nvmem_validate_keepouts failed, nvmem's memory will be freed and return, but nobody will free memory for nvmem->dev.kobj.name, there will be memleak, so moving nvmem_validate_keepouts() after device_register() and let the device core deal with cleaning name in error cases.	N/A	More Details
CVE-2022-50733	In the Linux kernel, the following vulnerability has been resolved: usb: idmouse: fix an uninit-value in idmouse_open In idmouse_create_image, if any ftip_command fails, it will go to the reset label. However, this leads to the data in bulk_in_buffer[HEADER..IMGSIZE] uninitialized. And the check for valid image incurs an uninitialized dereference. Fix this by moving the check before reset label since this check only be valid if the data after bulk_in_buffer[HEADER] has concrete data. Note that this is found by KMSAN, so only kernel compilation is tested.	N/A	More Details
CVE-2025-68742	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix invalid prog->stats access when update_effective_progs fails Syzkaller triggers an invalid memory access issue following fault injection in update_effective_progs. The issue can be described as follows: __cgroup_bpf_detach update_effective_progs compute_effective_progs bpf_prog_array_alloc <-- fault inject purge_effective_progs /* change to dummy_bpf_prog */ array->items[index] = &dummy_bpf_prog.prog ---softirq start--- __do_softirq ... __cgroup_bpf_run_filter_skb __bpf_prog_run_save_cb bpf_prog_run stats = this_cpu_ptr(prog->stats) /* invalid memory access */ flags = u64_stats_update_begin_irqsave(&stats->syncp) ---softirq end--- static_branch_dec(&cgroup_bpf_enabled_key[atype]) The reason is that fault injection caused update_effective_progs to fail and then changed the original prog into dummy_bpf_prog.prog in purge_effective_progs. Then a softirq came, and accessing the members of dummy_bpf_prog.prog in the softirq triggers invalid mem access. To fix it, skip updating stats when stats is NULL.	N/A	More Details
CVE-2022-50732	In the Linux kernel, the following vulnerability has been resolved: staging: rtl8192u: Fix use after free in ieee80211_rx() We cannot dereference the "skb" pointer after calling ieee80211_monitor_rx(), because it is a use after free.	N/A	More Details
CVE-2022-50730	In the Linux kernel, the following vulnerability has been resolved: ext4: silence the warning when evicting inode with dioread_nolock When evicting an inode with default dioread_nolock, it could be raced by the unwritten extents converting kworker after writeback some new allocated dirty blocks. It convert unwritten extents to written, the extents could be merged to upper level and free extent blocks, so it could mark the inode dirty again even this inode has been marked I_FREEING. But the inode->i_io_list check and warning in ext4_evict_inode() missing this corner case. Fortunately, ext4_evict_inode() will wait all extents converting finished before this check, so it will not lead to inode use-after-free problem, every thing is OK besides this warning. The WARN_ON_ONCE was originally designed for finding inode use-after-free issues in advance, but if we add current dioread_nolock case in, it will become not quite useful, so fix this warning by just remove this check. ====== WARNING: CPU: 7 PID: 1092 at fs/ext4/inode.c:227 ext4_evict_inode+0x875/0xc60 ... RIP: 0010:ext4_evict_inode+0x875/0xc60 ... Call Trace: <TASK> evict+0x11c/0x2b0 iput+0x236/0x3a0 do_unlinkat+0x1b4/0x490 __x64_sys_unlinkat+0x4c/0xb0 do_syscall_64+0x3b/0x90 entry_SYSCALL_64_after_hwframe+0x46/0xb0 RIP: 0033:0x7fa933c1115b ====== rm kworker ext4_end_io_end() vfs_unlink() ext4_unlink() ext4_convert_unwritten_io_end_vec() ext4_convert_unwritten_extents() ext4_map_blocks() ext4_ext_map_blocks() ext4_ext_try_to_merge_up() __mark_inode_dirty() check !I_FREEING locked_inode_to_wb_and_lock_list() iput() iput_final() evict() ext4_evict_inode() truncate_inode_pages_final() //wait release io_end inode_io_list_move_locked() ext4_release_io_end() trigger WARN_ON_ONCE()	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: landlock: Fix handling of disconnected directories Disconnected		

CVE-2025-68736	files or directories can appear when they are visible and opened from a bind mount, but have been renamed or moved from the source of the bind mount in a way that makes them inaccessible from the mount point (i.e. out of scope). Previously, access rights tied to files or directories opened through a disconnected directory were collected by walking the related hierarchy down to the root of the filesystem, without taking into account the mount point because it couldn't be found. This could lead to inconsistent access results, potential access right widening, and hard-to-debug renames, especially since such paths cannot be printed. For a sandboxed task to create a disconnected directory, it needs to have write access (i.e. FS_MAKE_REG, FS_REMOVE_FILE, and FS_REFER) to the underlying source of the bind mount, and read access to the related mount point. Because a sandboxed task cannot acquire more access rights than those defined by its Landlock domain, this could lead to inconsistent access rights due to missing permissions that should be inherited from the mount point hierarchy, while inheriting permissions from the filesystem hierarchy hidden by this mount point instead. Landlock now handles files and directories opened from disconnected directories by taking into account the filesystem hierarchy when the mount point is not found in the hierarchy walk, and also always taking into account the mount point from which these disconnected directories were opened. This ensures that a rename is not allowed if it would widen access rights [1]. The rationale is that, even if disconnected hierarchies might not be visible or accessible to a sandboxed task, relying on the collected access rights from them improves the guarantee that access rights will not be widened during a rename because of the access right comparison between the source and the destination (see LANDLOCK_ACCESS_FS_REFER). It may look like this would grant more access on disconnected files and directories, but the security policies are always enforced for all the evaluated hierarchies. This new behavior should be less surprising to users and safer from an access control perspective. Remove a wrong WARN_ON_ONCE() canary in collect_domain_accesses() and fix the related comment. Because opened files have their access rights stored in the related file security properties, there is no impact for disconnected or unlinked files.	N/A	More Details
CVE-2025-68741	In the Linux kernel, the following vulnerability has been resolved: scsi: qla2xxx: Fix improper freeing of purex item In qla2xxx_process_purls_iocb(), an item is allocated via qla27xx_copy_multiple_pkt(), which internally calls qla24xx_alloc_purex_item(). The qla24xx_alloc_purex_item() function may return a pre-allocated item from a per-adapter pool for small allocations, instead of dynamically allocating memory with kzalloc(). An error handling path in qla2xxx_process_purls_iocb() incorrectly uses kfree() to release the item. If the item was from the pre-allocated pool, calling kfree() on it is a bug that can lead to memory corruption. Fix this by using the correct deallocation function, qla24xx_free_purex_item(), which properly handles both dynamically allocated and pre-allocated items.	N/A	More Details
CVE-2022-50729	In the Linux kernel, the following vulnerability has been resolved: ksmbd: Fix resource leak in ksmbd_session_rpc_open() When ksmbd_rpc_open() fails then it must call ksmbd_rpc_id_free() to undo the result of ksmbd_ipc_id_alloc().	N/A	More Details
CVE-2022-50728	In the Linux kernel, the following vulnerability has been resolved: s390/lcs: Fix return type of lcs_start_xmit() With clang's kernel control flow integrity (kCFI, CONFIG_CFI_CLANG), indirect call targets are validated against the expected function pointer prototype to make sure the call target is valid to help mitigate ROP attacks. If they are not identical, there is a failure at run time, which manifests as either a kernel panic or thread getting killed. A proposed warning in clang aims to catch these at compile time, which reveals: drivers/s390/net/lcs.c:2090:21: error: incompatible function pointer types initializing 'netdev_tx_t (*)(struct sk_buff *, struct net_device *)' (aka 'enum netdev_tx (*)(struct sk_buff *, struct net_device *)') with an expression of type 'int (struct sk_buff *, struct net_device *)' [-Werror,-Wincompatible-function-pointer-types-strict] .ndo_start_xmit = lcs_start_xmit, ^~~~~~~~~~~~~~~ drivers/s390/net/lcs.c:2097:21: error: incompatible function pointer types initializing 'netdev_tx_t (*)(struct sk_buff *, struct net_device *)' (aka 'enum netdev_tx (*)(struct sk_buff *, struct net_device *)') with an expression of type 'int (struct sk_buff *, struct net_device *)' [-Werror,-Wincompatible-function-pointer-types-strict] .ndo_start_xmit = lcs_start_xmit, ^~~~~~~~~~~~~~~ ->ndo_start_xmit() in 'struct net_device_ops' expects a return type of 'netdev_tx_t', not 'int'. Adjust the return type of lcs_start_xmit() to match the prototype's to resolve the warning and potential CFI failure, should s390 select ARCH_SUPPORTS_CFI_CLANG in the future.	N/A	More Details
CVE-2025-68740	In the Linux kernel, the following vulnerability has been resolved: ima: Handle error code returned by ima_filter_rule_match() In ima_match_rules(), if ima_filter_rule_match() returns -ENOENT due to the rule being NULL, the function incorrectly skips the 'if (!rc)' check and sets 'result = true'. The LSM rule is considered a match, causing extra files to be measured by IMA. This issue can be reproduced in the following scenario: After unloading the SELinux policy module via 'semodule -d', if an IMA measurement is triggered before ima_lsm_rules is updated, in ima_match_rules(), the first call to ima_filter_rule_match() returns -ESTALE. This causes the code to enter the 'if (rc == -ESTALE && !rule_reinitialized)' block, perform ima_lsm_copy_rule() and retry. In ima_lsm_copy_rule(), since the SELinux module has been removed, the rule becomes NULL, and the second call to ima_filter_rule_match() returns -ENOENT. This bypasses the 'if (!rc)' check and results in a false match. Call trace: selinux_audit_rule_match+0x310/0x3b8 security_audit_rule_match+0x60/0xa0 ima_match_rules+0x2e4/0x4a0 ima_match_policy+0x9c/0x1e8 ima_get_action+0x48/0x60 process_measurement+0xf8/0xa98 ima_bprm_check+0x98/0xd8 security_bprm_check+0x5c/0x78 search_binary_handler+0x6c/0x318 exec_binprm+0x58/0x1b8 bprm_execve+0xb8/0x130 do_execveat_common.isra.0+0x1a8/0x258 __arm64_sys_execve+0x48/0x68 invoke_syscall+0x50/0x128 el0_svc_common.constprop.0+0xc8/0xf0 do_el0_svc+0x24/0x38 el0_svc+0x44/0x200 el0t_64_sync_handler+0x100/0x130 el0t_64_sync+0x3c8/0x3d0 Fix this by changing 'if (!rc)' to 'if (rc <= 0)' to ensure that error codes like -ENOENT do not bypass the check and accidentally result in a successful match.	N/A	More Details
CVE-2022-50727	In the Linux kernel, the following vulnerability has been resolved: scsi: efct: Fix possible memleak in efct_device_init() In efct_device_init(), when efct_scsi_reg_fc_transport() fails, efct_scsi_tgt_driver_exit() is not called to release memory for efct_scsi_tgt_driver_init() and causes memleak: unreferenced object 0xffff8881020ce000 (size 2048): comm "modprobe", pid 465, jiffies 4294928222 (age 55.872s) backtrace: [<0000000021a1ef1b>] kmalloc_trace+0x27/0x110 [<000000004c3ed51c>] target_register_template+0x4fd/0x7b0 [target_core_mod] [<00000000f3393296>] efct_scsi_tgt_driver_init+0x18/0x50 [efct] [<00000000115de533>] 0xffffffffc0d90011 [<00000000d608f646>] do_one_initcall+0xd0/0x4e0 [<0000000067828cf1>] do_init_module+0x1cc/0x6a0 ...	N/A	More Details
CVE-2025-68739	In the Linux kernel, the following vulnerability has been resolved: PM / devfreq: hisi: Fix potential UAF in OPP handling Ensure all required data is acquired before calling dev_pm_opp_put(opp) to maintain correct resource acquisition and release order.	N/A	More Details
CVE-2025-68738	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7996: fix null pointer deref in mt7996_conf_tx() If a link does not have an assigned channel yet, mt7996_vif_link returns NULL. We still need to store the updated queue settings in that case, and apply them later. Move the location of the queue params to within struct mt7996_vif_link.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Fix possible use-after-free in async command interface mlx5_cmd_cleanup_async_ctx should return only after all its callback handlers were completed. Before this patch, the below race between mlx5_cmd_cleanup_async_ctx and mlx5_cmd_exec_cb_handler was possible and lead to a use-after-free: 1. mlx5_cmd_cleanup_async_ctx is called while num_inflight is 2 (i.e. elevated by 1, a single inflight callback). 2. mlx5_cmd_cleanup_async_ctx decreases num_inflight to 1. 3. mlx5_cmd_exec_cb_handler is called, decreases num_inflight to 0 and		

CVE-2022-50726	is about to call wake_up(). 4. mlx5_cmd_cleanup_async_ctx calls wait_event, which returns immediately as the condition (num_inflight == 0) holds. 5. mlx5_cmd_cleanup_async_ctx returns. 6. The caller of mlx5_cmd_cleanup_async_ctx frees the mlx5_async_ctx object. 7. mlx5_cmd_exec_cb_handler goes on and calls wake_up() on the freed object. Fix it by syncing using a completion object. Mark it completed when num_inflight reaches 0. Trace: BUG: KASAN: use-after-free in do_raw_spin_lock+0x23d/0x270 Read of size 4 at addr ffff888139cd12f4 by task swapper/5/0 CPU: 5 PID: 0 Comm: swapper/5 Not tainted 6.0.0-rc3_for_upstream_debug_2022_08_30_13_10 #1 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 Call Trace: <IRQ> dump_stack_lvl+0x57/0x7d print_report.cold+0x2d5/0x684 ? do_raw_spin_lock+0x23d/0x270 kasan_report+0xb1/0x1a0 ? do_raw_spin_lock+0x23d/0x270 do_raw_spin_lock+0x23d/0x270 ? rwlock_bug.part.0+0x90/0x90 ? __delete_object+0xb8/0x100 ? lock_downgrade+0x6e0/0x6e0 _raw_spin_lock_irqsave+0x43/0x60 ? __wake_up_common_lock+0xb9/0x140 __wake_up_common_lock+0xb9/0x140 ? __wake_up_common+0x650/0x650 ? destroy_tis_callback+0x53/0x70 [mlx5_core] ? kasan_set_track+0x21/0x30 ? destroy_tis_callback+0x53/0x70 [mlx5_core] ? kfree+0x1ba/0x520 ? do_raw_spin_unlock+0x54/0x220 mlx5_cmd_exec_cb_handler+0x136/0x1a0 [mlx5_core] ? mlx5_cmd_cleanup_async_ctx+0x220/0x220 [mlx5_core] ? mlx5_cmd_cleanup_async_ctx+0x220/0x220 [mlx5_core] mlx5_cmd_comp_handler+0x65a/0x12b0 [mlx5_core] ? dump_command+0xcc0/0xcc0 [mlx5_core] ? lockdep_hardirqs_on_prepare+0x400/0x400 ? cmd_comp_notifier+0x7e/0xb0 [mlx5_core] cmd_comp_notifier+0x7e/0xb0 [mlx5_core] atomic_notifier_call_chain+0xd7/0x1d0 mlx5_eq_async_int+0x3ce/0xa20 [mlx5_core] atomic_notifier_call_chain+0xd7/0x1d0 ? irq_release+0x140/0x140 [mlx5_core] irq_int_handler+0x19/0x30 [mlx5_core] __handle_irq_event_percpu+0x1f2/0x620 handle_irq_event+0xb2/0x1d0 handle_edge_irq+0x21e/0xb00 __common_interrupt+0x79/0x1a0 common_interrupt+0x78/0xa0 </IRQ> <TASK> asm_common_interrupt+0x22/0x40 RIP: 0010:default_idle+0x42/0x60 Code: c1 83 e0 07 48 c1 e9 03 83 c0 03 0f b6 14 11 38 d0 7c 04 84 d2 75 14 8b 05 eb 47 22 02 85 c0 7e 07 0f 00 2d e0 9f 48 00 fb f4 <c3> 48 c7 c7 80 08 7f 85 e8 d1 d3 3e fe eb de 66 66 2e 0f 1f 84 00 RSP: 0018:ffff888100dbfdf0 EFLAGS: 00000242 RAX: 0000000000000001 RBX: ffffffff84ecbd48 RCX: 1ffffffff0afe110 RDX: 0000000000000004 RSI: 0000000000000000 RDI: ffffffff835cc9bc RBP: 0000000000000005 R08: 0000000000000001 R09: ffff88881dec4ac3 R10: ffffed1103bd8958 R11: 0000017d0ca571c9 R12: 0000000000000005 R13: ffffffff84f024e0 R14: 0000000000000000 R15: dfffffc0000000000 ? default_idle_call+0xcc/0x450 default_idle_call+0xec/0x450 do_idle+0x394/0x450 ? arch_cpu_idle_exit+0x40/0x40 ? do_idle+0x17/0x450 cpu_startup_entry+0x19/0x20 start_secondary+0x221/0x2b0 ? set_cpu_sibling_map+0x2070/0x2070 secondary_startup_64_no_verify+0xcd/0xdb </TASK> Allocated by task 49502: kasan_save_stack+0x1e/0x40 __kasan_kmalloc+0x81/0xa0 kvmalloc_node+0x48/0xe0 mlx5e_bulk_async_init+0x35/0x110 [mlx5_core] mlx5e_tls_priv_tx_list_cleanup+0x84/0x3e0 [mlx5_core] mlx5e_ktls_cleanup_tx+0x38f/0x760 [mlx5_core] mlx5e_cleanup_nic_tx+0xa7/0x100 [mlx5_core] mlx5e_detach_netdev+0x1c ---truncated---	N/A	More Details
CVE-2022-50737	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Validate index root when initialize NTFS security This enhances the sanity check for $SDH and $SII while initializing NTFS security, guarantees these index root are legit. [162.459513] BUG: KASAN: use-after-free in hdr_find_e.isra.0+0x10c/0x320 [162.460176] Read of size 2 at addr ffff8880037bca99 by task mount/243 [162.460851] [162.461252] CPU: 0 PID: 243 Comm: mount Not tainted 6.0.0-rc7 #42 [162.461744] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org 04/01/2014 [162.462609] Call Trace: [162.462954] <TASK> [162.463276] dump_stack_lvl+0x49/0x63 [162.463822] print_report.cold+0xf5/0x689 [162.464608] ? unwind_get_return_address+0x3a/0x60 [162.465766] ? hdr_find_e.isra.0+0x10c/0x320 [162.466975] kasan_report+0xa7/0x130 [162.467506] ? _raw_spin_lock_irq+0xc0/0xf0 [162.467998] ? hdr_find_e.isra.0+0x10c/0x320 [162.468536] __asan_load2+0x68/0x90 [162.468923] hdr_find_e.isra.0+0x10c/0x320 [162.469282] ? cmp_uints+0xe0/0xe0 [162.469557] ? cmp_sdh+0x90/0x90 [162.469864] ? ni_find_attr+0x214/0x300 [162.470217] ? ni_load_mi+0x80/0x80 [162.470479] ? entry_SYSCALL_64_after_hwframe+0x63/0xcd [162.470931] ? ntfs_bread_run+0x190/0x190 [162.471307] ? indx_get_root+0xe4/0x190 [162.471556] ? indx_get_root+0x140/0x190 [162.471833] ? indx_init+0x1e0/0x1e0 [162.472069] ? fnd_clear+0x115/0x140 [162.472363] ? _raw_spin_lock_irqsave+0x100/0x100 [162.472731] indx_find+0x184/0x470 [162.473461] ? sysvec_apic_timer_interrupt+0x57/0xc0 [162.474429] ? indx_find_buffer+0x2d0/0x2d0 [162.474704] ? do_syscall_64+0x3b/0x90 [162.474962] dir_search_u+0x196/0x2f0 [162.475381] ? ntfs_nls_to_utf16+0x450/0x450 [162.475661] ? ntfs_security_init+0x3d6/0x440 [162.475906] ? is_sd_valid+0x180/0x180 [162.476191] ntfs_extend_init+0x13f/0x2c0 [162.476496] ? ntfs_fix_post_read+0x130/0x130 [162.476861] ? iput.part.0+0x286/0x320 [162.477325] ntfs_fill_super+0x11e0/0x1b50 [162.477709] ? put_ntfs+0x1d0/0x1d0 [162.477970] ? vsprintf+0x20/0x20 [162.478258] ? set_blocksize+0x95/0x150 [162.478538] get_tree_bdev+0x232/0x370 [162.478789] ? put_ntfs+0x1d0/0x1d0 [162.479038] ntfs_fs_get_tree+0x15/0x20 [162.479374] vfs_get_tree+0x4c/0x130 [162.479729] path_mount+0x654/0xfe0 [162.480124] ? putname+0x80/0xa0 [162.480484] ? finish_automount+0x2e0/0x2e0 [162.480894] ? putname+0x80/0xa0 [162.481467] ? kmem_cache_free+0x1c4/0x440 [162.482280] ? putname+0x80/0xa0 [162.482714] do_mount+0xd6/0xf0 [162.483264] ? path_mount+0xfe0/0xfe0 [162.484782] ? __kasan_check_write+0x14/0x20 [162.485593] __x64_sys_mount+0xca/0x110 [162.486024] do_syscall_64+0x3b/0x90 [162.486543] entry_SYSCALL_64_after_hwframe+0x63/0xcd [162.487141] RIP: 0033:0x7f9d374e948a [162.488324] Code: 48 8b 0d 11 fa 2a 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 49 89 ca b8 a5 00 00 008 [162.489728] RSP: 002b:00007ffe30e73d18 EFLAGS: 00000206 ORIG_RAX: 00000000000000a5 [162.490971] RAX: ffffffffffffffda RBX: 0000561cdb43a060 RCX: 00007f9d374e948a [162.491669] RDX: 0000561cdb43a260 RSI: 0000561cdb43a2e0 RDI: 0000561cdb442af0 [162.492050] RBP: 0000000000000000 R08: 0000561cdb43a280 R09: 0000000000000020 [162.492459] R10: 00000000c0ed0000 R11: 0000000000000206 R12: 0000561cdb442af0 [162.493183] R13: 0000561cdb43a260 R14: 0000000000000000 R15: 00000000ffffffff [162.493644] </TASK> [162.493908] [162.494214] The buggy address belongs to the physical page: [162.494761] page:000000003e38a3d5 refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x37bc [162.496064] flags: 0xfffffc0000000(node=0|zone=1|lastcpupid=0x1fffff) [162.497278] raw: 000fffffc0000000 ffffea00000df1c8 ffffea00000df008 0000000000000000 [162.498928] raw: 0000000000000000 0000000000240000 0 ---truncated---	N/A	More Details
CVE-2025-68745	In the Linux kernel, the following vulnerability has been resolved: scsi: qla2xxx: Clear cmds after chip reset Commit aefed3e5548f ("scsi: qla2xxx: target: Fix offline port handling and host reset handling") caused two problems: 1. Commands sent to FW, after chip reset got stuck and never freed as FW is not going to respond to them anymore. 2. BUG_ON(cmd->sg_mapped) in qlt_free_cmd(). Commit 26f9ce53817a ("scsi: qla2xxx: Fix missed DMA unmap for aborted commands") attempted to fix this, but introduced another bug under different circumstances when two different CPUs were racing to call qlt_unmap_sg() at the same time: BUG_ON(!valid_dma_direction(dir)) in dma_unmap_sg_attrs(). So revert "scsi: qla2xxx: Fix missed DMA unmap for aborted commands" and partially revert "scsi: qla2xxx: target: Fix offline port handling and host reset handling" at __qla2x00_abort_all_cmds.	N/A	More Details
CVE-2025-68746	In the Linux kernel, the following vulnerability has been resolved: spi: tegra210-quad: Fix timeout handling When the CPU that the QSPI interrupt handler runs on (typically CPU 0) is excessively busy, it can lead to rare cases of the IRQ thread not running before the transfer timeout is reached. While handling the timeouts, any pending transfers are cleaned up and the message that they correspond to is marked as failed, which leaves the curr_xfer field pointing at stale memory. To avoid this, clear curr_xfer to NULL upon timeout and check for this condition when the IRQ thread is finally run. While at it, also make sure to clear interrupts on failure so that new interrupts can be run. A better, more involved, fix would move the interrupt clearing into a hard IRQ handler. Ideally we	N/A	More Details

	would also want to signal that the IRQ thread no longer needs to be run after the timeout is hit to avoid the extra check for a valid transfer.		
CVE-2025-68747	In the Linux kernel, the following vulnerability has been resolved: drm/panthor: Fix UAF on kernel BO VA nodes If the MMU is down, panthor_vm_unmap_range() might return an error. We expect the page table to be updated still, and if the MMU is blocked, the rest of the GPU should be blocked too, so no risk of accessing physical memory returned to the system (which the current code doesn't cover for anyway). Proceed with the rest of the cleanup instead of bailing out and leaving the va_node inserted in the drm_mm, which leads to UAF when other adjacent nodes are removed from the drm_mm tree.	N/A	More Details
CVE-2023-54071	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw88: use work to update rate to avoid RCU warning The ieee80211_ops::sta_rc_update must be atomic, because ieee80211_chan_bw_change() holds rcu_read lock while calling drv_sta_rc_update(), so create a work to do original things. Voluntary context switch within RCU read-side critical section! WARNING: CPU: 0 PID: 4621 at kernel/rcu/tree_plugin.h:318 rcu_note_context_switch+0x571/0x5d0 CPU: 0 PID: 4621 Comm: kworker/u16:2 Tainted: G W OE Workqueue: phy3 ieee80211_chswitch_work [mac80211] RIP: 0010:rcu_note_context_switch+0x571/0x5d0 Call Trace: <TASK> __schedule+0xb0/0x1460 ? __mod_timer+0x116/0x360 schedule+0x5a/0xc0 schedule_timeout+0x87/0x150 ? trace_raw_output_tick_stop+0x60/0x60 wait_for_completion_timeout+0x7b/0x140 usb_start_wait_urb+0x82/0x160 [usbcore usb_control_msg+0xe3/0x140 [usbcore rtw_usb_read+0x88/0xe0 [rtw_usb rtw_usb_read8+0xf/0x10 [rtw_usb rtw_fw_send_h2c_command+0xa0/0x170 [rtw_core rtw_fw_send_ra_info+0xc9/0xf0 [rtw_core drv_sta_rc_update+0x7c/0x160 [mac80211 ieee80211_chan_bw_change+0xfb/0x110 [mac80211 ieee80211_change_chanctx+0x38/0x130 [mac80211 ieee80211_vif_use_reserved_switch+0x34e/0x900 [mac80211 ieee80211_link_use_reserved_context+0x88/0xe0 [mac80211 ieee80211_chswitch_work+0x95/0x170 [mac80211 process_one_work+0x201/0x410 worker_thread+0x4a/0x3b0 ? process_one_work+0x410/0x410 kthread+0xe1/0x110 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork+0x1f/0x30 </TASK>	N/A	More Details
CVE-2022-50740	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: hif_usb: fix memory leak of urbs in ath9k_hif_usb_dealloc_tx_urbs() Syzkaller reports a long-known leak of urbs in ath9k_hif_usb_dealloc_tx_urbs(). The cause of the leak is that usb_get_urb() is called but usb_free_urb() (or usb_put_urb()) is not called inside usb_kill_urb() as urb->dev or urb->ep fields have not been initialized and usb_kill_urb() returns immediately. The patch removes trying to kill urbs located in hif_dev->tx.tx_buf because hif_dev->tx.tx_buf is not supposed to contain urbs which are in pending state (the pending urbs are stored in hif_dev->tx.tx_pending). The tx.tx_lock is acquired so there should not be any changes in the list. Found by Linux Verification Center (linuxtesting.org) with Syzkaller.	N/A	More Details
CVE-2023-54072	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcm: Fix potential data race at PCM memory allocation helpers The PCM memory allocation helpers have a sanity check against too many buffer allocations. However, the check is performed without a proper lock and the allocation isn't serialized; this allows user to allocate more memories than predefined max size. Practically seen, this isn't really a big problem, as it's more or less some "soft limit" as a sanity check, and it's not possible to allocate unlimitedly. But it's still better to address this for more consistent behavior. The patch covers the size check in do_alloc_pages() with the card->memory_mutex, and increases the allocated size there for preventing the further overflow. When the actual allocation fails, the size is decreased accordingly.	N/A	More Details
CVE-2023-54073	In the Linux kernel, the following vulnerability has been resolved: tpm: Add !tpm_amd_is_rng_defective() to the hwrng_unregister() call site The following crash was reported: [1950.279393] list_del corruption, ffff99560d485790->next is NULL [1950.279400] ----------[cut here]---------- [1950.279401] kernel BUG at lib/list_debug.c:49! [1950.279405] invalid opcode: 0000 [#1] PREEMPT SMP NOPTI [1950.279407] CPU: 11 PID: 5886 Comm: modprobe Tainted: G O 6.2.8_1 #1 [1950.279409] Hardware name: Gigabyte Technology Co., Ltd. B550M AORUS PRO-P/B550M AORUS PRO-P, BIOS F15c 05/11/2022 [1950.279410] RIP: 0010:__list_del_entry_valid+0x59/0xc0 [1950.279415] Code: 48 8b 01 48 39 f8 75 5a 48 8b 72 08 48 39 c6 75 65 b8 01 00 00 00 c3 cc cc cc cc 48 89 fe 48 c7 c7 08 a8 13 9e e8 b7 0a bc ff <0f> 0b 48 89 fe 48 c7 c7 38 a8 13 9e e8 a6 0a bc ff 0f 0b 48 89 fe [1950.279416] RSP: 0018:ffffa96d05647e08 EFLAGS: 00010246 [1950.279418] RAX: 0000000000000033 RBX: ffff99560d485750 RCX: 0000000000000000 [1950.279419] RDX: 0000000000000000 RSI: ffffffff9e107c59 RDI: 00000000ffffffff [1950.279420] RBP: ffffffffc19c5168 R08: 0000000000000000 R09: ffffa96d05647cc8 [1950.279421] R10: 0000000000000003 R11: ffffffff9ea2a568 R12: 0000000000000000 [1950.279422] R13: ffff99560140a2e0 R14: ffff99560127d2e0 R15: 0000000000000000 [1950.279422] FS: 00007f67da795380(0000) GS:ffff995d1f0c0000(0000) knlGS:0000000000000000 [1950.279424] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [1950.279424] CR2: 00007f67da7e65c0 CR3: 00000001feed2000 CR4: 0000000000750ee0 [1950.279426] PKRU: 55555554 [1950.279426] Call Trace: [1950.279428] <TASK> [1950.279430] hwrng_unregister+0x28/0xe0 [rng_core] [1950.279436] tpm_chip_unregister+0xd5/0xf0 [tpm] Add the forgotten !tpm_amd_is_rng_defective() invariant to the hwrng_unregister() call site inside tpm_chip_unregister().	N/A	More Details
CVE-2022-50739	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Add null pointer check for inode operations This adds a sanity check for the i_op pointer of the inode which is returned after reading Root directory MFT record. We should check the i_op is valid before trying to create the root dentry, otherwise we may encounter a NPD while mounting a image with a funny Root directory MFT record. [114.484325] BUG: kernel NULL pointer dereference, address: 0000000000000008 [114.484811] #PF: supervisor read access in kernel mode [114.485084] #PF: error_code(0x0000) - not-present page [114.485606] PGD 0 P4D 0 [114.485975] Oops: 0000 [#1] PREEMPT SMP KASAN NOPTI [114.486570] CPU: 0 PID: 237 Comm: mount Tainted: G B 6.0.0-rc4 #28 [114.486977] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org 04/01/2014 [114.488169] RIP: 0010:d_flags_for_inode+0xe0/0x110 [114.488816] Code: 24 f7 ff 49 83 3e 00 74 41 41 83 cd 02 66 44 89 6b 02 eb 92 48 8d 7b 20 e8 6d 24 f7 ff 4c 8b 73 20 49 8d 7e 08 e8 60 241 [114.490326] RSP: 0018:ffff8880065e7aa8 EFLAGS: 00000296 [114.490695] RAX: 0000000000000001 RBX: ffff888008ccd750 RCX: ffffffff84af2aea [114.490986] RDX: 0000000000000001 RSI: 0000000000000008 RDI: ffffffff87abd020 [114.491364] RBP: ffff8880065e7ac8 R08: 0000000000000001 R09: fffffbfff0f57a05 [114.491675] R10: ffffffff87abd027 R11: fffffbfff0f57a04 R12: 0000000000000000 [114.491954] R13: 0000000000000008 R14: 0000000000000000 R15: ffff888008ccd750 [114.492397] FS: 00007fdc8a627e40(0000) GS:ffff888058200000(0000) knlGS:0000000000000000 [114.492797] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [114.493150] CR2: 0000000000000008 CR3: 00000000013ba000 CR4: 00000000000006f0 [114.493671] Call Trace: [114.493890] <TASK> [114.494075] __d_instantiate+0x24/0x1c0 [114.494505] d_instantiate.part.0+0x35/0x50 [114.494754] d_make_root+0x53/0x80 [114.494998] ntfs_fill_super+0x1232/0x1b50 [114.495260] ? put_ntfs+0x1d0/0x1d0 [114.495499] ? vsprintf+0x20/0x20 [114.495723] ? set_blocksize+0x95/0x150 [114.495964] get_tree_bdev+0x232/0x370 [114.496272] ? put_ntfs+0x1d0/0x1d0 [114.496502] ntfs_fs_get_tree+0x15/0x20 [114.496859] vfs_get_tree+0x4c/0x130 [114.497099] path_mount+0x654/0xfe0 [114.497507] ? putname+0x80/0xa0 [114.497933] ? finish_automount+0x2e0/0x2e0 [114.498362] ? putname+0x80/0xa0 [114.498571] ? kmem_cache_free+0x1c4/0x440 [114.498819] ? putname+0x80/0xa0 [114.499069] do_mount+0xd6/0xf0 [114.499343] ? path_mount+0xfe0/0xfe0 [114.499683] ? __kasan_check_write+0x14/0x20 [114.500133] __x64_sys_mount+0xca/0x110 [114.500592] do_syscall_64+0x3b/0x90 [114.500930] entry_SYSCALL_64_after_hwframe+0x63/0xcd [114.501294] RIP: 0033:0x7fdc898e948a [114.501542] Code: 48 8b 0d 11 fa 2a 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 49 89 ca b8 a5 00 00 008 [114.502716] RSP: 002b:00007ffd793e58f8 EFLAGS: 00000202 ORIG_RAX: 00000000000000a5 [114.503175] RAX: ffffffffffffffda RBX:	N/A	More Details

	0000564b2228f060 RCX: 00007fdc898e948a [114.503588] RDX: 0000564b2228f260 RSI: 0000564b2228f2e0 RDI: 0000564b22297ce0 [114.504925] RBP: 0000000000000000 R08: 0000564b2228f280 R09: 0000000000000020 [114.505484] R10: 00000000c0ed0000 R11: 0000000000000202 R12: 0000564b22297ce0 [114.505823] R13: 0000564b2228f260 R14: 0000000000000000 R15: 00000000ffffffff [114.506562] </TASK> [114.506887] Modules linked in: [114.507648] CR2: 0000000000000008 [114.508884] ---[end trace 0000000000000000]--- [114.509675] RIP: 0010:d_flags_for_inode+0xe0/0x110 [114.510140] Code: 24 f7 ff 49 83 3e 00 74 41 41 83 cd 02 66 44 89 6b 02 eb 92 48 8d 7b 20 e8 6d 24 f7 ff 4c 8b 73 20 49 8d 7e 08 e8 60 241 [114.511762] RSP: 0018:ffff8880065e7aa8 EFLAGS: 00000296 [114.512401] RAX: 0000000000000001 RBX: ffff888008ccd750 RCX: ffffffff84af2aea [114.51 ---truncated---		
CVE-2025-14177	In PHP versions:8.1.* before 8.1.34, 8.2.* before 8.2.30, 8.3.* before 8.3.29, 8.4.* before 8.4.16, 8.5.* before 8.5.1, the getimagesize() function may leak uninitialized heap memory into the APPn segments (e.g., APP1) when reading images in multi-chunk mode (such as via php://filter). This occurs due to a bug in php_read_stream_all_chunks() that overwrites the buffer without advancing the pointer, leaving tail bytes uninitialized. This may lead to information disclosure of sensitive heap data and affect the confidentiality of the target server.	N/A	More Details
CVE-2022-50738	In the Linux kernel, the following vulnerability has been resolved: vhost-vdpa: fix an iotlb memory leak Before commit 3d5698793897 ("vhost-vdpa: introduce asid based IOTLB") we called vhost_vdpa_iotlb_unmap(v, iotlb, 0ULL, 0ULL - 1) during release to free all the resources allocated when processing user IOTLB messages through vhost_vdpa_process_iotlb_update(). That commit changed the handling of IOTLB a bit, and we accidentally removed some code called during the release. We partially fixed this with commit 037d4305569a ("vhost-vdpa: call vhost_vdpa_cleanup during the release") but a potential memory leak is still there as showed by kmemleak if the application does not send VHOST_IOTLB_INVALIDATE or crashes: unreferenced object 0xffff888007fbaa30 (size 16): comm "blkio-bench", pid 914, jiffies 4294993521 (age 885.500s) hex dump (first 16 bytes): 40 73 41 07 80 88 ff ff 00 00 00 00 00 00 00 00 @sA..... backtrace: [<0000000087736d2a>] kmem_cache_alloc_trace+0x142/0x1c0 [<0000000060740f50>] vhost_vdpa_process_iotlb_msg+0x68c/0x901 [vhost_vdpa] [<0000000083e8e205>] vhost_chr_write_iter+0xc0/0x4a0 [vhost] [<000000008f2f414a>] vhost_vdpa_chr_write_iter+0x18/0x20 [vhost_vdpa] [<00000000de1cd4a0>] vfs_write+0x216/0x4b0 [<00000000a2850200>] ksys_write+0x71/0xf0 [<00000000de8e720b>] __x64_sys_write+0x19/0x20 [<0000000018b12cbb>] do_syscall_64+0x3f/0x90 [<00000000986ec465>] entry_SYSCALL_64_after_hwframe+0x63/0xcd Let's fix this calling vhost_vdpa_iotlb_unmap() on the whole range in vhost_vdpa_remove_as(). We move that call before vhost_dev_cleanup() since we need a valid v->vdev.mm in vhost_vdpa_pa_unmap(). vhost_iotlb_reset() call can be removed, since vhost_vdpa_iotlb_unmap() on the whole range removes all the entries. The kmemleak log reported was observed with a vDPA device that has `use_va` set to true (e.g. VDUSE). This patch has been tested with both types of devices.	N/A	More Details
CVE-2025-14180	In PHP versions 8.1.* before 8.1.34, 8.2.* before 8.2.30, 8.3.* before 8.3.29, 8.4.* before 8.4.16, 8.5.* before 8.5.1 when using the PDO PostgreSQL driver with PDO::ATTR_EMULATE_PREPARES enabled, an invalid character sequence (such as \x99) in a prepared statement parameter may cause the quoting function PQescapeStringConn to return NULL, leading to a null pointer dereference in pdo_parse_params() function. This may lead to crashes (segmentation fault) and affect the availability of the target server.	N/A	More Details
CVE-2025-68750	In the Linux kernel, the following vulnerability has been resolved: usb: potential integer overflow in usbg_make_tpg() The variable tpgt in usbg_make_tpg() is defined as unsigned long and is assigned to tpgt->tport_tpgt, which is defined as u16. This may cause an integer overflow when tpgt is greater than USHRT_MAX (65535). I haven't tried to trigger it myself, but it is possible to trigger it by calling usbg_make_tpg() with a large value for tpgt. I modified the type of tpgt to match tpgt->tport_tpgt and adjusted the relevant code accordingly. This patch is similar to commit 59c816c1f24d ("vhost/scsi: potential memory corruption").	N/A	More Details
CVE-2025-43876	Under certain circumstances a successful exploitation could result in access to the device.	N/A	More Details
CVE-2025-43875	Under certain circumstances a successful exploitation could result in access to the device.	N/A	More Details
CVE-2023-54074	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Use correct encap attribute during invalidation With introduction of post action infrastructure most of the users of encap attribute had been modified in order to obtain the correct attribute by calling mlx5e_tc_get_encap_attr() helper instead of assuming encap action is always on default attribute. However, the cited commit didn't modify mlx5e_invalidate_encap() which prevents it from destroying correct modify header action which leads to a warning [0]. Fix the issue by using correct attribute. [0]: Feb 21 09:47:35 c-237-177-40-045 kernel: WARNING: CPU: 17 PID: 654 at drivers/net/ethernet/mellanox/mlx5/core/en_tc.c:684 mlx5e_tc_attach_mod_hdr+0x1cc/0x230 [mlx5_core] Feb 21 09:47:35 c-237-177-40-045 kernel: RIP: 0010:mlx5e_tc_attach_mod_hdr+0x1cc/0x230 [mlx5_core] Feb 21 09:47:35 c-237-177-40-045 kernel: Call Trace: Feb 21 09:47:35 c-237-177-40-045 kernel: <TASK> Feb 21 09:47:35 c-237-177-40-045 kernel: mlx5e_tc_fib_event_work+0x8e3/0x1f60 [mlx5_core] Feb 21 09:47:35 c-237-177-40-045 kernel: ? mlx5e_take_all_encap_flows+0xe0/0xe0 [mlx5_core] Feb 21 09:47:35 c-237-177-40-045 kernel: ? lock_downgrade+0x6d0/0x6d0 Feb 21 09:47:35 c-237-177-40-045 kernel: ? lockdep_hardirqs_on_prepare+0x273/0x3f0 Feb 21 09:47:35 c-237-177-40-045 kernel: ? lockdep_hardirqs_on_prepare+0x273/0x3f0 Feb 21 09:47:35 c-237-177-40-045 kernel: process_one_work+0x7c2/0x1310 Feb 21 09:47:35 c-237-177-40-045 kernel: ? lockdep_hardirqs_on_prepare+0x3f0/0x3f0 Feb 21 09:47:35 c-237-177-40-045 kernel: ? pwq_dec_nr_in_flight+0x230/0x230 Feb 21 09:47:35 c-237-177-40-045 kernel: ? rwlock_bug.part.0+0x90/0x90 Feb 21 09:47:35 c-237-177-40-045 kernel: worker_thread+0x59d/0xec0 Feb 21 09:47:35 c-237-177-40-045 kernel: ? __kthread_parkme+0xd9/0x1d0	N/A	More Details
CVE-2023-54075	In the Linux kernel, the following vulnerability has been resolved: ASoC: mediatek: common: Fix refcount leak in parse_dai_link_info Add missing of_node_put()s before the returns to balance of_node_get()s and of_node_put()s, which may get unbalanced in case the for loop 'for_each_available_child_of_node' returns early.	N/A	More Details
CVE-2023-54076	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix missed ses refcounting Use new cifs_smb_ses_inc_refcount() helper to get an active reference of @ses and @ses->dfs_root_ses (if set). This will prevent @ses->dfs_root_ses of being put in the next call to cifs_put_smb_ses() and thus potentially causing an use-after-free bug.	N/A	More Details
CVE-2025-5448	Rejected reason: This CVE id was assigned but later discarded.	N/A	More Details
CVE-2025-	In the Linux kernel, the following vulnerability has been resolved: accel/ivpu: Fix race condition when unbinding BOs Fix 'Memory manager not clean during takedown' warning that occurs when ivpu_gem_bo_free() removes the BO from the BOs list before it gets unmapped. Then file_priv_unbind() triggers a warning in drm_mm_takedown() during context teardown. Protect the unmapping	N/A	More

68749	sequence with bo_list_lock to ensure the BO is always fully unmapped when removed from the list. This ensures the BO is either fully unmapped at context teardown time or present on the list and unmapped by file_priv_unbind().		Details
CVE-2025-68748	In the Linux kernel, the following vulnerability has been resolved: drm/panthor: Fix UAF race between device unplug and FW event processing The function panthor_fw_unplug() will free the FW memory sections. The problem is that there could still be pending FW events which are yet not handled at this point. process_fw_events_work() can in this case try to access said freed memory. Simply call disable_work_sync() to both drain and prevent future invocation of process_fw_events_work().	N/A	More Details
CVE-2025-68737	In the Linux kernel, the following vulnerability has been resolved: arm64/pageattr: Propagate return value from __change_memory_common The rodata=on security measure requires that any code path which does vmalloc -> set_memory_ro/set_memory_rox must protect the linear map alias too. Therefore, if such a call fails, we must abort set_memory_* and caller must take appropriate action; currently we are suppressing the error, and there is a real chance of such an error arising post commit a166563e7ec3 ("arm64: mm: support large block mapping when rodata=full"). Therefore, propagate any error to the caller.	N/A	More Details
CVE-2025-68735	In the Linux kernel, the following vulnerability has been resolved: drm/panthor: Prevent potential UAF in group creation This commit prevents the possibility of a use after free issue in the GROUP_CREATE ioctl function, which arose as pointer to the group is accessed in that ioctl function after storing it in the Xarray. A malicious userspace can second guess the handle of a group and try to call GROUP_DESTROY ioctl from another thread around the same time as GROUP_CREATE ioctl. To prevent the use after free exploit, this commit uses a mark on an entry of group pool Xarray which is added just before returning from the GROUP_CREATE ioctl function. The mark is checked for all ioctls that specify the group handle and so userspace won't be abe to delete a group that isn't marked yet. v2: Add R-bs and fixes tags	N/A	More Details
CVE-2025-68948	SiYuan is self-hosted, open source personal knowledge management software. In versions 3.5.1 and prior, the SiYuan Note application utilizes a hardcoded cryptographic secret for its session store. This unsafe practice renders the session encryption ineffective. Since the sensitive AccessAuthCode is stored within the session cookie, an attacker who intercepts or obtains a user's encrypted session cookie (e.g., via session hijacking) can locally decrypt it using the public key. Once decrypted, the attacker can retrieve the AccessAuthCode in plain text and use it to authenticate or take over the session.	N/A	More Details
CVE-2025-68733	In the Linux kernel, the following vulnerability has been resolved: smack: fix bug: unprivileged task can create labels If an unprivileged task is allowed to relabel itself (/smack/relabel-self is not empty), it can freely create new labels by writing their names into own /proc/PID/attr/smack/current This occurs because do_setattr() imports the provided label in advance, before checking "relabel-self" list. This change ensures that the "relabel-self" list is checked before importing the label.	N/A	More Details
CVE-2022-50713	In the Linux kernel, the following vulnerability has been resolved: clk: visconti: Fix memory leak in visconti_register_pll() @pll->rate_table has allocated memory by kmemdup(), if clk_hw_register() fails, it should be freed, otherwise it will cause memory leak issue, this patch fixes it.	N/A	More Details
CVE-2023-54087	In the Linux kernel, the following vulnerability has been resolved: ubi: Fix possible null-ptr-deref in ubi_free_volume() It willl cause null-ptr-deref in the following case: uif_init() ubi_add_volume() cdev_add() -> if it fails, call kill_volumes() device_register() kill_volumes() -> if ubi_add_volume() fails call this function ubi_free_volume() cdev_del() device_unregister() -> trying to delete a not added device, it causes null-ptr-deref So in ubi_free_volume(), it delete devices whether they are added or not, it will causes null-ptr-deref. Handle the error case whlie calling ubi_add_volume() to fix this problem. If add volume fails, set the corresponding vol to null, so it can not be accessed in kill_volumes() and release the resource in ubi_add_volume() error path.	N/A	More Details
CVE-2022-50712	In the Linux kernel, the following vulnerability has been resolved: devlink: hold region lock when flushing snapshots Netdevsim triggers a splat on reload, when it destroys regions with snapshots pending: WARNING: CPU: 1 PID: 787 at net/core/devlink.c:6291 devlink_region_snapshot_del+0x12e/0x140 CPU: 1 PID: 787 Comm: devlink Not tainted 6.1.0-07460-g7ae9888d6e1c #580 RIP: 0010:devlink_region_snapshot_del+0x12e/0x140 Call Trace: <TASK> devl_region_destroy+0x70/0x140 nsim_dev_reload_down+0x2f/0x60 [netdevsim] devlink_reload+0x1f7/0x360 devlink_nl_cmd_reload+0x6ce/0x860 genl_family_rcv_msg_doit.isra.0+0x145/0x1c0 This is the locking assert in devlink_region_snapshot_del(), we're supposed to be holding the region->snapshot_lock here.	N/A	More Details
CVE-2025-68734	In the Linux kernel, the following vulnerability has been resolved: isdn: mISDN: hfcsusb: fix memory leak in hfcsusb_probe() In hfcsusb_probe(), the memory allocated for ctrl_urb gets leaked when setup_instance() fails with an error code. Fix that by freeing the urb before freeing the hw structure. Also change the error paths to use the goto ladder style. Compile tested only. Issue found using a prototype static analysis tool.	N/A	More Details
CVE-2023-54088	In the Linux kernel, the following vulnerability has been resolved: blk-cgroup: hold queue_lock when removing blkg->q_node When blkg is removed from q->blkg_list from blkg_free_workfn(), queue_lock has to be held, otherwise, all kinds of bugs(list corruption, hard lockup, ..) can be triggered from blkg_destroy_all().	N/A	More Details
CVE-2023-54089	In the Linux kernel, the following vulnerability has been resolved: virtio_pmem: add the missing REQ_OP_WRITE for flush bio When doing mkfs.xfs on a pmem device, the following warning was ------------[cut here]------------ WARNING: CPU: 2 PID: 384 at block/blk-core.c:751 submit_bio_noacct Modules linked in: CPU: 2 PID: 384 Comm: mkfs.xfs Not tainted 6.4.0-rc7+ #154 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996) RIP: 0010:submit_bio_noacct+0x340/0x520 ..... Call Trace: <TASK> ? submit_bio_noacct+0xd5/0x520 submit_bio+0x37/0x60 async_pmem_flush+0x79/0xa0 nvdimm_flush+0x17/0x40 pmem_submit_bio+0x370/0x390 __submit_bio+0xbc/0x190 submit_bio_noacct_nocheck+0x14d/0x370 submit_bio_noacct+0x1ef/0x520 submit_bio+0x55/0x60 submit_bio_wait+0x5a/0xc0 blkdev_issue_flush+0x44/0x60 The root cause is that submit_bio_noacct() needs bio_op() is either WRITE or ZONE_APPEND for flush bio and async_pmem_flush() doesn't assign REQ_OP_WRITE when allocating flush bio, so submit_bio_noacct just fail the flush bio. Simply fix it by adding the missing REQ_OP_WRITE for flush bio. And we could fix the flush order issue and do flush optimization later.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ixgbe: Fix panic during XDP_TX with > 64 CPUs Commit 4fe815850bdc ("ixgbe: let the xdpdrv work with more than 64 cpus") adds support to allow XDP programs to run on systems with more than 64 CPUs by locking the XDP TX rings and indexing them using cpu % 64 (IXGBE_MAX_XDP_QS). Upon trying this out patch on a system with more than 64 cores, the kernel paniced with an array-index-out-of-bounds at the return in ixgbe_determine_xdp_ring in ixgbe.h, which means ixgbe_determine_xdp_q_idx was just returning the cpu instead of cpu % IXGBE_MAX_XDP_QS. An example splat: == UBSAN: array-index-out-of-bounds in /var/lib/dkms/ixgbe/5.18.6+focal-1/build/src/ixgbe.h:1147:26 index 65 is out of range for type 'ixgbe_ring *[64]' == BUG:		

CVE-2023-54090	kernel NULL pointer dereference, address: 0000000000000058 #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: 0000 [#1] SMP NOPTI CPU: 65 PID: 408 Comm: ksoftirqd/65 Tainted: G IOE 5.15.0-48-generic #54~20.04.1-Ubuntu Hardware name: Dell Inc. PowerEdge R640/0W23H8, BIOS 2.5.4 01/13/2020 RIP: 0010:ixgbe_xmit_xdp_ring+0x1b/0x1c0 [ixgbe] Code: 3b 52 d4 cf e9 42 f2 ff ff 66 0f 1f 44 00 00 0f 1f 44 00 00 55 b9 00 00 00 00 48 89 e5 41 57 41 56 41 55 41 54 53 48 83 ec 08 <44> 0f b7 47 58 0f b7 47 5a 0f b7 57 54 44 0f b7 76 08 66 41 39 c0 RSP: 0018:ffffbc3fcd88fcb0 EFLAGS: 00010282 RAX: ffff92a253260980 RBX: ffffbc3fe68b00a0 RCX: 0000000000000000 RDX: ffff928b5f659000 RSI: ffff928b5f659000 RDI: 0000000000000000 RBP: ffffbc3fcd88fce0 R08: ffff92b9dfc20580 R09: 0000000000000001 R10: 3d3d3d3d3d3d3d3d R11: 3d3d3d3d3d3d3d3d R12: 0000000000000000 R13: ffff928b2f0fa8c0 R14: ffff928b9be20050 R15: 000000000000003c FS: 0000000000000000(0000) GS:ffff92b9dfc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000058 CR3: 000000011dd6a002 CR4: 00000000007706e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 55555554 Call Trace: <TASK> ixgbe_poll+0x103e/0x1280 [ixgbe] ? sched_clock_cpu+0x12/0xe0 __napi_poll+0x30/0x160 net_rx_action+0x11c/0x270 __do_softirq+0xda/0x2ee run_ksoftirqd+0x2f/0x50 smpboot_thread_fn+0xb7/0x150 ? sort_range+0x30/0x30 kthread+0x127/0x150 ? set_kthread_struct+0x50/0x50 ret_from_fork+0x1f/0x30 </TASK> I think this is how it happens: Upon loading the first XDP program on a system with more than 64 CPUs, ixgbe_xdp_locking_key is incremented in ixgbe_xdp_setup. However, immediately after this, the rings are reconfigured by ixgbe_setup_tc. ixgbe_setup_tc calls ixgbe_clear_interrupt_scheme which calls ixgbe_free_q_vectors which calls ixgbe_free_q_vector in a loop. ixgbe_free_q_vector decrements ixgbe_xdp_locking_key once per call if it is non-zero. Commenting out the decrement in ixgbe_free_q_vector stopped my system from panicing. I suspect to make the original patch work, I would need to load an XDP program and then replace it in order to get ixgbe_xdp_locking_key back above 0 since ixgbe_setup_tc is only called when transitioning between XDP and non-XDP ring configurations, while ixgbe_xdp_locking_key is incremented every time ixgbe_xdp_setup is called. Also, ixgbe_setup_tc can be called via ethtool --set-channels, so this becomes another path to decrement ixgbe_xdp_locking_key to 0 on systems with more than 64 CPUs. Since ixgbe_xdp_locking_key only protects the XDP_TX path and is tied to the number of CPUs present, there is no reason to disable it upon unloading an XDP program. To avoid confusion, I have moved enabling ixgbe_xdp_locking_key into ixgbe_sw_init, which is part of the probe path.	N/A	More Details
CVE-2023-54091	In the Linux kernel, the following vulnerability has been resolved: drm/client: Fix memory leak in drm_client_target_cloned dmt_mode is allocated and never freed in this function. It was found with the ast driver, but most drivers using generic fbdev setup are probably affected. This fixes the following kmemleak report: backtrace: [<00000000b391296d>] drm_mode_duplicate+0x45/0x220 [drm] [<00000000e45bb5b3>] drm_client_target_cloned.constprop.0+0x27b/0x480 [drm] [<00000000ed2d3a37>] drm_client_modeset_probe+0x6bd/0xf50 [drm] [<0000000010e5cc9d>] __drm_fb_helper_initial_config_and_unlock+0xb4/0x2c0 [drm_kms_helper] [<00000000909f82ca>] drm_fbdev_client_hotplug+0x2bc/0x4d0 [drm_kms_helper] [<00000000063a69aa>] drm_client_register+0x169/0x240 [drm] [<00000000a8c61525>] ast_pci_probe+0x142/0x190 [ast] [<00000000987f19bb>] local_pci_probe+0xdc/0x180 [<000000004fca231b>] work_for_cpu_fn+0x4e/0xa0 [<0000000000b85301>] process_one_work+0x8b7/0x1540 [<000000003375b17c>] worker_thread+0x70a/0xed0 [<00000000b0d43cd9>] kthread+0x29f/0x340 [<000000008d770833>] ret_from_fork+0x1f/0x30 unreferenced object 0xff11000333089a00 (size 128):	N/A	More Details
CVE-2025-68732	In the Linux kernel, the following vulnerability has been resolved: gpu: host1x: Fix race in syncpt alloc/free Fix race condition between host1x_syncpt_alloc() and host1x_syncpt_put() by using kref_put_mutex() instead of kref_put() + manual mutex locking. This ensures no thread can acquire the syncpt_mutex after the refcount drops to zero but before syncpt_release acquires it. This prevents races where syncpoints could be allocated while still being cleaned up from a previous release. Remove explicit mutex locking in syncpt_release as kref_put_mutex() handles this atomically.	N/A	More Details
CVE-2022-50725	In the Linux kernel, the following vulnerability has been resolved: media: vidtv: Fix use-after-free in vidtv_bridge_dvb_init() KASAN reports a use-after-free: BUG: KASAN: use-after-free in dvb_dmxdev_release+0x4d5/0x5d0 [dvb_core] Call Trace: ... dvb_dmxdev_release+0x4d5/0x5d0 [dvb_core] vidtv_bridge_probe+0x7bf/0xa40 [dvb_vidtv_bridge] platform_probe+0xb6/0x170 ... Allocated by task 1238: ... dvb_register_device+0x1a7/0xa70 [dvb_core] dvb_dmxdev_init+0x2af/0x4a0 [dvb_core] vidtv_bridge_probe+0x766/0xa40 [dvb_vidtv_bridge] ... Freed by task 1238: dvb_register_device+0x6d2/0xa70 [dvb_core] dvb_dmxdev_init+0x2af/0x4a0 [dvb_core] vidtv_bridge_probe+0x766/0xa40 [dvb_vidtv_bridge] ... It is because the error handling in vidtv_bridge_dvb_init() is wrong. First, vidtv_bridge_dmx(dev)_init() will clean themselves when fail, but goto fail_dmx(_dev): calls release functions again, which causes use-after-free. Also, in fail_fe, fail_tuner_probe and fail_demod_probe, j = i will cause out-of-bound when i finished its loop (i == NUM_FE). And the loop releasing is wrong, although now NUM_FE is 1 so it won't cause problem. Fix this by correctly releasing everything.	N/A	More Details
CVE-2025-68731	In the Linux kernel, the following vulnerability has been resolved: accel/amdxdna: Fix an integer overflow in aie2_query_ctx_status_array() The unpublished smatch static checker reported a warning. drivers/accel/amdxdna/aie2_pci.c:904 aie2_query_ctx_status_array() warn: potential user controlled sizeof overflow 'args->num_element * args->element_size' '1-u32max(user) * 1-u32max(user)' Even this will not cause a real issue, it is better to put a reasonable limitation for element_size and num_element. Add condition to make sure the input element_size <= 4K and num_element <= 1K.	N/A	More Details
CVE-2025-68730	In the Linux kernel, the following vulnerability has been resolved: accel/ivpu: Fix page fault in ivpu_bo_unbind_all_bos_from_context() Don't add BO to the vdev->bo_list in ivpu_gem_create_object(). When failure happens inside drm_gem_shmem_create(), the BO is not fully created and ivpu_gem_bo_free() callback will not be called causing a deleted BO to be left on the list.	N/A	More Details
CVE-2025-68729	In the Linux kernel, the following vulnerability has been resolved: wifi: ath12k: Fix MSDU buffer types handling in RX error path Currently, packets received on the REO exception ring from unassociated peers are of MSDU buffer type, while the driver expects link descriptor type packets. These packets are not parsed further due to a return check on packet type in ath12k_hal_desc_reo_parse_err(), but the associated skb is not freed. This may lead to kernel crashes and buffer leaks. Hence to fix, update the RX error handler to explicitly drop MSDU buffer type packets received on the REO exception ring. This prevents further processing of invalid packets and ensures stability in the RX error handling path. Tested-on: QCN9274 hw2.0 PCI WLAN.WBE.1.4.1-00199-QCAHKSWPL_SILICONZ-1	N/A	More Details
CVE-2025-68728	In the Linux kernel, the following vulnerability has been resolved: ntfs3: fix uninit memory after failed mi_read in mi_format_new Fix a KMSAN un-init bug found by syzkaller. ntfs_get_bh() expects a buffer from sb_getblk(), that buffer may not be uptodate. We do not bring the buffer uptodate before setting it as uptodate. If the buffer were to not be uptodate, it could mean adding a buffer with un-init data to the mi record. Attempting to load that record will trigger KMSAN. Avoid this by setting the buffer as uptodate, if it’s not already, by overwriting it.	N/A	More Details
CVE-2025-68727	In the Linux kernel, the following vulnerability has been resolved: ntfs3: Fix uninit buffer allocated by __getname() Fix uninit errors caused after buffer allocation given to 'de'; by initializing the buffer with zeroes. The fix was found by using KMSAN.	N/A	More Details

CVE-2025-68726	In the Linux kernel, the following vulnerability has been resolved: crypto: aead - Fix reqsize handling Commit afddce13ce81d ("crypto: api - Add reqsize to crypto_alg") introduced cra_reqsize field in crypto_alg struct to replace type specific reqsize fields. It looks like this was introduced specifically for ahash and acomp from the commit description as subsequent commits add necessary changes in these alg frameworks. However, this is being recommended for use in all crypto algs instead of setting reqsize using crypto_*_set_reqsize(). Using cra_reqsize in aead algorithms, hence, causes memory corruptions and crashes as the underlying functions in the algorithm framework have not been updated to set the reqsize properly from cra_reqsize. [1] Add proper set_reqsize calls in the aead init function to properly initialize reqsize for these algorithms in the framework. [1]: https://gist.github.com/Pratham-T/24247446f1faf4b7843e4014d5089f6b	N/A	More Details
CVE-2025-68725	In the Linux kernel, the following vulnerability has been resolved: bpf: Do not let BPF test infra emit invalid GSO types to stack Yinhao et al. reported that their fuzzer tool was able to trigger a skb_warn_bad_offload() from netif_skb_features() -> gso_features_check(). When a BPF program - triggered via BPF test infra - pushes the packet to the loopback device via bpf_clone_redirect() then mentioned offload warning can be seen. GSO-related features are then rightfully disabled. We get into this situation due to convert___skb_to_skb() setting gso_segs and gso_size but not gso_type. Technically, it makes sense that this warning triggers since the GSO properties are malformed due to the gso_type. Potentially, the gso_type could be marked non-trustworthy through setting it at least to SKB_GSO_DODGY without any other specific assumptions, but that also feels wrong given we should not go further into the GSO engine in the first place. The checks were added in 121d57af308d ("gso: validate gso_type in GSO handlers") because there were malicious (syzbot) senders that combine a protocol with a non-matching gso_type. If we would want to drop such packets, gso_features_check() currently only returns feature flags via netif_skb_features(), so one location for potentially dropping such skbs could be validate_xmit_unreadable_skb(), but then otoh it would be an additional check in the fast-path for a very corner case. Given bpf_clone_redirect() is the only place where BPF test infra could emit such packets, lets reject them right there.	N/A	More Details
CVE-2025-68724	In the Linux kernel, the following vulnerability has been resolved: crypto: asymmetric_keys - prevent overflow in asymmetric_key_generate_id Use check_add_overflow() to guard against potential integer overflows when adding the binary blob lengths and the size of an asymmetric_key_id structure and return ERR_PTR(-EOVERFLOW) accordingly. This prevents a possible buffer overflow when copying data from potentially malicious X.509 certificate fields that can be arbitrarily large, such as ASN.1 INTEGER serial numbers, issuer names, etc.	N/A	More Details
CVE-2023-54086	In the Linux kernel, the following vulnerability has been resolved: bpf: Add preempt_count_{sub,add} into btf id deny list The recursion check in __bpf_prog_enter* and __bpf_prog_exit* leave preempt_count_{sub,add} unprotected. When attaching trampoline to them we get panic as follows, [867.843050] BUG: TASK stack guard page was hit at 0000000009d325cf (stack is 0000000046a46a15..00000000537e7b28) [867.843064] stack guard page: 0000 [#1] PREEMPT SMP NOPTI [867.843067] CPU: 8 PID: 11009 Comm: trace Kdump: loaded Not tainted 6.2.0+ #4 [867.843100] Call Trace: [867.843101] <TASK> [867.843104] asm_exc_int3+0x3a/0x40 [867.843108] RIP: 0010:preempt_count_sub+0x1/0xa0 [867.843135] __bpf_prog_enter_recur+0x17/0x90 [867.843148] bpf_trampoline_6442468108_0+0x2e/0x1000 [867.843154] ? preempt_count_sub+0x1/0xa0 [867.843157] preempt_count_sub+0x5/0xa0 [867.843159] ? migrate_enable+0xac/0xf0 [867.843164] __bpf_prog_exit_recur+0x2d/0x40 [867.843168] bpf_trampoline_6442468108_0+0x55/0x1000 ... [867.843788] preempt_count_sub+0x5/0xa0 [867.843793] ? migrate_enable+0xac/0xf0 [867.843829] __bpf_prog_exit_recur+0x2d/0x40 [867.843837] BUG: IRQ stack guard page was hit at 0000000099bd8228 (stack is 00000000b23e2bc4..000000006d95af35) [867.843841] BUG: IRQ stack guard page was hit at 000000005ae07924 (stack is 00000000ffd69623..0000000014eb594c) [867.843843] BUG: IRQ stack guard page was hit at 00000000028320f0 (stack is 00000000034b6438..0000000078d1bcec) [867.843842] bpf_trampoline_6442468108_0+0x55/0x1000 ... That is because in __bpf_prog_exit_recur, the preempt_count_{sub,add} are called after prog->active is decreased. Fixing this by adding these two functions into btf ids deny list.	N/A	More Details
CVE-2023-54085	In the Linux kernel, the following vulnerability has been resolved: mptcp: fix NULL pointer dereference on fastopen early fallback In case of early fallback to TCP, subflow_syn_recv_sock() deletes the subflow context before returning the newly allocated sock to the caller. The fastopen path does not cope with the above unconditionally dereferencing the subflow context.	N/A	More Details
CVE-2023-54084	In the Linux kernel, the following vulnerability has been resolved: ALSA: firewire-digi00x: prevent potential use after free This code was supposed to return an error code if init_stream() failed, but it instead freed dg00x->rx_stream and returned success. This potentially leads to a use after free.	N/A	More Details
CVE-2023-54083	In the Linux kernel, the following vulnerability has been resolved: phy: tegra: xusb: Clear the driver reference in usb-phy dev For the dual-role port, it will assign the phy dev to usb-phy dev and use the port dev driver as the dev driver of usb-phy. When we try to destroy the port dev, it will destroy its dev driver as well. But we did not remove the reference from usb-phy dev. This might cause the use-after-free issue in KASAN.	N/A	More Details
CVE-2023-54077	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix memory leak if ntfs_read_mft failed Label ATTR_ROOT in ntfs_read_mft() sets is_root = true and ni->ni_flags |= NI_FLAG_DIR, then next attr will goto label ATTR_ALLOC and alloc ni->dir.alloc_run. However two states are not always consistent and can make memory leak. 1) attr_name in ATTR_ROOT does not fit the condition it will set is_root = true but NI_FLAG_DIR is not set. 2) next attr_name in ATTR_ALLOC fits the condition and alloc ni->dir.alloc_run 3) in cleanup function ni_clear(), when NI_FLAG_DIR is set, it frees ni->dir.alloc_run, otherwise it frees ni->file.run 4) because NI_FLAG_DIR is not set in this case, ni->dir.alloc_run is leaked as kmemleak reported: unreferenced object 0xffff888003bc5480 (size 64): backtrace: [<000000003d42e6b0>] __kmalloc_node+0x4e/0x1c0 [<00000000d8e19b8a>] kvmalloc_node+0x39/0x1f0 [<00000000fc3eb5b8>] run_add_entry+0x18a/0xa40 [ntfs3] [<0000000011c9f978>] run_unpack+0x75d/0x8e0 [ntfs3] [<00000000e7cf1819>] run_unpack_ex+0xbc/0x500 [ntfs3] [<00000000bbf0a43d>] ntfs_iget5+0xb25/0x2dd0 [ntfs3] [<00000000a6e50693>] ntfs_fill_super+0x218d/0x3580 [ntfs3] [<00000000b9170608>] get_tree_bdev+0x3fb/0x710 [<000000004833798a>] vfs_get_tree+0x8e/0x280 [<000000006e20b8e6>] path_mount+0xf3c/0x1930 [<000000007bf15a5f>] do_mount+0xf3/0x110 ... Fix this by always setting is_root and NI_FLAG_DIR together.	N/A	More Details
CVE-2023-54078	In the Linux kernel, the following vulnerability has been resolved: media: max9286: Free control handler The control handler is leaked in some probe-time error paths, as well as in the remove path. Fix it.	N/A	More Details
CVE-2023-54079	In the Linux kernel, the following vulnerability has been resolved: power: supply: bq27xxx: Fix poll_interval handling and races on remove Before this patch bq27xxx_battery_teardown() was setting poll_interval = 0 to avoid bq27xxx_battery_update() requeuing the delayed_work item. There are 2 problems with this: 1. If the driver is unbound through sysfs, rather then the module being rmmod-ed, this changes poll_interval unexpectedly 2. This is racy, after it being set poll_interval could be changed before bq27xxx_battery_update() checks it through /sys/module/bq27xxx_battery/parameters/poll_interval Fix this by added a removed attribute to struct bq27xxx_device_info and using that instead of setting poll_interval to 0. There also is another poll_interval related	N/A	More Details

	race on remove(), writing /sys/module/bq27xxx_battery/parameters/poll_interval will requeue the delayed_work item for all devices on the bq27xxx_battery_devices list and the device being removed was only removed from that list after cancelling the delayed_work item. Fix this by moving the removal from the bq27xxx_battery_devices list to before cancelling the delayed_work item.		
CVE-2023-54080	In the Linux kernel, the following vulnerability has been resolved: btrfs: zoned: skip splitting and logical rewriting on pre-alloc write When doing a relocation, there is a chance that at the time of btrfs_reloc_clone_csums(), there is no checksum for the corresponding region. In this case, btrfs_finish_ordered_zoned()'s sum points to an invalid item and so ordered_extent's logical is set to some invalid value. Then, btrfs_lookup_block_group() in btrfs_zone_finish_endio() failed to find a block group and will hit an assert or a null pointer dereference as following. This can be reprodcued by running btrfs/028 several times (e.g, 4 to 16 times) with a null_blk setup. The device's zone size and capacity is set to 32 MB and the storage size is set to 5 GB on my setup. KASAN: null-ptr-deref in range [0x0000000000000088-0x000000000000008f] CPU: 6 PID: 3105720 Comm: kworker/u16:13 Tainted: G W 6.5.0-rc6-kts+ #1 Hardware name: Supermicro Super Server/X10SRL-F, BIOS 2.0 12/17/2015 Workqueue: btrfs-endio-write btrfs_work_helper [btrfs] RIP: 0010:btrfs_zone_finish_endio.part.0+0x34/0x160 [btrfs] Code: 41 54 49 89 fc 55 48 89 f5 53 e8 57 7d fc ff 48 8d b8 88 00 00 00 48 89 c3 48 b8 00 00 00 00 00 > 3c 02 00 0f 85 02 01 00 00 f6 83 88 00 00 00 01 0f 84 a8 00 00 RSP: 0018:ffff88833cf87b08 EFLAGS: 00010206 RAX: dfffc0000000000 RBX: 0000000000000000 RCX: 0000000000000000 RDX: 0000000000000011 RSI: 0000000000000004 RDI: 0000000000000088 RBP: 0000000000000002 R08: 0000000000000001 R09: ffffed102877b827 R10: ffff888143bdc13b R11: ffff888125b1cbc0 R12: ffff888143bdc000 R13: 0000000000007000 R14: ffff888125b1cba8 R15: 0000000000000000 FS: 0000000000000000(0000) GS:ffff88881e500000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007f3ed85223d5 CR3: 00000001519b4005 CR4: 00000000001706e0 Call Trace: <TASK> ? die_addr+0x3c/0xa0 ? exc_general_protection+0x148/0x220 ? asm_exc_general_protection+0x22/0x30 ? btrfs_zone_finish_endio.part.0+0x34/0x160 [btrfs] ? btrfs_zone_finish_endio.part.0+0x19/0x160 [btrfs] btrfs_finish_one_ordered+0x7b8/0x1de0 [btrfs] ? rcu_is_watching+0x11/0xb0 ? lock_release+0x47a/0x620 ? btrfs_finish_ordered_zoned+0x59b/0x800 [btrfs] ? __pfx_btrfs_finish_one_ordered+0x10/0x10 [btrfs] ? btrfs_finish_ordered_zoned+0x358/0x800 [btrfs] ? __smp_call_single_queue+0x124/0x350 ? rcu_is_watching+0x11/0xb0 btrfs_work_helper+0x19f/0xc60 [btrfs] ? __pfx_try_to_wake_up+0x10/0x10 ? _raw_spin_unlock_irq+0x24/0x50 ? rcu_is_watching+0x11/0xb0 process_one_work+0x8c1/0x1430 ? __pfx_lock_acquire+0x10/0x10 ? __pfx_process_one_work+0x10/0x10 ? __pfx_do_raw_spin_lock+0x10/0x10 ? _raw_spin_lock_irq+0x52/0x60 worker_thread+0x100/0x12c0 ? __kthread_parkme+0xc1/0x1f0 ? __pfx_worker_thread+0x10/0x10 kthread+0x2ea/0x3c0 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x30/0x70 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1b/0x30 </TASK> On the zoned mode, writing to pre-allocated region means data relocation write. Such write always uses WRITE command so there is no need of splitting and rewriting logical address. Thus, we can just skip the function for the case.	N/A	More Details
CVE-2022-50724	In the Linux kernel, the following vulnerability has been resolved: regulator: core: fix resource leak in regulator_register() I got some resource leak reports while doing fault injection test: OF: ERROR: memory leak, expected refcount 1 instead of 100, of_node_get()/of_node_put() unbalanced - destroy cset entry: attach overlay node /i2c/pmic@64/regulators/buck1 unreferenced object 0xffff88810deea000 (size 512): comm "490-i2c-rt5190a", pid 253, jiffies 4294859840 (age 5061.046s) hex dump (first 32 bytes): 00 00 00 00 ad 4e ad de ff ff ff ff 00 00 00 00 .....N..... ff ff ff ff ff ff ff ff a0 1e 00 a1 ff ff ff ff ..... backtrace: [<00000000d78541e2>] kmalloc_trace+0x21/0x110 [<00000000b343d153>] device_private_init+0x32/0xd0 [<00000000be1f0c70>] device_add+0xb2d/0x1030 [<00000000e3e6344d>] regulator_register+0xaf2/0x12a0 [<00000000e2f5e754>] devm_regulator_register+0x57/0xb0 [<000000008b898197>] rt5190a_probe+0x52a/0x861 [rt5190a_regulator] unreferenced object 0xffff88810b617b80 (size 32): comm "490-i2c-rt5190a", pid 253, jiffies 4294859904 (age 5060.983s) hex dump (first 32 bytes): 72 65 67 75 6c 61 74 6f 72 2e 32 38 36 38 2d 53 regulator.2868-S 55 50 50 4c 59 00 ff ff 29 00 00 00 2b 00 00 00 UPPLY...)...+... backtrace: [<000000009da9280d>] __kmalloc_node_track_caller+0x44/0x1b0 [<0000000025c6a4e5>] kstrdup+0x3a/0x70 [<00000000790efb69>] create_regulator+0xc0/0x4e0 [<0000000005ed203a>] regulator_resolve_supply+0x2d4/0x440 [<0000000045796214>] regulator_register+0x10b3/0x12a0 [<00000000e2f5e754>] devm_regulator_register+0x57/0xb0 [<000000008b898197>] rt5190a_probe+0x52a/0x861 [rt5190a_regulator] After calling regulator_resolve_supply(), the 'rdev->supply' is set by set_supply(), after this set, in the error path, the resources need be released, so call regulator_put() to avoid the leaks.	N/A	More Details
CVE-2022-50723	In the Linux kernel, the following vulnerability has been resolved: bnxt_en: fix memory leak in bnxt_nvm_test() Free the kzalloc'ed buffer before returning in the success path.	N/A	More Details
CVE-2022-50722	In the Linux kernel, the following vulnerability has been resolved: media: ipu3-imgu: Fix NULL pointer dereference in active selection access What the IMGU driver did was that it first acquired the pointers to active and try V4L2 subdev state, and only then figured out which one to use. The problem with that approach and a later patch (see Fixes: tag) is that as sd_state argument to v4l2_subdev_get_try_crop() et al is NULL, there is now an attempt to dereference that. Fix this. Also rewrap lines a little.	N/A	More Details
CVE-2022-50721	In the Linux kernel, the following vulnerability has been resolved: dmaengine: qcom-adm: fix wrong calling convention for prep_slave_sg The calling convention for pre_slave_sg is to return NULL on error and provide an error log to the system. Qcom-adm instead provide error pointer when an error occur. This indirectly cause kernel panic for example for the nandc driver that checks only if the pointer returned by device_prep_slave_sg is not NULL. Returning an error pointer makes nandc think the device_prep_slave_sg function correctly completed and makes the kernel panics later in the code. While nandc is the one that makes the kernel crash, it was pointed out that the real problem is qcom-adm not following calling convention for that function. To fix this, drop returning error pointer and return NULL with an error log.	N/A	More Details
CVE-2023-54081	In the Linux kernel, the following vulnerability has been resolved: xen: speed up grant-table reclaim When a grant entry is still in use by the remote domain, Linux must put it on a deferred list. Normally, this list is very short, because the PV network and block protocols expect the backend to unmap the grant first. However, Qubes OS's GUI protocol is subject to the constraints of the X Window System, and as such winds up with the frontend unmapping the window first. As a result, the list can grow very large, resulting in a massive memory leak and eventual VM freeze. To partially solve this problem, make the number of entries that the VM will attempt to free at each iteration tunable. The default is still 10, but it can be overridden via a module parameter. This is Cc: stable because (when combined with appropriate userspace changes) it fixes a severe performance and stability problem for Qubes OS users.	N/A	More Details
CVE-2022-50720	In the Linux kernel, the following vulnerability has been resolved: x86/apic: Don't disable x2APIC if locked The APIC supports two modes, legacy APIC (or xAPIC), and Extended APIC (or x2APIC). X2APIC mode is mostly compatible with legacy APIC, but it disables the memory-mapped APIC interface in favor of one that uses MSRs. The APIC mode is controlled by the EXT bit in the APIC MSR. The MMIO/xAPIC interface has some problems, most notably the APIC LEAK [1]. This bug allows an attacker to use the APIC MMIO interface to extract data from the SGX enclave. Introduce support for a new feature that will allow the BIOS to lock the APIC in x2APIC mode. If the APIC is locked in x2APIC mode and the kernel tries to disable the APIC or revert to legacy APIC mode a GP fault	N/A	More Details

	will occur. Introduce support for a new MSR (IA32_XAPIC_DISABLE_STATUS) and handle the new locked mode when the LEGACY_XAPIC_DISABLED bit is set by preventing the kernel from trying to disable the x2APIC. On platforms with the IA32_XAPIC_DISABLE_STATUS MSR, if SGX or TDX are enabled the LEGACY_XAPIC_DISABLED will be set by the BIOS. If legacy APIC is required, then it SGX and TDX need to be disabled in the BIOS. [1]: https://aepicleak.com/aepicleak.pdf		
CVE-2022-50719	In the Linux kernel, the following vulnerability has been resolved: ALSA: line6: fix stack overflow in line6_midi_transmit Correctly calculate available space including the size of the chunk buffer. This fixes a buffer overflow when multiple MIDI sysex messages are sent to a PODxt device.	N/A	More Details
CVE-2022-50718	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix pci device refcount leak As comment of pci_get_domain_bus_and_slot() says, it returns a pci device with refcount increment, when finish using it, the caller must decrement the reference count by calling pci_dev_put(). So before returning from amdgpu_device_resume|suspend_display_audio(), pci_dev_put() is called to avoid refcount leak.	N/A	More Details
CVE-2022-50717	In the Linux kernel, the following vulnerability has been resolved: nvmet-tcp: add bounds check on Transfer Tag ttag is used as an index to get cmd in nvmet_tcp_handle_h2c_data_pdu(), add a bounds check to avoid out-of-bounds access.	N/A	More Details
CVE-2022-50716	In the Linux kernel, the following vulnerability has been resolved: wifi: ar5523: Fix use-after-free on ar5523_cmd() timed out syzkaller reported use-after-free with the stack trace like below [1]: [38.960489][C3] == [38.963216][C3] BUG: KASAN: use-after-free in ar5523_cmd_tx_cb+0x220/0x240 [38.964950][C3] Read of size 8 at addr ffff888048e03450 by task swapper/3/0 [38.966363][C3] [38.967053][C3] CPU: 3 PID: 0 Comm: swapper/3 Not tainted 6.0.0-09039-ga6afa4199d3d-dirty #18 [38.968464][C3] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.16.0-1.fc36 04/01/2014 [38.969959][C3] Call Trace: [38.970841][C3] <IRQ> [38.971663][C3] dump_stack_lvl+0xfc/0x174 [38.972620][C3] print_report.cold+0x2c3/0x752 [38.973626][C3] ? ar5523_cmd_tx_cb+0x220/0x240 [38.974644][C3] kasan_report+0xb1/0x1d0 [38.975720][C3] ? ar5523_cmd_tx_cb+0x220/0x240 [38.976831][C3] ar5523_cmd_tx_cb+0x220/0x240 [38.978412][C3] __usb_hcd_giveback_urb+0x353/0x5b0 [38.979755][C3] usb_hcd_giveback_urb+0x385/0x430 [38.981266][C3] dummy_timer+0x140c/0x34e0 [38.982925][C3] ? notifier_call_chain+0xb5/0x1e0 [38.984761][C3] ? rcu_read_lock_sched_held+0xb/0x60 [38.986242][C3] ? lock_release+0x51c/0x790 [38.987323][C3] ? _raw_read_unlock_irqrestore+0x37/0x70 [38.988483][C3] ? __wake_up_common_lock+0xde/0x130 [38.989621][C3] ? reacquire_held_locks+0x4a0/0x4a0 [38.990777][C3] ? lock_acquire+0x472/0x550 [38.991919][C3] ? rcu_read_lock_sched_held+0xb/0x60 [38.993138][C3] ? lock_acquire+0x472/0x550 [38.994890][C3] ? dummy_urb_enqueue+0x860/0x860 [38.996266][C3] ? do_raw_spin_unlock+0x16f/0x230 [38.997670][C3] ? dummy_urb_enqueue+0x860/0x860 [38.999116][C3] call_timer_fn+0x1a0/0x6a0 [39.000668][C3] ? add_timer_on+0x4a0/0x4a0 [39.002137][C3] ? reacquire_held_locks+0x4a0/0x4a0 [39.003809][C3] ? __next_timer_interrupt+0x226/0x2a0 [39.005509][C3] __run_timers.part.0+0x69a/0xac0 [39.007025][C3] ? dummy_urb_enqueue+0x860/0x860 [39.008716][C3] ? call_timer_fn+0x6a0/0x6a0 [39.010254][C3] ? cpuacct_percpu_seq_show+0x10/0x10 [39.011795][C3] ? kvm_sched_clock_read+0x14/0x40 [39.013277][C3] ? sched_clock_cpu+0x69/0x2b0 [39.014724][C3] run_timer_softirq+0xb6/0x1d0 [39.016196][C3] __do_softirq+0x1d2/0x9be [39.017616][C3] __irq_exit_rcu+0xeb/0x190 [39.019004][C3] irq_exit_rcu+0x5/0x20 [39.020361][C3] sysvec_apic_timer_interrupt+0x8f/0xb0 [39.021965][C3] </IRQ> [39.023237][C3] <TASK> In ar5523_probe(), ar5523_host_available() calls ar5523_cmd() as below (there are other functions which finally call ar5523_cmd()): ar5523_probe() -> ar5523_host_available() -> ar5523_cmd_read() -> ar5523_cmd() If ar5523_cmd() timed out, then ar5523_host_available() failed and ar5523_probe() freed the device structure. So, ar5523_cmd_tx_cb() might touch the freed structure. This patch fixes this issue by canceling in-flight tx cmd if submitted urb timed out.	N/A	More Details
CVE-2022-50715	In the Linux kernel, the following vulnerability has been resolved: md/raid1: stop mdx_raid1 thread when raid1 array run failed fail run raid1 array when we assemble array with the inactive disk only, but the mdx_raid1 thread were not stop, Even if the associated resources have been released. it will caused a NULL dereference when we do poweroff. This causes the following Oops: [287.587787] BUG: kernel NULL pointer dereference, address: 0000000000000070 [287.594762] #PF: supervisor read access in kernel mode [287.599912] #PF: error_code(0x0000) - not-present page [287.605061] PGD 0 P4D 0 [287.607612] Oops: 0000 [#1] SMP NOPTI [287.611287] CPU: 3 PID: 5265 Comm: md0_raid1 Tainted: G U 5.10.146 #0 [287.619029] Hardware name: xxxxxxx/To be filled by O.E.M, BIOS 5.19 06/16/2022 [287.626775] RIP: 0010:md_check_recovery+0x57/0x500 [md_mod] [287.632357] Code: fe 01 00 00 48 83 bb 10 03 00 00 00 74 08 48 89 ..... [287.651118] RSP: 0018:ffffc90000433d78 EFLAGS: 00010202 [287.656347] RAX: 0000000000000000 RBX: ffff888105986800 RCX: 0000000000000000 [287.663491] RDX: ffffc90000433bb0 RSI: 00000000ffffefff RDI: ffff888105986800 [287.670634] RBP: ffffc90000433da0 R08: 0000000000000000 R09: c0000000ffffefff [287.677771] R10: 0000000000000001 R11: ffffc90000433ba8 R12: ffff888105986800 [287.684907] R13: 0000000000000000 R14: fffffffffffffe00 R15: ffff888100b6b500 [287.692052] FS: 0000000000000000(0000) GS:ffff888277f80000(0000) knlGS:0000000000000000 [287.700149] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [287.705897] CR2: 0000000000000070 CR3: 000000000320a000 CR4: 0000000000350ee0 [287.713033] Call Trace: [287.715498] raid1d+0x6c/0xbbb [raid1] [287.719256] ? __schedule+0x1ff/0x760 [287.722930] ? schedule+0x3b/0xb0 [287.726260] ? schedule_timeout+0x1ed/0x290 [287.730456] ? __switch_to+0x11f/0x400 [287.734219] md_thread+0xe9/0x140 [md_mod] [287.738328] ? md_thread+0xe9/0x140 [md_mod] [287.742601] ? wait_woken+0x80/0x80 [287.746097] ? md_register_thread+0xe0/0xe0 [md_mod] [287.751064] kthread+0x11a/0x140 [287.754300] ? kthread_park+0x90/0x90 [287.757974] ret_from_fork+0x1f/0x30 In fact, when raid1 array run fail, we need to do md_unregister_thread() before raid1_free().	N/A	More Details
CVE-2022-50714	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7921e: fix rmmod crash in driver reload test In insmod/rmmod stress test, the following crash dump shows up immediately. The problem is caused by missing mt76_dev in mt7921_pci_remove(). We should make sure the drvdata is ready before probe() finished. [168.862789] == [168.862797] BUG: KASAN: user-memory-access in try_to_grab_pending+0x59/0x480 [168.862805] Write of size 8 at addr 0000000000006df0 by task rmmod/5361 [168.862812] CPU: 7 PID: 5361 Comm: rmmod Tainted: G OE 5.19.0-rc6 #1 [168.862816] Hardware name: Intel(R) Client Systems NUC8i7BEH/NUC8BEB, 05/04/2020 [168.862820] Call Trace: [168.862822] <TASK> [168.862825] dump_stack_lvl+0x49/0x63 [168.862832] print_report.cold+0x493/0x6b7 [168.862845] kasan_report+0xa7/0x120 [168.862857] kasan_check_range+0x163/0x200 [168.862861] __kasan_check_write+0x14/0x20 [168.862866] try_to_grab_pending+0x59/0x480 [168.862870] __cancel_work_timer+0xbb/0x340 [168.862898] cancel_work_sync+0x10/0x20 [168.862902] mt7921_pci_remove+0x61/0x1c0 [mt7921e] [168.862909] pci_device_remove+0xa3/0x1d0 [168.862914] device_remove+0xc4/0x170 [168.862920] device_release_driver_internal+0x163/0x300 [168.862925] driver_detach+0xc7/0x1a0 [168.862930] bus_remove_driver+0xeb/0x2d0 [168.862935] driver_unregister+0x71/0xb0 [168.862939] pci_unregister_driver+0x30/0x230 [168.862944] mt7921_pci_driver_exit+0x10/0x1b [mt7921e] [168.862949] __x64_sys_delete_module+0x2f9/0x4b0 [168.862968] do_syscall_64+0x38/0x90 [168.862973] entry_SYSCALL_64_after_hwframe+0x63/0xcd Test steps: 1. insmode 2. do not ifup 3. rmmod quickly (within 1 second)	N/A	More Details

CVE-2023-54082	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2025-68952	Eigent is a multi-agent Workforce. In version 0.0.60, a 1-click Remote Code Execution (RCE) vulnerability has been identified in Eigent. This vulnerability allows an attacker to execute arbitrary code on the victim's machine or server through a specific interaction (1-click). This issue has been patched in version 0.0.61.	N/A	More Details
CVE-2025-68927	Libredesk is a self-hosted customer support desk. Prior to version 0.8.6-beta, LibreDesk is vulnerable to stored HTML injection in the contact notes feature. When adding notes via POST /api/v1/contacts/{id}/notes, the backend automatically wraps user input in <p> tags. However, by intercepting the request and removing the <p> tag, an attacker can inject arbitrary HTML elements such as forms and images, which are then stored and rendered without proper sanitization. This can lead to phishing, CSRF-style forced actions, and UI redress attacks. This issue has been patched in version 0.8.6-beta.	N/A	More Details
CVE-2025-68379	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix null deref on srq->rq.queue after resize failure A NULL pointer dereference can occur in rxe_srq_chk_attr() when ibv_modify_srq() is invoked twice in succession under certain error conditions. The first call may fail in rxe_queue_resize(), which leads rxe_srq_from_attr() to set srq->rq.queue = NULL. The second call then triggers a crash (null deref) when accessing srq->rq.queue->buf->index_mask. Call Trace: <TASK> rxe_modify_srq+0x170/0x480 [rdma_rxe] ? __pfx_rxe_modify_srq+0x10/0x10 [rdma_rxe] ? uverbs_try_lock_object+0x4f/0xa0 [ib_uverbs] ? rdma_lookup_get_uobject+0x1f0/0x380 [ib_uverbs] ib_uverbs_modify_srq+0x204/0x290 [ib_uverbs] ? __pfx_ib_uverbs_modify_srq+0x10/0x10 [ib_uverbs] ? tryinc_node_nr_active+0xe6/0x150 ? uverbs_fill_udata+0xed/0x4f0 [ib_uverbs] ib_uverbs_handler_UVERBS_METHOD_INVOKE_WRITE+0x2c0/0x470 [ib_uverbs] ? __pfx_ib_uverbs_handler_UVERBS_METHOD_INVOKE_WRITE+0x10/0x10 [ib_uverbs] ? uverbs_fill_udata+0xed/0x4f0 [ib_uverbs] ib_uverbs_run_method+0x55a/0x6e0 [ib_uverbs] ? __pfx_ib_uverbs_handler_UVERBS_METHOD_INVOKE_WRITE+0x10/0x10 [ib_uverbs] ib_uverbs_cmd_verbs+0x54d/0x800 [ib_uverbs] ? __pfx_ib_uverbs_cmd_verbs+0x10/0x10 [ib_uverbs] ? __pfx___raw_spin_lock_irqsave+0x10/0x10 ? __pfx_do_vfs_ioctl+0x10/0x10 ? ioctl_has_perm.constprop.0.isra.0+0x2c7/0x4c0 ? __pfx_ioctl_has_perm.constprop.0.isra.0+0x10/0x10 ib_uverbs_ioctl+0x13e/0x220 [ib_uverbs] ? __pfx_ib_uverbs_ioctl+0x10/0x10 [ib_uverbs] __x64_sys_ioctl+0x138/0x1c0 do_syscall_64+0x82/0x250 ? fdget_pos+0x58/0x4c0 ? ksys_write+0xf3/0x1c0 ? __pfx_ksys_write+0x10/0x10 ? do_syscall_64+0xc8/0x250 ? __pfx_vm_mmap_pgoff+0x10/0x10 ? fget+0x173/0x230 ? fput+0x2a/0x80 ? ksys_mmap_pgoff+0x224/0x4c0 ? do_syscall_64+0xc8/0x250 ? do_user_addr_fault+0x37b/0xfe0 ? clear_bhb_loop+0x50/0xa0 ? clear_bhb_loop+0x50/0xa0 ? clear_bhb_loop+0x50/0xa0 entry_SYSCALL_64_after_hwframe+0x76/0x7e	N/A	More Details
CVE-2022-50763	In the Linux kernel, the following vulnerability has been resolved: crypto: marvell/octeontx - prevent integer overflows The "code_length" value comes from the firmware file. If your firmware is untrusted realistically there is probably very little you can do to protect yourself. Still we try to limit the damage as much as possible. Also Smatch marks any data read from the filesystem as untrusted and prints warnings if it not capped correctly. The "code_length * 2" can overflow. The round_up(ucode_size, 16) + sizeof() expression can overflow too. Prevent these overflows.	N/A	More Details
CVE-2022-50767	In the Linux kernel, the following vulnerability has been resolved: fbdev: smscufx: Fix several use-after-free bugs Several types of UAFs can occur when physically removing a USB device. Adds ufx_ops_destroy() function to .fb_destroy of fb_ops, and in this function, there is kref_put() that finally calls ufx_free(). This fix prevents multiple UAFs.	N/A	More Details
CVE-2025-14715	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2025-14820	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2022-50766	In the Linux kernel, the following vulnerability has been resolved: btrfs: set generation before calling btrfs_clean_tree_block in btrfs_init_new_buffer syzbot is reporting uninit-value in btrfs_clean_tree_block() [1], for commit bc877d285ca3dba2 ("btrfs: Deduplicate extent_buffer init code") missed that btrfs_set_header_generation() in btrfs_init_new_buffer() must not be moved to after clean_tree_block() because clean_tree_block() is calling btrfs_header_generation() since commit 55c69072d6bd5be1 ("Btrfs: Fix extent_buffer usage when nodesize != leafsize"). Since memzero_extent_buffer() will reset "struct btrfs_header" part, we can't move btrfs_set_header_generation() to before memzero_extent_buffer(). Just re-add btrfs_set_header_generation() before btrfs_clean_tree_block().	N/A	More Details
CVE-2022-50765	In the Linux kernel, the following vulnerability has been resolved: RISC-V: kexec: Fix memory leak of elf header buffer This is reported by kmemleak detector: unreferenced object 0xff2000000403d000 (size 4096): comm "kexec", pid 146, jiffies 4294900633 (age 64.792s) hex dump (first 32 bytes): 7f 45 4c 46 02 01 01 00 00 00 00 00 00 00 00 00 .ELF..... 04 00 f3 00 01 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<00000000566ca97c>] kmemleak_vmalloc+0x3c/0xbe [<00000000979283d8>] __vmalloc_node_range+0x3ac/0x560 [<00000000b4b3712a>] __vmalloc_node+0x56/0x62 [<00000000854f75e2>] vzalloc+0x2c/0x34 [<00000000e9a00db9>] crash_prepare_elf64_headers+0x80/0x30c [<0000000067e8bf48>] elf_kexec_load+0x3e8/0x4ec [<0000000036548e09>] kexec_image_load_default+0x40/0x4c [<0000000079fbe1b4>] sys_kexec_file_load+0x1c4/0x322 [<0000000040c62c03>] ret_from_syscall+0x0/0x2 In elf_kexec_load(), a buffer is allocated via vzalloc() to store elf headers. While it's not freed back to system when kdump kernel is reloaded or unloaded, or when image->elf_header is successfully set and then fails to load kdump kernel for some reason. Fix it by freeing the buffer in arch_kimage_file_post_load_cleanup().	N/A	More Details
CVE-2023-54050	In the Linux kernel, the following vulnerability has been resolved: ubifs: Fix memleak when insert_old_idx() failed Following process will cause a memleak for copied up znode: dirty_cow_znode zn = copy_znode(c, znode); err = insert_old_idx(c, zbr->lnum, zbr->offs); if (unlikely(err)) return ERR_PTR(err); // No one refers to zn. Fetch a reproducer in [Link]. Function copy_znode() is split into 2 parts: resource allocation and znode replacement, insert_old_idx() is split in similar way, so resource cleanup could be done in error handling path without corrupting metadata(mem & disk). It's okay that old index inserting is put behind of add_idx_dirt(), old index is used in layout_leb_in_gaps(), so the two processes do not depend on each other.	N/A	More Details
CVE-2022-50764	In the Linux kernel, the following vulnerability has been resolved: ipv6/sit: use DEV_STATS_INC() to avoid data-races syzbot/KCSAN reported that multiple cpus are updating dev->stats.tx_error concurrently. This is because sit tunnels are NETIF_F_LLTX, meaning their ndo_start_xmit() is not protected by a spinlock. While original KCSAN report was about tx path, rx path has the same issue.	N/A	More Details
CVE-	Forgejo before 13.0.2 allows attackers to write to unintended files, and possibly obtain server shell access, because of mishandling		More

2025-68937	of out-of-repository symlink destinations for template repositories. This is also fixed for 11 LTS in 11.0.7 and later.	N/A	Details
CVE-2023-54051	In the Linux kernel, the following vulnerability has been resolved: net: do not allow gso_size to be set to GSO_BY_FRAGS One missing check in virtio_net_hdr_to_skb() allowed syzbot to crash kernels again [1] Do not allow gso_size to be set to GSO_BY_FRAGS (0xffff), because this magic value is used by the kernel. [1] general protection fault, probably for non-canonical address 0xdfffffc000000000e: 0000 [#1] PREEMPT SMP KASAN KASAN: null-ptr-deref in range [0x0000000000000070-0x0000000000000077] CPU: 0 PID: 5039 Comm: syz-executor401 Not tainted 6.5.0-rc5-next-20230809-syzkaller #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 07/26/2023 RIP: 0010:skb_segment+0x1a52/0x3ef0 net/core/skbuff.c:4500 Code: 00 00 00 e9 ab eb ff ff e8 6b 96 5d f9 48 8b 84 24 00 01 00 00 48 8d 78 70 48 b8 00 00 00 00 00 fc ff df 48 89 fa 48 c1 ea 03 <0f> b6 04 02 84 c0 74 08 3c 03 0f 8e ea 21 00 00 48 8b 84 24 00 01 RSP: 0018:ffffc90003d3f1c8 EFLAGS: 00010202 RAX: dffffc0000000000 RBX: 000000000001fffe RCX: 0000000000000000 RDX: 000000000000000e RSI: ffffffff882a3115 RDI: 0000000000000070 RBP: ffffc90003d3f378 R08: 0000000000000005 R09: 000000000000ffff R10: 000000000000ffff R11: 5ee4a93e456187d6 R12: 000000000001ffc6 R13: dffffc0000000000 R14: 0000000000000008 R15: 000000000000ffff FS: 00005555563f2380(0000) GS:ffff8880b9800000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000020020000 CR3: 000000001626d000 CR4: 00000000003506f0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> udp6_ufo_fragment+0x9d2/0xd50 net/ipv6/udp_offload.c:109 ipv6_gso_segment+0x5c4/0x17b0 net/ipv6/ip6_offload.c:120 skb_mac_gso_segment+0x292/0x610 net/core/gso.c:53 __skb_gso_segment+0x339/0x710 net/core/gso.c:124 skb_gso_segment include/net/gso.h:83 [inline] validate_xmit_skb+0x3a5/0xf10 net/core/dev.c:3625 __dev_queue_xmit+0x8f0/0x3d60 net/core/dev.c:4329 dev_queue_xmit include/linux/netdevice.h:3082 [inline] packet_xmit+0x257/0x380 net/packet/af_packet.c:276 packet_snd net/packet/af_packet.c:3087 [inline] packet_sendmsg+0x24c7/0x5570 net/packet/af_packet.c:3119 sock_sendmsg_nosec net/socket.c:727 [inline] sock_sendmsg+0xd9/0x180 net/socket.c:750 ____sys_sendmsg+0x6ac/0x940 net/socket.c:2496 ___sys_sendmsg+0x135/0x1d0 net/socket.c:2550 __sys_sendmsg+0x117/0x1e0 net/socket.c:2579 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x38/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7ff27cdb34d9	N/A	More Details
CVE-2022-50759	In the Linux kernel, the following vulnerability has been resolved: media: i2c: ov5648: Free V4L2 fwnode data on unbind The V4L2 fwnode data structure doesn't get freed on unbind, which leads to a memleak.	N/A	More Details
CVE-2023-54052	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7921: fix skb leak by txs missing in AMSDU txs may be dropped if the frame is aggregated in AMSDU. When the problem shows up, some SKBs would be hold in driver to cause network stopped temporarily. Even if the problem can be recovered by txs timeout handling, mt7921 still need to disable txs in AMSDU to avoid this issue.	N/A	More Details
CVE-2023-54053	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: pcie: fix possible NULL pointer dereference It is possible that iwl_pci_probe() will fail and free the trans, then afterwards iwl_pci_remove() will be called and crash by trying to access trans which is already freed, fix it. iwlwifi 0000:01:00.0: Detected crf-id 0xa5a5a5a2, cnv-id 0xa5a5a5a2 wfpm id 0xa5a5a5a2 iwlwifi 0000:01:00.0: Can't find a correct rfid for crf id 0x5a2 ... BUG: kernel NULL pointer dereference, address: 0000000000000028 ... RIP: 0010:iwl_pci_remove+0x12/0x30 [iwlwifi] pci_device_remove+0x3e/0xb0 device_release_driver_internal+0x103/0x1f0 driver_detach+0x4c/0x90 bus_remove_driver+0x5c/0xd0 driver_unregister+0x31/0x50 pci_unregister_driver+0x40/0x90 iwl_pci_unregister_driver+0x15/0x20 [iwlwifi] __exit_compat+0x9/0x98 [iwlwifi] __x64_sys_delete_module+0x147/0x260	N/A	More Details
CVE-2023-54054	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2022-50762	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Avoid UBSAN error on true_sectors_per_clst() syzbot reported UBSAN error as below: [76.901829][T6677] == [76.903908][T6677] UBSAN: shift-out-of-bounds in fs/ntfs3/super.c:675:13 [76.905363][T6677] shift exponent -247 is negative This patch avoid this error.	N/A	More Details
CVE-2022-50761	In the Linux kernel, the following vulnerability has been resolved: x86/xen: Fix memory leak in xen_init_lock_cpu() In xen_init_lock_cpu(), the @name has allocated new string by kasprintf(), if bind_ipi_to_irqhandler() fails, it should be freed, otherwise may lead to a memory leak issue, fix it.	N/A	More Details
CVE-2022-50760	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix PCI device refcount leak in amdgpu_atrm_get_bios() As comment of pci_get_class() says, it returns a pci_device with its refcount increased and decreased the refcount for the input parameter @from if it is not NULL. If we break the loop in amdgpu_atrm_get_bios() with 'pdev' not NULL, we need to call pci_dev_put() to decrease the refcount. Add the missing pci_dev_put() to avoid refcount leak.	N/A	More Details
CVE-2023-54055	In the Linux kernel, the following vulnerability has been resolved: RDMA/irdma: Fix memory leak of PBLE objects On rmmod of irdma, the PBLE object memory is not being freed. PBLE object memory are not statically pre-allocated at function initialization time unlike other HMC objects. PBLEs objects and the Segment Descriptors (SD) for it can be dynamically allocated during scale up and SD's remain allocated till function deinitialization. Fix this leak by adding IRDMA_HMC_IW_PBLE to the iw_hmc_obj_types[] table and skip pbles in irdma_create_hmc_obj but not in irdma_del_hmc_objects().	N/A	More Details
CVE-2023-54056	In the Linux kernel, the following vulnerability has been resolved: kheaders: Use array declaration instead of char Under CONFIG_FORTIFY_SOURCE, memcpy() will check the size of destination and source buffers. Defining kernel_headers_data as "char" would trip this check. Since these addresses are treated as byte arrays, define them as arrays (as done everywhere else). This was seen with: $ cat /sys/kernel/kheaders.tar.xz >> /dev/null detected buffer overflow in memcpy kernel BUG at lib/string_helpers.c:1027! ... RIP: 0010:fortify_panic+0xf/0x20 [...] Call Trace: <TASK> ikheaders_read+0x45/0x50 [kheaders] kernfs_fop_read_iter+0x1a4/0x2f0 ...	N/A	More Details
CVE-2023-54049	In the Linux kernel, the following vulnerability has been resolved: rpmsg: glink: Add check for kstrdup Add check for the return value of kstrdup() and return the error if it fails in order to avoid NULL pointer dereference.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: RDMA/bnxt_re: Prevent handling any completions after qp destroy HW may generate completions that indicates QP is destroyed. Driver should not be scheduling any more completion handlers for this QP, after the QP is destroyed. Since CQs are active during the QP destroy, driver may still schedule completion handlers. This can		

CVE-2023-54048	cause a race where the destroy_cq and poll_cq running simultaneously. Snippet of kernel panic while doing bnxt_re driver load unload in loop. This indicates a poll after the CQ is freed. [77786.481636] Call Trace: [77786.481640] <TASK> [77786.481644] bnxt_re_poll_cq+0x14a/0x620 [bnxt_re] [77786.481658] ? kvm_clock_read+0x14/0x30 [77786.481693] __ib_process_cq+0x57/0x190 [ib_core] [77786.481728] ib_cq_poll_work+0x26/0x80 [ib_core] [77786.481761] process_one_work+0x1e5/0x3f0 [77786.481768] worker_thread+0x50/0x3a0 [77786.481785] ? __pfx_worker_thread+0x10/0x10 [77786.481790] kthread+0xe2/0x110 [77786.481794] ? __pfx_kthread+0x10/0x10 [77786.481797] ret_from_fork+0x2c/0x50 To avoid this, complete all completion handlers before returning the destroy QP. If free_cq is called soon after destroy_qp, IB stack will cancel the CQ work before invoking the destroy_cq verb and this will prevent any race mentioned.	N/A	More Details
CVE-2022-50768	In the Linux kernel, the following vulnerability has been resolved: scsi: smartpqi: Correct device removal for multi-actuator devices Correct device count for multi-actuator drives which can cause kernel panics.	N/A	More Details
CVE-2022-50769	In the Linux kernel, the following vulnerability has been resolved: mmc: mxcmmc: fix return value check of mmc_add_host() mmc_add_host() may return error, if we ignore its return value, the memory that allocated in mmc_alloc_host() will be leaked and it will lead a kernel crash because of deleting not added device in the remove path. So fix this by checking the return value and goto error path which will call mmc_free_host().	N/A	More Details
CVE-2022-50781	In the Linux kernel, the following vulnerability has been resolved: amdgpu/pm: prevent array underflow in vega20_odn_edit_dpm_table() In the PP_OD_EDIT_VDDC_CURVE case the "input_index" variable is capped at 2 but not checked for negative values so it results in an out of bounds read. This value comes from the user via sysfs.	N/A	More Details
CVE-2022-50780	In the Linux kernel, the following vulnerability has been resolved: net: fix UAF issue in nfqnl_nf_hook_drop() when ops_init() failed When the ops_init() interface is invoked to initialize the net, but ops->init() fails, data is released. However, the ptr pointer in net->gen is invalid. In this case, when nfqnl_nf_hook_drop() is invoked to release the net, invalid address access occurs. The process is as follows: setup_net() ops_init() data = kzalloc(...) ---> alloc "data" net_assign_generic() ---> assign "date" to ptr in net->gen ... ops->init() ---> failed ... kfree(data); ---> ptr in net->gen is invalid ... ops_exit_list() ... nfqnl_nf_hook_drop() *q = nfnl_queue_pernet(net) ---> q is invalid The following is the Call Trace information: BUG: KASAN: use-after-free in nfqnl_nf_hook_drop+0x264/0x280 Read of size 8 at addr ffff88810396b240 by task ip/15855 Call Trace: <TASK> dump_stack_lvl+0x8e/0xd1 print_report+0x155/0x454 kasan_report+0xba/0x1f0 nfqnl_nf_hook_drop+0x264/0x280 nf_queue_nf_hook_drop+0x8b/0x1b0 __nf_unregister_net_hook+0x1ae/0x5a0 nf_unregister_net_hooks+0xde/0x130 ops_exit_list+0xb0/0x170 setup_net+0x7ac/0xbd0 copy_net_ns+0x2e6/0x6b0 create_new_namespaces+0x382/0xa50 unshare_nsproxy_namespaces+0xa6/0x1c0 ksys_unshare+0x3a4/0x7e0 __x64_sys_unshare+0x2d/0x40 do_syscall_64+0x35/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0 </TASK> Allocated by task 15855: kasan_save_stack+0x1e/0x40 kasan_set_track+0x21/0x30 __kasan_kmalloc+0xa1/0xb0 __kmalloc+0x49/0xb0 ops_init+0xe7/0x410 setup_net+0x5aa/0xbd0 copy_net_ns+0x2e6/0x6b0 create_new_namespaces+0x382/0xa50 unshare_nsproxy_namespaces+0xa6/0x1c0 ksys_unshare+0x3a4/0x7e0 __x64_sys_unshare+0x2d/0x40 do_syscall_64+0x35/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0 Freed by task 15855: kasan_save_stack+0x1e/0x40 kasan_set_track+0x21/0x30 kasan_save_free_info+0x2a/0x40 ____kasan_slab_free+0x155/0x1b0 slab_free_freelist_hook+0x11b/0x220 __kmem_cache_free+0xa4/0x360 ops_init+0xb9/0x410 setup_net+0x5aa/0xbd0 copy_net_ns+0x2e6/0x6b0 create_new_namespaces+0x382/0xa50 unshare_nsproxy_namespaces+0xa6/0x1c0 ksys_unshare+0x3a4/0x7e0 __x64_sys_unshare+0x2d/0x40 do_syscall_64+0x35/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0	N/A	More Details
CVE-2022-50779	In the Linux kernel, the following vulnerability has been resolved: orangefs: Fix kmemleak in orangefs_prepare_debugfs_help_string() When insert and remove the orangefs module, then debug_help_string will be leaked: unreferenced object 0xffff8881652ba000 (size 4096): comm "insmod", pid 1701, jiffies 4294893639 (age 13218.530s) hex dump (first 32 bytes): 43 6c 69 65 6e 74 20 44 65 62 75 67 20 4b 65 79 Client Debug Key 77 6f 72 64 73 20 61 72 65 20 75 6e 6b 6e 6f 77 words are unknow backtrace: [<0000000004e6f8e3>] kmalloc_trace+0x27/0xa0 [<0000000006f75d85>] orangefs_prepare_debugfs_help_string+0x5e/0x480 [orangefs] [<0000000091270a2a>] _sub_I_65535_1+0x57/0xf70 [crc_itu_t] [<000000004b1ee1a3>] do_one_initcall+0x87/0x2a0 [<000000001d0614ae>] do_init_module+0xdf/0x320 [<00000000efef068c>] load_module+0x2f98/0x3330 [<000000006533b44d>] __do_sys_finit_module+0x113/0x1b0 [<00000000a0da6f99>] do_syscall_64+0x35/0x80 [<0000000007790b19b>] entry_SYSCALL_64_after_hwframe+0x46/0xb0 When remove the module, should always free debug_help_string. Should always free the allocated buffer when change the free_debug_help_string.	N/A	More Details
CVE-2022-50778	In the Linux kernel, the following vulnerability has been resolved: fortify: Fix __compiletime_strlen() under UBSAN_BOUNDS_LOCAL With CONFIG_FORTIFY=y and CONFIG_UBSAN_LOCAL_BOUNDS=y enabled, we observe a runtime panic while running Android's Compatibility Test Suite's (CTS) android.hardware.input.cts.tests. This is stemming from a strlen() call in hidinput_allocate(). __compiletime_strlen() is implemented in terms of __builtin_object_size(), then does an array access to check for NUL-termination. A quirk of __builtin_object_size() is that for strings whose values are runtime dependent, __builtin_object_size(str, 1 or 0) returns the maximum size of possible values when those sizes are determinable at compile time. Example: static const char *v = "FOO BAR"; static const char *y = "FOO BA"; unsigned long x (int z) { // Returns 8, which is: // max(__builtin_object_size(v, 1), __builtin_object_size(y, 1)) return __builtin_object_size(z ? v : y, 1); } So when FORTIFY_SOURCE is enabled, the current implementation of __compiletime_strlen() will try to access beyond the end of y at runtime using the size of v. Mixed with UBSAN_LOCAL_BOUNDS we get a fault. hidinput_allocate() has a local C string whose value is control flow dependent on a switch statement, so __builtin_object_size(str, 1) evaluates to the maximum string length, making all other cases fault on the last character check. hidinput_allocate() could be cleaned up to avoid runtime calls to strlen() since the local variable can only have literal values, so there's no benefit to trying to fortify the strlen call site there. Perform a __builtin_constant_p() check against index 0 earlier in the macro to filter out the control-flow-dependant case. Add a KUnit test for checking the expected behavioral characteristics of FORTIFY_SOURCE internals.	N/A	More Details
CVE-2023-54043	In the Linux kernel, the following vulnerability has been resolved: iommufd: Do not add the same hwpt to the ioas->hwpt_list twice The hwpt is added to the hwpt_list only during its creation, it is never added again. This hunk is some missed leftover from rework. Adding it twice will corrupt the linked list in some cases. It effects HWPT specific attachment, which is something the test suite cannot cover until we can create a legitimate struct device with a non-system iommu "driver" (ie we need the bus removed from the iommu code)	N/A	More Details
CVE-2022-50777	In the Linux kernel, the following vulnerability has been resolved: net: phy: xgmiitorgmii: Fix refcount leak in xgmiitorgmii_probe of_phy_find_device() return device node with refcount incremented. Call put_device() to relese it when not needed anymore.	N/A	More Details
CVE-2022-50776	In the Linux kernel, the following vulnerability has been resolved: clk: st: Fix memory leak in st_of_quadfs_setup() If st_clk_register_quadfs_pll() fails, @lock should be freed before goto @err_exit, otherwise will cause meory leak issue, fix it.	N/A	More Details

CVE-2023-54044	In the Linux kernel, the following vulnerability has been resolved: spmi: Add a check for remove callback when removing a SPMI driver When removing a SPMI driver, there can be a crash due to NULL pointer dereference if it does not have a remove callback defined. This is one such call trace observed when removing the QCOM SPMI PMIC driver: dump_backtrace.cfi_jt+0x0/0x8 dump_stack_lvl+0xd8/0x16c panic+0x188/0x498 __cfi_slowpath+0x0/0x214 __cfi_slowpath+0x1dc/0x214 spmi_drv_remove+0x16c/0x1e0 device_release_driver_internal+0x468/0x79c driver_detach+0x11c/0x1a0 bus_remove_driver+0xc4/0x124 driver_unregister+0x58/0x84 cleanup_module+0x1c/0xc24 [qcom_spmi_pmic] __do_sys_delete_module+0x3ec/0x53c __arm64_sys_delete_module+0x18/0x28 el0_svc_common+0xdc/0x294 el0_svc+0x38/0x9c el0_sync_handler+0x8c/0xf0 el0_sync+0x1b4/0x1c0 If a driver has all its resources allocated through devm_() APIs and does not need any other explicit cleanup, it would not require a remove callback to be defined. Hence, add a check for remove callback presence before calling it when removing a SPMI driver.	N/A	More Details
CVE-2022-50775	In the Linux kernel, the following vulnerability has been resolved: RDMA/hns: Fix refcount leak in hns_roce_mmap rdma_user_mmap_entry_get_pgoff() takes the reference. Add missing rdma_user_mmap_entry_put() to release the reference. Acked-by Haoyue Xu <xuhaoyue1@hisilicon.com>	N/A	More Details
CVE-2022-50774	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - fix DMA transfer direction When CONFIG_DMA_API_DEBUG is selected, while running the crypto self test on the QAT crypto algorithms, the function add_dma_entry() reports a warning similar to the one below, saying that overlapping mappings are not supported. This occurs in tests where the input and the output scatter list point to the same buffers (i.e. two different scatter lists which point to the same chunks of memory). The logic that implements the mapping uses the flag DMA_BIDIRECTIONAL for both the input and the output scatter lists which leads to overlapped write mappings. These are not supported by the DMA layer. Fix by specifying the correct DMA transfer directions when mapping buffers. For in-place operations where the input scatter list matches the output scatter list, buffers are mapped once with DMA_BIDIRECTIONAL, otherwise input buffers are mapped using the flag DMA_TO_DEVICE and output buffers are mapped with DMA_FROM_DEVICE. Overlapping a read mapping with a write mapping is a valid case in dma-coherent devices like QAT. The function that frees and unmaps the buffers, qat_alg_free_bufl() has been changed accordingly to the changes to the mapping function. DMA-API: 4xxx 0000:06:00.0: cacheline tracking EEXIST, overlapping mappings aren't supported WARNING: CPU: 53 PID: 4362 at kernel/dma/debug.c:570 add_dma_entry+0x1e9/0x270 ... Call Trace: dma_map_page_attrs+0x82/0x2d0 ? preempt_count_add+0x6a/0xa0 qat_alg_sgl_to_bufl+0x45b/0x990 [intel_qat] qat_alg_aead_dec+0x71/0x250 [intel_qat] crypto_aead_decrypt+0x3d/0x70 test_aead_vec_cfg+0x649/0x810 ? number+0x310/0x3a0 ? vsnprintf+0x2a3/0x550 ? scnprintf+0x42/0x70 ? valid_sg_divisions.constprop.0+0x86/0xa0 ? test_aead_vec+0xdf/0x120 test_aead_vec+0xdf/0x120 alg_test_aead+0x185/0x400 alg_test+0x3d8/0x500 ? crypto_acomp_scomp_free_ctx+0x30/0x30 ? __schedule+0x32a/0x12a0 ? ttwu_queue_wakelist+0xbf/0x110 ? _raw_spin_unlock_irqrestore+0x23/0x40 ? try_to_wake_up+0x83/0x570 ? _raw_spin_unlock_irqrestore+0x23/0x40 ? __set_cpus_allowed_ptr_locked+0xea/0x1b0 ? crypto_acomp_scomp_free_ctx+0x30/0x30 cryptomgr_test+0x27/0x50 kthread+0xe6/0x110 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork+0x1f/0x30	N/A	More Details
CVE-2022-50773	In the Linux kernel, the following vulnerability has been resolved: ALSA: mts64: fix possible null-ptr-defer in snd_mts64_interrupt I got a null-ptr-defer error report when I do the following tests on the qemu platform: make defconfig and CONFIG_PARPORT=m, CONFIG_PARPORT_PC=m, CONFIG_SND_MTS64=m Then making test scripts: cat>test_mod1.sh<<EOF modprobe snd-mts64 modprobe snd-mts64 EOF Executing the script, perhaps several times, we will get a null-ptr-defer report, as follow: syzkaller:~# ./test_mod.sh snd_mts64: probe of snd_mts64.0 failed with error -5 modprobe: ERROR: could not insert 'snd_mts64': No such device BUG: kernel NULL pointer dereference, address: 0000000000000000 #PF: supervisor write access in kernel mode #PF: error_code(0x0002) - not-present page PGD 0 P4D 0 Oops: 0002 [#1] PREEMPT SMP PTI CPU: 0 PID: 205 Comm: modprobe Not tainted 6.1.0-rc8-00588-g76dcd734eca2 #6 Call Trace: <IRQ> snd_mts64_interrupt+0x24/0xa0 [snd_mts64] parport_irq_handler+0x37/0x50 [parport] __handle_irq_event_percpu+0x39/0x190 handle_irq_event_percpu+0xa/0x30 handle_irq_event+0x2f/0x50 handle_edge_irq+0x99/0x1b0 __common_interrupt+0x5d/0x100 common_interrupt+0xa0/0xc0 </IRQ> <TASK> asm_common_interrupt+0x22/0x40 RIP: 0010:_raw_write_unlock_irqrestore+0x11/0x30 parport_claim+0xbd/0x230 [parport] snd_mts64_probe+0x14a/0x465 [snd_mts64] platform_probe+0x3f/0xa0 really_probe+0x129/0x2c0 __driver_probe_device+0x6d/0xc0 driver_probe_device+0x1a/0xa0 __device_attach_driver+0x7a/0xb0 bus_for_each_drv+0x62/0xb0 __device_attach+0xe4/0x180 bus_probe_device+0x82/0xa0 device_add+0x550/0x920 platform_device_add+0x106/0x220 snd_mts64_attach+0x2e/0x80 [snd_mts64] port_check+0x14/0x20 [parport] bus_for_each_dev+0x6e/0xc0 __parport_register_driver+0x7c/0xb0 [parport] snd_mts64_module_init+0x31/0x1000 [snd_mts64] do_one_initcall+0x3c/0x1f0 do_init_module+0x46/0x1c6 load_module+0x1d8d/0x1e10 __do_sys_finit_module+0xa2/0xf0 do_syscall_64+0x37/0x90 entry_SYSCALL_64_after_hwframe+0x63/0xcd </TASK> Kernel panic - not syncing: Fatal exception in interrupt Rebooting in 1 seconds.. The mts wa not initialized during interrupt, we add check for mts to fix this bug.	N/A	More Details
CVE-2022-50772	In the Linux kernel, the following vulnerability has been resolved: netdevsim: fix memory leak in nsim_bus_dev_new() If device_register() failed in nsim_bus_dev_new(), the value of reference in nsim_bus_dev->dev is 1. obj->name in nsim_bus_dev->dev will not be released. unreferenced object 0xffff88810352c480 (size 16): comm "echo", pid 5691, jiffies 4294945921 (age 133.270s) hex dump (first 16 bytes): 6e 65 74 64 65 76 73 69 6d 31 00 00 00 00 00 00 netdevsim1..... backtrace: [<000000005e2e5e26>] __kmalloc_node_track_caller+0x3a/0xb0 [<0000000094ca4fc8>] kvasprintf+0xc3/0x160 [<00000000aad09bcc>] kvasprintf_const+0x55/0x180 [<000000009bac868d>] kobject_set_name_vargs+0x56/0x150 [<000000007c1a5d70>] dev_set_name+0xbb/0xf0 [<00000000ad0d126b>] device_add+0x1f8/0x1cb0 [<00000000c222ae24>] new_device_store+0x3b6/0x5e0 [<0000000043593421>] bus_attr_store+0x72/0xa0 [<00000000cbb1833a>] sysfs_kf_write+0x106/0x160 [<00000000d0dedb8a>] kernfs_fop_write_iter+0x3a8/0x5a0 [<00000000770b66e2>] vfs_write+0x8f0/0xc80 [<0000000078bb39be>] ksys_write+0x106/0x210 [<00000000005e55a4>] do_syscall_64+0x35/0x80 [<00000000eaa40bbc>] entry_SYSCALL_64_after_hwframe+0x46/0xb0	N/A	More Details
CVE-2022-50771	In the Linux kernel, the following vulnerability has been resolved: rcu: Fix __this_cpu_read() lockdep warning in rcu_force_quiescent_state() Running rcutorture with non-zero fqs_duration module parameter in a kernel built with CONFIG_PREEMPTION=y results in the following splat: BUG: using __this_cpu_read() in preemptible [00000000] code: rcu_torture_fqs/398 caller is __this_cpu_preempt_check+0x13/0x20 CPU: 3 PID: 398 Comm: rcu_torture_fqs Not tainted 6.0.0-rc1-yoctodev-standard+ Call Trace: <TASK> dump_stack_lvl+0x5b/0x86 dump_stack+0x10/0x16 check_preemption_disabled+0xe5/0xf0 __this_cpu_preempt_check+0x13/0x20 rcu_force_quiescent_state.part.0+0x1c/0x170 rcu_force_quiescent_state+0x1e/0x30 rcu_torture_fqs+0xca/0x160 ? rcu_torture_boost+0x430/0x430 kthread+0x192/0x1d0 ? kthread_complete_and_exit+0x30/0x30 ret_from_fork+0x22/0x30 </TASK> The problem is that rcu_force_quiescent_state() uses __this_cpu_read() in preemptible code instead of the proper raw_cpu_read(). This commit therefore changes __this_cpu_read() to raw_cpu_read().	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: ocfs2: fix memory leak in ocfs2_mount_volume() There is a memory leak reported by kmemleak: unreferenced object 0xffff88810cc65e60 (size 32): comm "mount.ocfs2", pid 23753, jiffies 4302528942 (age 34735.105s) hex dump (first 32 bytes): 10 00 00 00 00 00 00 00 00 00 01 01 01 01 01 01 ..... 01 01 01 01		

CVE-2022-50770	01 01 01 01 00 00 00 00 00 00 00 00 ..... backtrace: [<ffffffff8170f73d>] __kmalloc+0x4d/0x150 [<ffffffffa0ac3f51>] ocfs2_compute_replay_slots+0x121/0x330 [ocfs2] [<ffffffffa0b65165>] ocfs2_check_volume+0x485/0x900 [ocfs2] [<ffffffffa0b68129>] ocfs2_mount_volume.isra.0+0x1e9/0x650 [ocfs2] [<ffffffffa0b7160b>] ocfs2_fill_super+0xe0b/0x1740 [ocfs2] [<ffffffff818e1fe2>] mount_bdev+0x312/0x400 [<ffffffff819a086d>] legacy_get_tree+0xed/0x1d0 [<ffffffff818de82d>] vfs_get_tree+0x7d/0x230 [<ffffffff81957f92>] path_mount+0xd62/0x1760 [<ffffffff81958a5a>] do_mount+0xca/0xe0 [<ffffffff81958d3c>] __x64_sys_mount+0x12c/0x1a0 [<ffffffff82f26f15>] do_syscall_64+0x35/0x80 [<ffffffff8300006a>] entry_SYSCALL_64_after_hwframe+0x46/0xb0 This call stack is related to two problems. Firstly, the ocfs2 super uses "replay_map" to trace online/offline slots, in order to recover offline slots during recovery and mount. But when ocfs2_truncate_log_init() returns an error in ocfs2_mount_volume(), the memory of "replay_map" will not be freed in error handling path. Secondly, the memory of "replay_map" will not be freed if d_make_root() returns an error in ocfs2_fill_super(). But the memory of "replay_map" will be freed normally when completing recovery and mount in ocfs2_complete_mount_recovery(). Fix the first problem by adding error handling path to free "replay_map" when ocfs2_truncate_log_init() fails. And fix the second problem by calling ocfs2_free_replay_slots(osb) in the error handling path "out_dismount". In addition, since ocfs2_free_replay_slots() is static, it is necessary to remove its static attribute and declare it in header file.	N/A	More Details
CVE-2023-54045	In the Linux kernel, the following vulnerability has been resolved: audit: fix possible soft lockup in __audit_inode_child() Tracefs or debugfs maybe cause hundreds to thousands of PATH records, too many PATH records maybe cause soft lockup. For example: 1. CONFIG_KASAN=y && CONFIG_PREEMPTION=n 2. auditctl -a exit,always -S open -k key 3. sysctl -w kernel.watchdog_thresh=5 4. mkdir /sys/kernel/debug/tracing/instances/test There may be a soft lockup as follows: watchdog: BUG: soft lockup - CPU#45 stuck for 7s! [mkdir:15498] Kernel panic - not syncing: softlockup: hung tasks Call trace: dump_backtrace+0x0/0x30c show_stack+0x20/0x30 dump_stack+0x11c/0x174 panic+0x27c/0x494 watchdog_timer_fn+0x2bc/0x390 __run_hrtimer+0x148/0x4fc __hrtimer_run_queues+0x154/0x210 hrtimer_interrupt+0x2c4/0x760 arch_timer_handler_phys+0x48/0x60 handle_percpu_devid_irq+0xe0/0x340 __handle_domain_irq+0xbc/0x130 gic_handle_irq+0x78/0x460 el1_irq+0xb8/0x140 __audit_inode_child+0x240/0x7bc tracefs_create_file+0x1b8/0x2a0 trace_create_file+0x18/0x50 event_create_dir+0x204/0x30c __trace_add_new_event+0xac/0x100 event_trace_add_tracer+0xa0/0x130 trace_array_create_dir+0x60/0x140 trace_array_create+0x1e0/0x370 instance_mkdir+0x90/0xd0 tracefs_syscall_mkdir+0x68/0xa0 vfs_mkdir+0x21c/0x34c do_mkdirat+0x1b4/0x1d4 __arm64_sys_mkdirat+0x4c/0x60 el0_svc_common.constprop.0+0xa8/0x240 do_el0_svc+0x8c/0xc0 el0_svc+0x20/0x30 el0_sync_handler+0xb0/0xb4 el0_sync+0x160/0x180 Therefore, we add cond_resched() to __audit_inode_child() to fix it.	N/A	More Details
CVE-2023-54046	In the Linux kernel, the following vulnerability has been resolved: crypto: essiv - Handle EBUSY correctly As it is essiv only handles the special return value of EINPROGERSS, which means that in all other cases it will free data related to the request. However, as the caller of essiv may specify MAY_BACKLOG, we also need to expect EBUSY and treat it in the same way. Otherwise backlogged requests will trigger a use-after-free.	N/A	More Details
CVE-2023-54047	In the Linux kernel, the following vulnerability has been resolved: drm/rockchip: dw_hdmi: cleanup drm encoder during unbind This fixes a use-after-free crash during rmmod. The DRM encoder is embedded inside the larger rockchip_hdmi, which is allocated with the component. The component memory gets freed before the main drm device is destroyed. Fix it by running encoder cleanup before tearing down its container. [moved encoder cleanup above clk_disable, similar to bind-error-path]	N/A	More Details
CVE-2023-54057	In the Linux kernel, the following vulnerability has been resolved: iommu/amd: Add a length limitation for the ivrs_acpihid command-line parameter The 'acpiid' buffer in the parse_ivrs_acpihid function may overflow, because the string specifier in the format string sscanf() has no width limitation. Found by InfoTeCS on behalf of Linux Verification Center (linuxtesting.org) with SVACE.	N/A	More Details
CVE-2023-54058	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_ffa: Check if ffa_driver remove is present before executing Currently ffa_drv->remove() is called unconditionally from ffa_device_remove(). Since the driver registration doesn't check for it and allows it to be registered without .remove callback, we need to check for the presence of it before executing it from ffa_device_remove() to above a NULL pointer dereference like the one below: | Unable to handle kernel NULL pointer dereference at virtual address 0000000000000000 | Mem abort info: | ESR = 0x0000000086000004 | EC = 0x21: IABT (current EL), IL = 32 bits | SET = 0, FnV = 0 | EA = 0, S1PTW = 0 | FSC = 0x04: level 0 translation fault | user pgtable: 4k pages, 48-bit VAs, pgdp=0000000881cc8000 | [0000000000000000] pgd=0000000000000000, p4d=0000000000000000 | Internal error: Oops: 0000000086000004 [#1] PREEMPT SMP | CPU: 3 PID: 130 Comm: rmmod Not tainted 6.3.0-rc7 #6 | Hardware name: FVP Base RevC (DT) | pstate: 63402809 (nZCv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=-c) | pc : 0x0 | lr : ffa_device_remove+0x20/0x2c | Call trace: | 0x0 | device_release_driver_internal+0x16c/0x260 | driver_detach+0x90/0xd0 | bus_remove_driver+0xdc/0x11c | driver_unregister+0x30/0x54 | ffa_driver_unregister+0x14/0x20 | cleanup_module+0x18/0xeec | __arm64_sys_delete_module+0x234/0x378 | invoke_syscall+0x40/0x108 | el0_svc_common+0xb4/0xf0 | do_el0_svc+0x30/0xa4 | el0_svc+0x2c/0x7c | el0t_64_sync_handler+0x84/0xf0 | el0t_64_sync+0x190/0x194	N/A	More Details
CVE-2022-50741	In the Linux kernel, the following vulnerability has been resolved: media: imx-jpeg: Disable useless interrupt to avoid kernel panic There is a hardware bug that the interrupt STMBUF_HALF may be triggered after or when disable interrupt. It may led to unexpected kernel panic. And interrupt STMBUF_HALF and STMBUF_RTND have no other effect. So disable them and the unused interrupts. meanwhile clear the interrupt status when disable interrupt.	N/A	More Details
CVE-2022-50746	In the Linux kernel, the following vulnerability has been resolved: erofs: validate the extent length for uncompressed pclusters syzkaller reported a KASAN use-after-free: https://syzkaller.appspot.com/bug?extid=2ae90e873e97f1faf6f2 The referenced fuzzed image actually has two issues: - m_pa == 0 as a non-inlined pcluster; - The logical length is longer than its physical length. The first issue has already been addressed. This patch addresses the second issue by checking the extent length validity.	N/A	More Details
CVE-2022-50751	In the Linux kernel, the following vulnerability has been resolved: configfs: fix possible memory leak in configfs_create_dir() kmemleak reported memory leaks in configfs_create_dir(): unreferenced object 0xffff888009f6af00 (size 192): comm "modprobe", pid 3777, jiffies 4295537735 (age 233.784s) backtrace: kmem_cache_alloc (mm/slub.c:3250 mm/slub.c:3256 mm/slub.c:3263 mm/slub.c:3273) new_fragment (./include/linux/slab.h:600 fs/configfs/dir.c:163) configfs_register_subsystem (fs/configfs/dir.c:1857) basic_write (drivers/hwtracing/stm/p_basic.c:14) stm_p_basic do_one_initcall (init/main.c:1296) do_init_module (kernel/module/main.c:2455) ... unreferenced object 0xffff888003ba7180 (size 96): comm "modprobe", pid 3777, jiffies 4295537735 (age 233.784s) backtrace: kmem_cache_alloc (mm/slub.c:3250 mm/slub.c:3256 mm/slub.c:3263 mm/slub.c:3273) configfs_new_dirent (./include/linux/slab.h:723 fs/configfs/dir.c:194) configfs_make_dirent (fs/configfs/dir.c:248) configfs_create_dir (fs/configfs/dir.c:296) configfs_attach_group.isra.28 (fs/configfs/dir.c:816 fs/configfs/dir.c:852) configfs_register_subsystem (fs/configfs/dir.c:1881) basic_write (drivers/hwtracing/stm/p_basic.c:14) stm_p_basic do_one_initcall (init/main.c:1296) do_init_module (kernel/module/main.c:2455) ... This is because the refcount is not correct in configfs_make_dirent(). For normal stage, the refcount is changing as: configfs_register_subsystem() configfs_create_dir() configfs_make_dirent() configfs_new_dirent() # set s_count = 1 dentry->d_fsdata = configfs_get(sd); # s_count = 2 ... configfs_unregister_subsystem() configfs_remove_dir() remove_dir() configfs_remove_dirent() # s_count = 1 dput() ... *dentry_unlink_inode()* configfs_d_iput() # s_count = 0, release	N/A	More Details

	However, if we failed in configfs_create(): configfs_register_subsystem() configfs_create_dir() configfs_make_dirent() # s_count = 2 ... configfs_create() # fail ->out_remove: configfs_remove_dirent(dentry) configfs_put(sd) # s_count = 1 return PTR_ERR(inode); There is no inode in the error path, so the configfs_d_iput() is lost and makes sd and fragment memory leaked. To fix this, when we failed in configfs_create(), manually call configfs_put(sd) to keep the refcount correct.		
CVE-2023-54067	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix race when deleting free space root from the dirty cow roots list When deleting the free space tree we are deleting the free space root from the list fs_info->dirty_cowonly_roots without taking the lock that protects it, which is struct btrfs_fs_info::trans_lock. This unsynchronized list manipulation may cause chaos if there's another concurrent manipulation of this list, such as when adding a root to it with ctree.c:add_root_to_dirty_list(). This can result in all sorts of weird failures caused by a race, such as the following crash: [337571.278245] general protection fault, probably for non-canonical address 0xdead000000000108: 0000 [#1] PREEMPT SMP PTI [337571.278933] CPU: 1 PID: 115447 Comm: btrfs Tainted: G W 6.4.0-rc6-btrfs-next-134+ #1 [337571.279153] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org 04/01/2014 [337571.279572] RIP: 0010:commit_cowonly_roots+0x11f/0x250 [btrfs] [337571.279928] Code: 85 38 06 00 (...) [337571.280363] RSP: 0018:ffff9f63446efba0 EFLAGS: 00010206 [337571.280582] RAX: ffff942d98ec2638 RBX: ffff9430b82b4c30 RCX: 0000000449e1c000 [337571.280798] RDX: dead000000000100 RSI: ffff9430021e4900 RDI: 0000000000036070 [337571.281015] RBP: ffff942d98ec2000 R08: ffff942d98ec2000 R09: 000000000000015b [337571.281254] R10: 0000000000000009 R11: 0000000000000001 R12: ffff942fe8fbf600 [337571.281476] R13: ffff942dabe23040 R14: ffff942dabe20800 R15: ffff942d92cf3b48 [337571.281723] FS: 00007f478adb7340(0000) GS:ffff94349fa40000(0000) knlGS:0000000000000000 [337571.281950] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [337571.282184] CR2: 00007f478ab9a3d5 CR3: 000000001e02c001 CR4: 0000000000370ee0 [337571.282416] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [337571.282647] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [337571.282874] Call Trace: [337571.283101] <TASK> [337571.283327] ? __die_body+0x1b/0x60 [337571.283570] ? die_addr+0x39/0x60 [337571.283796] ? exc_general_protection+0x22e/0x430 [337571.284022] ? asm_exc_general_protection+0x22/0x30 [337571.284251] ? commit_cowonly_roots+0x11f/0x250 [btrfs] [337571.284531] btrfs_commit_transaction+0x42e/0xf90 [btrfs] [337571.284803] ? _raw_spin_unlock+0x15/0x30 [337571.285031] ? release_extent_buffer+0x103/0x130 [btrfs] [337571.285305] reset_balance_state+0x152/0x1b0 [btrfs] [337571.285578] btrfs_balance+0xa50/0x11e0 [btrfs] [337571.285864] ? __kmem_cache_alloc_node+0x14a/0x410 [337571.286086] btrfs_ioctl+0x249a/0x3320 [btrfs] [337571.286358] ? mod_objcg_state+0xd2/0x360 [337571.286577] ? refill_obj_stock+0xb0/0x160 [337571.286798] ? seq_release+0x25/0x30 [337571.287016] ? __rseq_handle_notify_resume+0x3ba/0x4b0 [337571.287235] ? percpu_counter_add_batch+0x2e/0xa0 [337571.287455] ? __x64_sys_ioctl+0x88/0xc0 [337571.287675] __x64_sys_ioctl+0x88/0xc0 [337571.287901] do_syscall_64+0x38/0x90 [337571.288126] entry_SYSCALL_64_after_hwframe+0x72/0xdc [337571.288352] RIP: 0033:0x7f478aaffe9b So fix this by locking struct btrfs_fs_info::trans_lock before deleting the free space root from that list.	N/A	More Details
CVE-2023-54068	In the Linux kernel, the following vulnerability has been resolved: f2fs: compress: fix to call f2fs_wait_on_page_writeback() in f2fs_write_raw_pages() BUG_ON() will be triggered when writing files concurrently, because the same page is writtenback multiple times. 1597 void folio_end_writeback(struct folio *folio) 1598 { ..... 1618 if (!__folio_end_writeback(folio)) 1619 BUG(); ..... 1625 } kernel BUG at mm/filemap.c:1619! Call Trace: <TASK> f2fs_write_end_io+0x1a0/0x370 blk_update_request+0x6c/0x410 blk_mq_end_request+0x15/0x130 blk_complete_reqs+0x3c/0x50 __do_softirq+0xb8/0x29b ? sort_range+0x20/0x20 run_ksoftirqd+0x19/0x20 smpboot_thread_fn+0x10b/0x1d0 kthread+0xde/0x110 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork+0x22/0x30 </TASK> Below is the concurrency scenario: [Process A] [Process B] [Process C] f2fs_write_raw_pages() - redirty_page_for_writepage() - unlock page() f2fs_do_write_data_page() - lock_page() - clear_page_dirty_for_io() - set_page_writeback() [1st writeback] ..... - unlock page() generic_perform_write() - f2fs_write_begin() - wait_for_stable_page() - f2fs_write_end() - set_page_dirty() - lock_page() - f2fs_do_write_data_page() - set_page_writeback() [2st writeback] This problem was introduced by the previous commit 7377e853967b ("f2fs: compress: fix potential deadlock of compress file"). All pagelocks were released in f2fs_write_raw_pages(), but whether the page was in the writeback state was ignored in the subsequent writing process. Let's fix it by waiting for the page to writeback before writing.	N/A	More Details
CVE-2025-13158	Prototype pollution vulnerability in apidoc-core versions 0.2.0 and all subsequent versions allows remote attackers to modify JavaScript object prototypes via malformed data structures, including the “define” property processed by the application, potentially leading to denial of service or unintended behavior in applications relying on the integrity of prototype chains. This affects the preProcess() function in api_group.js, api_param_title.js, api_use.js, and api_permission.js worker modules.	N/A	More Details
CVE-2022-50750	In the Linux kernel, the following vulnerability has been resolved: drm/panel/panel-sitronix-st7701: Remove panel on DSI attach failure In case mipi_dsi_attach() fails, call drm_panel_remove() to avoid memory leak.	N/A	More Details
CVE-2022-50749	In the Linux kernel, the following vulnerability has been resolved: acct: fix potential integer overflow in encode_comp_t() The integer overflow is descripted with following codes: > 317 static comp_t encode_comp_t(u64 value) > 318 { > 319 int exp, rnd; ..... > 341 exp <<= MANTSIZE; > 342 exp += value; > 343 return exp; > 344 } Currently comp_t is defined as type of '__u16', but the variable 'exp' is type of 'int', so overflow would happen when variable 'exp' in line 343 is greater than 65535.	N/A	More Details
CVE-2022-50748	In the Linux kernel, the following vulnerability has been resolved: ipc: mqueue: fix possible memory leak in init_mqueue_fs() commit db7cfc380900 ("ipc: Free mq_sysctls if ipc namespace creation failed") Here's a similar memory leak to the one fixed by the patch above. retire_mq_sysctls need to be called when init_mqueue_fs fails after setup_mq_sysctls.	N/A	More Details
CVE-2022-50747	In the Linux kernel, the following vulnerability has been resolved: hfs: Fix OOB Write in hfs_asc2mac Syzbot reported a OOB Write bug: loop0: detected capacity change from 0 to 64 == BUG: KASAN: slab-out-of-bounds in hfs_asc2mac+0x467/0x9a0 fs/hfs/trans.c:133 Write of size 1 at addr ffff888018483l4e by task syz-executor391/3632 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x1b1/0x28e lib/dump_stack.c:106 print_address_description+0x74/0x340 mm/kasan/report.c:284 print_report+0x107/0x1f0 mm/kasan/report.c:395 kasan_report+0xcd/0x100 mm/kasan/report.c:495 hfs_asc2mac+0x467/0x9a0 fs/hfs/trans.c:133 hfs_cat_build_key+0x92/0x170 fs/hfs/catalog.c:28 hfs_lookup+0x1ab/0x2c0 fs/hfs/dir.c:31 lookup_open fs/namei.c:3391 [inline] open_last_lookups fs/namei.c:3481 [inline] path_openat+0x10e6/0x2df0 fs/namei.c:3710 do_filp_open+0x264/0x4f0 fs/namei.c:3740 If in->len is much larger than HFS_NAMELEN(31) which is the maximum length of an HFS filename, a OOB write could occur in hfs_asc2mac(). In that case, when the dst reaches the boundary, the srclen is still greater than 0, which causes a OOB write. Fix this by adding a check on dstlen in while() before writing to dst address.	N/A	More Details
CVE-2022-50745	In the Linux kernel, the following vulnerability has been resolved: staging: media: tegra-video: fix device_node use after free At probe time this code path is followed: * tegra_csi_init * tegra_csi_channels_alloc * for_each_child_of_node(node, channel) -- iterates over channels * automatically gets 'channel' * tegra_csi_channel_alloc() * saves into chan->of_node a pointer to the channel OF node * automatically gets and puts 'channel' * now the node saved in chan->of_node has refcount 0, can disappear * tegra_csi_channels_init * iterates over channels * tegra_csi_channel_init -- uses chan->of_node After that, chan->of_node keeps	N/A	More Details

	storing the node until the device is removed. of_node_get() the node and of_node_put() it during teardown to avoid any risk.		
CVE-2018-25153	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority as the reported issue does not constitute a security vulnerability and represents a minor, non-exploitable memory leak.	N/A	More Details
CVE-2023-54069	In the Linux kernel, the following vulnerability has been resolved: ext4: fix BUG in ext4_mb_new_inode_pa() due to overflow When we calculate the end position of ext4_free_extent, this position may be exactly where ext4_lblk_t (i.e. uint) overflows. For example, if ac_g_ex.fe_logical is 4294965248 and ac_orig_goal_len is 2048, then the computed end is 0x100000000, which is 0. If ac->ac_o_ex.fe_logical is not the first case of adjusting the best extent, that is, new_bex_end > 0, the following BUG_ON will be triggered: ========== kernel BUG at fs/ext4/mballoc.c:5116! invalid opcode: 0000 [#1] PREEMPT SMP PTI CPU: 3 PID: 673 Comm: xfs_io Tainted: G E 6.5.0-rc1+ #279 RIP: 0010:ext4_mb_new_inode_pa+0xc5/0x430 Call Trace: <TASK> ext4_mb_use_best_found+0x203/0x2f0 ext4_mb_try_best_found+0x163/0x240 ext4_mb_regular_allocator+0x158/0x1550 ext4_mb_new_blocks+0x86a/0xe10 ext4_ext_map_blocks+0xb0c/0x13a0 ext4_map_blocks+0x2cd/0x8f0 ext4_iomap_begin+0x27b/0x400 iomap_iter+0x222/0x3d0 __iomap_dio_rw+0x243/0xcb0 iomap_dio_rw+0x16/0x80 ========== A simple reproducer demonstrating the problem: mkfs.ext4 -F /dev/sda -b 4096 100M mount /dev/sda /tmp/test fallocate -l1M /tmp/test/tmp fallocate -l10M /tmp/test/file fallocate -i -o 1M -l16777203M /tmp/test/file fsstress -d /tmp/test -l 0 -n 100000 -p 8 & sleep 10 && killall -9 fsstress rm -f /tmp/test/tmp xfs_io -c "open -ad /tmp/test/file" -c "pwrite -S 0xff 0 8192" We simply refactor the logic for adjusting the best extent by adding a temporary ext4_free_extent ex and use extent_logical_end() to avoid overflow, which also simplifies the code.	N/A	More Details
CVE-2022-50744	In the Linux kernel, the following vulnerability has been resolved: scsi: lpfc: Fix hard lockup when reading the rx_monitor from debugfs During I/O and simultaneous cat of /sys/kernel/debug/lpfc/fnX/rx_monitor, a hard lockup similar to the call trace below may occur. The spin_lock_bh in lpfc_rx_monitor_report is not protecting from timer interrupts as expected, so change the strength of the spin lock to _irq. Kernel panic - not syncing: Hard LOCKUP CPU: 3 PID: 110402 Comm: cat Kdump: loaded exception RIP: native_queued_spin_lock_slowpath+91 [IRQ stack] native_queued_spin_lock_slowpath at ffffffffb814e30b _raw_spin_lock at ffffffffb89a667a lpfc_rx_monitor_record at ffffffffc0a73a36 [lpfc] lpfc_cmf_timer at ffffffffc0abbc67 [lpfc] __hrtimer_run_queues at ffffffffb8184250 hrtimer_interrupt at ffffffffb8184ab0 smp_apic_timer_interrupt at ffffffffb8a026ba apic_timer_interrupt at ffffffffb8a01c4f [End of IRQ stack] apic_timer_interrupt at ffffffffb8a01c4f lpfc_rx_monitor_report at ffffffffc0a73c80 [lpfc] lpfc_rx_monitor_read at ffffffffc0addde1 [lpfc] full_proxy_read at ffffffffb83e7fc3 vfs_read at ffffffffb833fe71 ksys_read at ffffffffb83402af do_syscall_64 at ffffffffb800430b entry_SYSCALL_64_after_hwframe at ffffffffb8a000ad	N/A	More Details
CVE-2022-50743	In the Linux kernel, the following vulnerability has been resolved: erofs: Fix pcluster memleak when its block address is zero syzkaller reported a memleak: https://syzkaller.appspot.com/bug?id=62f37ff612f0021641eda5b17f056f1668aa9aed unreferenced object 0xffff88811009c7f8 (size 136): ... backtrace: [<ffffffff821db19b>] z_erofs_do_read_page+0x99b/0x1740 [<ffffffff821dee9e>] z_erofs_readahead+0x24e/0x580 [<ffffffff814bc0d6>] read_pages+0x86/0x3d0 ... syzkaller constructed a case: in z_erofs_register_pcluster(), ztailpacking = false and map->m_pa = zero. This makes pcl->obj.index be zero although pcl is not a inline pcluster. Then following path adds refcount for grp, but the refcount won't be put because pcl is inline. z_erofs_readahead() z_erofs_do_read_page() # for another page z_erofs_collector_begin() erofs_find_workgroup() erofs_workgroup_get() Since it's illegal for the block address of a non-inlined pcluster to be zero, add check here to avoid registering the pcluster which would be leaked.	N/A	More Details
CVE-2022-50742	In the Linux kernel, the following vulnerability has been resolved: misc: ocxl: fix possible refcount leak in afu_ioctl() eventfd_ctx_put need to be called to put the refcount that gotten by eventfd_ctx_fdget when ocxl_irq_set_handler fails.	N/A	More Details
CVE-2023-54070	In the Linux kernel, the following vulnerability has been resolved: igb: clean up in all error paths when enabling SR-IOV After commit 50f303496d92 ("igb: Enable SR-IOV after reinit"), removing the igb module could hang or crash (depending on the machine) when the module has been loaded with the max_vfs parameter set to some value != 0. In case of one test machine with a dual port 82580, this hang occurred: [232.480687] igb 0000:41:00.1: removed PHC on enp65s0f1 [233.093257] igb 0000:41:00.1: IOV Disabled [233.329969] pcieport 0000:40:01.0: AER: Multiple Uncorrected (Non-Fatal) err0 [233.340302] igb 0000:41:00.0: PCIe Bus Error: severity=Uncorrected (Non-Fata) [233.352248] igb 0000:41:00.0: device [8086:1516] error status/mask=00100000 [233.361088] igb 0000:41:00.0: [20] UnsupReq (First) [233.368183] igb 0000:41:00.0: AER: TLP Header: 40000001 0000040f cdbfc00c c [233.376846] igb 0000:41:00.1: PCIe Bus Error: severity=Uncorrected (Non-Fata) [233.388779] igb 0000:41:00.1: device [8086:1516] error status/mask=00100000 [233.397629] igb 0000:41:00.1: [20] UnsupReq (First) [233.404736] igb 0000:41:00.1: AER: TLP Header: 40000001 0000040f cdbfc00c c [233.538214] pci 0000:41:00.1: AER: can't recover (no error_detected callback) [233.538401] igb 0000:41:00.0: removed PHC on enp65s0f0 [233.546197] pcieport 0000:40:01.0: AER: device recovery failed [234.157244] igb 0000:41:00.0: IOV Disabled [371.619705] INFO: task irq/35-aerdrv:257 blocked for more than 122 seconds. [371.627489] Not tainted 6.4.0-dirty #2 [371.632257] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this. [371.641000] task:irq/35-aerdrv state:D stack:0 pid:257 ppid:2 f0 [371.650330] Call Trace: [371.653061] <TASK> [371.655407] __schedule+0x20e/0x660 [371.659313] schedule+0x5a/0xd0 [371.662824] schedule_preempt_disabled+0x11/0x20 [371.667983] __mutex_lock.constprop.0+0x372/0x6c0 [371.673237] ? __pfx_aer_root_reset+0x10/0x10 [371.678105] report_error_detected+0x25/0x1c0 [371.682974] ? __pfx_report_normal_detected+0x10/0x10 [371.688618] pci_walk_bus+0x72/0x90 [371.692519] pcie_do_recovery+0xb2/0x330 [371.696899] aer_process_err_devices+0x117/0x170 [371.702055] aer_isr+0x1c0/0x1e0 [371.705661] ? __set_cpus_allowed_ptr+0x54/0xa0 [371.710723] ? __pfx_irq_thread_fn+0x10/0x10 [371.715496] irq_thread_fn+0x20/0x60 [371.719491] irq_thread+0xe6/0x1b0 [371.723291] ? __pfx_irq_thread_dtor+0x10/0x10 [371.728255] ? __pfx_irq_thread+0x10/0x10 [371.732731] kthread+0xe2/0x110 [371.736243] ? __pfx_kthread+0x10/0x10 [371.740430] ret_from_fork+0x2c/0x50 [371.744428] </TASK> The reproducer was a simple script: #!/bin/sh for i in `seq 1 5`; do modprobe -rv igb modprobe -v igb max_vfs=1 sleep 1 modprobe -rv igb done It turned out that this could only be reproduce on 82580 (quad and dual-port), but not on 82576, i350 and i210. Further debugging showed that igb_enable_sriov()'s call to pci_enable_sriov() is failing, because dev->is_physfn is 0 on 82580. Prior to commit 50f303496d92 ("igb: Enable SR-IOV after reinit"), igb_enable_sriov() jumped into the "err_out" cleanup branch. After this commit it only returned the error code. So the cleanup didn't take place, and the incorrect VF setup in the igb_adapter structure fooled the igb driver into assuming that VFs have been set up where no VF actually existed. Fix this problem by cleaning up again if pci_enable_sriov() fails.	N/A	More Details
CVE-2025-68473	ESF-IDF is the Espressif Internet of Things (IOT) Development Framework. In versions 5.5.1, 5.4.3, 5.3.4, 5.2.6, 5.1.6, and earlier, in the ESP-IDF Bluetooth host stack (BlueDroid), the function bta_dm_sdp_result() used a fixed-size array uuid_list[32][MAX_UUID_SIZE] to store discovered service UUIDs during the SDP (Service Discovery Protocol) process. On modern Bluetooth devices, it is possible for the number of available services to exceed this fixed limit (32). In such cases, if more than 32 services are discovered, subsequent writes to uuid_list could exceed the bounds of the array, resulting in a potential out-of-bounds write condition.	N/A	More Details
	ESF-IDF is the Espressif Internet of Things (IOT) Development Framework. In versions 5.5.1, 5.4.3, 5.3.4, 5.2.6, 5.1.6, and earlier, in the avrc_vendor_msg() function of the ESP-IDF BlueDroid AVRCP stack, the allocated buffer size was validated using		

CVE-2025-68474	AVRC_MIN_CMD_LEN (20 bytes). However, the actual fixed header data written before the vendor payload exceeds this value. This totals 29 bytes written before p_msg->p_vendor_data is copied. Using the old AVRC_MIN_CMD_LEN could allow an out-of-bounds write if vendor_len approaches the buffer limit. For commands where vendor_len is large, the original buffer allocation may be insufficient, causing writes beyond the allocated memory. This can lead to memory corruption, crashes, or other undefined behavior. The overflow could be larger when assertions are disabled.	N/A	More Details
CVE-2025-68932	FreshRSS is a free, self-hostable RSS aggregator. Prior to version 1.28.0, FreshRSS uses cryptographically weak random number generators (mt_rand() and uniqid()) to generate remember-me authentication tokens and challenge-response nonces. This allows attackers to predict valid session tokens, leading to account takeover through persistent session hijacking. The remember-me tokens provide permanent authentication and are the sole credential for "keep me logged in" functionality. This issue has been patched in version 1.28.0.	N/A	More Details
CVE-2023-54066	In the Linux kernel, the following vulnerability has been resolved: media: dvb-usb-v2: gl861: Fix null-ptr-deref in gl861_i2c_master_xfer In gl861_i2c_master_xfer, msg is controlled by user. When msg[i].buf is null and msg[i].len is zero, former checks on msg[i].buf would be passed. Malicious data finally reach gl861_i2c_master_xfer. If accessing msg[i].buf[0] without sanity check, null ptr deref would happen. We add check on msg[i].len to prevent crash. Similar commit: commit 0ed554fd769a ("media: dvb-usb: az6027: fix null-ptr-deref in az6027_i2c_xfer()")	N/A	More Details
CVE-2022-50752	In the Linux kernel, the following vulnerability has been resolved: md/raid5: Remove unnecessary bio_put() in raid5_read_one_chunk() When running chunk-sized reads on disks with badblocks duplicate bio free/puts are observed: == BUG bio-200 (Not tainted): Object already free ---------- Allocated in mempool_alloc_slab+0x17/0x20 age=3 cpu=2 pid=7504 __slab_alloc.constprop.0+0x5a/0xb0 kmem_cache_alloc+0x31e/0x330 mempool_alloc_slab+0x17/0x20 mempool_alloc+0x100/0x2b0 bio_alloc_bioset+0x181/0x460 do_mpage_readpage+0x776/0xd00 mpage_readahead+0x166/0x320 blkdev_readahead+0x15/0x20 read_pages+0x13f/0x5f0 page_cache_ra_unbounded+0x18d/0x220 force_page_cache_ra+0x181/0x1c0 page_cache_sync_ra+0x65/0xb0 filemap_get_pages+0x1df/0xaf0 filemap_read+0x1e1/0x700 blkdev_read_iter+0x1e5/0x330 vfs_read+0x42a/0x570 Freed in mempool_free_slab+0x17/0x20 age=3 cpu=2 pid=7504 kmem_cache_free+0x46d/0x490 mempool_free_slab+0x17/0x20 mempool_free+0x66/0x190 bio_free+0x78/0x90 bio_put+0x100/0x1a0 raid5_make_request+0x2259/0x2450 md_handle_request+0x402/0x600 md_submit_bio+0xd9/0x120 __submit_bio+0x11f/0x1b0 submit_bio_noacct_nocheck+0x204/0x480 submit_bio_noacct+0x32e/0xc70 submit_bio+0x98/0x1a0 mpage_readahead+0x250/0x320 blkdev_readahead+0x15/0x20 read_pages+0x13f/0x5f0 page_cache_ra_unbounded+0x18d/0x220 Slab 0xffffea000481b600 objects=21 used=0 fp=0xffff8881206d8940 flags=0x17ffffc0010201(locked|slab|head|node=0|zone=2|lastcpupid=0x1fffff) CPU: 0 PID: 34525 Comm: kworker/u24:2 Not tainted 6.0.0-rc2-localyes-265166-gf11c5343fa3f #143 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.13.0-1ubuntu1.1 04/01/2014 Workqueue: raid5wq raid5_do_work Call Trace: <TASK> dump_stack_lvl+0x5a/0x78 dump_stack+0x10/0x16 print_trailer+0x158/0x165 object_err+0x35/0x50 free_debug_processing.cold+0xb7/0xbe __slab_free+0x1ae/0x330 kmem_cache_free+0x46d/0x490 mempool_free_slab+0x17/0x20 mempool_free+0x66/0x190 bio_free+0x78/0x90 bio_put+0x100/0x1a0 mpage_end_io+0x36/0x150 bio_endio+0x2fd/0x360 md_end_io_acct+0x7e/0x90 bio_endio+0x2fd/0x360 handle_failed_stripe+0x960/0xb80 handle_stripe+0x1348/0x3760 handle_active_stripes.constprop.0+0x72a/0xaf0 raid5_do_work+0x177/0x330 process_one_work+0x616/0xb20 worker_thread+0x2bd/0x6f0 kthread+0x179/0x1b0 ret_from_fork+0x22/0x30 </TASK> The double free is caused by an unnecessary bio_put() in the if(is_badblock(...)) error path in raid5_read_one_chunk(). The error path was moved ahead of bio_alloc_clone() in c82aa1b76787c ("md/raid5: move checking badblock before clone bio in raid5_read_one_chunk"). The previous code checked and freed align_bio which required a bio_put. After the move that is no longer needed as raid_bio is returned to the control of the common io path which performs its own endio resulting in a double free on bad device blocks.	N/A	More Details
CVE-2023-54065	In the Linux kernel, the following vulnerability has been resolved: net: dsa: realtek: fix out-of-bounds access The probe function sets priv->chip_data to (void *)priv + sizeof(*priv) with the expectation that priv has enough trailing space. However, only realtek-smi actually allocated this chip_data space. Do likewise in realtek-mdio to fix out-of-bounds accesses. These accesses likely went unnoticed so far, because of an (unused) buf[4096] member in struct realtek_priv, which caused kmalloc to round up the allocated buffer to a big enough size, so nothing of value was overwritten. With a different allocator (like in the barebox bootloader port of the driver) or with KASAN, the memory corruption becomes quickly apparent.	N/A	More Details
CVE-2023-54064	In the Linux kernel, the following vulnerability has been resolved: ipmi:ssif: Fix a memory leak when scanning for an adapter The adapter scan ssif_info_find() sets info->adapter_name if the adapter info came from SMBIOS, as it's not set in that case. However, this function can be called more than once, and it will leak the adapter name if it had already been set. So check for NULL before setting it.	N/A	More Details
CVE-2023-54059	In the Linux kernel, the following vulnerability has been resolved: soc: mediatek: mtk-svs: Enable the IRQ later If the system does not come from reset (like when is booted via kexec()), the peripheral might triger an IRQ before the data structures are initialised. [0.227710] Unable to handle kernel NULL pointer dereference at virtual address 0000000000000f08 [0.227913] Call trace: [0.227918] svs_isr+0x8c/0x538	N/A	More Details
CVE-2025-52598	Cybersecurity Nozomi Networks Labs, a specialized security company focused on Industrial Control Systems (ICS) and OT/IoT security, has found a flaw that camera's client service does not perform certificate validation. The manufacturer has released patch firmware for the flaw, please refer to the manufacturer's report for details and workarounds.	N/A	More Details
CVE-2025-52599	Cybersecurity Nozomi Networks Labs, a specialized security company focused on Industrial Control Systems (ICS) and OT/IoT security, has discovered Inadequate of permission management for camera guest account. The manufacturer has released patch firmware for the flaw, please refer to the manufacturer's report for details and workarounds.	N/A	More Details
CVE-2025-52600	Cybersecurity Nozomi Networks Labs, a specialized security company focused on Industrial Control Systems (ICS) and OT/IoT security, has discovered a vulnerability in camera video analytics that Improper input validation. This vulnerability could allow an attacker to execute specific commands on the user's host PC.The manufacturer has released patch firmware for the flaw, please refer to the manufacturer's report for details and workarounds.	N/A	More Details
CVE-2025-52601	Cybersecurity Nozomi Networks Labs, a specialized security company focused on Industrial Control Systems (ICS) and OT/IoT security, has discovered a vulnerability in Device Manager that a hardcoded encryption key for sensitive information. An attacker can use key to decrypt sensitive information. The manufacturer has released patch firmware for the flaw, please refer to the manufacturer's report for details and workarounds.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: iommufd: Set end correctly when doing batch carry Even though the test suite covers this it somehow became obscured that this wasn't working. The test		

CVE-2023-54060	iommufd_ioas.mock_domain.access_domain_destory would blow up rarely. end should be set to 1 because this just pushed an item, the carry, to the pfns list. Sometimes the test would blow up with: BUG: kernel NULL pointer dereference, address: 0000000000000000 #PF: supervisor read access in kernel mode #PF: error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: 0000 [#1] SMP CPU: 5 PID: 584 Comm: iommufd Not tainted 6.5.0-rc1-dirty #1236 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 RIP: 0010:batch_unpin+0xa2/0x100 [iommufd] Code: 17 48 81 fe ff ff 07 00 77 70 48 8b 15 b7 be 97 e2 48 85 d2 74 14 48 8b 14 fa 48 85 d2 74 0b 40 0f b6 f6 48 c1 e6 04 48 01 f2 <48> 8b 3a 48 c1 e0 06 89 ca 48 89 de 48 83 e7 f0 48 01 c7 e8 96 dc RSP: 0018:ffffc90001677a58 EFLAGS: 00010246 RAX: 00007f7e2646f000 RBX: 0000000000000000 RCX: 0000000000000001 RDX: 0000000000000000 RSI: 00000000fefc4c8d RDI: 0000000000fefc4c RBP: ffffc90001677a80 R08: 0000000000000048 R09: 0000000000000200 R10: 0000000000030b98 R11: ffffffff81f3bb40 R12: 0000000000000001 R13: ffff888101f75800 R14: ffffc90001677ad0 R15: 00000000000001fe FS: 00007f9323679740(0000) GS:ffff8881ba540000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000000 CR3: 0000000105ede003 CR4: 00000000003706a0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> ? show_regs+0x5c/0x70 ? __die+0x1f/0x60 ? page_fault_oops+0x15d/0x440 ? lock_release+0xbc/0x240 ? exc_page_fault+0x4a4/0x970 ? asm_exc_page_fault+0x27/0x30 ? batch_unpin+0xa2/0x100 [iommufd] ? batch_unpin+0xba/0x100 [iommufd] __iopt_area_unfill_domain+0x198/0x430 [iommufd] ? __mutex_lock+0x8c/0xb80 ? __mutex_lock+0x6aa/0xb80 ? xa_erase+0x28/0x30 ? iopt_table_remove_domain+0x162/0x320 [iommufd] ? lock_release+0xbc/0x240 iopt_area_unfill_domain+0xd/0x10 [iommufd] iopt_table_remove_domain+0x195/0x320 [iommufd] iommufd_hw_pagetable_destroy+0xb3/0x110 [iommufd] iommufd_object_destroy_user+0x8e/0xf0 [iommufd] iommufd_device_detach+0xc5/0x140 [iommufd] iommufd_selftest_destroy+0x1f/0x70 [iommufd] iommufd_object_destroy_user+0x8e/0xf0 [iommufd] iommufd_destroy+0x3a/0x50 [iommufd] iommufd_fops_ioctl+0xfb/0x170 [iommufd] __x64_sys_ioctl+0x40d/0x9a0 do_syscall_64+0x3c/0x80 entry_SYSCALL_64_after_hwframe+0x46/0xb0	N/A	More Details
CVE-2025-8075	Cybersecurity Nozomi Networks Labs, a specialized security company focused on Industrial Control Systems (ICS) and OT/IoT security, has discovered that validation of incoming XML format request messages is inadequate. This vulnerability could allow an attacker to XSS on the user's browser. The manufacturer has released patch firmware for the flaw, please refer to the manufacturer's report for details and workarounds.	N/A	More Details
CVE-2025-62578	DVP-12SE - Modbus/TCP Cleartext Transmission of Sensitive Information	N/A	More Details
CVE-2022-50758	In the Linux kernel, the following vulnerability has been resolved: staging: vt6655: fix potential memory leak In function device_init_td0_ring, memory is allocated for member td_info of priv->apTD0Rings[i], with i increasing from 0. In case of allocation failure, the memory is freed in reversed order, with i decreasing to 0. However, the case i=0 is left out and thus memory is leaked. Modify the memory freeing loop to include the case i=0.	N/A	More Details
CVE-2022-50757	In the Linux kernel, the following vulnerability has been resolved: media: camss: Clean up received buffers on failed start of streaming It is required to return the received buffers, if streaming can not be started. For instance media_pipeline_start() may fail with EPIPE, if a link validation between entities is not passed, and in such a case a user gets a kernel warning: WARNING: CPU: 1 PID: 520 at drivers/media/common/videobuf2/videobuf2-core.c:1592 vb2_start_streaming+0xec/0x160 <snip> Call trace: vb2_start_streaming+0xec/0x160 vb2_core_streamon+0x9c/0x1a0 vb2_ioctl_streamon+0x68/0xbc v4l_streamon+0x30/0x3c __video_do_ioctl+0x184/0x3e0 video_usercopy+0x37c/0x7b0 video_ioctl2+0x24/0x40 v4l2_ioctl+0x4c/0x70 The fix is to correct the error path in video_start_streaming() of camss.	N/A	More Details
CVE-2022-50756	In the Linux kernel, the following vulnerability has been resolved: nvme-pci: fix mempool alloc size Convert the max size to bytes to match the units of the divisor that calculates the worst-case number of PRP entries. The result is used to determine how many PRP Lists are required. The code was previously rounding this to 1 list, but we can require 2 in the worst case. In that scenario, the driver would corrupt memory beyond the size provided by the mempool. While unlikely to occur (you'd need a 4MB in exactly 127 phys segments on a queue that doesn't support SGLs), this memory corruption has been observed by kfence.	N/A	More Details
CVE-2022-50755	In the Linux kernel, the following vulnerability has been resolved: udf: Avoid double brelse() in udf_rename() syzbot reported a warning like below [1]: VFS: brelse: Trying to free free buffer WARNING: CPU: 2 PID: 7301 at fs/buffer.c:1145 __brelse+0x67/0xa0 ... Call Trace: <TASK> invalidate_bh_lru+0x99/0x150 smp_call_function_many_cond+0xe2a/0x10c0 ? generic_remap_file_range_prep+0x50/0x50 ? __brelse+0xa0/0xa0 ? __mutex_lock+0x21c/0x12d0 ? smp_call_on_cpu+0x250/0x250 ? rcu_read_lock_sched_held+0xb/0x60 ? lock_release+0x587/0x810 ? __brelse+0xa0/0xa0 ? generic_remap_file_range_prep+0x50/0x50 on_each_cpu_cond_mask+0x3c/0x80 blkdev_flush_mapping+0x13a/0x2f0 blkdev_put_whole+0xd3/0xf0 blkdev_put+0x222/0x760 deactivate_locked_super+0x96/0x160 deactivate_super+0xda/0x100 cleanup_mnt+0x222/0x3d0 task_work_run+0x149/0x240 ? task_work_cancel+0x30/0x30 do_exit+0xb29/0x2a40 ? reacquire_held_locks+0x4a0/0x4a0 ? do_raw_spin_lock+0x12a/0x2b0 ? mm_update_next_owner+0x7c0/0x7c0 ? rwlock_bug.part.0+0x90/0x90 ? zap_other_threads+0x234/0x2d0 do_group_exit+0xd0/0x2a0 __x64_sys_exit_group+0x3a/0x50 do_syscall_64+0x34/0xb0 entry_SYSCALL_64_after_hwframe+0x63/0xcd The cause of the issue is that brelse() is called on both ofibh.sbh and ofibh.ebh by udf_find_entry() when it returns NULL. However, brelse() is called by udf_rename(), too. So, b_count on buffer_head becomes unbalanced. This patch fixes the issue by not calling brelse() by udf_rename() when udf_find_entry() returns NULL.	N/A	More Details
CVE-2023-54061	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2023-54062	In the Linux kernel, the following vulnerability has been resolved: ext4: fix invalid free tracking in ext4_xattr_move_to_block() In ext4_xattr_move_to_block(), the value of the extended attribute which we need to move to an external block may be allocated by kvmalloc() if the value is stored in an external inode. So at the end of the function the code tried to check if this was the case by testing entry->e_value_inum. However, at this point, the pointer to the xattr entry is no longer valid, because it was removed from the original location where it had been stored. So we could end up calling kvfree() on a pointer which was not allocated by kvmalloc(); or we could also potentially leak memory by not freeing the buffer when it should be freed. Fix this by storing whether it should be freed in a separate variable.	N/A	More Details
CVE-2022-50754	In the Linux kernel, the following vulnerability has been resolved: apparmor: fix a memleak in multi_transaction_new() In multi_transaction_new(), the variable t is not freed or passed out on the failure of copy_from_user(t->data, buf, size), which could lead to a memleak. Fix this bug by adding a put_multi_transaction(t) in the error path.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to do sanity check on summary info As Wenqing Liu		

CVE-2022-50753	reported in bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=216456 BUG: KASAN: use-after-free in recover_data+0x63ae/0x6ae0 [f2fs] Read of size 4 at addr ffff8881464dcd80 by task mount/1013 CPU: 3 PID: 1013 Comm: mount Tainted: G W 6.0.0-rc4 #1 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.15.0-1 04/01/2014 Call Trace: dump_stack_lvl+0x45/0x5e print_report.cold+0xf3/0x68d kasan_report+0xa8/0x130 recover_data+0x63ae/0x6ae0 [f2fs] f2fs_recover_fsync_data+0x120d/0x1fc0 [f2fs] f2fs_fill_super+0x4665/0x61e0 [f2fs] mount_bdev+0x2cf/0x3b0 legacy_get_tree+0xed/0x1d0 vfs_get_tree+0x81/0x2b0 path_mount+0x47e/0x19d0 do_mount+0xce/0xf0 __x64_sys_mount+0x12c/0x1a0 do_syscall_64+0x38/0x90 entry_SYSCALL_64_after_hwframe+0x63/0xcd The root cause is: in fuzzed image, SSA table is corrupted: ofs_in_node is larger than ADDRS_PER_PAGE(), result in out-of-range access on 4k-size page. - recover_data - do_recover_data - check_index_in_prev_nodes - f2fs_data_blkaddr This patch adds sanity check on summary info in recovery and GC flow in where the flows rely on them. After patch: [29.310883] F2FS-fs (loop0): Inconsistent ofs_in_node:65286 in summary, ino:0, nid:6, max:1018	N/A	More Details
CVE-2023-54063	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix OOB read in indx_insert_into_buffer Syzbot reported a OOB read bug: BUG: KASAN: slab-out-of-bounds in indx_insert_into_buffer+0xaa3/0x13b0 fs/ntfs3/index.c:1755 Read of size 17168 at addr ffff8880255e06c0 by task syz-executor308/3630 Call Trace: <TASK> memmove+0x25/0x60 mm/kasan/shadow.c:54 indx_insert_into_buffer+0xaa3/0x13b0 fs/ntfs3/index.c:1755 indx_insert_entry+0x446/0x6b0 fs/ntfs3/index.c:1863 ntfs_create_inode+0x1d3f/0x35c0 fs/ntfs3/inode.c:1548 ntfs_create+0x3e/0x60 fs/ntfs3/namei.c:100 lookup_open fs/namei.c:3413 [inline] If the member struct INDEX_BUFFER *index of struct indx_node is incorrect, that is, the value of __le32 used is greater than the value of __le32 total in struct INDEX_HDR. Therefore, OOB read occurs when memmove is called in indx_insert_into_buffer(). Fix this by adding a check in hdr_find_e().	N/A	More Details
CVE-2025-68380	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: fix peer HE MCS assignment In ath11k_wmi_send_peer_assoc_cmd(), peer's transmit MCS is sent to firmware as receive MCS while peer's receive MCS sent as transmit MCS, which goes against firmwire's definition. While connecting to a misbehaved AP that advertises 0xffff (meaning not supported) for 160 MHz transmit MCS map, firmware crashes due to 0xffff is assigned to he_mcs->rx_mcs_set field. Ext Tag: HE Capabilities [...] Supported HE-MCS and NSS Set [...] Rx and Tx MCS Maps 160 MHz [...] Tx HE-MCS Map 160 MHz: 0xffff Swap the assignment to fix this issue. As the HE rate control mask is meant to limit our own transmit MCS, it needs to go via he_mcs->rx_mcs_set field. With the aforementioned swapping done, change is needed as well to apply it to the peer's receive MCS. Tested-on: WCN6855 hw2.1 PCI WLAN.HSP.1.1-03125-QCAHSPSWPL_V1_V2_SILICONZ_LITE-3.6510.41 Tested-on: QCN9274 hw2.0 PCI WLAN.WBE.1.4.1-00199-QCAHKSWPL_SILICONZ-1	N/A	More Details
CVE-2023-54092	In the Linux kernel, the following vulnerability has been resolved: KVM: s390: pv: fix index value of replaced ASCE The index field of the struct page corresponding to a guest ASCE should be 0. When replacing the ASCE in s390_replace_asce(), the index of the new ASCE should also be set to 0. Having the wrong index might lead to the wrong addresses being passed around when notifying pte invalidations, and eventually to validity intercepts (VM crash) if the prefix gets unmapped and the notifier gets called with the wrong address.	N/A	More Details
CVE-2022-50597	Rejected reason: ** REJECT ** DO NOT USE THIS CVE RECORD. ConsultIDs: none. Reason: This record was in a CNA pool that was not assigned to any issues during 2022. Notes: none.	N/A	More Details
CVE-2023-54012	In the Linux kernel, the following vulnerability has been resolved: net: fix stack overflow when LRO is disabled for virtual interfaces When the virtual interface's feature is updated, it synchronizes the updated feature for its own lower interface. This propagation logic should be worked as the iteration, not recursively. But it works recursively due to the netdev notification unexpectedly. This problem occurs when it disables LRO only for the team and bonding interface type. team0 | +------+------+-----+-----+ | | | | team1 team2 team3 ... team200 If team0's LRO feature is updated, it generates the NETDEV_FEAT_CHANGE event to its own lower interfaces(team1 ~ team200). It is worked by netdev_sync_lower_features(). So, the NETDEV_FEAT_CHANGE notification logic of each lower interface work iteratively. But generated NETDEV_FEAT_CHANGE event is also sent to the upper interface too. upper interface(team0) generates the NETDEV_FEAT_CHANGE event for its own lower interfaces again. lower and upper interfaces receive this event and generate this event again and again. So, the stack overflow occurs. But it is not the infinite loop issue. Because the netdev_sync_lower_features() updates features before generating the NETDEV_FEAT_CHANGE event. Already synchronized lower interfaces skip notification logic. So, it is just the problem that iteration logic is changed to the recursive unexpectedly due to the notification mechanism. Reproducer: ip link add team0 type team ethtool -K team0 lro on for i in {1..200} do ip link add team$i master team0 type team ethtool -K team$i lro on done ethtool -K team0 lro off In order to fix it, the notifier_ctx member of bonding/team is introduced.	N/A	More Details
CVE-2023-54016	In the Linux kernel, the following vulnerability has been resolved: wifi: ath12k: Fix memory leak in rx_desc and tx_desc Currently when ath12k_dp_cc_desc_init() is called we allocate memory to rx_descs and tx_descs. In ath12k_dp_cc_cleanup(), during descriptor cleanup rx_descs and tx_descs memory is not freed. This is cause of memory leak. These allocated memory should be freed in ath12k_dp_cc_cleanup. In ath12k_dp_cc_desc_init(), we can save base address of rx_descs and tx_descs. In ath12k_dp_cc_cleanup(), we can free rx_descs and tx_descs memory using their base address. Tested-on: QCN9274 hw2.0 PCI WLAN.WBE.1.0.1-00029-QCAHKSWPL_SILICONZ-1	N/A	More Details
CVE-2023-54160	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_sdei: Fix sleep from invalid context BUG Running a preempt-rt (v6.2-rc3-rt1) based kernel on an Ampere Altra triggers: BUG: sleeping function called from invalid context at kernel/locking/spinlock_rt.c:46 in_atomic(): 0, irqs_disabled(): 128, non_block: 0, pid: 24, name: cpuhp/0 preempt_count: 0, expected: 0 RCU nest depth: 0, expected: 0 3 locks held by cpuhp/0/24: #0: ffffda30217c70d0 (cpu_hotplug_lock){++++}-{0:0}, at: cpuhp_thread_fun+0x5c/0x248 #1: ffffda30217c7120 (cpuhp_state-up){+.+.}-{0:0}, at: cpuhp_thread_fun+0x5c/0x248 #2: ffffda3021c711f0 (sdei_list_lock){....}-{3:3}, at: sdei_cpuhp_up+0x3c/0x130 irq event stamp: 36 hardirqs last enabled at (35): [<ffffda301e85b7bc>] finish_task_switch+0xb4/0x2b0 hardirqs last disabled at (36): [<ffffda301e812fec>] cpuhp_thread_fun+0x21c/0x248 softirqs last enabled at (0): [<ffffda301e80b184>] copy_process+0x63c/0x1ac0 softirqs last disabled at (0): [<0000000000000000>] 0x0 CPU: 0 PID: 24 Comm: cpuhp/0 Not tainted 5.19.0-rc3-rt5-[...] Hardware name: WIWYNN Mt.Jade Server [...] Call trace: dump_backtrace+0x114/0x120 show_stack+0x20/0x70 dump_stack_lvl+0x9c/0xd8 dump_stack+0x18/0x34 __might_resched+0x188/0x228 rt_spin_lock+0x70/0x120 sdei_cpuhp_up+0x3c/0x130 cpuhp_invoke_callback+0x250/0xf08 cpuhp_thread_fun+0x120/0x248 smpboot_thread_fn+0x280/0x320 kthread+0x130/0x140 ret_from_fork+0x10/0x20 sdei_cpuhp_up() is called in the STARTING hotplug section, which runs with interrupts disabled. Use a CPUHP_AP_ONLINE_DYN entry instead to execute the cpuhp cb later, with preemption enabled. SDEI originally got its own cpuhp slot to allow interacting with perf. It got superseded by pNMI and this early slot is not relevant anymore. [1] Some SDEI calls (e.g. SDEI_1_0_FN_SDEI_PE_MASK) take actions on the calling CPU. It is checked that preemption is disabled for them. _ONLINE cpuhp cb are executed in the 'per CPU hotplug thread'. Preemption is enabled in those threads, but their cpumask is limited to 1 CPU. Move 'WARN_ON_ONCE(preemptible())' statements so that SDEI cpuhp cb don't trigger them. Also add a check for the SDEI_1_0_FN_SDEI_PRIVATE_RESET SDEI call which acts on the calling CPU. [1]: https://lore.kernel.org/all/5813b8c5-ae3e-87fd-fccc-	N/A	More Details

	94c9cd08816d@arm.com/		
CVE-2023-54159	In the Linux kernel, the following vulnerability has been resolved: usb: mtu3: fix kernel panic at qmu transfer done irq handler When handle qmu transfer irq, it will unlock @mtu->lock before give back request, if another thread handle disconnect event at the same time, and try to disable ep, it may lock @mtu->lock and free qmu ring, then qmu irq hanlder may get a NULL gpd, avoid the KE by checking gpd's value before handling it. e.g. qmu done irq on cpu0 thread running on cpu1 qmu_done_tx() handle gpd [0] mtu3_requ_complete() mtu3_gadget_ep_disable() unlock @mtu->lock give back request lock @mtu->lock mtu3_ep_disable() mtu3_gpd_ring_free() unlock @mtu->lock lock @mtu->lock get next gpd [1] [1]: goto [0] to handle next gpd, and next gpd may be NULL.	N/A	More Details
CVE-2023-54015	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Devcom, fix error flow in mlx5_devcom_register_device In case devcom allocation is failed, mlx5 is always freeing the priv. However, this priv might have been allocated by a different thread, and freeing it might lead to use-after-free bugs. Fix it by freeing the priv only in case it was allocated by the running thread.	N/A	More Details
CVE-2023-54014	In the Linux kernel, the following vulnerability has been resolved: scsi: qla2xxx: Check valid rport returned by fc_bsg_to_rport() Klocwork reported warning of rport maybe NULL and will be dereferenced. rport returned by call to fc_bsg_to_rport() could be NULL and dereferenced. Check valid rport returned by fc_bsg_to_rport().	N/A	More Details
CVE-2023-54158	In the Linux kernel, the following vulnerability has been resolved: btrfs: don't free qgroup space unless specified Boris noticed in his simple quotas testing that he was getting a leak with Sweet Tea's change to subvol create that stopped doing a transaction commit. This was just a side effect of that change. In the delayed inode code we have an optimization that will free extra reservations if we think we can pack a dir item into an already modified leaf. Previously this wouldn't be triggered in the subvolume create case because we'd commit the transaction, it was still possible but much harder to trigger. It could actually be triggered if we did a mkdir && subvol create with qgroups enabled. This occurs because in btrfs_insert_delayed_dir_index(), which gets called when we're adding the dir item, we do the following: btrfs_block_rsv_release(fs_info, trans->block_rsv, bytes, NULL); if we're able to skip reserving space. The problem here is that trans->block_rsv points at the temporary block rsv for the subvolume create, which has qgroup reservations in the block rsv. This is a problem because btrfs_block_rsv_release() will do the following: if (block_rsv->qgroup_rsv_reserved >= block_rsv->qgroup_rsv_size) { qgroup_to_release = block_rsv->qgroup_rsv_reserved - block_rsv->qgroup_rsv_size; block_rsv->qgroup_rsv_reserved = block_rsv->qgroup_rsv_size; } The temporary block rsv just has ->qgroup_rsv_reserved set, ->qgroup_rsv_size == 0. The optimization in btrfs_insert_delayed_dir_index() sets ->qgroup_rsv_reserved = 0. Then later on when we call btrfs_subvolume_release_metadata() which has btrfs_block_rsv_release(fs_info, rsv, (u64)-1, &qgroup_to_release); btrfs_qgroup_convert_reserved_meta(root, qgroup_to_release); qgroup_to_release is set to 0, and we do not convert the reserved metadata space. The problem here is that the block rsv code has been unconditionally messing with ->qgroup_rsv_reserved, because the main place this is used is delalloc, and any time we call btrfs_block_rsv_release() we do it with qgroup_to_release set, and thus do the proper accounting. The subvolume code is the only other code that uses the qgroup reservation stuff, but it's intermingled with the above optimization, and thus was getting its reservation freed out from underneath it and thus leaking the reserved space. The solution is to simply not mess with the qgroup reservations if we don't have qgroup_to_release set. This works with the existing code as anything that messes with the delalloc reservations always have qgroup_to_release set. This fixes the leak that Boris was observing.	N/A	More Details
CVE-2023-54157	In the Linux kernel, the following vulnerability has been resolved: binder: fix UAF of alloc->vma in race with munmap() [cmllamas: clean forward port from commit 015ac18be7de ("binder: fix UAF of alloc->vma in race with munmap()") in 5.10 stable. It is needed in mainline after the revert of commit a43cfc87caaf ("android: binder: stop saving a pointer to the VMA") as pointed out by Liam. The commit log and tags have been tweaked to reflect this.] In commit 720c24192404 ("ANDROID: binder: change down_write to down_read") binder assumed the mmap read lock is sufficient to protect alloc->vma inside binder_update_page_range(). This used to be accurate until commit dd2283f2605e ("mm: mmap: zap pages with read mmap_sem in munmap"), which now downgrades the mmap_lock after detaching the vma from the rbtree in munmap(). Then it proceeds to teardown and free the vma with only the read lock held. This means that accesses to alloc->vma in binder_update_page_range() now will race with vm_area_free() in munmap() and can cause a UAF as shown in the following KASAN trace: == BUG: KASAN: use-after-free in vm_insert_page+0x7c/0x1f0 Read of size 8 at addr ffff16204ad00600 by task server/558 CPU: 3 PID: 558 Comm: server Not tainted 5.10.150-00001-gdc8dcf942daa #1 Hardware name: linux,dummy-virt (DT) Call trace: dump_backtrace+0x0/0x2a0 show_stack+0x18/0x2c dump_stack+0xf8/0x164 print_address_description.constprop.0+0x9c/0x538 kasan_report+0x120/0x200 __asan_load8+0xa0/0xc4 vm_insert_page+0x7c/0x1f0 binder_update_page_range+0x278/0x50c binder_alloc_new_buf+0x3f0/0xba0 binder_transaction+0x64c/0x3040 binder_thread_write+0x924/0x2020 binder_ioctl+0x1610/0x2e5c __arm64_sys_ioctl+0xd4/0x120 el0_svc_common.constprop.0+0xac/0x270 do_el0_svc+0x38/0xa0 el0_svc+0x1c/0x2c el0_sync_handler+0xe8/0x114 el0_sync+0x180/0x1c0 Allocated by task 559: kasan_save_stack+0x38/0x6c __kasan_kmalloc.constprop.0+0xe4/0xf0 kasan_slab_alloc+0x18/0x2c kmem_cache_alloc+0x1b0/0x2d0 vm_area_alloc+0x28/0x94 mmap_region+0x378/0x920 do_mmap+0x3f0/0x600 vm_mmap_pgoff+0x150/0x17c ksys_mmap_pgoff+0x284/0x2dc __arm64_sys_mmap+0x84/0xa4 el0_svc_common.constprop.0+0xac/0x270 do_el0_svc+0x38/0xa0 el0_svc+0x1c/0x2c el0_sync_handler+0xe8/0x114 el0_sync+0x180/0x1c0 Freed by task 560: kasan_save_stack+0x38/0x6c kasan_set_track+0x28/0x40 kasan_set_free_info+0x24/0x4c __kasan_slab_free+0x100/0x164 kasan_slab_free+0x14/0x20 kmem_cache_free+0xc4/0x34c vm_area_free+0x1c/0x2c remove_vma+0x7c/0x94 __do_munmap+0x358/0x710 __vm_munmap+0xbc/0x130 __arm64_sys_munmap+0x4c/0x64 el0_svc_common.constprop.0+0xac/0x270 do_el0_svc+0x38/0xa0 el0_svc+0x1c/0x2c el0_sync_handler+0xe8/0x114 el0_sync+0x180/0x1c0 [...] == To prevent the race above, revert back to taking the mmap write lock inside binder_update_page_range(). One might expect an increase of mmap lock contention. However, binder already serializes these calls via top level alloc->mutex. Also, there was no performance impact shown when running the binder benchmark tests.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: interconnect: Fix locking for runpm vs reclaim For cases where icc_bw_set() can be called in callbaths that could deadlock against shrinker/reclaim, such as runpm resume, we need to decouple the icc locking. Introduce a new icc_bw_lock for cases where we need to serialize bw aggregation and update to decouple that from paths that require memory allocation such as node/link creation/ destruction. Fixes this lockdep splat: == WARNING: possible circular locking dependency detected 6.2.0-rc8-debug+ #554 Not tainted -- ring0/132 is trying to acquire lock: ffffff80871916d0 (&gmu->lock){+.+.}-{3:3}, at: a6xx_pm_resume+0xf0/0x234 but task is already holding lock: ffffffdb5aee57e8 (dma_fence_map){++++}-{0:0}, at: msm_job_run+0x68/0x150 which lock already depends on the new lock. the existing dependency chain (in reverse order) is: -> #4 (dma_fence_map){++++}-{0:0}: __dma_fence_might_wait+0x74/0xc0 dma_resv_lockdep+0x1f4/0x2f4 do_one_initcall+0x104/0x2bc kernel_init_freeable+0x344/0x34c kernel_init+0x30/0x134 ret_from_fork+0x10/0x20 -> #3 (mmu_notifier_invalidate_range_start){+.+.}-{0:0}: fs_reclaim_acquire+0x80/0xa8 slab_pre_alloc_hook.constprop.0+0x40/0x25c __kmem_cache_alloc_node+0x60/0x1cc __kmalloc+0xd8/0x100		

CVE-2023-54013	topology_parse_cpu_capacity+0x8c/0x178 get_cpu_for_node+0x88/0xc4 parse_cluster+0x1b0/0x28c parse_cluster+0x8c/0x28c init_cpu_topology+0x168/0x188 smp_prepare_cpus+0x24/0xf8 kernel_init_freeable+0x18c/0x34c kernel_init+0x30/0x134 ret_from_fork+0x10/0x20 -> #2 (fs_reclaim){+.+.}-{0:0}: __fs_reclaim_acquire+0x3c/0x48 fs_reclaim_acquire+0x54/0xa8 slab_pre_alloc_hook.constprop.0+0x40/0x25c __kmem_cache_alloc_node+0x60/0x1cc __kmalloc+0xd8/0x100 kzalloc.constprop.0+0x14/0x20 icc_node_create_nolock+0x4c/0xc4 icc_node_create+0x38/0x58 qcom_icc_rpmh_probe+0x1b8/0x248 platform_probe+0x70/0xc4 really_probe+0x158/0x290 __driver_probe_device+0xc8/0xe0 driver_probe_device+0x44/0x100 __driver_attach+0xf8/0x108 bus_for_each_dev+0x78/0xc4 driver_attach+0x2c/0x38 bus_add_driver+0xd0/0x1d8 driver_register+0xbc/0xf8 __platform_driver_register+0x30/0x3c qnoc_driver_init+0x24/0x30 do_one_initcall+0x104/0x2bc kernel_init_freeable+0x344/0x34c kernel_init+0x30/0x134 ret_from_fork+0x10/0x20 -> #1 (icc_lock) {+.+.}-{3:3}: __mutex_lock+0xcc/0x3c8 mutex_lock_nested+0x30/0x44 icc_set_bw+0x88/0x2b4 _set_opp_bw+0x8c/0xd8 _set_opp+0x19c/0x300 dev_pm_opp_set_opp+0x84/0x94 a6xx_gmu_resume+0x18c/0x804 a6xx_pm_resume+0xf8/0x234 adreno_runtime_resume+0x2c/0x38 pm_generic_runtime_resume+0x30/0x44 __rpm_callback+0x15c/0x174 rpm_callback+0x78/0x7c rpm_resume+0x318/0x524 __pm_runtime_resume+0x78/0xbc adreno_load_gpu+0xc4/0x17c msm_open+0x50/0x120 drm_file_alloc+0x17c/0x228 drm_open_helper+0x74/0x118 drm_open+0xa0/0x144 drm_stub_open+0xd4/0xe4 chrdev_open+0x1b8/0x1e4 do_dentry_open+0x2f8/0x38c vfs_open+0x34/0x40 path_openat+0x64c/0x7b4 do_filp_open+0x54/0xc4 do_sys_openat2+0x9c/0x100 do_sys_open+0x50/0x7c __arm64_sys_openat+0x28/0x34 invoke_syscall+0x8c/0x128 el0_svc_common.constprop.0+0xa0/0x11c do_el0_ ---truncated---	N/A	More Details
CVE-2023-54011	In the Linux kernel, the following vulnerability has been resolved: scsi: mpi3mr: Fix an issue found by KASAN Write only correct size (32 instead of 64 bytes).	N/A	More Details
CVE-2023-54152	In the Linux kernel, the following vulnerability has been resolved: can: j1939: prevent deadlock by moving j1939_sk_errqueue() This commit addresses a deadlock situation that can occur in certain scenarios, such as when running data TP/ETP transfer and subscribing to the error queue while receiving a net down event. The deadlock involves locks in the following order: 3 j1939_session_list_lock -> active_session_list_lock j1939_session_activate ... j1939_sk_queue_activate_next -> sk_session_queue_lock ... j1939_xtp_rx_eoma_one 2 j1939_sk_queue_drop_all -> sk_session_queue_lock ... j1939_sk_netdev_event_netdown -> j1939_socks_lock j1939_netdev_notify 1 j1939_sk_errqueue -> j1939_socks_lock __j1939_session_cancel -> active_session_list_lock j1939_tp_rxtimer CPU0 CPU1 ---- ---- lock(&priv->active_session_list_lock); lock(&jsk->sk_session_queue_lock); lock(&priv->active_session_list_lock); lock(&priv->j1939_socks_lock); The solution implemented in this commit is to move the j1939_sk_errqueue() call out of the active_session_list_lock context, thus preventing the deadlock situation.	N/A	More Details
CVE-2023-54010	In the Linux kernel, the following vulnerability has been resolved: ACPICA: ACPICA: check null return of ACPI_ALLOCATE_ZEROED in acpi_db_display_objects ACPICA commit 0d5f467d6a0ba852ea3aad68663cbcbd43300fd4 ACPI_ALLOCATE_ZEROED may fails, object_info might be null and will cause null pointer dereference later.	N/A	More Details
CVE-2023-54009	In the Linux kernel, the following vulnerability has been resolved: i2c: cadence: cdns_i2c_master_xfer(): Fix runtime PM leak on error path The cdns_i2c_master_xfer() function gets a runtime PM reference when the function is entered. This reference is released when the function is exited. There is currently one error path where the function exits directly, which leads to a leak of the runtime PM reference. Make sure that this error path also releases the runtime PM reference.	N/A	More Details
CVE-2023-54008	In the Linux kernel, the following vulnerability has been resolved: virtio_vdpa: build affinity masks conditionally We try to build affinity mask via create_affinity_masks() unconditionally which may lead several issues: - the affinity mask is not used for parent without affinity support (only VDUSE support the affinity now) - the logic of create_affinity_masks() might not work for devices other than block. For example it's not rare in the networking device where the number of queues could exceed the number of CPUs. Such case breaks the current affinity logic which is based on group_cpus_evenly() who assumes the number of CPUs are not less than the number of groups. This can trigger a warning[1]: if (ret >= 0) WARN_ON(nr_present + nr_others < numgrps); Fixing this by only build the affinity masks only when - Driver passes affinity descriptor, driver like virtio-blk can make sure to limit the number of queues when it exceeds the number of CPUs - Parent support affinity setting config ops This help to avoid the warning. More optimizations could be done on top. [1] [682.146655] WARNING: CPU: 6 PID: 1550 at lib/group_cpus.c:400 group_cpus_evenly+0x1aa/0x1c0 [682.146668] CPU: 6 PID: 1550 Comm: vdpa Not tainted 6.5.0-rc5jason+ #79 [682.146671] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.2-0-gea1b7a073390-prebuilt.qemu.org 04/01/2014 [682.146673] RIP: 0010:group_cpus_evenly+0x1aa/0x1c0 [682.146676] Code: 4c 89 e0 5b 5d 41 5c 41 5d 41 5e c3 cc cc cc cc e8 1b c4 74 ff 48 89 ef e8 13 ac 98 ff 4c 89 e7 45 31 e4 e8 08 ac 98 ff eb c2 <0f> 0b eb b6 e8 fd 05 c3 00 45 31 e4 eb e5 cc cc cc cc cc cc cc cc cc [682.146679] RSP: 0018:ffffc9000215f498 EFLAGS: 00010293 [682.146682] RAX: 0000000000001f1e0 RBX: 0000000000000041 RCX: 0000000000000000 [682.146684] RDX: ffff888109922058 RSI: 0000000000000041 RDI: 0000000000000030 [682.146686] RBP: ffff888109922058 R08: ffffc9000215f498 R09: ffffc9000215f4a0 [682.146687] R10: 00000000000198d0 R11: 0000000000000030 R12: ffff888107e02800 [682.146689] R13: 0000000000000030 R14: 0000000000000030 R15: 0000000000000041 [682.146692] FS: 00007fef52315740(0000) GS:ffff888237380000(0000) knlGS:0000000000000000 [682.146695] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [682.146696] CR2: 00007fef52509000 CR3: 0000000110dbc004 CR4: 0000000000370ee0 [682.146698] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [682.146700] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [682.146701] Call Trace: [682.146703] <TASK> [682.146705] ? __warn+0x7b/0x130 [682.146709] ? group_cpus_evenly+0x1aa/0x1c0 [682.146712] ? report_bug+0x1c8/0x1e0 [682.146717] ? handle_bug+0x3c/0x70 [682.146721] ? exc_invalid_op+0x14/0x70 [682.146723] ? asm_exc_invalid_op+0x16/0x20 [682.146727] ? group_cpus_evenly+0x1aa/0x1c0 [682.146729] ? group_cpus_evenly+0x15c/0x1c0 [682.146731] create_affinity_masks+0xaf/0x1a0 [682.146735] virtio_vdpa_find_vqs+0x83/0x1d0 [682.146738] ? __pfx_default_calc_sets+0x10/0x10 [682.146742] virtnet_find_vqs+0x1f0/0x370 [682.146747] virtnet_probe+0x501/0xcd0 [682.146749] ? vp_modern_get_status+0x12/0x20 [682.146751] ? get_cap_addr.isra.0+0x10/0xc0 [682.146754] virtio_dev_probe+0x1af/0x260 [682.146759] really_probe+0x1a5/0x410	N/A	More Details
CVE-2023-54156	In the Linux kernel, the following vulnerability has been resolved: sfc: fix crash when reading stats while NIC is resetting efx_net_stats() (.ndo_get_stats64) can be called during an ethtool selftest, during which time nic_data->mc_stats is NULL as the NIC has been fini'd. In this case do not attempt to fetch the latest stats from the hardware, else we will crash on a NULL dereference: BUG: kernel NULL pointer dereference, address: 0000000000000038 RIP efx_nic_update_stats abridged calltrace: efx_ef10_update_stats_pf efx_net_stats dev_get_stats dev_seq_printf_stats Skipping the read is safe, we will simply give out stale stats. To ensure that the free in efx_ef10_fini_nic() does not race against efx_ef10_update_stats_pf(), which could cause a TOCTTOU bug, take the efx->stats_lock in fini_nic (it is already held across update_stats).	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: vmci_host: fix a race condition in vmci_host_poll() causing GPF During fuzzing, a general protection fault is observed in vmci_host_poll(). general protection fault, probably for non-canonical address 0xdffffc0000000019: 0000 [#1] PREEMPT SMP KASAN KASAN: null-ptr-deref in range [0x00000000000000c8-		

CVE-2023-54007	0x00000000000000cf] RIP: 0010:__lock_acquire+0xf3/0x5e00 kernel/locking/lockdep.c:4926 <- omitting registers -> Call Trace: <TASK> lock_acquire+0x1a4/0x4a0 kernel/locking/lockdep.c:5672 __raw_spin_lock_irqsave include/linux/spinlock_api_smp.h:110 [inline] _raw_spin_lock_irqsave+0xb3/0x100 kernel/locking/spinlock.c:162 add_wait_queue+0x3d/0x260 kernel/sched/wait.c:22 poll_wait include/linux/poll.h:49 [inline] vmci_host_poll+0xf8/0x2b0 drivers/misc/vmw_vmci/vmci_host.c:174 vfs_poll include/linux/poll.h:88 [inline] do_pollfd fs/select.c:873 [inline] do_poll fs/select.c:921 [inline] do_sys_poll+0xc7c/0x1aa0 fs/select.c:1015 __do_sys_ppoll fs/select.c:1121 [inline] __se_sys_ppoll+0x2cc/0x330 fs/select.c:1101 do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_64+0x4e/0xa0 arch/x86/entry/common.c:82 entry_SYSCALL_64_after_hwframe+0x46/0xb0 Example thread interleaving that causes the general protection fault is as follows: CPU1 (vmci_host_poll) CPU2 (vmci_host_do_init_context) ----- ----- // Read uninitialized context context = vmci_host_dev->context; // Initialize context vmci_host_dev->context = vmci_ctx_create(); vmci_host_dev->ct_type = VMCIOBJ_CONTEXT; if (vmci_host_dev->ct_type == VMCIOBJ_CONTEXT) { // Dereferencing the wrong pointer poll_wait(..., &context->host_context); } In this scenario, vmci_host_poll() reads vmci_host_dev->context first, and then reads vmci_host_dev->ct_type to check that vmci_host_dev->context is initialized. However, since these two reads are not atomically executed, there is a chance of a race condition as described above. To fix this race condition, read vmci_host_dev->context after checking the value of vmci_host_dev->ct_type so that vmci_host_poll() always reads an initialized context.	N/A	More Details
CVE-2023-54155	In the Linux kernel, the following vulnerability has been resolved: net: core: remove unnecessary frame_sz check in bpf_xdp_adjust_tail() Syzkaller reported the following issue: ========== Too BIG xdp->frame_sz = 131072 WARNING: CPU: 0 PID: 5020 at net/core/filter.c:4121 ____bpf_xdp_adjust_tail net/core/filter.c:4121 [inline] WARNING: CPU: 0 PID: 5020 at net/core/filter.c:4121 bpf_xdp_adjust_tail+0x466/0xa10 net/core/filter.c:4103 ... Call Trace: <TASK> bpf_prog_4add87e5301a4105+0x1a/0x1c __bpf_prog_run include/linux/filter.h:600 [inline] bpf_prog_run_xdp include/linux/filter.h:775 [inline] bpf_prog_run_generic_xdp+0x57e/0x11e0 net/core/dev.c:4721 netif_receive_generic_xdp net/core/dev.c:4807 [inline] do_xdp_generic+0x35c/0x770 net/core/dev.c:4866 tun_get_user+0x2340/0x3ca0 drivers/net/tun.c:1919 tun_chr_write_iter+0xe8/0x210 drivers/net/tun.c:2043 call_write_iter include/linux/fs.h:1871 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_write+0x650/0xe40 fs/read_write.c:584 ksys_write+0x12f/0x250 fs/read_write.c:637 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x38/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd xdp->frame_sz > PAGE_SIZE check was introduced in commit c8741e2bfe87 ("xdp: Allow bpf_xdp_adjust_tail() to grow packet size"). But Jesper Dangaard Brouer <jbrouer@redhat.com> noted that after introducing the xdp_init_buff() which all XDP driver use - it's safe to remove this check. The original intend was to catch cases where XDP drivers have not been updated to use xdp.frame_sz, but that is not longer a concern (since xdp_init_buff). Running the initial syzkaller repro it was discovered that the contiguous physical memory allocation is used for both xdp paths in tun_get_user(), e.g. tun_build_skb() and tun_alloc_skb(). It was also stated by Jesper Dangaard Brouer <jbrouer@redhat.com> that XDP can work on higher order pages, as long as this is contiguous physical memory (e.g. a page).	N/A	More Details
CVE-2023-54154	In the Linux kernel, the following vulnerability has been resolved: scsi: target: core: Fix target_cmd_counter leak The target_cmd_counter struct allocated via target_alloc_cmd_counter() is never freed, resulting in leaks across various transport types, e.g.: unreferenced object 0xffff88801f920120 (size 96): comm "sh", pid 102, jiffies 4294892535 (age 713.412s) hex dump (first 32 bytes): 07 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 38 01 92 1f 80 88 ff ff ....8..... backtrace: [<00000000e58a6252>] kmalloc_trace+0x11/0x20 [<0000000043af4b2f>] target_alloc_cmd_counter+0x17/0x90 [target_core_mod] [<000000007da2dfa7>] target_setup_session+0x2d/0x140 [target_core_mod] [<0000000068feef86>] tcm_loop_tpg_nexus_store+0x19b/0x350 [tcm_loop] [<000000006a80e021>] configfs_write_iter+0xb1/0x120 [<00000000e9f4d860>] vfs_write+0x2e4/0x3c0 [<000000008143433b>] ksys_write+0x80/0xb0 [<00000000a7df29b2>] do_syscall_64+0x42/0x90 [<0000000053f45fb8>] entry_SYSCALL_64_after_hwframe+0x6e/0xd8 Free the structure alongside the corresponding iscsit_conn / se_sess parent.	N/A	More Details
CVE-2023-54153	In the Linux kernel, the following vulnerability has been resolved: ext4: turn quotas off if mount failed after enabling quotas Yi found during a review of the patch "ext4: don't BUG on inconsistent journal feature" that when ext4_mark_recovery_complete() returns an error value, the error handling path does not turn off the enabled quotas, which triggers the following kmemleak: ========== unreferenced object 0xffff8cf68678e7c0 (size 64): comm "mount", pid 746, jiffies 4294871231 (age 11.540s) hex dump (first 32 bytes): 00 90 ef 82 f6 8c ff ff 00 00 00 00 41 01 00 00 ..... A... c7 00 00 00 bd 00 00 00 0a 00 00 00 48 00 00 00 ..... H... backtrace: [<00000000c561ef24>] __kmem_cache_alloc_node+0x4d4/0x880 [<00000000d4e621d7>] kmalloc_trace+0x39/0x140 [<00000000837eee74>] v2_read_file_info+0x18a/0x3a0 [<0000000088f6c877>] dquot_load_quota_sb+0x2ed/0x770 [<00000000340a4782>] dquot_load_quota_inode+0xc6/0x1c0 [<0000000089a18bd5>] ext4_enable_quotas+0x17e/0x3a0 [ext4] [<000000003a0268fa>] __ext4_fill_super+0x3448/0x3910 [ext4] [<00000000b0f2a8a8>] ext4_fill_super+0x13d/0x340 [ext4] [<000000004a9489c4>] get_tree_bdev+0x1dc/0x370 [<000000006e723bf1>] ext4_get_tree+0x1d/0x30 [ext4] [<00000000c7cb663d>] vfs_get_tree+0x31/0x160 [<00000000320e1bed>] do_new_mount+0x1d5/0x480 [<00000000c074654c>] path_mount+0x22e/0xbe0 [<0000000003e97a8e>] do_mount+0x95/0xc0 [<000000002f3d3736>] __x64_sys_mount+0xc4/0x160 [<0000000027d2140c>] do_syscall_64+0x3f/0x90 ========== To solve this problem, we add a "failed_mount10" tag, and call ext4_quota_off_umount() in this tag to release the enabled qoutas.	N/A	More Details
CVE-2023-54161	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2023-54017	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: fix possible memory leak in ibmebus_bus_init() If device_register() returns error in ibmebus_bus_init(), name of kobject which is allocated in dev_set_name() called in device_add() is leaked. As comment of device_add() says, it should call put_device() to drop the reference count that was set in device_initialize() when it fails, so the name can be freed in kobject_cleanup().	N/A	More Details
CVE-2023-54119	In the Linux kernel, the following vulnerability has been resolved: inotify: Avoid reporting event with invalid wd When inotify_freeing_mark() races with inotify_handle_inode_event() it can happen that inotify_handle_inode_event() sees that i_mark->wd got already reset to -1 and reports this value to userspace which can confuse the inotify listener. Avoid the problem by validating that wd is sensible (and pretend the mark got removed before the event got generated otherwise).	N/A	More Details
CVE-2023-54018	In the Linux kernel, the following vulnerability has been resolved: drm/msm/hdmi: Add missing check for alloc_ordered_workqueue Add check for the return value of alloc_ordered_workqueue as it may return NULL pointer and cause NULL pointer dereference in `hdmi_hdcp.c` and `hdmi_hpd.c`. Patchwork: https://patchwork.freedesktop.org/patch/517211/	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: iio: core: Prevent invalid memory access when there is no parent Commit 813665564b3d ("iio: core: Convert to use firmware node handle instead of OF node") switched the kind of nodes to use for label retrieval in device registration. Probably an unwanted change in that commit was that if the device has no parent then NULL		

CVE-2023-54027	pointer is accessed. This is what happens in the stock IIO dummy driver when a new entry is created in configfs: # mkdir /sys/kernel/config/iio/devices/dummy/foo BUG: kernel NULL pointer dereference, address: ... ... Call Trace: __iio_device_register iio_dummy_probe Since there seems to be no reason to make a parent device of an IIO dummy device mandatory, let’s prevent the invalid memory access in __iio_device_register when the parent device is NULL. With this change, the IIO dummy driver works fine with configfs.	N/A	More Details
CVE-2023-54026	In the Linux kernel, the following vulnerability has been resolved: opp: Fix use-after-free in lazy_opp_tables after probe deferral When dev_pm_opp_of_find_icc_paths() in _allocate_opp_table() returns -EPROBE_DEFER, the opp_table is freed again, to wait until all the interconnect paths are available. However, if the OPP table is using required-opps then it may already have been added to the global lazy_opp_tables list. The error path does not remove the opp_table from the list again. This can cause crashes later when the provider of the required-opps is added, since we will iterate over OPP tables that have already been freed. E.g.: Unable to handle kernel NULL pointer dereference when read CPU: 0 PID: 7 Comm: kworker/0:0 Not tainted 6.4.0-rc3 PC is at _of_add_opp_table_v2 (include/linux/of.h:949 drivers/opp/of.c:98 drivers/opp/of.c:344 drivers/opp/of.c:404 drivers/opp/of.c:1032) -> lazy_link_required_opp_table() Fix this by calling _of_clear_opp_table() to remove the opp_table from the list and clear other allocated resources. While at it, also add the missing mutex_destroy() calls in the error path.	N/A	More Details
CVE-2023-54025	In the Linux kernel, the following vulnerability has been resolved: wifi: rsi: Do not configure WoWlan in shutdown hook if not enabled In case WoWlan was never configured during the operation of the system, the hw->wiphy->wowlan_config will be NULL. rsi_config_wowlan() checks whether wowlan_config is non-NULL and if it is not, then WARNs about it. The warning is valid, as during normal operation the rsi_config_wowlan() should only ever be called with non-NULL wowlan_config. In shutdown this rsi_config_wowlan() should only ever be called if WoWlan was configured before by the user. Add checks for non-NULL wowlan_config into the shutdown hook. While at it, check whether the wiphy is also non-NULL before accessing wowlan_config . Drop the single-use wowlan_config variable, just inline it into function call.	N/A	More Details
CVE-2023-54024	In the Linux kernel, the following vulnerability has been resolved: KVM: Destroy target device if coalesced MMIO unregistration fails Destroy and free the target coalesced MMIO device if unregistering said device fails. As clearly noted in the code, kvm_io_bus_unregister_dev() does not destroy the target device. BUG: memory leak unreferenced object 0xffff888112a54880 (size 64): comm "syz-executor.2", pid 5258, jiffies 4297861402 (age 14.129s) hex dump (first 32 bytes): 38 c7 67 15 00 c9 ff ff 38 c7 67 15 00 c9 ff ff 8.g.....8.g..... e0 c7 e1 83 ff ff ff ff 00 30 67 15 00 c9 ff ff .....0g..... backtrace: [<0000000006995a8a>] kmalloc include/linux/slab.h:556 [inline] [<0000000006995a8a>] kzalloc include/linux/slab.h:690 [inline] [<0000000006995a8a>] kvm_vm_ioctl_register_coalesced_mmio+0x8e/0x3d0 arch/x86/kvm/../../../virt/kvm/coalesced_mmio.c:150 [<00000000022550c2>] kvm_vm_ioctl+0x47d/0x1600 arch/x86/kvm/../../../virt/kvm/kvm_main.c:3323 [<000000008a75102f>] vfs_ioctl fs/ioctl.c:46 [inline] [<000000008a75102f>] file_ioctl fs/ioctl.c:509 [inline] [<000000008a75102f>] do_vfs_ioctl+0xbab/0x1160 fs/ioctl.c:696 [<0000000080e3f669>] ksys_ioctl+0x76/0xa0 fs/ioctl.c:713 [<0000000059ef4888>] __do_sys_ioctl fs/ioctl.c:720 [inline] [<0000000059ef4888>] __se_sys_ioctl fs/ioctl.c:718 [inline] [<0000000059ef4888>] __x64_sys_ioctl+0x6f/0xb0 fs/ioctl.c:718 [<000000006444fa05>] do_syscall_64+0x9f/0x4e0 arch/x86/entry/common.c:290 [<000000009a4ed50b>] entry_SYSCALL_64_after_hwframe+0x49/0xbe BUG: leak checking failed	N/A	More Details
CVE-2023-54111	In the Linux kernel, the following vulnerability has been resolved: pinctrl: rockchip: Fix refcount leak in rockchip_pinctrl_parse_groups of_find_node_by_phandle() returns a node pointer with refcount incremented, We should use of_node_put() on it when not needed anymore. Add missing of_node_put() to avoid refcount leak.	N/A	More Details
CVE-2023-54112	In the Linux kernel, the following vulnerability has been resolved: kcm: Fix memory leak in error path of kcm_sendmsg() syzbot reported a memory leak like below: BUG: memory leak unreferenced object 0xffff88810b088c00 (size 240): comm "syz-executor186", pid 5012, jiffies 4294943306 (age 13.680s) hex dump (first 32 bytes): 00 89 08 0b 81 88 ff ff 00 00 00 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<ffffffff83e5d5ff>] __alloc_skb+0x1ef/0x230 net/core/skbuff.c:634 [<ffffffff84606e59>] alloc_skb include/linux/skbuff.h:1289 [inline] [<ffffffff84606e59>] kcm_sendmsg+0x269/0x1050 net/kcm/kcmsock.c:815 [<ffffffff83e479c6>] sock_sendmsg_nosec net/socket.c:725 [inline] [<ffffffff83e479c6>] sock_sendmsg+0x56/0xb0 net/socket.c:748 [<ffffffff83e47f55>] ____sys_sendmsg+0x365/0x470 net/socket.c:2494 [<ffffffff83e4c389>] ___sys_sendmsg+0xc9/0x130 net/socket.c:2548 [<ffffffff83e4c536>] __sys_sendmsg+0xa6/0x120 net/socket.c:2577 [<ffffffff84ad7bb8>] do_syscall_x64 arch/x86/entry/common.c:50 [inline] [<ffffffff84ad7bb8>] do_syscall_64+0x38/0xb0 arch/x86/entry/common.c:80 [<ffffffff84c0008b>] entry_SYSCALL_64_after_hwframe+0x63/0xcd In kcm_sendmsg(), kcm_tx_msg(head)->last_skb is used as a cursor to append newly allocated skbs to 'head'. If some bytes are copied, an error occurred, and jumped to out_error label, 'last_skb' is left unmodified. A later kcm_sendmsg() will use an obsoleted 'last_skb' reference, corrupting the 'head' frag_list and causing the leak. This patch fixes this issue by properly updating the last allocated skb in 'last_skb'.	N/A	More Details
CVE-2023-54113	In the Linux kernel, the following vulnerability has been resolved: rcu: dump vmalloc memory info safely Currently, for double invoke call_rcu(), will dump rcu_head objects memory info, if the objects is not allocated from the slab allocator, the vmalloc_dump_obj() will be invoke and the vmap_area_lock spinlock need to be held, since the call_rcu() can be invoked in interrupt context, therefore, there is a possibility of spinlock deadlock scenarios. And in Preempt-RT kernel, the rcutorture test also trigger the following lockdep warning: BUG: sleeping function called from invalid context at kernel/locking/spinlock_rt.c:48 in_atomic(): 1, irqs_disabled(): 1, non_block: 0, pid: 1, name: swapper/0 preempt_count: 1, expected: 0 RCU nest depth: 1, expected: 1 3 locks held by swapper/0/1: #0: ffffffffb534ee80 (fullstop_mutex){+.+.}-{4:4}, at: torture_init_begin+0x24/0xa0 #1: ffffffffb5307940 (rcu_read_lock){....}-{1:3}, at: rcu_torture_init+0x1ec7/0x2370 #2: ffffffffb536af40 (vmap_area_lock){+.+.}-{3:3}, at: find_vmap_area+0x1f/0x70 irq event stamp: 565512 hardirqs last enabled at (565511): [<ffffffffb379b138>] __call_rcu_common+0x218/0x940 hardirqs last disabled at (565512): [<ffffffffb5804262>] rcu_torture_init+0x20b2/0x2370 softirqs last enabled at (399112): [<ffffffffb36b2586>] __local_bh_enable_ip+0x126/0x170 softirqs last disabled at (399106): [<ffffffffb43fef59>] inet_register_protosw+0x9/0x1d0 Preemption disabled at: [<ffffffffb58040c3>] rcu_torture_init+0x1f13/0x2370 CPU: 0 PID: 1 Comm: swapper/0 Tainted: G W 6.5.0-rc4-rt2-yocto-preempt-rt+ #15 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.16.2-0-gea1b7a073390-prebuilt.qemu.org 04/01/2014 Call Trace: <TASK> dump_stack_lvl+0x68/0xb0 dump_stack+0x14/0x20 __might_resched+0x1aa/0x280 ? __pfx_rcu_torture_err_cb+0x10/0x10 rt_spin_lock+0x53/0x130 ? find_vmap_area+0x1f/0x70 find_vmap_area+0x1f/0x70 vmalloc_dump_obj+0x20/0x60 mem_dump_obj+0x22/0x90 __call_rcu_common+0x5bf/0x940 ? debug_smp_processor_id+0x1b/0x30 call_rcu_hurry+0x14/0x20 rcu_torture_init+0x1f82/0x2370 ? __pfx_rcu_torture_leak_cb+0x10/0x10 ? __pfx_rcu_torture_leak_cb+0x10/0x10 ? __pfx_rcu_torture_init+0x10/0x10 do_one_initcall+0x6c/0x300 ? debug_smp_processor_id+0x1b/0x30 kernel_init_freeable+0x2b9/0x540 ? __pfx_kernel_init+0x10/0x10 kernel_init+0x1f/0x150 ret_from_fork+0x40/0x50 ? __pfx_kernel_init+0x10/0x10 ret_from_fork_asm+0x1b/0x30 </TASK> The previous patch fixes this by using the deadlock-safe best-effort version of find_vm_area. However, in case of failure print the fact that the pointer was a vmalloc pointer so that we print at least something.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: net: nsh: Use correct mac_offset to unwind gso skb in nsh_gso_segment() As the call trace shows, skb_panic was caused by wrong skb->mac_header in nsh_gso_segment(): invalid		

CVE-2023-54114	opcode: 0000 [#1] PREEMPT SMP KASAN PTI CPU: 3 PID: 2737 Comm: syz Not tainted 6.3.0-next-20230505 #1 RIP: 0010:skb_panic+0xda/0xe0 call Trace: skb_push+0x91/0xa0 nsh_gso_segment+0x4f3/0x570 skb_mac_gso_segment+0x19e/0x270 __skb_gso_segment+0x1e8/0x3c0 validate_xmit_skb+0x452/0x890 validate_xmit_skb_list+0x99/0xd0 sch_direct_xmit+0x294/0x7c0 __dev_queue_xmit+0x16f0/0x1d70 packet_xmit+0x185/0x210 packet_snd+0xc15/0x1170 packet_sendmsg+0x7b/0xa0 sock_sendmsg+0x14f/0x160 The root cause is: nsh_gso_segment() use skb->network_header - nhoff to reset mac_header in skb_gso_error_unwind() if inner-layer protocol gso fails. However, skb->network_header may be reset by inner-layer protocol gso function e.g. mpls_gso_segment. skb->mac_header reset by the inaccurate network_header will be larger than skb headroom. nsh_gso_segment nhoff = skb->network_header - skb->mac_header; __skb_pull(skb,nsh_len) skb_mac_gso_segment mpls_gso_segment skb_reset_network_header(skb);//skb->network_header+=nsh_len return -EINVAL; skb_gso_error_unwind skb_push(skb, nsh_len); skb->mac_header = skb->network_header - nhoff; // skb->mac_header > skb->headroom, cause skb_push panic Use correct mac_offset to restore mac_header and get rid of nhoff.	N/A	More Details
CVE-2023-54115	In the Linux kernel, the following vulnerability has been resolved: pcmcia: rsrc_nonstatic: Fix memory leak in nonstatic_release_resource_db() When nonstatic_release_resource_db() frees all resources associated with an PCMCIA socket, it forgets to free socket_data too, causing a memory leak observable with kmemleak: unreferenced object 0xc28d1000 (size 64): comm "systemd-udevd", pid 297, jiffies 4294898478 (age 194.484s) hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 f0 85 0e c3 00 00 00 00 ..... 00 00 00 00 0c 10 8d c2 00 00 00 00 00 00 00 00 ..... backtrace: [<ffda4245>] __kmem_cache_alloc_node+0x2d7/0x4a0 [<7e51f0c8>] kmalloc_trace+0x31/0xa4 [<d52b4ca0>] nonstatic_init+0x24/0x1a4 [pcmcia_rsrc] [<a2f13e08>] pcmcia_register_socket+0x200/0x35c [pcmcia_core] [<a728be1b>] yenta_probe+0x4d8/0xa70 [yenta_socket] [<c48fac39>] pci_device_probe+0x99/0x194 [<84b7c690>] really_probe+0x181/0x45c [<8060fe6e>] __driver_probe_device+0x75/0x1f4 [<b9b76f43>] driver_probe_device+0x28/0xac [<648b766f>] __driver_attach+0xeb/0x1e4 [<6e9659eb>] bus_for_each_dev+0x61/0xb4 [<25a669f3>] driver_attach+0x1e/0x28 [<d8671d6b>] bus_add_driver+0x102/0x20c [<df0d323c>] driver_register+0x5b/0x120 [<942cd8a4>] __pci_register_driver+0x44/0x4c [<e536027e>] __UNIQUE_ID___addressable_cleanup_module188+0x1c/0xfffff000 [iTCO_vendor_support] Fix this by freeing socket_data too. Tested on a Acer Travelmate 4002WLMi by manually binding/unbinding the yenta_cardbus driver (yenta_socket).	N/A	More Details
CVE-2023-54116	In the Linux kernel, the following vulnerability has been resolved: drm/fbdev-generic: prohibit potential out-of-bounds access The fbdev test of IGT may write after EOF, which lead to out-of-bound access for drm drivers with fbdev-generic. For example, run fbdev test on a x86+ast2400 platform, with 1680x1050 resolution, will cause the linux kernel hang with the following call trace: Oops: 0000 [#1] PREEMPT SMP PTI [IGT] fbdev: starting subtest eof Workqueue: events drm_fb_helper_damage_work [drm_kms_helper] [IGT] fbdev: starting subtest nullptr RIP: 0010:memcpy_erms+0xa/0x20 RSP: 0018:ffffa17d40167d98 EFLAGS: 00010246 RAX: ffffa17d4eb7fa80 RBX: ffffa17d40e0aa80 RCX: 00000000000014c0 RDX: 0000000000001a40 RSI: ffffa17d40e0b000 RDI: ffffa17d4eb80000 RBP: ffffa17d40167e20 R08: 0000000000000000 R09: ffff89522ecff8c0 R10: ffffa17d4e4c5000 R11: 0000000000000000 R12: ffffa17d4eb7fa80 R13: 0000000000001a40 R14: 000000000000041a R15: ffffa17d40167e30 FS: 0000000000000000(0000) GS:ffff895257380000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: ffffa17d40e0b000 CR3: 00000001eaeca006 CR4: 00000000001706e0 Call Trace: <TASK> ? drm_fbdev_generic_helper_fb_dirty+0x207/0x330 [drm_kms_helper] drm_fb_helper_damage_work+0x8f/0x170 [drm_kms_helper] process_one_work+0x21f/0x430 worker_thread+0x4e/0x3c0 ? __pfx_worker_thread+0x10/0x10 kthread+0xf4/0x120 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x2c/0x50 </TASK> CR2: ffffa17d40e0b000 ---[end trace 0000000000000000]--- The is because damage rectangles computed by drm_fb_helper_memory_range_to_clip() function is not guaranteed to be bound in the screen's active display area. Possible reasons are: 1) Buffers are allocated in the granularity of page size, for mmap system call support. The shadow screen buffer consumed by fbdev emulation may also choosed be page size aligned. 2) The DIV_ROUND_UP() used in drm_fb_helper_memory_range_to_clip() will introduce off-by-one error. For example, on a 16KB page size system, in order to store a 1920x1080 XRGB framebuffer, we need allocate 507 pages. Unfortunately, the size 1920*1080*4 can not be divided exactly by 16KB. 1920 * 1080 * 4 = 8294400 bytes 506 * 16 * 1024 = 8290304 bytes 507 * 16 * 1024 = 8306688 bytes line_length = 1920*4 = 7680 bytes 507 * 16 * 1024 / 7680 = 1081.6 off / line_length = 507 * 16 * 1024 / 7680 = 1081 DIV_ROUND_UP(507 * 16 * 1024, 7680) will yeild 1082 memcpy_toio() typically issue the copy line by line, when copy the last line, out-of-bound access will be happen. Because: 1082 * line_length = 1082 * 7680 = 8309760, and 8309760 > 8306688 Note that userspace may still write to the invisiable area if a larger buffer than width x stride is exposed. But it is not a big issue as long as there still have memory resolve the access if not drafting so far. - Also limit the y1 (Daniel) - keep fix patch it to minimal (Daniel) - screen_size is page size aligned because of it need mmap (Thomas) - Adding fixes tag (Thomas)	N/A	More Details
CVE-2023-54117	In the Linux kernel, the following vulnerability has been resolved: s390/dcssblk: fix kernel crash with list_add corruption Commit fb08a1908cb1 ("dax: simplify the dax_device <-> gendisk association") introduced new logic for gendisk association, requiring drivers to explicitly call dax_add_host() and dax_remove_host(). For dcssblk driver, some dax_remove_host() calls were missing, e.g. in device remove path. The commit also broke error handling for out_dax case in device add path, resulting in an extra put_device() w/o the previous get_device() in that case. This lead to stale xarray entries after device add / remove cycles. In the case when a previously used struct gendisk pointer (xarray index) would be used again, because blk_alloc_disk() happened to return such a pointer, the xa_insert() in dax_add_host() would fail and go to out_dax, doing the extra put_device() in the error path. In combination with an already flawed error handling in dcssblk (device_register() cleanup), which needs to be addressed in a separate patch, this resulted in a missing device_del() / klist_del(), and eventually in the kernel crash with list_add corruption on a subsequent device_add() / klist_add(). Fix this by adding the missing dax_remove_host() calls, and also move the put_device() in the error path to restore the previous logic.	N/A	More Details
CVE-2023-54023	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix race between balance and cancel/pause Syzbot reported a panic that looks like this: assertion failed: fs_info->exclusive_operation == BTRFS_EXCLOP_BALANCE_PAUSED, in fs/btrfs/ioctl.c:465 ------------[cut here]------------ kernel BUG at fs/btrfs/messages.c:259! RIP: 0010:btrfs_assertfail+0x2c/0x30 fs/btrfs/messages.c:259 Call Trace: <TASK> btrfs_exclop_balance fs/btrfs/ioctl.c:465 [inline] btrfs_ioctl_balance fs/btrfs/ioctl.c:3564 [inline] btrfs_ioctl+0x531e/0x5b30 fs/btrfs/ioctl.c:4632 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:870 [inline] __se_sys_ioctl fs/ioctl.c:856 [inline] __x64_sys_ioctl+0x197/0x210 fs/ioctl.c:856 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x39/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd The reproducer is running a balance and a cancel or pause in parallel. The way balance finishes is a bit wonky, if we were paused we need to save the balance_ctl in the fs_info, but clear it otherwise and cleanup. However we rely on the return values being specific errors, or having a cancel request or no pause request. If balance completes and returns 0, but we have a pause or cancel request we won't do the appropriate cleanup, and then the next time we try to start a balance we'll trip this ASSERT. The error handling is just wrong here, we always want to clean up, unless we got -ECANCELLED and we set the appropriate pause flag in the exclusive op. With this patch the reproducer ran for an hour without tripping, previously it would trip in less than a few minutes.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: ALSA: usb-audio: Fix potential memory leaks at error path for UMP open The allocation and initialization errors at alloc_midi_urbs() that is called at MIDI 2.0 / UMP device are supposed to be handled at the caller side by invoking free_midi_urbs(). However, free_midi_urbs() loops only for ep->num_urbs entries, and since ep->num_entries wasn't updated yet at the allocation / init error in alloc_midi_urbs(), this entry won't be released. The intention of	N/A	More Details

54022	free_midi_urbs() is to release the whole elements, so change the loop size to NUM_URBS to scan over all elements for fixing the missed releases. Also, the call of free_midi_urbs() is missing at snd_usb_midi_v2_open(). Although it'll be released later at reopen/close or disconnection, it's better to release immediately at the error path.		
CVE-2023-54021	In the Linux kernel, the following vulnerability has been resolved: ext4: set goal start correctly in ext4_mb_normalize_request We need to set ac_g_ex to notify the goal start used in ext4_mb_find_by_goal. Set ac_g_ex instead of ac_f_ex in ext4_mb_normalize_request. Besides we should assure goal start is in range [first_data_block, blocks_count) as ext4_mb_initialize_context does. [Added a check to make sure size is less than ar->pright; otherwise we could end up passing an underflowed value of ar->pright - size to ext4_get_group_no_and_offset(), which will trigger a BUG_ON later on. - TYT]	N/A	More Details
CVE-2023-54020	In the Linux kernel, the following vulnerability has been resolved: dmaengine: sf-pdma: pdma_desc memory leak fix Commit b2cc5c465c2c ("dmaengine: sf-pdma: Add multithread support for a DMA channel") changed sf_pdma_prep_dma_memcpy() to unconditionally allocate a new sf_pdma_desc each time it is called. The driver previously recycled descs, by checking the in_use flag, only allocating additional descs if the existing one was in use. This logic was removed in commit b2cc5c465c2c ("dmaengine: sf-pdma: Add multithread support for a DMA channel"), but sf_pdma_free_desc() was not changed to handle the new behaviour. As a result, each time sf_pdma_prep_dma_memcpy() is called, the previous descriptor is leaked, over time leading to memory starvation: unreferenced object 0xffffffe008447300 (size 192): comm "irq/39-mchp_dsc", pid 343, jiffies 4294906910 (age 981.200s) hex dump (first 32 bytes): 00 00 00 ff 00 00 00 00 b8 c1 00 00 00 00 00 00 ..... 00 00 70 08 10 00 00 00 00 00 c0 00 00 00 00 00 ..p..... backtrace: [<00000000064a04f4>] kmemleak_alloc+0x1e/0x28 [<00000000018927a7>] kmem_cache_alloc+0x11e/0x178 [<000000002aea8d16>] sf_pdma_prep_dma_memcpy+0x40/0x112 Add the missing kfree() to sf_pdma_free_desc(), and remove the redundant in_use flag.	N/A	More Details
CVE-2023-54118	In the Linux kernel, the following vulnerability has been resolved: serial: sc16is7xx: setup GPIO controller later in probe The GPIO controller component of the sc16is7xx driver is setup too early, which can result in a race condition where another device tries to utilise the GPIO lines before the sc16is7xx device has finished initialising. This issue manifests itself as an Oops when the GPIO lines are configured: Unable to handle kernel read from unreadable memory at virtual address ... pc : sc16is7xx_gpio_direction_output+0x68/0x108 [sc16is7xx] lr : sc16is7xx_gpio_direction_output+0x4c/0x108 [sc16is7xx] ... Call trace: sc16is7xx_gpio_direction_output+0x68/0x108 [sc16is7xx] gpiod_direction_output_raw_commit+0x64/0x318 gpiod_direction_output+0xb0/0x170 create_gpio_led+0xec/0x198 gpio_led_probe+0x16c/0x4f0 platform_drv_probe+0x5c/0xb0 really_probe+0xe8/0x448 driver_probe_device+0xe8/0x138 __device_attach_driver+0x94/0x118 bus_for_each_drv+0x8c/0xe0 __device_attach+0x100/0x1b8 device_initial_probe+0x28/0x38 bus_probe_device+0xa4/0xb0 deferred_probe_work_func+0x90/0xe0 process_one_work+0x1c4/0x480 worker_thread+0x54/0x430 kthread+0x138/0x150 ret_from_fork+0x10/0x1c This patch moves the setup of the GPIO controller functions to later in the probe function, ensuring the sc16is7xx device has already finished initialising by the time other devices try to make use of the GPIO lines. The error handling has also been reordered to reflect the new initialisation order.	N/A	More Details
CVE-2023-54019	In the Linux kernel, the following vulnerability has been resolved: sched/psi: use kernfs polling functions for PSI trigger polling Destroying psi trigger in cgroup_file_release causes UAF issues when a cgroup is removed from under a polling process. This is happening because cgroup removal causes a call to cgroup_file_release while the actual file is still alive. Destroying the trigger at this point would also destroy its waitqueue head and if there is still a polling process on that file accessing the waitqueue, it will step on the freed pointer: do_select vfs_poll do_rmdir cgroup_rmdir kernfs_drain_open_files cgroup_file_release cgroup_pressure_release psi_trigger_destroy wake_up_pollfree(&t->event_wait) // vfs_poll is unblocked synchronize_rcu kfree(t) poll_freewait -> UAF access to the trigger's waitqueue head Patch [1] fixed this issue for epoll() case using wake_up_pollfree(), however the same issue exists for synchronous poll() case. The root cause of this issue is that the lifecycles of the psi trigger's waitqueue and of the file associated with the trigger are different. Fix this by using kernfs_generic_poll function when polling on cgroup-specific psi triggers. It internally uses kernfs_open_node->poll waitqueue head with its lifecycle tied to the file's lifecycle. This also renders the fix in [1] obsolete, so revert it. [1] commit c2dbe32d5db5 ("sched/psi: Fix use-after-free in ep_remove_wait_queue()")	N/A	More Details
CVE-2023-54006	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix data-race around unix_tot_inflight. unix_tot_inflight is changed under spin_lock(unix_gc_lock), but unix_release_sock() reads it locklessly. Let's use READ_ONCE() for unix_tot_inflight. Note that the writer side was marked by commit 9d6d7f1cb67c ("af_unix: annote lockless accesses to unix_tot_inflight & gc_in_progress") BUG: KCSAN: data-race in unix_inflight / unix_release_sock write (marked) to 0xffffffff871852b8 of 4 bytes by task 123 on cpu 1: unix_inflight+0x130/0x180 net/unix/scm.c:64 unix_attach_fds+0x137/0x1b0 net/unix/scm.c:123 unix_scm_to_skb net/unix/af_unix.c:1832 [inline] unix_dgram_sendmsg+0x46a/0x14f0 net/unix/af_unix.c:1955 sock_sendmsg_nosec net/socket.c:724 [inline] sock_sendmsg+0x148/0x160 net/socket.c:747 ____sys_sendmsg+0x4e4/0x610 net/socket.c:2493 ___sys_sendmsg+0xc6/0x140 net/socket.c:2547 __sys_sendmsg+0x94/0x140 net/socket.c:2576 __do_sys_sendmsg net/socket.c:2585 [inline] __se_sys_sendmsg net/socket.c:2583 [inline] __x64_sys_sendmsg+0x45/0x50 net/socket.c:2583 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3b/0x90 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x72/0xdc read to 0xffffffff871852b8 of 4 bytes by task 4891 on cpu 0: unix_release_sock+0x608/0x910 net/unix/af_unix.c:671 unix_release+0x59/0x80 net/unix/af_unix.c:1058 __sock_release+0x7d/0x170 net/socket.c:653 sock_close+0x19/0x30 net/socket.c:1385 __fput+0x179/0x5e0 fs/file_table.c:321 ____fput+0x15/0x20 fs/file_table.c:349 task_work_run+0x116/0x1a0 kernel/task_work.c:179 resume_user_mode_work include/linux/resume_user_mode.h:49 [inline] exit_to_user_mode_loop kernel/entry/common.c:171 [inline] exit_to_user_mode_prepare+0x174/0x180 kernel/entry/common.c:204 __syscall_exit_to_user_mode_work kernel/entry/common.c:286 [inline] syscall_exit_to_user_mode+0x1a/0x30 kernel/entry/common.c:297 do_syscall_64+0x4b/0x90 arch/x86/entry/common.c:86 entry_SYSCALL_64_after_hwframe+0x72/0xdc value changed: 0x00000000 -> 0x00000001 Reported by Kernel Concurrency Sanitizer on: CPU: 0 PID: 4891 Comm: systemd-coredum Not tainted 6.4.0-rc5-01219-gfa0e21fa4443 #5 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014	N/A	More Details
CVE-2023-54151	In the Linux kernel, the following vulnerability has been resolved: f2fs: Fix system crash due to lack of free space in LFS When f2fs tries to checkpoint during foreground gc in LFS mode, system crash occurs due to lack of free space if the amount of dirty node and dentry pages generated by data migration exceeds free space. The reproduction sequence is as follows. - 20GiB capacity block device (null_blk) - format and mount with LFS mode - create a file and write 20,000MiB - 4k random write on full range of the file RIP: 0010:new_curseg+0x48a/0x510 [f2fs] Code: 55 e7 f5 89 c0 48 0f af c3 48 8b 5d c0 48 c1 e8 20 83 c0 01 89 43 6c 48 83 c4 28 5b 41 5c 41 5d 41 5e 41 5f 5d c3 cc cc cc cc <0f> 0b f0 41 80 4f 48 04 45 85 f6 0f 84 ba fd ff ff e9 ef fe ff ff RSP: 0018:ffff977bc397b218 EFLAGS: 00010246 RAX: 0000000000027b9 RBX: 0000000000000000 RCX: 00000000000027c0 RDX: 0000000000000000 RSI: 00000000000027b9 RDI: ffff8c25ab4e74f8 RBP: ffff977bc397b268 R08: 0000000000027b9 R09: ffff8c29e4a34b40 R10: 0000000000000001 R11: ffff977bc397b0d8 R12: 0000000000000000 R13: ffff8c25b4dd81a0 R14: 0000000000000000 R15: ffff8c2f667f9000 FS: 0000000000000000(0000) GS:ffff8c344ec80000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 000000c00055d000 CR3: 0000000e30810003 CR4: 00000000003706e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> allocate_segment_by_default+0x9c/0x110 [f2fs] f2fs_allocate_data_block+0x243/0xa30 [f2fs] ?	N/A	More Details

	__mod_lruvec_page_state+0xa0/0x150 do_write_page+0x80/0x160 [f2fs] f2fs_do_write_node_page+0x32/0x50 [f2fs] __write_node_page+0x339/0x730 [f2fs] f2fs_sync_node_pages+0x5a6/0x780 [f2fs] block_operations+0x257/0x340 [f2fs] f2fs_write_checkpoint+0x102/0x1050 [f2fs] f2fs_gc+0x27c/0x630 [f2fs] ? folio_mark_dirty+0x36/0x70 f2fs_balance_fs+0x16f/0x180 [f2fs] This patch adds checking whether free sections are enough before checkpoint during gc. [Jaegeuk Kim: code clean-up]		
CVE-2025-69234	Whale browser before 4.35.351.12 allows an attacker to escape the iframe sandbox in a sidebar environment.	N/A	More Details
CVE-2025-69091	Missing Authorization vulnerability in Kraft Plugins Demo Importer Plus demo-importer-plus allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Demo Importer Plus: from n/a through <= 2.0.8.	N/A	More Details
CVE-2025-69029	Authorization Bypass Through User-Controlled Key vulnerability in Select-Themes Struktur struktur allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Struktur: from n/a through <= 2.5.1.	N/A	More Details
CVE-2025-69030	Authorization Bypass Through User-Controlled Key vulnerability in Mikado-Themes Backpack Traveler backpacktraveler allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Backpack Traveler: from n/a through <= 2.10.3.	N/A	More Details
CVE-2025-69031	Missing Authorization vulnerability in Skywarrior Arcane arcane allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Arcane: from n/a through <= 3.6.6.	N/A	More Details
CVE-2025-69032	Authorization Bypass Through User-Controlled Key vulnerability in Mikado-Themes FiveStar fivestar allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects FiveStar: from n/a through <= 1.7.	N/A	More Details
CVE-2025-69033	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in A WP Life Blog Filter blog-filter allows DOM-Based XSS.This issue affects Blog Filter: from n/a through <= 1.7.3.	N/A	More Details
CVE-2025-69034	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion') vulnerability in Mikado-Themes Lekker lekker allows PHP Local File Inclusion.This issue affects Lekker: from n/a through <= 1.8.	N/A	More Details
CVE-2025-69088	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Vidish Combo Offers WooCommerce woo-combo-offers allows DOM-Based XSS.This issue affects Combo Offers WooCommerce: from n/a through <= 4.2.	N/A	More Details
CVE-2025-69089	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in autolistings Auto Listings auto-listings allows Stored XSS.This issue affects Auto Listings: from n/a through <= 2.7.1.	N/A	More Details
CVE-2025-69092	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in WPDeveloper Essential Addons for Elementor essential-addons-for-elementor-lite allows DOM-Based XSS.This issue affects Essential Addons for Elementor: from n/a through <= 6.5.3.	N/A	More Details
CVE-2023-54150	In the Linux kernel, the following vulnerability has been resolved: drm/amd: Fix an out of bounds error in BIOS parser The array is hardcoded to 8 in atomfirmware.h, but firmware provides a bigger one sometimes. Deferencing the larger array causes an out of bounds error. commit 4fc1ba4aa589 ("drm/amd/display: fix array index out of bound error in bios parser") fixed some of this, but there are two other cases not covered by it. Fix those as well.	N/A	More Details
CVE-2025-69093	Missing Authorization vulnerability in wpdesk ShopMagic shopmagic-for-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects ShopMagic: from n/a through <= 4.7.2.	N/A	More Details
CVE-2022-50784	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mei: fix potential NULL-ptr deref after clone If cloning the SKB fails, don't try to use it, but rather return as if we should pass it. Coverity CID: 1503456	N/A	More Details
CVE-2022-50785	In the Linux kernel, the following vulnerability has been resolved: fsi: occ: Prevent use after free Use get_device and put_device in the open and close functions to make sure the device doesn't get freed while a file descriptor is open. Also, lock around the freeing of the device buffer and check the buffer before using it in the submit function.	N/A	More Details
CVE-2022-50786	In the Linux kernel, the following vulnerability has been resolved: media: s5p-mfc: Clear workbit to handle error condition During error on CLOSE_INSTANCE command, ctx_work_bits was not getting cleared. During consequent mfc execution NULL pointer dereferencing of this context led to kernel panic. This patch fixes this issue by making sure to clear ctx_work_bits always.	N/A	More Details
CVE-2022-	In the Linux kernel, the following vulnerability has been resolved: ext4: fix bug_on in __es_tree_search caused by bad quota inode We got a issue as fllows: == kernel BUG at fs/ext4/extents_status.c:202! invalid opcode: 0000 [#1] PREEMPT SMP CPU: 1 PID: 810 Comm: mount Not tainted 6.1.0-rc1-next-g9631525255e3 #352 RIP: 0010:__es_tree_search.isra.0+0xb8/0xe0 RSP: 0018:ffffc90001227900 EFLAGS: 00010202 RAX: 0000000000000000 RBX: 0000000077512a0f RCX: 0000000000000000 RDX: 0000000000000002 RSI: 0000000000002a10 RDI: ffff8881004cd0c8 RBP: ffff888177512ac8 R08: 47ffffffffffffff R09: 0000000000000001 R10: 0000000000000001 R11: 00000000000679af R12: 0000000000002a10 R13: ffff888177512d88 R14: 0000000077512a10 R15: 0000000000000000 FS: 00007f4bd76dbc40(0000)GS:ffff88842fd00000(0000)knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00005653bf993cf8 CR3: 000000017bfdf000 CR4: 00000000000006e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> ext4_es_cache_extent+0xe2/0x210 ext4_cache_extents+0xd2/0x110 ext4_find_extent+0x5d5/0x8c0 ext4_ext_map_blocks+0x9c/0x1d30 ext4_map_blocks+0x431/0xa50 ext4_getblk+0x82/0x340 ext4_bread+0x14/0x110 ext4_quota_read+0xf0/0x180 v2_read_header+0x24/0x90 v2_check_quota_file+0x2f/0xa0 dquot_load_quota_sb+0x26c/0x760 dquot_load_quota_inode+0xa5/0x190 ext4_enable_quotas+0x14c/0x300 __ext4_fill_super+0x31cc/0x32c0	N/A	More Details

50782	ext4_fill_super+0x115/0x2d0 get_tree_bdev+0x1d2/0x360 ext4_get_tree+0x19/0x30 vfs_get_tree+0x26/0xe0 path_mount+0x81d/0xfc0 do_mount+0x8d/0xc0 __x64_sys_mount+0xc0/0x160 do_syscall_64+0x35/0x80 entry_SYSCALL_64_after_hwframe+0x63/0xcd </TASK> ========== Above issue may happen as follows: ---------- ext4_fill_super ext4_orphan_cleanup ext4_enable_quotas ext4_quota_enable ext4_iget --> get error inode <5> ext4_ext_check_inode --> Wrong imode makes it escape inspection make_bad_inode(inode) --> EXT4_BOOT_LOADER_INO set imode dquot_load_quota_inode vfs_setup_quota_inode --> check pass dquot_load_quota_sb v2_check_quota_file v2_read_header ext4_quota_read ext4_bread ext4_getblk ext4_map_blocks ext4_ext_map_blocks ext4_find_extent ext4_cache_extents ext4_es_cache_extent __es_tree_search.isra.0 ext4_es_end --> Wrong extents trigger BUG_ON In the above issue, s_usr_quota_inum is set to 5, but inode<5> contains incorrect imode and disordered extents. Because 5 is EXT4_BOOT_LOADER_INO, the ext4_ext_check_inode check in the ext4_iget function can be bypassed, finally, the extents that are not checked trigger the BUG_ON in the __es_tree_search function. To solve this issue, check whether the inode is bad_inode in vfs_setup_quota_inode().		
CVE-2023-54001	In the Linux kernel, the following vulnerability has been resolved: staging: r8712: Fix memory leak in _r8712_init_xmit_priv() In the above mentioned routine, memory is allocated in several places. If the first succeeds and a later one fails, the routine will leak memory. This patch fixes commit 2865d42c78a9 ("staging: r8712u: Add the new driver to the mainline kernel"). A potential memory leak in r8712_xmit_resource_alloc() is also addressed.	N/A	More Details
CVE-2023-54000	In the Linux kernel, the following vulnerability has been resolved: net: hns3: fix deadlock issue when externel_lb and reset are executed together When externel_lb and reset are executed together, a deadlock may occur: [3147.217009] INFO: task kworker/u321:0:7 blocked for more than 120 seconds. [3147.230483] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message. [3147.238999] task:kworker/u321:0 state:D stack: 0 pid: 7 ppid: 2 flags:0x00000008 [3147.248045] Workqueue: hclge hclge_service_task [hclge] [3147.253957] Call trace: [3147.257093] __switch_to+0x7c/0xbc [3147.261183] __schedule+0x338/0x6f0 [3147.265357] schedule+0x50/0xe0 [3147.269185] schedule_preempt_disabled+0x18/0x24 [3147.274488] __mutex_lock.constprop.0+0x1d4/0x5dc [3147.279880] __mutex_lock_slowpath+0x1c/0x30 [3147.284839] mutex_lock+0x50/0x60 [3147.288841] rtnl_lock+0x20/0x2c [3147.292759] hclge_reset_prepare+0x68/0x90 [hclge] [3147.298239] hclge_reset_subtask+0x88/0xe0 [hclge] [3147.303718] hclge_reset_service_task+0x84/0x120 [hclge] [3147.309718] hclge_service_task+0x2c/0x70 [hclge] [3147.315109] process_one_work+0x1d0/0x490 [3147.319805] worker_thread+0x158/0x3d0 [3147.324240] kthread+0x108/0x13c [3147.328154] ret_from_fork+0x10/0x18 In externel_lb process, the hns3 driver call napi_disable() first, then the reset happen, then the restore process of the externel_lb will fail, and will not call napi_enable(). When doing externel_lb again, napi_disable() will be double call, cause a deadlock of rtnl_lock(). This patch use the HNS3_NIC_STATE_DOWN state to protect the calling of napi_disable() and napi_enable() in externel_lb process, just as the usage in ndo_stop() and ndo_start().	N/A	More Details
CVE-2023-53999	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: TC, Fix internal port memory leak The flow rule can be splited, and the extra post_act rules are added to post_act table. It's possible to trigger memleak when the rule forwards packets from internal port and over tunnel, in the case that, for example, CT 'new' state offload is allowed. As int_port object is assigned to the flow attribute of post_act rule, and its refcnt is incremented by mlx5e_tc_int_port_get(), but mlx5e_tc_int_port_put() is not called, the refcnt is never decremented, then int_port is never freed. The kmemleak reports the following error: unreferenced object 0xffff888128204b80 (size 64): comm "handler20", pid 50121, jiffies 4296973009 (age 642.932s) hex dump (first 32 bytes): 01 00 00 00 19 00 00 00 03 f0 00 00 04 00 00 00 ..... 98 77 67 41 81 88 ff ff 98 77 67 41 81 88 ff ff .wgA.....wgA.... backtrace: [<00000000e992680d>] kmalloc_trace+0x27/0x120 [<000000009e945a98>] mlx5e_tc_int_port_get+0x3f3/0xe20 [mlx5_core] [<0000000035a537f0>] mlx5e_tc_add_fdb_flow+0x473/0xcf0 [mlx5_core] [<0000000070c2cec6>] __mlx5e_add_fdb_flow+0x7cf/0xe90 [mlx5_core] [<000000005cc84048>] mlx5e_configure_flower+0xd40/0x4c40 [mlx5_core] [<000000004f8a2031>] mlx5e_rep_indr_offload.isra.0+0x10e/0x1c0 [mlx5_core] [<000000007df797dc>] mlx5e_rep_indr_setup_tc_cb+0x90/0x130 [mlx5_core] [<0000000016c15cc3>] tc_setup_cb_add+0x1cf/0x410 [<00000000a63305b4>] fl_hw_replace_filter+0x38f/0x670 [cls_flower] [<000000008bc9e77c>] fl_change+0x1fd5/0x4430 [cls_flower] [<00000000e7f766e4>] tc_new_tfilter+0x867/0x2010 [<00000000e101c0ef>] rtnetlink_rcv_msg+0x6fc/0x9f0 [<00000000e1111d44>] netlink_rcv_skb+0x12c/0x360 [<0000000082dd6c8b>] netlink_unicast+0x438/0x710 [<00000000fc568f70>] netlink_sendmsg+0x794/0xc50 [<0000000016e92590>] sock_sendmsg+0xc5/0x190 So fix this by moving int_port cleanup code to the flow attribute free helper, which is used by all the attribute free cases.	N/A	More Details
CVE-2025-69028	Missing Authorization vulnerability in BoldGrid weForms weforms allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects weForms: from n/a through <= 1.6.25.	N/A	More Details
CVE-2025-69027	Missing Authorization vulnerability in tychesoftwares Product Delivery Date for WooCommerce – Lite product-delivery-date-for-woocommerce-lite allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Product Delivery Date for WooCommerce – Lite: from n/a through <= 3.2.0.	N/A	More Details
CVE-2025-69026	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in Roxnor PopupKit popup-builder-block allows Retrieve Embedded Sensitive Data.This issue affects PopupKit: from n/a through <= 2.1.5.	N/A	More Details
CVE-2025-69025	Exposure of Sensitive System Information to an Unauthorized Control Sphere vulnerability in Aethonic Poptics: AI-Powered Popup Builder for Lead Generation, Conversions, Exit-Intent, Email Opt-ins & WooCommerce Sales poptics allows Retrieve Embedded Sensitive Data.This issue affects Poptics: AI-Powered Popup Builder for Lead Generation, Conversions, Exit-Intent, Email Opt-ins & WooCommerce Sales: from n/a through <= 1.0.20.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: net: dsa: avoid suspicious RCU usage for synced VLAN-aware MAC addresses When using the felix driver (the only one which supports UC filtering and MC filtering) as a DSA master for a random other DSA switch, one can see the following stack trace when the downstream switch ports join a VLAN-aware bridge: ========== WARNING: suspicious RCU usage ---------- net/8021q/vlan_core.c:238 suspicious rcu_dereference_protected() usage! stack backtrace: Workqueue: dsa_ordered dsa_slave_switchdev_event_work Call trace: lockdep_rcu_suspicious+0x170/0x210 vlan_for_each+0x8c/0x188 dsa_slave_sync_uc+0x128/0x178 __hw_addr_sync_dev+0x138/0x158 dsa_slave_set_rx_mode+0x58/0x70 __dev_set_rx_mode+0x88/0xa8 dev_uc_add+0x74/0xa0 dsa_port_bridge_host_fdb_add+0xec/0x180 dsa_slave_switchdev_event_work+0x7c/0x1c8 process_one_work+0x290/0x568 What it's saying is that vlan_for_each() expects rtnl_lock() context and it's not getting it, when it's called from the DSA master's ndo_set_rx_mode(). The caller of that - dsa_slave_set_rx_mode() - is the slave DSA interface's dsa_port_bridge_host_fdb_add() which comes from the deferred dsa_slave_switchdev_event_work(). We went to great lengths to avoid the rtnl_lock() context in that call	N/A	More

54149	path in commit 0faf890fc519 ("net: dsa: drop rtnl_lock from dsa_slave_switchdev_event_work"), and calling rtnl_lock() is simply not an option due to the possibility of deadlocking when calling dsa_flush_workqueue() from the call paths that do hold rtnl_lock() - basically all of them. So, when the DSA master calls vlan_for_each() from its ndo_set_rx_mode(), the state of the 8021q driver on this device is really not protected from concurrent access by anything. Looking at net/8021q/, I don't think that vlan_info->vid_list was particularly designed with RCU traversal in mind, so introducing an RCU read-side form of vlan_for_each() - vlan_for_each_rcu() - won't be so easy, and it also wouldn't be exactly what we need anyway. In general I believe that the solution isn't in net/8021q/ anyway; vlan_for_each() is not cut out for this task. DSA doesn't need rtnl_lock() to be held per se - since it's not a netdev state change that we're blocking, but rather, just concurrent additions/removals to a VLAN list. We don't even need sleepable context - the callback of vlan_for_each() just schedules deferred work. The proposed escape is to remove the dependency on vlan_for_each() and to open-code a non-sleepable, rtnl-free alternative to that, based on copies of the VLAN list modified from .ndo_vlan_rx_add_vid() and .ndo_vlan_rx_kill_vid().		Details
CVE-2023-54148	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Move representor neigh cleanup to profile cleanup_tx For IP tunnel encapsulation in ECMP (Equal-Cost Multipath) mode, as the flow is duplicated to the peer eswitch, the related neighbour information on the peer uplink representor is created as well. In the cited commit, eswitch devcom unpair is moved to uplink unload API, specifically the profile->cleanup_tx. If there is a encap rule offloaded in ECMP mode, when one eswitch does unpair (because of unloading the driver, for instance), and the peer rule from the peer eswitch is going to be deleted, the use-after-free error is triggered while accessing neigh info, as it is already cleaned up in uplink's profile->disable, which is before its profile->cleanup_tx. To fix this issue, move the neigh cleanup to profile's cleanup_tx callback, and after mlx5e_cleanup_uplink_rep_tx is called. The neigh init is moved to init_tx for symmeter. [2453.376299] BUG: KASAN: slab-use-after-free in mlx5e_rep_neigh_entry_release+0x109/0x3a0 [mlx5_core] [2453.379125] Read of size 4 at addr ffff888127af9008 by task modprobe/2496 [2453.381542] CPU: 7 PID: 2496 Comm: modprobe Tainted: G B 6.4.0-rc7+ #15 [2453.383386] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qemu.org 04/01/2014 [2453.384335] Call Trace: [2453.384625] <TASK> [2453.384891] dump_stack_lvl+0x33/0x50 [2453.385285] print_report+0xc2/0x610 [2453.385667] ? __virt_addr_valid+0xb1/0x130 [2453.386091] ? mlx5e_rep_neigh_entry_release+0x109/0x3a0 [mlx5_core] [2453.386757] kasan_report+0xae/0xe0 [2453.387123] ? mlx5e_rep_neigh_entry_release+0x109/0x3a0 [mlx5_core] [2453.387798] mlx5e_rep_neigh_entry_release+0x109/0x3a0 [mlx5_core] [2453.388465] mlx5e_rep_encap_entry_detach+0xa6/0xe0 [mlx5_core] [2453.389111] mlx5e_encap_dealloc+0xa7/0x100 [mlx5_core] [2453.389706] mlx5e_tc_tun_encap_dests_unset+0x61/0xb0 [mlx5_core] [2453.390361] mlx5_free_flow_attr_actions+0x11e/0x340 [mlx5_core] [2453.391015] ? complete_all+0x43/0xd0 [2453.391398] ? free_flow_post_acts+0x38/0x120 [mlx5_core] [2453.392004] mlx5e_tc_del_fdb_flow+0x4ae/0x690 [mlx5_core] [2453.392618] mlx5e_tc_del_fdb_peers_flow+0x308/0x370 [mlx5_core] [2453.393276] mlx5e_tc_clean_fdb_peer_flows+0xf5/0x140 [mlx5_core] [2453.393925] mlx5_esw_offloads_unpair+0x86/0x540 [mlx5_core] [2453.394546] ? mlx5_esw_offloads_set_ns_peer.isra.0+0x180/0x180 [mlx5_core] [2453.395268] ? down_write+0xaa/0x100 [2453.395652] mlx5_esw_offloads_devcom_event+0x203/0x530 [mlx5_core] [2453.396317] mlx5_devcom_send_event+0xbb/0x190 [mlx5_core] [2453.396917] mlx5_esw_offloads_devcom_cleanup+0xb0/0xd0 [mlx5_core] [2453.397582] mlx5e_tc_esw_cleanup+0x42/0x120 [mlx5_core] [2453.398182] mlx5e_rep_tc_cleanup+0x15/0x30 [mlx5_core] [2453.398768] mlx5e_cleanup_rep_tx+0x6c/0x80 [mlx5_core] [2453.399367] mlx5e_detach_netdev+0xee/0x120 [mlx5_core] [2453.399957] mlx5e_netdev_change_profile+0x84/0x170 [mlx5_core] [2453.400598] mlx5e_vport_rep_unload+0xe0/0xf0 [mlx5_core] [2453.403781] mlx5_eswitch_unregister_vport_reps+0x15e/0x190 [mlx5_core] [2453.404479] ? mlx5_eswitch_register_vport_reps+0x200/0x200 [mlx5_core] [2453.405170] ? up_write+0x39/0x60 [2453.405529] ? kernfs_remove_by_name_ns+0xb7/0xe0 [2453.405985] auxiliary_bus_remove+0x2e/0x40 [2453.406405] device_release_driver_internal+0x243/0x2d0 [2453.406900] ? kobject_put+0x42/0x2d0 [2453.407284] bus_remove_device+0x128/0x1d0 [2453.407687] device_del+0x240/0x550 [2453.408053] ? waiting_for_supplier_show+0xe0/0xe0 [2453.408511] ? kobject_put+0xfa/0x2d0 [2453.408889] ? __kmem_cache_free+0x14d/0x280 [2453.409310] mlx5_rescan_drivers_locked.part.0+0xcd/0x2b0 [mlx5_core] [2453.409973] mlx5_unregister_device+0x40/0x50 [mlx5_core] [2453.410561] mlx5_uninit_one+0x3d/0x110 [mlx5_core] [2453.411111] remove_one+0x89/0x130 [mlx5_core] [24 ---truncated---	N/A	More Details
CVE-2023-54147	In the Linux kernel, the following vulnerability has been resolved: media: platform: mtk-mdp3: Add missing check and free for ida_alloc Add the check for the return value of the ida_alloc in order to avoid NULL pointer dereference. Moreover, free allocated "ctx->id" if mdp_m2m_open fails later in order to avoid memory leak.	N/A	More Details
CVE-2023-54146	In the Linux kernel, the following vulnerability has been resolved: x86/kexec: Fix double-free of elf header buffer After b3e34a47f989 ("x86/kexec: fix memory leak of elf header buffer"), freeing image->elf_headers in the error path of crash_load_segments() is not needed because kimage_file_post_load_cleanup() will take care of that later. And not clearing it could result in a double-free. Drop the superfluous vfree() call at the error path of crash_load_segments().	N/A	More Details
CVE-2023-54145	In the Linux kernel, the following vulnerability has been resolved: bpf: drop unnecessary user-triggerable WARN_ONCE in verifierl log It's trivial for user to trigger "verifier log line truncated" warning, as verifier has a fixed-sized buffer of 1024 bytes (as of now), and there are at least two pieces of user-provided information that can be output through this buffer, and both can be arbitrarily sized by user: - BTF names; - BTF.ext source code lines strings. Verifier log buffer should be properly sized for typical verifier state output. But it's sort-of expected that this buffer won't be long enough in some circumstances. So let's drop the check. In any case code will work correctly, at worst truncating a part of a single line output.	N/A	More Details
CVE-2023-54144	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Fix kernel warning during topology setup This patch fixes the following kernel warning seen during driver load by correctly initializing the p2plink attr before creating the sysfs file: [+0.002865] ------------[cut here]------------ [+0.002327] kobject: '(null)' (0000000056260cfb): is not initialized, yet kobject_put() is being called. [+0.004780] WARNING: CPU: 32 PID: 1006 at lib/kobject.c:718 kobject_put+0xaa/0x1c0 [+0.001361] Call Trace: [+0.001234] <TASK> [+0.001067] kfd_remove_sysfs_node_entry+0x24a/0x2d0 [amdgpu] [+0.003147] kfd_topology_update_sysfs+0x3d/0x750 [amdgpu] [+0.002890] kfd_topology_add_device+0xbd7/0xc70 [amdgpu] [+0.002844] ? lock_release+0x13c/0x2e0 [+0.001936] ? smu_cmn_send_smc_msg_with_param+0x1e8/0x2d0 [amdgpu] [+0.003313] ? amdgpu_dpm_get_mclk+0x54/0x60 [amdgpu] [+0.002703] kgd2kfd_device_init.cold+0x39f/0x4ed [amdgpu] [+0.002930] amdgpu_amdkfd_device_init+0x13d/0x1f0 [amdgpu] [+0.002944] amdgpu_device_init.cold+0x1464/0x17b4 [amdgpu] [+0.002970] ? pci_bus_read_config_word+0x43/0x80 [+0.002380] amdgpu_driver_load_kms+0x15/0x100 [amdgpu] [+0.002744] amdgpu_pci_probe+0x147/0x370 [amdgpu] [+0.002522] local_pci_probe+0x40/0x80 [+0.001896] work_for_cpu_fn+0x10/0x20 [+0.001892] process_one_work+0x26e/0x5a0 [+0.002029] worker_thread+0x1fd/0x3e0 [+0.001890] ? process_one_work+0x5a0/0x5a0 [+0.002115] kthread+0xea/0x110 [+0.001618] ? kthread_complete_and_exit+0x20/0x20 [+0.002422] ret_from_fork+0x1f/0x30 [+0.001808] </TASK> [+0.001103] irq event stamp: 59837 [+0.001718] hardirqs last enabled at (59849): [<ffffffffb30fab12>] __up_console_sem+0x52/0x60 [+0.004414] hardirqs last disabled at (59860): [<ffffffffb30faaf7>] __up_console_sem+0x37/0x60 [+0.004414] softirqs last enabled at (59654): [<ffffffffb307d9c7>] irq_exit_rcu+0xd7/0x130 [+0.004205] softirqs last disabled at (59649): [<ffffffffb307d9c7>] irq_exit_rcu+0xd7/0x130 [+0.004203] ---[end trace 0000000000000000]---	N/A	More Details

CVE-2023-54143	In the Linux kernel, the following vulnerability has been resolved: media: mediatek: vcodec: fix resource leaks in vdec_msg_queue_init() If we encounter any error in the vdec_msg_queue_init() then we need to set "msg_queue->wdma_addr.size = 0;". Normally, this is done inside the vdec_msg_queue_deinit() function. However, if the first call to allocate &msg_queue->wdma_addr fails, then the vdec_msg_queue_deinit() function is a no-op. For that situation, just set the size to zero explicitly and return. There were two other error paths which did not clean up before returning. Change those error paths to goto mem_alloc_err.	N/A	More Details
CVE-2023-54142	In the Linux kernel, the following vulnerability has been resolved: gtp: Fix use-after-free in __gtp_encap_destroy(). syzkaller reported use-after-free in __gtp_encap_destroy(). [0] It shows the same process freed sk and touched it illegally. Commit e198987e7dd7 ("gtp: fix suspicious RCU usage") added lock_sock() and release_sock() in __gtp_encap_destroy() to protect sk->sk_user_data, but release_sock() is called after sock_put() releases the last refcnt. [0]: BUG: KASAN: slab-use-after-free in instrument_atomic_read_write include/linux/instrumented.h:96 [inline] BUG: KASAN: slab-use-after-free in atomic_try_cmpxchg_acquire include/linux/atomic/atomic-instrumented.h:541 [inline] BUG: KASAN: slab-use-after-free in queued_spin_lock include/asm-generic/qspinlock.h:111 [inline] BUG: KASAN: slab-use-after-free in do_raw_spin_lock include/linux/spinlock.h:186 [inline] BUG: KASAN: slab-use-after-free in __raw_spin_lock_bh include/linux/spinlock_api_smp.h:127 [inline] BUG: KASAN: slab-use-after-free in _raw_spin_lock_bh+0x75/0xe0 kernel/locking/spinlock.c:178 Write of size 4 at addr ffff88800dbef398 by task syz-executor.2/2401 CPU: 1 PID: 2401 Comm: syz-executor.2 Not tainted 6.4.0-rc5-01219-gfa0e21fa4443 #2 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x72/0xa0 lib/dump_stack.c:106 print_address_description mm/kasan/report.c:351 [inline] print_report+0xcc/0x620 mm/kasan/report.c:462 kasan_report+0xb2/0xe0 mm/kasan/report.c:572 check_region_inline mm/kasan/generic.c:181 [inline] kasan_check_range+0x39/0x1c0 mm/kasan/generic.c:187 instrument_atomic_read_write include/linux/instrumented.h:96 [inline] atomic_try_cmpxchg_acquire include/linux/atomic/atomic-instrumented.h:541 [inline] queued_spin_lock include/asm-generic/qspinlock.h:111 [inline] do_raw_spin_lock include/linux/spinlock.h:186 [inline] __raw_spin_lock_bh include/linux/spinlock_api_smp.h:127 [inline] _raw_spin_lock_bh+0x75/0xe0 kernel/locking/spinlock.c:178 spin_lock_bh include/linux/spinlock.h:355 [inline] release_sock+0x1f/0x1a0 net/core/sock.c:3526 gtp_encap_disable_sock drivers/net/gtp.c:651 [inline] gtp_encap_disable+0xb9/0x220 drivers/net/gtp.c:664 gtp_dev_uninit+0x19/0x50 drivers/net/gtp.c:728 unregister_netdevice_many_notify+0x97e/0x1520 net/core/dev.c:10841 rtnl_delete_link net/core/rtnetlink.c:3216 [inline] rtnl_dellink+0x3c0/0xb30 net/core/rtnetlink.c:3268 rtnetlink_rcv_msg+0x450/0xb10 net/core/rtnetlink.c:6423 netlink_rcv_skb+0x15d/0x450 net/netlink/af_netlink.c:2548 netlink_unicast_kernel net/netlink/af_netlink.c:1339 [inline] netlink_unicast+0x700/0x930 net/netlink/af_netlink.c:1365 netlink_sendmsg+0x91c/0xe30 net/netlink/af_netlink.c:1913 sock_sendmsg_nosec net/socket.c:724 [inline] sock_sendmsg+0x1b7/0x200 net/socket.c:747 ____sys_sendmsg+0x75a/0x990 net/socket.c:2493 ___sys_sendmsg+0x11d/0x1c0 net/socket.c:2547 __sys_sendmsg+0xfe/0x1d0 net/socket.c:2576 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x3f/0x90 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x72/0xdc RIP: 0033:0x7f1168b1fe5d Code: ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 1e fa 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 8b 0d 73 9f 1b 00 f7 d8 64 89 01 48 RSP: 002b:00007f1167edccc8 EFLAGS: 00000246 ORIG_RAX: 000000000000002e RAX: ffffffffffffffda RBX: 00000000004bbf80 RCX: 00007f1168b1fe5d RDX: 0000000000000000 RSI: 00000000200002c0 RDI: 0000000000000003 RBP: 00000000004bbf80 R08: 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000000 R13: 000000000000000b R14: 00007f1168b80530 R15: 0000000000000000 </TASK> Allocated by task 1483: kasan_save_stack+0x22/0x50 mm/kasan/common.c:45 kasan_set_track+0x25/0x30 mm/kasan/common.c:52 __kasan_slab_alloc+0x ---truncated---	N/A	More Details
CVE-2023-54141	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: Add missing hw_ops->get_ring_selector() for IPQ5018 During sending data after clients connected, hw_ops->get_ring_selector() will be called. But for IPQ5018, this member isn't set, and the following NULL pointer exception will be occurred: [38.840478] 8<--- cut here --- [38.840517] Unable to handle kernel NULL pointer dereference at virtual address 00000000 ... [38.923161] PC is at 0x0 [38.927930] LR is at ath11k_dp_tx+0x70/0x730 [ath11k] ... [39.063264] Process hostapd (pid: 1034, stack limit = 0x801ceb3d) [39.068994] Stack: (0x856a9a68 to 0x856aa000) ... [39.438467] [<7f323804>] (ath11k_dp_tx [ath11k]) from [<7f314e6c>] (ath11k_mac_op_tx+0x80/0x190 [ath11k]) [39.446607] [<7f314e6c>] (ath11k_mac_op_tx [ath11k]) from [<7f17dbe0>] (ieee80211_handle_wake_tx_queue+0x7c/0xc0 [mac80211]) [39.456162] [<7f17dbe0>] (ieee80211_handle_wake_tx_queue [mac80211]) from [<7f174450>] (ieee80211_probereq_get+0x584/0x704 [mac80211]) [39.467443] [<7f174450>] (ieee80211_probereq_get [mac80211]) from [<7f178c40>] (ieee80211_tx_prepare_skb+0x1f8/0x248 [mac80211]) [39.479334] [<7f178c40>] (ieee80211_tx_prepare_skb [mac80211]) from [<7f179e28>] (__ieee80211_subif_start_xmit+0x32c/0x3d4 [mac80211]) [39.491053] [<7f179e28>] (__ieee80211_subif_start_xmit [mac80211]) from [<7f17af08>] (ieee80211_tx_control_port+0x19c/0x288 [mac80211]) [39.502946] [<7f17af08>] (ieee80211_tx_control_port [mac80211]) from [<7f0fc704>] (nl80211_tx_control_port+0x174/0x1d4 [cfg80211]) [39.515017] [<7f0fc704>] (nl80211_tx_control_port [cfg80211]) from [<808ceac4>] (genl_rcv_msg+0x154/0x340) [39.526814] [<808ceac4>] (genl_rcv_msg) from [<808cdb74>] (netlink_rcv_skb+0xb8/0x11c) [39.536446] [<808cdb74>] (netlink_rcv_skb) from [<808ce1d0>] (genl_rcv+0x28/0x34) [39.544344] [<808ce1d0>] (genl_rcv) from [<808cd234>] (netlink_unicast+0x174/0x274) [39.551895] [<808cd234>] (netlink_unicast) from [<808cd510>] (netlink_sendmsg+0x1dc/0x440) [39.559362] [<808cd510>] (netlink_sendmsg) from [<808596e0>] (____sys_sendmsg+0x1a8/0x1fc) [39.567697] [<808596e0>] (____sys_sendmsg) from [<8085b1a8>] (___sys_sendmsg+0xa4/0xdc) [39.575941] [<8085b1a8>] (___sys_sendmsg) from [<8085b310>] (sys_sendmsg+0x44/0x74) [39.583841] [<8085b310>] (sys_sendmsg) from [<80300060>] (ret_fast_syscall+0x0/0x40) ... [39.620734] Code: bad PC value [39.625869] ---[end trace 8aef983ad3cbc032]---	N/A	More Details
CVE-2023-54005	In the Linux kernel, the following vulnerability has been resolved: binder: fix memory leak in binder_init() In binder_init(), the destruction of binder_alloc_shrinker_init() is not performed in the wrong path, which will cause memory leaks. So this commit introduces binder_alloc_shrinker_exit() and calls it in the wrong path to fix that.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: udplite: Fix NULL pointer dereference in __sk_mem_raise_allocated(). syzbot reported [0] a null-ptr-deref in sk_get_rmem0() while using IPPROTO_UDPLITE (0x88): 14:25:52 executing program 1: r0 = socket$inet6(0xa, 0x80002, 0x88) We had a similar report [1] for probably sk_memory_allocated_add() in __sk_mem_raise_allocated(), and commit c915fe13cbaa ("udplite: fix NULL pointer dereference") fixed it by setting .memory_allocated for udplite_prot and udplitev6_prot. To fix the variant, we need to set either .sysctl_wmem_offset or .sysctl_rmem. Now UDP and UDPLITE share the same value for .memory_allocated, so we use the same .sysctl_wmem_offset for UDP and UDPLITE. [0]: general protection fault, probably for non-canonical address 0xdfffffc0000000000: 0000 [#1] PREEMPT SMP KASAN KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007] CPU: 0 PID: 6829 Comm: syz-executor.1 Not tainted 6.4.0-rc2-syzkaller #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 04/28/2023 RIP: 0010:sk_get_rmem0 include/net/sock.h:2907 [inline] RIP: 0010:__sk_mem_raise_allocated+0x806/0x17a0 net/core/sock.c:3006 Code: c1 ea 03 80 3c 02 00 0f 85 23 0f 00 00 48 8b 44 24 08 48 8b 98 38 01 00 00 48 b8 00 00 00 00 00 fc ff df 48 89 da 48 c1 ea 03 <0f> b6 14 02 48 89 d8 83 e0 07 83 c0 03 38 d0 0f 8d 6f 0a 00 00 8b RSP: 0018:ffffc90005d7f450 EFLAGS: 00010246 RAX: dfffffc0000000000 RBX: 0000000000000000 RCX: ffffc90004d92000 RDX: 0000000000000000 RSI: ffffffff88066482 RDI: ffffffff8e2ccbb8 RBP: ffff8880173f7000 R08: 0000000000000005 R09: 0000000000000000 R10: 0000000000000000 R11:		

CVE-2023-54004	0000000000000000 R12: 0000000000030000 R13: 0000000000000001 R14: 0000000000000340 R15: 0000000000000001 FS: 0000000000000000(0000) GS:ffff8880b9800000(0063) knlGS:00000000f7f1cb40 CS: 0010 DS: 002b ES: 002b CR0: 0000000080050033 CR2: 000000002e82f000 CR3: 0000000034ff0000 CR4: 00000000003506f0 Call Trace: <TASK> __sk_mem_schedule+0x6c/0xe0 net/core/sock.c:3077 udp_rmem_schedule net/ipv4/udp.c:1539 [inline] __udp_enqueue_schedule_skb+0x776/0xb30 net/ipv4/udp.c:1581 __udpv6_queue_rcv_skb net/ipv6/udp.c:666 [inline] udpv6_queue_rcv_one_skb+0xc39/0x16c0 net/ipv6/udp.c:775 udpv6_queue_rcv_skb+0x194/0xa10 net/ipv6/udp.c:793 __udp6_lib_mcast_deliver net/ipv6/udp.c:906 [inline] __udp6_lib_rcv+0x1bda/0x2bd0 net/ipv6/udp.c:1013 ip6_protocol_deliver_rcu+0x2e7/0x1250 net/ipv6/ip6_input.c:437 ip6_input_finish+0x150/0x2f0 net/ipv6/ip6_input.c:482 NF_HOOK include/linux/netfilter.h:303 [inline] NF_HOOK include/linux/netfilter.h:297 [inline] ip6_input+0xa0/0xd0 net/ipv6/ip6_input.c:491 ip6_mc_input+0x40b/0xf50 net/ipv6/ip6_input.c:585 dst_input include/net/dst.h:468 [inline] ip6_rcv_finish net/ipv6/ip6_input.c:79 [inline] NF_HOOK include/linux/netfilter.h:303 [inline] NF_HOOK include/linux/netfilter.h:297 [inline] ipv6_rcv+0x250/0x380 net/ipv6/ip6_input.c:309 __netif_receive_skb_one_core+0x114/0x180 net/core/dev.c:5491 __netif_receive_skb+0x1f/0x1c0 net/core/dev.c:5605 netif_receive_skb_internal net/core/dev.c:5691 [inline] netif_receive_skb+0x133/0x7a0 net/core/dev.c:5750 tun_rx_batched+0x4b3/0x7a0 drivers/net/tun.c:1553 tun_get_user+0x2452/0x39c0 drivers/net/tun.c:1989 tun_chr_write_iter+0xdf/0x200 drivers/net/tun.c:2035 call_write_iter include/linux/fs.h:1868 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_write+0x945/0xd50 fs/read_write.c:584 ksys_write+0x12b/0x250 fs/read_write.c:637 do_syscall_32_irqs_on arch/x86/entry/common.c:112 [inline] __do_fast_syscall_32+0x65/0xf0 arch/x86/entry/common.c:178 do_fast_syscall_32+0x33/0x70 arch/x86/entry/common.c:203 entry_SYSENTER_compat_after_hwframe+0x70/0x82 RIP: 0023:0xf7f21579 Code: b8 01 10 06 03 74 b4 01 10 07 03 74 b0 01 10 08 03 74 d8 01 00 00 00 00 00 00 00 00 00 00 00 00 51 52 55 89 e5 0f 34 cd 80 <5d> 5a 59 c3 90 90 90 90 8d b4 26 00 00 00 00 8d b4 26 00 00 00 00 ---truncated---	N/A	More Details
CVE-2023-54003	In the Linux kernel, the following vulnerability has been resolved: RDMA/core: Fix GID entry ref leak when create_ah fails If AH create request fails, release sgid_attr to avoid GID entry referrence leak reported while releasing GID table	N/A	More Details
CVE-2023-54002	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix assertion of exclop condition when starting balance Balance as exclusive state is compatible with paused balance and device add, which makes some things more complicated. The assertion of valid states when starting from paused balance needs to take into account two more states, the combinations can be hit when there are several threads racing to start balance and device add. This won't typically happen when the commands are started from command line. Scenario 1: With exclusive_operation state == BTRFS_EXCLOP_NONE. Concurrently adding multiple devices to the same mount point and btrfs_exclop_finish executed finishes before assertion in btrfs_exclop_balance, exclusive_operation will changed to BTRFS_EXCLOP_NONE state which lead to assertion failed: fs_info->exclusive_operation == BTRFS_EXCLOP_BALANCE || fs_info->exclusive_operation == BTRFS_EXCLOP_DEV_ADD, in fs/btrfs/ioctl.c:456 Call Trace: <TASK> btrfs_exclop_balance+0x13c/0x310 ? memdup_user+0xab/0xc0 ? PTR_ERR+0x17/0x20 btrfs_ioctl_add_dev+0x2ee/0x320 btrfs_ioctl+0x9d5/0x10d0 ? btrfs_ioctl_encoded_write+0xb80/0xb80 __x64_sys_ioctl+0x197/0x210 do_syscall_64+0x3c/0xb0 entry_SYSCALL_64_after_hwframe+0x63/0xcd Scenario 2: With exclusive_operation state == BTRFS_EXCLOP_BALANCE_PAUSED. Concurrently adding multiple devices to the same mount point and btrfs_exclop_balance executed finish before the latter thread execute assertion in btrfs_exclop_balance, exclusive_operation will changed to BTRFS_EXCLOP_BALANCE_PAUSED state which lead to assertion failed: fs_info->exclusive_operation == BTRFS_EXCLOP_BALANCE || fs_info->exclusive_operation == BTRFS_EXCLOP_DEV_ADD || fs_info->exclusive_operation == BTRFS_EXCLOP_NONE, fs/btrfs/ioctl.c:458 Call Trace: <TASK> btrfs_exclop_balance+0x240/0x410 ? memdup_user+0xab/0xc0 ? PTR_ERR+0x17/0x20 btrfs_ioctl_add_dev+0x2ee/0x320 btrfs_ioctl+0x9d5/0x10d0 ? btrfs_ioctl_encoded_write+0xb80/0xb80 __x64_sys_ioctl+0x197/0x210 do_syscall_64+0x3c/0xb0 entry_SYSCALL_64_after_hwframe+0x63/0xcd An example of the failed assertion is below, which shows that the paused balance is also needed to be checked. root@syzkaller:/home/xsk# ./repro Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 Failed to add device /dev/vda, errno 14 [416.611428][T7970] BTRFS info (device loop0): fs_info exclusive_operation: 0 Failed to add device /dev/vda, errno 14 [416.613973][T7971] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.615456][T7972] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.617528][T7973] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.618359][T7974] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.622589][T7975] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.624034][T7976] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.626420][T7977] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.627643][T7978] BTRFS info (device loop0): fs_info exclusive_operation: 3 Failed to add device /dev/vda, errno 14 [416.629006][T7979] BTRFS info (device loop0): fs_info exclusive_operation: 3 [416.630298][T7980] BTRFS info (device loop0): fs_info exclusive_operation: 3 Fai ---truncated---	N/A	More Details
CVE-2025-69021	Cross-Site Request Forgery (CSRF) vulnerability in Ays Pro Popup box ays-popup-box allows Cross Site Request Forgery.This issue affects Popup box: from n/a through <= 6.0.7.	N/A	More Details
CVE-2025-69022	Missing Authorization vulnerability in Weblizar - WordPress Themes & Plugin HR Management Lite hr-management-lite allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects HR Management Lite: from n/a through <= 3.5.	N/A	More Details
CVE-2025-69023	Missing Authorization vulnerability in Marketing Fire Discussion Board wp-discussion-board allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects Discussion Board: from n/a through <= 2.5.7.	N/A	More Details
CVE-2025-69024	Missing Authorization vulnerability in bizswoop BizPrint print-google-cloud-print-gcp-woocommerce allows Exploiting Incorrectly Configured Access Control Security Levels.This issue affects BizPrint: from n/a through <= 4.6.7.	N/A	More Details
CVE-2025-69235	Whale browser before 4.35.351.12 allows an attacker to bypass the Same-Origin Policy in a sidebar environment.	N/A	More Details
CVE-2023-54028	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix the error "trying to register non-static key in rxe_cleanup_task" In the function rxe_create_qp(), rxe_qp_from_init() is called to initialize qp, internally things like rxe_init_task are not setup until rxe_qp_init_req(). If an error occurred before this point then the unwind will call rxe_cleanup() and eventually to rxe_qp_do_cleanup()/rxe_cleanup_task() which will oops when trying to access the uninitialized spinlock. If rxe_init_task is not executed, rxe_cleanup_task will not be called.	N/A	More Details

CVE-2025-68378	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix stackmap overflow check in __bpf_get_stackid() Syzkaller reported a KASAN slab-out-of-bounds write in __bpf_get_stackid() when copying stack trace data. The issue occurs when the perf trace contains more stack entries than the stack map bucket can hold, leading to an out-of-bounds write in the bucket's data array.	N/A	More Details
CVE-2025-68363	In the Linux kernel, the following vulnerability has been resolved: bpf: Check skb->transport_header is set in bpf_skb_check_mtu The bpf_skb_check_mtu helper needs to use skb->transport_header when the BPF_MTU_CHK_SEGS flag is used: bpf_skb_check_mtu(skb, ifindex, &mtu_len, 0, BPF_MTU_CHK_SEGS) The transport_header is not always set. There is a WARN_ON_ONCE report when CONFIG_DEBUG_NET is enabled + skb->gso_size is set + bpf_prog_test_run is used: WARNING: CPU: 1 PID: 2216 at ./include/linux/skbuff.h:3071 skb_gso_validate_network_len bpf_skb_check_mtu bpf_prog_3920e25740a41171_tc_chk_segs_flag # A test in the next patch bpf_test_run bpf_prog_test_run_skb For a normal ingress skb (not test_run), skb_reset_transport_header is performed but there is plan to avoid setting it as described in commit 2170a1f09148 ("net: no longer reset transport_header in __netif_receive_skb_core()"). This patch fixes the bpf helper by checking skb_transport_header_was_set(). The check is done just before skb->transport_header is used, to avoid breaking the existing bpf prog. The WARN_ON_ONCE is limited to bpf_prog_test_run, so targeting bpf-next.	N/A	More Details
CVE-2025-69211	Nest is a framework for building scalable Node.js server-side applications. Versions prior to 11.1.11 have a Fastify URL encoding middleware bypass. A NestJS application is vulnerable if it uses `@nestjs/platform-fastify`; relies on `NestMiddleware` (via `MiddlewareConsumer`) for security checks (authentication, authorization, etc.), or through `app.use()`; and applies middleware to specific routes using string paths or controllers (e.g., `.forRoutes('admin')`). Exploitation can result in unauthenticated users accessing protected routes, restricted administrative endpoints becoming accessible to lower-privileged users, and/or middleware performing sanitization or validation being skipped. This issue is patched in `@nestjs/platform-fastify@11.1.11`.	N/A	More Details
CVE-2025-68366	In the Linux kernel, the following vulnerability has been resolved: nbd: defer config unlock in nbd_genl_connect There is one use-after-free warning when running NBD_CMD_CONNECT and NBD_CLEAR_SOCK: nbd_genl_connect nbd_alloc_and_init_config // config_refs=1 nbd_start_device // config_refs=2 set NBD_RT_HAS_CONFIG_REF open nbd // config_refs=3 recv_work done // config_refs=2 NBD_CLEAR_SOCK // config_refs=1 close nbd // config_refs=0 refcount_inc -> uaf ------------[cut here]------------ refcount_t: addition on 0; use-after-free. WARNING: CPU: 24 PID: 1014 at lib/refcount.c:25 refcount_warn_saturate+0x12e/0x290 nbd_genl_connect+0x16d0/0x1ab0 genl_family_rcv_msg_doit+0x1f3/0x310 genl_rcv_msg+0x44a/0x790 The issue can be easily reproduced by adding a small delay before refcount_inc(&nbd->config_refs) in nbd_genl_connect(): mutex_unlock(&nbd->config_lock); if (!ret) { set_bit(NBD_RT_HAS_CONFIG_REF, &config->runtime_flags); + printk("before sleep\n"); + mdelay(5 * 1000); + printk("after sleep\n"); refcount_inc(&nbd->config_refs); nbd_connect_reply(info, nbd->index); }	N/A	More Details
CVE-2023-54099	In the Linux kernel, the following vulnerability has been resolved: fs: Protect reconfiguration of sb read-write from racing writes The reconfigure / remount code takes a lot of effort to protect filesystem's reconfiguration code from racing writes on remounting read-only. However during remounting read-only filesystem to read-write mode userspace writes can start immediately once we clear SB_RDONLY flag. This is inconvenient for example for ext4 because we need to do some writes to the filesystem (such as preparation of quota files) before we can take userspace writes so we are clearing SB_RDONLY flag before we are fully ready to accept userpace writes and syzbot has found a way to exploit this [1]. Also as far as I'm reading the code the filesystem remount code was protected from racing writes in the legacy mount path by the mount's MNT_READONLY flag so this is relatively new problem. It is actually fairly easy to protect remount read-write from racing writes using sb->s_readonly_remount flag so let's just do that instead of having to workaround these races in the filesystem code. [1] https://lore.kernel.org/all/00000000000006a0df05f6667499@google.com/T/	N/A	More Details
CVE-2023-54100	In the Linux kernel, the following vulnerability has been resolved: scsi: qedi: Fix use after free bug in qedi_remove() In qedi_probe() we call __qedi_probe() which initializes &qedi->recovery_work with qedi_recovery_handler() and &qedi->board_disable_work with qedi_board_disable_work(). When qedi_schedule_recovery_handler() is called, schedule_delayed_work() will finally start the work. In qedi_remove(), which is called to remove the driver, the following sequence may be observed: Fix this by finishing the work before cleanup in qedi_remove(). CPU0 CPU1 |qedi_recovery_handler qedi_remove | __qedi_remove | iscsi_host_free | scsi_host_put | //free shost | |iscsi_host_for_each_session |//use qedi->shost Cancel recovery_work and board_disable_work in __qedi_remove().	N/A	More Details
CVE-2025-66861	An issue was discovered in function d_unqualified_name in file cp-demangle.c in BinUtils 2.26 allowing attackers to cause a denial of service via crafted PE file.	N/A	More Details
CVE-2025-68365	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Initialize allocated memory before use KMSAN reports: Multiple uninitialized values detected: - KMSAN: uninit-value in ntfs_read_hdr (3) - KMSAN: uninit-value in bcmp (3) Memory is allocated by __getname(), which is a wrapper for kmem_cache_alloc(). This memory is used before being properly cleared. Change kmem_cache_alloc() to kmem_cache_zalloc() to properly allocate and clear memory before use.	N/A	More Details
CVE-2025-68364	In the Linux kernel, the following vulnerability has been resolved: ocfs2: relax BUG() to ocfs2_error() in __ocfs2_move_extent() In '__ocfs2_move_extent()', relax 'BUG()' to 'ocfs2_error()' just to avoid crashing the whole kernel due to a filesystem corruption.	N/A	More Details
CVE-2025-66864	An issue was discovered in function d_print_comp_inner in file cp-demangle.c in BinUtils 2.26 allows attackers to cause a denial of service via crafted PE file.	N/A	More Details
CVE-2025-66866	An issue was discovered in function d_abi_tags in file cp-demangle.c in BinUtils 2.26 allows attackers to cause a denial of service via crafted PE file.	N/A	More Details
CVE-2025-68357	In the Linux kernel, the following vulnerability has been resolved: iomap: allocate s_dio_done_wq for async reads as well Since commit 222f2c7c6d14 ("iomap: always run error completions in user context"), read error completions are deferred to s_dio_done_wq. This means the workqueue also needs to be allocated for async reads.	N/A	More Details
CVE-2025-68362	In the Linux kernel, the following vulnerability has been resolved: wifi: rtl818x: rtl8187: Fix potential buffer underflow in rtl8187_rx_cb() The rtl8187_rx_cb() calculates the rx descriptor header address by subtracting its size from the skb tail pointer. However, it does not validate if the received packet (skb->len from urb->actual_length) is large enough to contain this header. If a truncated packet is received, this will lead to a buffer underflow, reading memory before the start of the skb data area, and causing a kernel panic. Add length checks for both rtl8187 and rtl8187b descriptor headers before attempting to access them, dropping the packet cleanly if the check fails.	N/A	More Details
CVE-	In the Linux kernel, the following vulnerability has been resolved: erofs: limit the level of fs stacking for file-backed mounts		More

2025-68361	Otherwise, it could cause potential kernel stack overflow (e.g., EROFS mounting itself).	N/A	Details
CVE-2025-68360	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: wed: use proper wed reference in mt76 wed driver callabacks MT7996 driver can use both wed and wed_hif2 devices to offload traffic from/to the wireless NIC. In the current codebase we assume to always use the primary wed device in wed callbacks resulting in the following crash if the hw runs wed_hif2 (e.g. 6GHz link). [297.455876] Unable to handle kernel read from unreadable memory at virtual address 000000000000080a [297.464928] Mem abort info: [297.467722] ESR = 0x0000000096000005 [297.471461] EC = 0x25: DABT (current EL), IL = 32 bits [297.476766] SET = 0, FnV = 0 [297.479809] EA = 0, S1PTW = 0 [297.482940] FSC = 0x05: level 1 translation fault [297.487809] Data abort info: [297.490679] ISV = 0, ISS = 0x00000005, ISS2 = 0x00000000 [297.496156] CM = 0, WnR = 0, TnD = 0, TagAccess = 0 [297.501196] GCS = 0, Overlay = 0, DirtyBit = 0, Xs = 0 [297.506500] user pgtable: 4k pages, 39-bit VAs, pgdp=0000000107480000 [297.512927] [000000000000080a] pgd=08000001097fb003, p4d=08000001097fb003, pud=08000001097fb003, pmd=0000000000000000 [297.523532] Internal error: Oops: 0000000096000005 [#1] SMP [297.715393] CPU: 2 UID: 0 PID: 45 Comm: kworker/u16:2 Tainted: G O 6.12.50 #0 [297.723908] Tainted: [O]=OOT_MODULE [297.727384] Hardware name: Banana Pi BPI-R4 (2x SFP+) (DT) [297.732857] Workqueue: nf_ft_offload_del nf_flow_rule_route_ipv6 [nf_flow_table] [297.740254] pstate: 60400005 (nZCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) [297.747205] pc : mt76_wed_offload_disable+0x64/0xa0 [mt76] [297.752688] lr : mtk_wed_flow_remove+0x58/0x80 [297.757126] sp : ffffffc080fe3ae0 [297.760430] x29: ffffffc080fe3ae0 x28: ffffffc080fe3be0 x27: 00000000deadbef7 [297.767557] x26: ffffff80c5ebca00 x25: 0000000000000001 x24: ffffff80c85f4c00 [297.774683] x23: ffffff80c1875b78 x22: ffffffc080d42cd0 x21: ffffffc080660018 [297.781809] x20: ffffff80c6a076d0 x19: ffffff80c6a043c8 x18: 0000000000000000 [297.788935] x17: 0000000000000000 x16: 0000000000000001 x15: 0000000000000000 [297.796060] x14: 0000000000000019 x13: ffffff80c0ad8ec0 x12: 00000000fa83b2da [297.803185] x11: ffffff80c02700c0 x10: ffffff80c0ad8ec0 x9 : ffffff81fef96200 [297.810311] x8 : ffffff80c02700c0 x7 : ffffff80c02700d0 x6 : 0000000000000002 [297.817435] x5 : 0000000000000400 x4 : 0000000000000000 x3 : 0000000000000000 [297.824561] x2 : 0000000000000001 x1 : 0000000000000800 x0 : ffffff80c6a063c8 [297.831686] Call trace: [297.834123] mt76_wed_offload_disable+0x64/0xa0 [mt76] [297.839254] mtk_wed_flow_remove+0x58/0x80 [297.843342] mtk_flow_offload_cmd+0x434/0x574 [297.847689] mtk_wed_setup_tc_block_cb+0x30/0x40 [297.852295] nf_flow_offload_ipv6_hook+0x7f4/0x964 [nf_flow_table] [297.858466] nf_flow_rule_route_ipv6+0x438/0x4a4 [nf_flow_table] [297.864463] process_one_work+0x174/0x300 [297.868465] worker_thread+0x278/0x430 [297.872204] kthread+0xd8/0xdc [297.875251] ret_from_fork+0x10/0x20 [297.878820] Code: 928b5ae0 8b000273 91400a60 f943fa61 (79401421) [297.884901] ---[end trace 0000000000000000]--- Fix the issue detecting the proper wed reference to use running wed callabacks.	N/A	More Details
CVE-2025-68359	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix double free of qgroup record after failure to add delayed ref head In the previous code it was possible to incur into a double kfree() scenario when calling add_delayed_ref_head(). This could happen if the record was reported to already exist in the btrfs_qgroup_trace_extent_nolock() call, but then there was an error later on add_delayed_ref_head(). In this case, since add_delayed_ref_head() returned an error, the caller went to free the record. Since add_delayed_ref_head() couldn't set this kfree'd pointer to NULL, then kfree() would have acted on a non-NULL 'record' object which was pointing to memory already freed by the callee. The problem comes from the fact that the responsibility to kfree the object is on both the caller and the callee at the same time. Hence, the fix for this is to shift the ownership of the 'qrecord' object out of the add_delayed_ref_head(). That is, we will never attempt to kfree() the given object inside of this function, and will expect the caller to act on the 'qrecord' object on its own. The only exception where the 'qrecord' object cannot be kfree'd is if it was inserted into the tracing logic, for which we already have the 'qrecord_inserted_ret' boolean to account for this. Hence, the caller has to kfree the object only if add_delayed_ref_head() reports not to have inserted it on the tracing logic. As a side-effect of the above, we must guarantee that 'qrecord_inserted_ret' is properly initialized at the start of the function, not at the end, and then set when an actual insert happens. This way we avoid 'qrecord_inserted_ret' having an invalid value on an early exit. The documentation from the add_delayed_ref_head() has also been updated to reflect on the exact ownership of the 'qrecord' object.	N/A	More Details
CVE-2023-54101	In the Linux kernel, the following vulnerability has been resolved: driver: soc: xilinx: use _safe loop iterator to avoid a use after free The hash_for_each_possible() loop dereferences "eve_data" to get the next item on the list. However the loop frees eve_data so it leads to a use after free. Use hash_for_each_possible_safe() instead.	N/A	More Details
CVE-2025-68358	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix racy bitfield write in btrfs_clear_space_info_full() From the memory-barriers.txt document regarding memory barrier ordering guarantees: (*) These guarantees do not apply to bitfields, because compilers often generate code to modify these using non-atomic read-modify-write sequences. Do not attempt to use bitfields to synchronize parallel algorithms. (*) Even in cases where bitfields are protected by locks, all fields in a given bitfield must be protected by one lock. If two fields in a given bitfield are protected by different locks, the compiler's non-atomic read-modify-write sequences can cause an update to one field to corrupt the value of an adjacent field. btrfs_space_info has a bitfield sharing an underlying word consisting of the fields full, chunk_alloc, and flush: struct btrfs_space_info { struct btrfs_fs_info * fs_info; /* 0 8 */ struct btrfs_space_info * parent; /* 8 8 */ ... int clamp; /* 172 4 */ unsigned int full:1; /* 176: 0 4 */ unsigned int chunk_alloc:1; /* 176: 1 4 */ unsigned int flush:1; /* 176: 2 4 */ ... Therefore, to be safe from parallel read-modify-writes losing a write to one of the bitfield members protected by a lock, all writes to all the bitfields must use the lock. They almost universally do, except for btrfs_clear_space_info_full() which iterates over the space_infos and writes out found->full = 0 without a lock. Imagine that we have one thread completing a transaction in which we finished deleting a block_group and are thus calling btrfs_clear_space_info_full() while simultaneously the data reclaim ticket infrastructure is running do_async_reclaim_data_space(): T1 T2 btrfs_commit_transaction btrfs_clear_space_info_full data_sinfo->full = 0 READ: full:0, chunk_alloc:0, flush:1 do_async_reclaim_data_space(data_sinfo) spin_lock(&space_info->lock); if(list_empty(tickets)) space_info->flush = 0; READ: full: 0, chunk_alloc:0, flush:1 MOD/WRITE: full: 0, chunk_alloc:0, flush:0 spin_unlock(&space_info->lock); return; MOD/WRITE: full:0, chunk_alloc:0, flush:1 and now data_sinfo->flush is 1 but the reclaim worker has exited. This breaks the invariant that flush is 0 iff there is no work queued or running. Once this invariant is violated, future allocations that go into __reserve_bytes() will add tickets to space_info->tickets but will see space_info->flush is set to 1 and not queue the work. After this, they will block forever on the resulting ticket, as it is now impossible to kick the worker again. I also confirmed by looking at the assembly of the affected kernel that it is doing RMW operations. For example, to set the flush (3rd) bit to 0, the assembly is: andb $0xfb,0x60(%rbx) and similarly for setting the full (1st) bit to 0: andb $0xfe,-0x20(%rax) So I think this is really a bug on practical systems. I have observed a number of systems in this exact state, but am currently unable to reproduce it. Rather than leaving this footgun lying around for the future, take advantage of the fact that there is room in the struct anyway, and that it is already quite large and simply change the three bitfield members to bools. This avoids writes to space_info->full having any effect on ---truncated---	N/A	More Details
CVE-2023-54102	In the Linux kernel, the following vulnerability has been resolved: scsi: lpfc: Prevent lpfc_debugfs_lockstat_write() buffer overflow A static code analysis tool flagged the possibility of buffer overflow when using copy_from_user() for a debugfs entry. Currently, it is possible that copy_from_user() copies more bytes than what would fit in the mybuf char array. Add a min() restriction check between sizeof(mybuf) - 1 and nbytes passed from the userspace buffer to protect against buffer overflow.	N/A	More Details

CVE-2023-54103	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.	N/A	More Details
CVE-2023-54098	In the Linux kernel, the following vulnerability has been resolved: drm/i915/gvt: fix gvt debugfs destroy When gvt debug fs is destroyed, need to have a sane check if drm minor's debugfs root is still available or not, otherwise in case like device remove through unbinding, drm minor's debugfs directory has already been removed, then intel_gvt_debugfs_clean() would act upon dangling pointer like below oops. i915 0000:00:02.0: Direct firmware load for i915/gvt/vid_0x8086_did_0x1926_rid_0x0a.golden_hw_state failed with error -2 i915 0000:00:02.0: MDEV: Registered Console: switching to colour dummy device 80x25 i915 0000:00:02.0: MDEV: Unregistering BUG: kernel NULL pointer dereference, address: 00000000000000a0 PGD 0 P4D 0 Oops: 0002 [#1] PREEMPT SMP PTI CPU: 2 PID: 2486 Comm: gfx-unbind.sh Tainted: G I 6.1.0-rc8+ #15 Hardware name: Dell Inc. XPS 13 9350/0JXC1H, BIOS 1.13.0 02/10/2020 RIP: 0010:down_write+0x1f/0x90 Code: 1d ff ff 0f 1f 84 00 00 00 00 00 0f 1f 44 00 00 53 48 89 fb e8 62 c0 ff ff bf 01 00 00 00 e8 28 5e 31 ff 31 c0 ba 01 00 00 00 <f0> 48 0f b1 13 75 33 65 48 8b 04 25 c0 bd 01 00 48 89 43 08 bf 01 RSP: 0018:ffff9eb3036ffcc8 EFLAGS: 00010246 RAX: 0000000000000000 RBX: 00000000000000a0 RCX: ffffff8100000000 RDX: 0000000000000001 RSI: 0000000000000064 RDI: ffffffffa48787a8 RBP: ffff9eb3036ffd30 R08: ffffeb1fc45a0608 R09: ffffeb1fc45a05c0 R10: 0000000000000002 R11: 0000000000000000 R12: 0000000000000000 R13: ffff91acc33fa328 R14: ffff91acc033f080 R15: ffff91acced533e0 FS: 00007f6947bba740(0000) GS:ffff91ae36d00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00000000000000a0 CR3: 00000001133a2002 CR4: 0000000003706e0 Call Trace: <TASK> simple_recursive_removal+0x9f/0x2a0 ? start_creating.part.0+0x120/0x120 ? _raw_spin_lock+0x13/0x40 debugfs_remove+0x40/0x60 intel_gvt_debugfs_clean+0x15/0x30 [kvmgt] intel_gvt_clean_device+0x49/0xe0 [kvmgt] intel_gvt_driver_remove+0x2f/0xb0 i915_driver_remove+0xa4/0xf0 i915_pci_remove+0x1a/0x30 pci_device_remove+0x33/0xa0 device_release_driver_internal+0x1b2/0x230 unbind_store+0xe0/0x110 kernfs_fop_write_iter+0x11b/0x1f0 vfs_write+0x203/0x3d0 ksys_write+0x63/0xe0 do_syscall_64+0x37/0x90 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7f6947cb5190 Code: 40 00 48 8b 15 71 9c 0d 00 f7 d8 64 89 02 48 c7 c0 ff ff ff ff eb b7 0f 1f 00 80 3d 51 24 0e 00 00 74 17 b8 01 00 00 00 0f 05 <48> 3d 00 f0 ff ff 77 58 c3 0f 1f 80 00 00 00 00 48 83 ec 28 48 89 RSP: 002b:00007ffcbac45a28 EFLAGS: 00000202 ORIG_RAX: 0000000000000001 RAX: ffffffffffffffda RBX: 000000000000000d RCX: 00007f6947cb5190 RDX: 000000000000000d RSI: 0000555e35c866a0 RDI: 0000000000000001 RBP: 0000555e35c866a0 R08: 0000000000000002 R09: 0000555e358cb97c R10: 0000000000000000 R11: 0000000000000202 R12: 0000000000000001 R13: 000000000000000d R14: 0000000000000000 R15: 0000555e358cb8e0 </TASK> Modules linked in: kvmgt CR2: 00000000000000a0 ---[end trace 0000000000000000]---	N/A	More Details
CVE-2025-69201	Tugtainer is a self-hosted app for automating updates of docker containers. In versions prior to 1.15.1, arbitary arguments can be injected in tugtainer-agent `POST api/command/run`. Version 1.15.1 fixes the issue.	N/A	More Details
CVE-2025-68367	In the Linux kernel, the following vulnerability has been resolved: macintosh/mac_hid: fix race condition in mac_hid_toggle_emumouse The following warning appears when running syzkaller, and this issue also exists in the mainline code. ----------[cut here]------------ list_add double add: new=ffffffffa57eee28, prev=ffffffffa57eee28, next=ffffffffa5e63100. WARNING: CPU: 0 PID: 1491 at lib/list_debug.c:35 __list_add_valid_or_report+0xf7/0x130 Modules linked in: CPU: 0 PID: 1491 Comm: syz.1.28 Not tainted 6.6.0+ #3 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 RIP: 0010:__list_add_valid_or_report+0xf7/0x130 RSP: 0018:ff1100010dfb7b78 EFLAGS: 00010282 RAX: 0000000000000000 RBX: ffffffffa57eee18 RCX: ffffffff97fc9817 RDX: 0000000000040000 RSI: ffa0000002383000 RDI: 0000000000000001 RBP: ffffffffa57eee28 R08: 0000000000000001 R09: ffe21c0021bf6f2c R10: 0000000000000001 R11: 6464615f7473696c R12: ffffffffa5e63100 R13: ffffffffa57eee28 R14: ffffffffa57eee28 R15: ff1100010dfb7d48 FS: 00007fb14398b640(0000) GS:ff11000119600000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000000 CR3: 000000010d096005 CR4: 0000000000773ef0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 80000000 Call Trace: <TASK> input_register_handler+0xb3/0x210 mac_hid_start_emulation+0x1c5/0x290 mac_hid_toggle_emumouse+0x20a/0x240 proc_sys_call_handler+0x4c2/0x6e0 new_sync_write+0x1b1/0x2d0 vfs_write+0x709/0x950 ksys_write+0x12a/0x250 do_syscall_64+0x5a/0x110 entry_SYSCALL_64_after_hwframe+0x78/0xe2 The WARNING occurs when two processes concurrently write to the mac-hid emulation sysctl, causing a race condition in mac_hid_toggle_emumouse(). Both processes read old_val=0, then both try to register the input handler, leading to a double list_add of the same handler. CPU0 CPU1 ---------- ---------- vfs_write() //write 1 vfs_write() //write 1 proc_sys_write() proc_sys_write() mac_hid_toggle_emumouse() mac_hid_toggle_emumouse() old_val = *valp // old_val=0 old_val = *valp // old_val=0 mutex_lock_killable() proc_dointvec() // *valp=1 mac_hid_start_emulation() input_register_handler() mutex_unlock() mutex_lock_killable() proc_dointvec() mac_hid_start_emulation() input_register_handler() //Trigger Warning mutex_unlock() Fix this by moving the old_val read inside the mutex lock region.	N/A	More Details
CVE-2023-54097	In the Linux kernel, the following vulnerability has been resolved: regulator: stm32-pwr: fix of_iomap leak Smatch reports: drivers/regulator/stm32-pwr.c:166 stm32_pwr_regulator_probe() warn: 'base' from of_iomap() not released on lines: 151,166. In stm32_pwr_regulator_probe(), base is not released when devm_kzalloc() fails to allocate memory or devm_regulator_register() fails to register a new regulator device, which may cause a leak. To fix this issue, replace of_iomap() with devm_platform_ioremap_resource(). devm_platform_ioremap_resource() is a specialized function for platform devices. It allows 'base' to be automatically released whether the probe function succeeds or fails. Besides, use IS_ERR(base) instead of !base as the return value of devm_platform_ioremap_resource() can either be a pointer to the remapped memory or an ERR_PTR() encoded error code if the operation fails.	N/A	More Details
CVE-2025-68377	In the Linux kernel, the following vulnerability has been resolved: ns: initialize ns_list_node for initial namespaces Make sure that the list is always initialized for initial namespaces.	N/A	More Details
CVE-2025-68376	In the Linux kernel, the following vulnerability has been resolved: coresight: ETR: Fix ETR buffer use-after-free issue When ETR is enabled as CS_MODE_SYSFS, if the buffer size is changed and enabled again, currently sysfs_buf will point to the newly allocated memory(buf_new) and free the old memory(buf_old). But the etr_buf that is being used by the ETR remains pointed to buf_old, not updated to buf_new. In this case, it will result in a memory use-after-free issue. Fix this by checking ETR's mode before updating and releasing buf_old, if the mode is CS_MODE_SYSFS, then skip updating and releasing it.	N/A	More Details
CVE-2025-57460	File upload vulnerability in machsol machpanel 8.0.32 allows attacker to gain a webshell.	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: perf/x86: Fix NULL event access and potential PEBS record loss		

CVE-2025-68375	When intel_pmu_drain_pebs_icl() is called to drain PEBS records, the perf_event_overflow() could be called to process the last PEBS record. While perf_event_overflow() could trigger the interrupt throttle and stop all events of the group, like what the below call-chain shows. perf_event_overflow() -> __perf_event_overflow() ->__perf_event_account_interrupt() -> perf_event_throttle_group() -> perf_event_throttle() -> event->pmu->stop() -> x86_pmu_stop() The side effect of stopping the events is that all corresponding event pointers in cpuc->events[] array are cleared to NULL. Assume there are two PEBS events (event a and event b) in a group. When intel_pmu_drain_pebs_icl() calls perf_event_overflow() to process the last PEBS record of PEBS event a, interrupt throttle is triggered and all pointers of event a and event b are cleared to NULL. Then intel_pmu_drain_pebs_icl() tries to process the last PEBS record of event b and encounters NULL pointer access. To avoid this issue, move cpuc->events[] clearing from x86_pmu_stop() to x86_pmu_del(). It's safe since cpuc->active_mask or cpuc->pebs_enabled is always checked before access the event pointer from cpuc->events[].	N/A	More Details
CVE-2025-68374	In the Linux kernel, the following vulnerability has been resolved: md: fix rcu protection in md_wakeup_thread We attempted to use RCU to protect the pointer 'thread', but directly passed the value when calling md_wakeup_thread(). This means that the RCU pointer has been acquired before rcu_read_lock(), which renders rcu_read_lock() ineffective and could lead to a use-after-free.	N/A	More Details
CVE-2023-54093	In the Linux kernel, the following vulnerability has been resolved: media: anysee: fix null-ptr-deref in anysee_master_xfer In anysee_master_xfer, msg is controlled by user. When msg[i].buf is null and msg[i].len is zero, former checks on msg[i].buf would be passed. Malicious data finally reach anysee_master_xfer. If accessing msg[i].buf[0] without sanity check, null ptr deref would happen. We add check on msg[i].len to prevent crash. Similar commit: commit 0ed554fd769a ("media: dvb-usb: az6027: fix null-ptr-deref in az6027_i2c_xfer()") [hverkuil: add spaces around +]	N/A	More Details
CVE-2025-68373	In the Linux kernel, the following vulnerability has been resolved: md: avoid repeated calls to del_gendisk There is a uaf problem which is found by case 23rdev-lifetime: Oops: general protection fault, probably for non-canonical address 0xdead000000000122 RIP: 0010:bdi_unregister+0x4b/0x170 Call Trace: <TASK> __del_gendisk+0x356/0x3e0 mddev_unlock+0x351/0x360 rdev_attr_store+0x217/0x280 kernfs_fop_write_iter+0x14a/0x210 vfs_write+0x29e/0x550 ksys_write+0x74/0xf0 do_syscall_64+0xbb/0x380 entry_SYSCALL_64_after_hwframe+0x77/0x7f RIP: 0033:0x7ff5250a177e The sequence is: 1. rdev remove path gets reconfig_mutex 2. rdev remove path release reconfig_mutex in mddev_unlock 3. md stop calls do_md_stop and sets MD_DELETED 4. rdev remove path calls del_gendisk because MD_DELETED is set 5. md stop path release reconfig_mutex and calls del_gendisk again So there is a race condition we should resolve. This patch adds a flag MD_DO_DELETE to avoid the race condition.	N/A	More Details
CVE-2023-54094	In the Linux kernel, the following vulnerability has been resolved: net: prevent skb corruption on frag list segmentation Ian reported several skb corruptions triggered by rx-gro-list, collecting different oops alike: [62.624003] BUG: kernel NULL pointer dereference, address: 00000000000000c0 [62.631083] #PF: supervisor read access in kernel mode [62.636312] #PF: error_code(0x0000) - not-present page [62.641541] PGD 0 P4D 0 [62.644174] Oops: 0000 [#1] PREEMPT SMP NOPTI [62.648629] CPU: 1 PID: 913 Comm: napi/eno2-79 Not tainted 6.4.0 #364 [62.655162] Hardware name: Supermicro Super Server/A2SDi-12C-HLN4F, BIOS 1.7a 10/13/2022 [62.663344] RIP: 0010:__udp_gso_segment (./include/linux/skbuff.h:2858 ./include/linux/udp.h:23 net/ipv4/udp_offload.c:228 net/ipv4/udp_offload.c:261 net/ipv4/udp_offload.c:277) [62.687193] RSP: 0018:ffffbd3a83b4f868 EFLAGS: 00010246 [62.692515] RAX: 00000000000000ce RBX: 0000000000000000 RCX: 0000000000000000 [62.699743] RDX: ffffa124def8a000 RSI: 0000000000000079 RDI: ffffa125952a14d4 [62.706970] RBP: ffffa124def8a000 R08: 0000000000000022 R09: 00002000001558c9 [62.714199] R10: 0000000000000000 R11: 00000000be554639 R12: 00000000000000e2 [62.721426] R13: ffffa125952a1400 R14: ffffa125952a1400 R15: 00002000001558c9 [62.728654] FS: 0000000000000000(0000) GS:ffffa127efa40000(0000) knlGS:0000000000000000 [62.736852] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [62.742702] CR2: 00000000000000c0 CR3: 00000001034b0000 CR4: 00000000003526e0 [62.749948] Call Trace: [62.752498] <TASK> [62.779267] inet_gso_segment (net/ipv4/af_inet.c:1398) [62.787605] skb_mac_gso_segment (net/core/gro.c:141) [62.791906] __skb_gso_segment (net/core/dev.c:3403 (discriminator 2)) [62.800492] validate_xmit_skb (./include/linux/netdevice.h:4862 net/core/dev.c:3659) [62.804695] validate_xmit_skb_list (net/core/dev.c:3710) [62.809158] sch_direct_xmit (net/sched/sch_generic.c:330) [62.813198] __dev_queue_xmit (net/core/dev.c:3805 net/core/dev.c:4210) net/netfilter/core.c:626) [62.821093] br_dev_queue_push_xmit (net/bridge/br_forward.c:55) [62.825652] maybe_deliver (net/bridge/br_forward.c:193) [62.829420] br_flood (net/bridge/br_forward.c:233) [62.832758] br_handle_frame_finish (net/bridge/br_input.c:215) [62.837403] br_handle_frame (net/bridge/br_input.c:298 net/bridge/br_input.c:416) [62.851417] __netif_receive_skb_core.constprop.0 (net/core/dev.c:5387) [62.866114] __netif_receive_skb_list_core (net/core/dev.c:5570) [62.871367] netif_receive_skb_list_internal (net/core/dev.c:5638 net/core/dev.c:5727) [62.876795] napi_complete_done (./include/linux/list.h:37 ./include/net/gro.h:434 ./include/net/gro.h:429 net/core/dev.c:6067) [62.881004] ixgbe_poll (drivers/net/ethernet/intel/ixgbe/ixgbe_main.c:3191) [62.893534] __napi_poll (net/core/dev.c:6498) [62.897133] napi_threaded_poll (./include/linux/netpoll.h:89 net/core/dev.c:6640) [62.905276] kthread (kernel/kthread.c:379) [62.913435] ret_from_fork (arch/x86/entry/entry_64.S:314) [62.917119] </TASK> In the critical scenario, rx-gro-list GRO-ed packets are fed, via a bridge, both to the local input path and to an egress device (tun). The segmentation of such packets unsafely writes to the cloned skbs with shared heads. This change addresses the issue by uncloning as needed the to-be-segmented skbs.	N/A	More Details
CVE-2025-65570	A type confusion in jsish 2.0 allows incorrect control flow during execution of the OP_NEXT opcode. When an “instanceof” expression uses an array element access as the left-hand operand inside a for-in loop, the instructions implementation leaves an additional array reference on the stack rather than consuming it during OP_INSTANCEOF. As a result, OP_NEXT interprets the array as an iterator object and reads the iterCmd function pointer from an invalid structure, potentially causing a crash or enabling code execution depending on heap layout.	N/A	More Details
CVE-2023-54095	In the Linux kernel, the following vulnerability has been resolved: powerpc/iommu: Fix notifiers being shared by PCI and VIO buses fail_iommu_setup() registers the fail_iommu_bus_notifier struct to both PCI and VIO buses. struct notifier_block is a linked list node, so this causes any notifiers later registered to either bus type to also be registered to the other since they share the same node. This causes issues in (at least) the vgaarb code, which registers a notifier for PCI buses. pci_notify() ends up being called on a vio device, converted with to_pci_dev() even though it's not a PCI device, and finally makes a bad access in vga_arbiter_add_pci_device() as discovered with KASAN: BUG: KASAN: slab-out-of-bounds in vga_arbiter_add_pci_device+0x60/0xe00 Read of size 4 at addr c000000264c26fdc by task swapper/0/1 Call Trace: dump_stack_lvl+0x1bc/0x2b8 (unreliable) print_report+0x3f4/0xc60 kasan_report+0x244/0x698 __asan_load4+0xe8/0x250 vga_arbiter_add_pci_device+0x60/0xe00 pci_notify+0x88/0x444 notifier_call_chain+0x104/0x320 blocking_notifier_call_chain+0xa0/0x140 device_add+0xac8/0x1d30 device_register+0x58/0x80 vio_register_device_node+0x9ac/0xce0 vio_bus_scan_register_devices+0xc4/0x13c __machine_initcall_pseries_vio_device_init+0x94/0xf0 do_one_initcall+0x12c/0xaa8 kernel_init_freeable+0xa48/0xba8 kernel_init+0x64/0x400 ret_from_kernel_thread+0x5c/0x64 Fix this by creating separate notifier_block structs for each bus type. [mpe: Add #ifdef to fix CONFIG_IBMVIO=n build]	N/A	More Details
	In the Linux kernel, the following vulnerability has been resolved: nbd: defer config put in recv_work There is one uaf issue in recv_work when running NBD_CLEAR_SOCK and NBD_CMD_RECONFIGURE: nbd_genl_connect // conf_ref=2 (connect and recv_work A) nbd_open // conf_ref=3 recv_work A done // conf_ref=2 NBD_CLEAR_SOCK // conf_ref=1 nbd_genl_reconfigure // conf_ref=2		

CVE-2025-68372	(trigger recv_work B) close nbd // conf_ref=1 recv_work B config_put // conf_ref=0 atomic_dec(&config->recv_threads); -> UAF Or only running NBD_CLEAR_SOCK: nbd_genl_connect // conf_ref=2 nbd_open // conf_ref=3 NBD_CLEAR_SOCK // conf_ref=2 close nbd nbd_release config_put // conf_ref=1 recv_work config_put // conf_ref=0 atomic_dec(&config->recv_threads); -> UAF Commit 87aac3a80af5 ("nbd: call nbd_config_put() before notifying the waiter") moved nbd_config_put() to run before waking up the waiter in recv_work, in order to ensure that nbd_start_device_ioctl() would not be woken up while nbd->task_recv was still uncleared. However, in nbd_start_device_ioctl(), after being woken up it explicitly calls flush_workqueue() to make sure all current works are finished. Therefore, there is no need to move the config put ahead of the wakeup. Move nbd_config_put() to the end of recv_work, so that the reference is held for the whole lifetime of the worker thread. This makes sure the config cannot be freed while recv_work is still running, even if clear + reconfigure interleave. In addition, we don't need to worry about recv_work dropping the last nbd_put (which causes deadlock): path A (netlink with NBD_CFLAG_DESTROY_ON_DISCONNECT): connect // nbd_refs=1 (trigger recv_work) open nbd // nbd_refs=2 NBD_CLEAR_SOCK close nbd nbd_release nbd_disconnect_and_put flush_workqueue // recv_work done nbd_config_put nbd_put // nbd_refs=1 nbd_put // nbd_refs=0 queue_work path B (netlink without NBD_CFLAG_DESTROY_ON_DISCONNECT): connect // nbd_refs=2 (trigger recv_work) open nbd // nbd_refs=3 NBD_CLEAR_SOCK // conf_refs=2 close nbd nbd_release nbd_config_put // conf_refs=1 nbd_put // nbd_refs=2 recv_work done // conf_refs=0, nbd_refs=1 rmmod // nbd_refs=0 Depends-on: e2daec488c57 ("nbd: Fix hungtask when nbd_config_put")	N/A	More Details
CVE-2025-56333	An issue in Fossorial fosrl/pangolin v.1.6.2 and before allows a remote attacker to escalate privileges via the 2FA component	N/A	More Details
CVE-2025-68371	In the Linux kernel, the following vulnerability has been resolved: scsi: smartpqi: Fix device resources accessed after device removal Correct possible race conditions during device removal. Previously, a scheduled work item to reset a LUN could still execute after the device was removed, leading to use-after-free and other resource access issues. This race condition occurs because the abort handler may schedule a LUN reset concurrently with device removal via sdev_destroy(), leading to use-after-free and improper access to freed resources. - Check in the device reset handler if the device is still present in the controller's SCSI device list before running; if not, the reset is skipped. - Cancel any pending TMF work that has not started in sdev_destroy(). - Ensure device freeing in sdev_destroy() is done while holding the LUN reset mutex to avoid races with ongoing resets.	N/A	More Details
CVE-2025-68370	In the Linux kernel, the following vulnerability has been resolved: coresight: tmc: add the handle of the event to the path The handle is essential for retrieving the AUX_EVENT of each CPU and is required in perf mode. It has been added to the coresight_path so that dependent devices can access it from the path when needed. The existing bug can be reproduced with: perf record -e cs_etm//k -C 0-9 dd if=/dev/zero of=/dev/null Showing an oops as follows: Unable to handle kernel paging request at virtual address 000f6e84934ed19e Call trace: tmc_etr_get_buffer+0x30/0x80 [coresight_tmc] (P) catu_enable_hw+0xbc/0x3d0 [coresight_catu] catu_enable+0x70/0xe0 [coresight_catu] coresight_enable_path+0xb0/0x258 [coresight]	N/A	More Details
CVE-2025-68369	In the Linux kernel, the following vulnerability has been resolved: ntfs3: init run lock for extend inode After setting the inode mode of $Extend to a regular file, executing the truncate system call will enter the do_truncate() routine, causing the run_lock uninitialized error reported by syzbot. Prior to patch 4e8011ffec79, if the inode mode of $Extend was not set to a regular file, the do_truncate() routine would not be entered. Add the run_lock initialization when loading $Extend. syzbot reported: INFO: trying to register non-static key. Call Trace: dump_stack_lvl+0x189/0x250 lib/dump_stack.c:120 assign_lock_key+0x133/0x150 kernel/locking/lockdep.c:984 register_lock_class+0x105/0x320 kernel/locking/lockdep.c:1299 __lock_acquire+0x99/0xd20 kernel/locking/lockdep.c:5112 lock_acquire+0x120/0x360 kernel/locking/lockdep.c:5868 down_write+0x96/0x1f0 kernel/locking/rwsem.c:1590 ntfs_set_size+0x140/0x200 fs/ntfs3/inode.c:860 ntfs_extend+0x1d9/0x970 fs/ntfs3/file.c:387 ntfs_setattr+0x2e8/0xbe0 fs/ntfs3/file.c:808	N/A	More Details
CVE-2025-68368	In the Linux kernel, the following vulnerability has been resolved: md: init bioset in mddev_init IO operations may be needed before md_run(), such as updating metadata after writing sysfs. Without bioset, this triggers a NULL pointer dereference as below: BUG: kernel NULL pointer dereference, address: 0000000000000020 Call Trace: md_update_sb+0x658/0xe00 new_level_store+0xc5/0x120 md_attr_store+0xc9/0x1e0 sysfs_kf_write+0x6f/0xa0 kernfs_fop_write_iter+0x141/0x2a0 vfs_write+0x1fc/0x5a0 ksys_write+0x79/0x180 __x64_sys_write+0x1d/0x30 x64_sys_call+0x2818/0x2880 do_syscall_64+0xa9/0x580 entry_SYSCALL_64_after_hwframe+0x4b/0x53 Reproducer ``` mdadm -CR /dev/md0 -l1 -n2 /dev/sd[cd] echo inactive > /sys/block/md0/md/array_state echo 10 > /sys/block/md0/md/new_level ``` mddev_init() can only be called once per mddev, no need to test if bioset has been initialized anymore.	N/A	More Details
CVE-2023-54096	In the Linux kernel, the following vulnerability has been resolved: soundwire: fix enumeration completion The soundwire subsystem uses two completion structures that allow drivers to wait for soundwire device to become enumerated on the bus and initialised by their drivers, respectively. The code implementing the signalling is currently broken as it does not signal all current and future waiters and also uses the wrong reinitialisation function, which can potentially lead to memory corruption if there are still waiters on the queue. Not signalling future waiters specifically breaks sound card probe deferrals as codec drivers can not tell that the soundwire device is already attached when being reprobed. Some codec runtime PM implementations suffer from similar problems as waiting for enumeration during resume can also timeout despite the device already having been enumerated.	N/A	More Details
CVE-2023-54104	In the Linux kernel, the following vulnerability has been resolved: mtd: rawnand: fsl_upm: Fix an off-by one test in fun_exec_op() 'op->cs' is copied in 'fun->mchip_number' which is used to access the 'mchip_offsets' and the 'rnb_gpio' arrays. These arrays have NAND_MAX_CHIPS elements, so the index must be below this limit. Fix the sanity check in order to avoid the NAND_MAX_CHIPS value. This would lead to out-of-bound accesses.	N/A	More Details
CVE-2025-68356	In the Linux kernel, the following vulnerability has been resolved: gfs2: Prevent recursive memory reclaim Function new_inode() returns a new inode with inode->i_mapping->gfp_mask set to GFP_HIGHUSER_MOVABLE. This value includes the __GFP_FS flag, so allocations in that address space can recurse into filesystem memory reclaim. We don't want that to happen because it can consume a significant amount of stack memory. Worse than that is that it can also deadlock: for example, in several places, gfs2_unstuff_dinode() is called inside filesystem transactions. This calls filemap_grab_folio(), which can allocate a new folio, which can trigger memory reclaim. If memory reclaim recurses into the filesystem and starts another transaction, a deadlock will ensue. To fix these kinds of problems, prevent memory reclaim from recursing into filesystem code by making sure that the gfp_mask of inode address spaces doesn't include __GFP_FS. The "meta" and resource group address spaces were already using GFP_NOFS as their gfp_mask (which doesn't include __GFP_FS). The default value of GFP_HIGHUSER_MOVABLE is less restrictive than GFP_NOFS, though. To avoid being overly limiting, use the default value and only knock off the __GFP_FS flag. I'm not sure if this will actually make a difference, but it also shouldn't hurt. This patch is loosely based on commit ad22c7a043c2 ("xfs: prevent stack overflows from page cache allocation"). Fixes xfstest generic/273.	N/A	More Details
CVE-2023-	In the Linux kernel, the following vulnerability has been resolved: usb: rndis_host: Secure rndis_query check against int overflow Variables off and len typed as uint32 in rndis_query function are controlled by incoming RNDIS response message thus their value may be manipulated. Setting off to a unexpectetly large value will cause the sum with len and 8 to overflow and pass the	N/A	More Details

54110	implemented validation step. Consequently the response pointer will be referring to a location past the expected buffer boundaries allowing information leakage e.g. via RNDIS_OID_802_3_PERMANENT_ADDRESS OID.		
CVE-2023-54037	In the Linux kernel, the following vulnerability has been resolved: ice: prevent NULL pointer deref during reload Calling ethtool during reload can lead to call trace, because VSI isn't configured for some time, but netdev is alive. To fix it add rtnl lock for VSI deconfig and config. Set ::num_q_vectors to 0 after freeing and add a check for ::tx/rx_rings in ring related ethtool ops. Add proper unroll of filters in ice_start_eth(). Reproduction: $watch -n 0.1 -d 'ethtool -g enp24s0f0np0' $devlink dev reload pci/0000:18:00.0 action driver_reinit Call trace before fix: [66303.926205] BUG: kernel NULL pointer dereference, address: 0000000000000000 [66303.926259] #PF: supervisor read access in kernel mode [66303.926286] #PF: error_code(0x0000) - not-present page [66303.926311] PGD 0 P4D 0 [66303.926332] Oops: 0000 [#1] PREEMPT SMP PTI [66303.926358] CPU: 4 PID: 933821 Comm: ethtool Kdump: loaded Tainted: G OE 6.4.0-rc5+ #1 [66303.926400] Hardware name: Intel Corporation S2600WFT/S2600WFT, BIOS SE5C620.86B.00.01.0014.070920180847 07/09/2018 [66303.926446] RIP: 0010:ice_get_ringparam+0x22/0x50 [ice] [66303.926649] Code: 90 90 90 90 90 90 90 90 f3 0f 1e fa 0f 1f 44 00 00 48 8b 87 c0 09 00 00 c7 46 04 e0 1f 00 00 c7 46 10 e0 1f 00 00 48 8b 50 20 <48> 8b 12 0f b7 52 3a 89 56 14 48 8b 40 28 48 8b 00 0f b7 40 58 48 [66303.926722] RSP: 0018:ffffad40472f39c8 EFLAGS: 00010246 [66303.926749] RAX: ffff98a8ada05828 RBX: ffff98a8c46dd060 RCX: ffffad40472f3b48 [66303.926781] RDX: 0000000000000000 RSI: ffff98a8c46dd068 RDI: ffff98a8b23c4000 [66303.926811] RBP: ffffad40472f3b48 R08: 00000000000337b0 R09: 0000000000000000 [66303.926843] R10: 0000000000000001 R11: 0000000000000100 R12: ffff98a8b23c4000 [66303.926874] R13: ffff98a8c46dd060 R14: 000000000000000f R15: ffffad40472f3a50 [66303.926906] FS: 00007f6397966740(0000) GS:ffff98b390900000(0000) knlGS:0000000000000000 [66303.926941] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [66303.926967] CR2: 0000000000000000 CR3: 000000011ac20002 CR4: 00000000007706e0 [66303.926999] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [66303.927029] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [66303.927060] PKRU: 55555554 [66303.927075] Call Trace: [66303.927094] <TASK> [66303.927111] ? __die+0x23/0x70 [66303.927140] ? page_fault_oops+0x171/0x4e0 [66303.927176] ? exc_page_fault+0x7f/0x180 [66303.927209] ? asm_exc_page_fault+0x26/0x30 [66303.927244] ? ice_get_ringparam+0x22/0x50 [ice] [66303.927433] rings_prepare_data+0x62/0x80 [66303.927469] ethnl_default_doit+0xe2/0x350 [66303.927501] genl_family_rcv_msg_doit.isra.0+0xe3/0x140 [66303.927538] genl_rcv_msg+0x1b1/0x2c0 [66303.927561] ? __pfx_ethnl_default_doit+0x10/0x10 [66303.927590] ? __pfx_genl_rcv_msg+0x10/0x10 [66303.927615] netlink_rcv_skb+0x58/0x110 [66303.927644] genl_rcv+0x28/0x40 [66303.927665] netlink_unicast+0x19e/0x290 [66303.927691] netlink_sendmsg+0x254/0x4d0 [66303.927717] sock_sendmsg+0x93/0xa0 [66303.927743] __sys_sendto+0x126/0x170 [66303.927780] __x64_sys_sendto+0x24/0x30 [66303.928593] do_syscall_64+0x5d/0x90 [66303.929370] ? __count_memcg_events+0x60/0xa0 [66303.930146] ? count_memcg_events.constprop.0+0x1a/0x30 [66303.930920] ? handle_mm_fault+0x9e/0x350 [66303.931688] ? do_user_addr_fault+0x258/0x740 [66303.932452] ? exc_page_fault+0x7f/0x180 [66303.933193] entry_SYSCALL_64_after_hwframe+0x72/0xdc	N/A	More Details
CVE-2025-68344	In the Linux kernel, the following vulnerability has been resolved: ALSA: wavefront: Fix integer overflow in sample size validation The wavefront_send_sample() function has an integer overflow issue when validating sample size. The header->size field is u32 but gets cast to int for comparison with dev->freemem Fix by using unsigned comparison to avoid integer overflow.	N/A	More Details
CVE-2023-54042	In the Linux kernel, the following vulnerability has been resolved: powerpc/64s: Fix VAS mm use after free The refcount on mm is dropped before the coprocessor is detached.	N/A	More Details
CVE-2023-54041	In the Linux kernel, the following vulnerability has been resolved: io_uring: fix memory leak when removing provided buffers When removing provided buffers, io_buffer structs are not being disposed of, leading to a memory leak. They can't be freed individually, because they are allocated in page-sized groups. They need to be added to some free list instead, such as io_buffers_cache. All callers already hold the lock protecting it, apart from when destroying buffers, so had to extend the lock there.	N/A	More Details
CVE-2023-54040	In the Linux kernel, the following vulnerability has been resolved: ice: fix wrong fallback logic for FDIR When adding a FDIR filter, if ice_vc_fdir_set_irq_ctx returns failure, the inserted fdir entry will not be removed and if ice_vc_fdir_write_fltr returns failure, the fdir context info for irq handler will not be cleared which may lead to inconsistent or memory leak issue. This patch refines failure cases to resolve this issue.	N/A	More Details
CVE-2023-54108	In the Linux kernel, the following vulnerability has been resolved: scsi: qla2xxx: Fix DMA-API call trace on NVMe LS requests The following message and call trace was seen with debug kernels: DMA-API: qla2xxx 0000:41:00.0: device driver failed to check map error [device address=0x00000002a3ff38d8] [size=1024 bytes] [mapped as single] WARNING: CPU: 0 PID: 2930 at kernel/dma/debug.c:1017 check_unmap+0xf42/0x1990 Call Trace: debug_dma_unmap_page+0xc9/0x100 qla_nvme_ls_unmap+0x141/0x210 [qla2xxx] Remove DMA mapping from the driver altogether, as it is already done by FC layer. This prevents the warning.	N/A	More Details
CVE-2023-54121	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix incorrect splitting in btrfs_drop_extent_map_range In production we were seeing a variety of WARN_ON()'s in the extent_map code, specifically in btrfs_drop_extent_map_range() when we have to call add_extent_mapping() for our second split. Consider the following extent map layout PINNED [0 16K) [32K, 48K) and then we call btrfs_drop_extent_map_range for [0, 36K), with skip_pinned == true. The initial loop will have start = 0 end = 36K len = 36K we will find the [0, 16k) extent, but since we are pinned we will skip it, which has this code start = em_end; if (end != (u64)-1) len = start + len - em_end; em_end here is 16K, so now the values are start = 16K len = 16K + 36K - 16K = 36K len should instead be 20K. This is a problem when we find the next extent at [32K, 48K), we need to split this extent to leave [36K, 48k), however the code for the split looks like this split->start = start + len; split->len = em_end - (start + len); In this case we have em_end = 48K split->start = 16K + 36K // this should be 16K + 20K split->len = 48K - (16K + 36K) // this overflows as 16K + 36K is 52K and now we have an invalid extent_map in the tree that potentially overlaps other entries in the extent map. Even in the non-overlapping case we will have split->start set improperly, which will cause problems with any block related calculations. We don't actually need len in this loop, we can simply use end as our end point, and only adjust start up when we find a pinned extent we need to skip. Adjust the logic to do this, which keeps us from inserting an invalid extent map. We only skip_pinned in the relocation case, so this is relatively rare, except in the case where you are running relocation a lot, which can happen with auto relocation on.	N/A	More Details
CVE-2023-54039	In the Linux kernel, the following vulnerability has been resolved: can: j1939: j1939_tp_tx_dat_new(): fix out-of-bounds memory access In the j1939_tp_tx_dat_new() function, an out-of-bounds memory access could occur during the memcpy() operation if the size of skb->cb is larger than the size of struct j1939_sk_buff_cb. This is because the memcpy() operation uses the size of skb->cb, leading to a read beyond the struct j1939_sk_buff_cb. Updated the memcpy() operation to use the size of struct j1939_sk_buff_cb instead of the size of skb->cb. This ensures that the memcpy() operation only reads the memory within the bounds of struct j1939_sk_buff_cb, preventing out-of-bounds memory access. Additionally, add a BUILD_BUG_ON() to check that the size of skb->cb is greater than or equal to the size of struct j1939_sk_buff_cb. This ensures that the skb->cb buffer is large enough to hold the j1939_sk_buff_cb structure. [mkl: rephrase commit message]	N/A	More Details

CVE-2023-54038	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_conn: return ERR_PTR instead of NULL when there is no link hci_connect_sco currently returns NULL when there is no link (i.e. when hci_conn_link() returns NULL). sco_connect() expects an ERR_PTR in case of any error (see line 266 in sco.c). Thus, hcon set as NULL passes through to sco_conn_add(), which tries to get hcon->hdev, resulting in dereferencing a NULL pointer as reported by syzkaller. The same issue exists for iso_connect_cis() calling hci_connect_cis(). Thus, make hci_connect_sco() and hci_connect_cis() return ERR_PTR instead of NULL.	N/A	More Details
CVE-2023-54036	In the Linux kernel, the following vulnerability has been resolved: wifi: rtl8xxxu: Fix memory leaks with RTL8723BU, RTL8192EU The wifi + bluetooth combo chip RTL8723BU can leak memory (especially?) when it's connected to a bluetooth audio device. The busy bluetooth traffic generates lots of C2H (card to host) messages, which are not freed correctly. To fix this, move the dev_kfree_skb() call in rtl8xxxu_c2hcmd_callback() inside the loop where skb_dequeue() is called. The RTL8192EU leaks memory because the C2H messages are added to the queue and left there forever. (This was fine in the past because it probably wasn't sending any C2H messages until commit e542e66b7c2e ("wifi: rtl8xxxu: gen2: Turn on the rate control"). Since that commit it sends a C2H message when the TX rate changes.) To fix this, delete the check for rf_paths > 1 and the goto. Let the function process the C2H messages from RTL8192EU like the ones from the other chips. Theoretically the RTL8188FU could also leak like RTL8723BU, but it most likely doesn't send C2H messages frequently enough. This change was tested with RTL8723BU by Erhard F. I tested it with RTL8188FU and RTL8192EU.	N/A	More Details
CVE-2025-68355	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix exclusive map memory leak When excl_prog_hash is 0 and excl_prog_hash_size is non-zero, the map also needs to be freed. Otherwise, the map memory will not be reclaimed, just like the memory leak problem reported by syzbot [1]. syzbot reported: BUG: memory leak backtrace (crc 7b9fb9b4): map_create+0x322/0x11e0 kernel/bpf/syscall.c:1512 __sys_bpf+0x3556/0x3610 kernel/bpf/syscall.c:6131	N/A	More Details
CVE-2023-54035	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: fix underflow in chain reference counter Set element addition error path decrements reference counter on chains twice: once on element release and again via nft_data_release(). Then, d6b478666ffa ("netfilter: nf_tables: fix underflow in object reference counter") incorrectly fixed this by removing the stateful object reference count decrement. Restore the stateful object decrement as in b91d90368837 ("netfilter: nf_tables: fix leaking object reference count") and let nft_data_release() decrement the chain reference counter, so this is done only once.	N/A	More Details
CVE-2023-54034	In the Linux kernel, the following vulnerability has been resolved: iommufd: Make sure to zero vfio_iommu_type1_info before copying to user Missed a zero initialization here. Most of the struct is filled with a copy_from_user(), however minsz for that copy is smaller than the actual struct by 8 bytes, thus we don't fill the padding.	N/A	More Details
CVE-2023-54033	In the Linux kernel, the following vulnerability has been resolved: bpf: fix a memory leak in the LRU and LRU_PERCPU hash maps The LRU and LRU_PERCPU maps allocate a new element on update before locking the target hash table bucket. Right after that the maps try to lock the bucket. If this fails, then maps return -EBUSY to the caller without releasing the allocated element. This makes the element untracked: it doesn't belong to either of free lists, and it doesn't belong to the hash table, so can't be re-used; this eventually leads to the permanent -ENOMEM on LRU map updates, which is unexpected. Fix this by returning the element to the local free list if bucket locking fails.	N/A	More Details
CVE-2023-54109	In the Linux kernel, the following vulnerability has been resolved: media: rcar_fdp1: Fix refcount leak in probe and remove function rcar_fcp_get() take reference, which should be balanced with rcar_fcp_put(). Add missing rcar_fcp_put() in fdp1_remove and the error paths of fdp1_probe() to fix this. [hverkuil: resolve merge conflict, remove() is now void]	N/A	More Details
CVE-2023-54032	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix race when deleting quota root from the dirty cow roots list When disabling quotas we are deleting the quota root from the list fs_info->dirty_cowonly_roots without taking the lock that protects it, which is struct btrfs_fs_info::trans_lock. This unsynchronized list manipulation may cause chaos if there's another concurrent manipulation of this list, such as when adding a root to it with ctree.c:add_root_to_dirty_list(). This can result in all sorts of weird failures caused by a race, such as the following crash: [337571.278245] general protection fault, probably for non-canonical address 0xdead000000000108: 0000 [#1] PREEMPT SMP PTI [337571.278933] CPU: 1 PID: 115447 Comm: btrfs Tainted: G W 6.4.0-rc6-btrfs-next-134+ #1 [337571.279153] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu.org 04/01/2014 [337571.279572] RIP: 0010:commit_cowonly_roots+0x11f/0x250 [btrfs] [337571.279928] Code: 85 38 06 00 (...) [337571.280363] RSP: 0018:ffff9f63446efba0 EFLAGS: 00010206 [337571.280582] RAX: ffff942d98ec2638 RBX: ffff9430b82b4c30 RCX: 0000000449e1c000 [337571.280798] RDX: dead000000000100 RSI: ffff9430021e4900 RDI: 0000000000036070 [337571.281015] RBP: ffff942d98ec2000 R08: ffff942d98ec2000 R09: 000000000000015b [337571.281254] R10: 0000000000000009 R11: 0000000000000001 R12: ffff942fe8fbf600 [337571.281476] R13: ffff942dabe23040 R14: ffff942dabe20800 R15: ffff942d92cf3b48 [337571.281723] FS: 00007f478adb7340(0000) GS:ffff94349fa40000(0000) knlGS:0000000000000000 [337571.281950] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [337571.282184] CR2: 00007f478ab9a3d5 CR3: 000000001e02c001 CR4: 0000000000370ee0 [337571.282416] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [337571.282647] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [337571.282874] Call Trace: [337571.283101] <TASK> [337571.283327] ? __die_body+0x1b/0x60 [337571.283570] ? die_addr+0x39/0x60 [337571.283796] ? exc_general_protection+0x22e/0x430 [337571.284022] ? asm_exc_general_protection+0x22/0x30 [337571.284251] ? commit_cowonly_roots+0x11f/0x250 [btrfs] [337571.284531] btrfs_commit_transaction+0x42e/0xf90 [btrfs] [337571.284803] ? _raw_spin_unlock+0x15/0x30 [337571.285031] ? release_extent_buffer+0x103/0x130 [btrfs] [337571.285305] reset_balance_state+0x152/0x1b0 [btrfs] [337571.285578] btrfs_balance+0xa50/0x11e0 [btrfs] [337571.285864] ? __kmem_cache_alloc_node+0x14a/0x410 [337571.286086] btrfs_ioctl+0x249a/0x3320 [btrfs] [337571.286358] ? mod_objcg_state+0xd2/0x360 [337571.286577] ? refill_obj_stock+0xb0/0x160 [337571.286798] ? seq_release+0x25/0x30 [337571.287016] ? __rseq_handle_notify_resume+0x3ba/0x4b0 [337571.287235] ? percpu_counter_add_batch+0x2e/0xa0 [337571.287455] ? __x64_sys_ioctl+0x88/0xc0 [337571.287675] __x64_sys_ioctl+0x88/0xc0 [337571.287901] do_syscall_64+0x38/0x90 [337571.288126] entry_SYSCALL_64_after_hwframe+0x72/0xdc [337571.288352] RIP: 0033:0x7f478aaffe9b So fix this by locking struct btrfs_fs_info::trans_lock before deleting the quota root from that list.	N/A	More Details
CVE-2023-54031	In the Linux kernel, the following vulnerability has been resolved: vdpa: Add queue index attr to vdpa_nl_policy for nlattr length check The vdpa_nl_policy structure is used to validate the nlattr when parsing the incoming nlmsg. It will ensure the attribute being described produces a valid nlattr pointer in info->attrs before entering into each handler in vdpa_nl_ops. That is to say, the missing part in vdpa_nl_policy may lead to illegal nlattr after parsing, which could lead to OOB read just like CVE-2023-3773. This patch adds the missing nla_policy for vdpa queue index attr to avoid such bugs.	N/A	More Details
CVE-2023-54030	In the Linux kernel, the following vulnerability has been resolved: io_uring/net: don't overflow multishot recv Don't allow overflowing multishot recv CQEs, it might get out of hand, hurt performance, and in the worst case scenario OOM the task.	N/A	More Details

CVE-2023-54029	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: fix iwl_mvm_max_amsdu_size() for MLO For MLO, we cannot use vif->bss_conf.chandef.chan->band, since that will lead to a NULL-ptr dereference as bss_conf isn't used. However, in case of real MLO, we also need to take both LMACs into account if they exist, since the station might be active on both LMACs at the same time.	N/A	More Details
CVE-2025-68345	In the Linux kernel, the following vulnerability has been resolved: ALSA: hda: cs35l41: Fix NULL pointer dereference in cs35l41_hda_read_acpi() The acpi_get_first_physical_node() function can return NULL, in which case the get_device() function also returns NULL, but this value is then dereferenced without checking,so add a check to prevent a crash. Found by Linux Verification Center (linuxtesting.org) with SVACE.	N/A	More Details
CVE-2025-68346	In the Linux kernel, the following vulnerability has been resolved: ALSA: dice: fix buffer overflow in detect_stream_formats() The function detect_stream_formats() reads the stream_count value directly from a FireWire device without validating it. This can lead to out-of-bounds writes when a malicious device provides a stream_count value greater than MAX_STREAMS. Fix by applying the same validation to both TX and RX stream counts in detect_stream_formats().	N/A	More Details
CVE-2023-54107	In the Linux kernel, the following vulnerability has been resolved: blk-cgroup: dropping parent refcount after pd_free_fn() is done Some cgroup policies will access parent pd through child pd even after pd_offline_fn() is done. If pd_free_fn() for parent is called before child, then UAF can be triggered. Hence it's better to guarantee the order of pd_free_fn(). Currently refcount of parent blkg is dropped in __blkg_release(), which is before pd_free_fn() is called in blkg_free_work_fn() while blkg_free_work_fn() is called asynchronously. This patch make sure pd_free_fn() called from removing cgroup is ordered by delaying dropping parent refcount after calling pd_free_fn() for child. BTW, pd_free_fn() will also be called from blkcg_deactivate_policy() from deleting device, and following patches will guarantee the order.	N/A	More Details
CVE-2025-68347	In the Linux kernel, the following vulnerability has been resolved: ALSA: firewire-motu: fix buffer overflow in hwdep read for DSP events The DSP event handling code in hwdep_read() could write more bytes to the user buffer than requested, when a user provides a buffer smaller than the event header size (8 bytes). Fix by using min_t() to clamp the copy size, This ensures we never copy more than the user requested.	N/A	More Details
CVE-2023-54105	In the Linux kernel, the following vulnerability has been resolved: can: isotp: check CAN address family in isotp_bind() Add missing check to block non-AF_CAN binds. Syzbot created some code which matched the right sockaddr struct size but used AF_XDP (0x2C) instead of AF_CAN (0x1D) in the address family field: bind$xdp(r2, &(0x7f0000000540)={0x2c, 0x0, r4, 0x0, r2}, 0x10) ^^^^ This has no funtional impact but the userspace should be notified about the wrong address family field content.	N/A	More Details
CVE-2025-68354	In the Linux kernel, the following vulnerability has been resolved: regulator: core: Protect regulator_supply_alias_list with regulator_list_mutex regulator_supply_alias_list was accessed without any locking in regulator_supply_alias(), regulator_register_supply_alias(), and regulator_unregister_supply_alias(). Concurrent registration, unregistration and lookups can race, leading to: 1 use-after-free if an alias entry is removed while being read, 2 duplicate entries when two threads register the same alias, 3 inconsistent alias mappings observed by consumers. Protect all traversals, insertions and deletions on regulator_supply_alias_list with the existing regulator_list_mutex.	N/A	More Details
CVE-2025-68353	In the Linux kernel, the following vulnerability has been resolved: net: vxlan: prevent NULL deref in vxlan_xmit_one Neither sock4 nor sock6 pointers are guaranteed to be non-NULL in vxlan_xmit_one, e.g. if the iface is brought down. This can lead to the following NULL dereference: BUG: kernel NULL pointer dereference, address: 0000000000000010 Oops: Oops: 0000 [#1] SMP NOPTI RIP: 0010:vxlan_xmit_one+0xbb3/0x1580 Call Trace: vxlan_xmit+0x429/0x610 dev_hard_start_xmit+0x55/0xa0 __dev_queue_xmit+0x6d0/0x7f0 ip_finish_output2+0x24b/0x590 ip_output+0x63/0x110 Mentioned commits changed the code path in vxlan_xmit_one and as a side effect the sock4/6 pointer validity checks in vxlan(6)_get_route were lost. Fix this by adding back checks. Since both commits being fixed were released in the same version (v6.7) and are strongly related, bundle the fixes in a single commit.	N/A	More Details
CVE-2023-54106	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: fix potential memory leak in mlx5e_init_rep_rx The memory pointed to by the priv->rx_res pointer is not freed in the error path of mlx5e_init_rep_rx, which can lead to a memory leak. Fix by freeing the memory in the error path, thereby making the error path identical to mlx5e_cleanup_rep_rx().	N/A	More Details
CVE-2025-67254	NagiosXI 2026R1.0.1 build 1762361101 is vulnerable to Directory Traversal in /admin/coreconfigsnapshots.php.	N/A	More Details
CVE-2025-67255	In NagiosXI 2026R1.0.1 build 1762361101, Dashboard parameters lack proper filtering, allowing any authenticated user to exploit a SQL Injection vulnerability.	N/A	More Details
CVE-2025-68352	In the Linux kernel, the following vulnerability has been resolved: spi: ch341: fix out-of-bounds memory access in ch341_transfer_one Discovered by Atuin - Automated Vulnerability Discovery Engine. The 'len' variable is calculated as 'min(32, trans->len + 1)', which includes the 1-byte command header. When copying data from 'trans->tx_buf' to 'ch341->tx_buf + 1', using 'len' as the length is incorrect because: 1. It causes an out-of-bounds read from 'trans->tx_buf' (which has size 'trans->len', i.e., 'len - 1' in this context). 2. It can cause an out-of-bounds write to 'ch341->tx_buf' if 'len' is CH341_PACKET_LENGTH (32). Writing 32 bytes to ch341->tx_buf + 1 overflows the buffer. Fix this by copying 'len - 1' bytes.	N/A	More Details
CVE-2025-68706	A stack-based buffer overflow exists in the GoAhead-Webs HTTP daemon on KuWFi 4G LTE AC900 devices with firmware 1.0.13. The /goform/formMultiApnSetting handler uses sprintf() to copy the user-supplied pincode parameter into a fixed 132-byte stack buffer with no bounds checks. This allows an attacker to corrupt adjacent stack memory, crash the web server, and (under certain conditions) may enable arbitrary code execution.	N/A	More Details
CVE-2024-30855	DedeCMS v5.7 was discovered to contain a Cross-Site Request Forgery (CSRF) vulnerability via /src/dede/makehtml_list_action.php.	N/A	More Details
CVE-2025-14175	A vulnerability in the SSH server of TP-Link TL-WR820N v2.80 allows the use of a weak cryptographic algorithm, enabling an adjacent attacker to intercept and decrypt SSH traffic. Exploitation may expose sensitive information and compromise confidentiality.	N/A	More Details
	Axios Cache Interceptor is a cache interceptor for axios. Prior to version 1.11.1, when a server calls an upstream service using different auth tokens, axios-cache-interceptor returns incorrect cached responses, leading to authorization bypass. The cache key is generated only from the URL, ignoring request headers like `Authorization`. When the server responds with `Vary: Authorization`		

CVE-2025-69202	(indicating the response varies by auth token), the library ignores this, causing all requests to share the same cache regardless of authorization. Server-side applications (APIs, proxies, backend services) that use axios-cache-interceptor to cache requests to upstream services, handle requests from multiple users with different auth tokens, and upstream services replies on `Vary` to differentiate caches are affected. Browser/client-side applications (single user per browser session) are not affected. Services using different auth tokens to call upstream services will return incorrect cached data, bypassing authorization checks and leaking user data across different authenticated sessions. After `v1.11.1`, automatic `Vary` header support is now enabled by default. When server responds with `Vary: Authorization`, cache keys now include the authorization header value. Each user gets their own cache.	N/A	More Details
CVE-2024-25182	givanz VvvebJs 1.7.2 suffers from a File Upload vulnerability via save.php.	N/A	More Details
CVE-2025-68351	In the Linux kernel, the following vulnerability has been resolved: exfat: fix refcount leak in exfat_find Fix refcount leaks in `exfat_find` related to `exfat_get_dentry_set`. Function `exfat_get_dentry_set` would increase the reference counter of `es->bh` on success. Therefore, `exfat_put_dentry_set` must be called after `exfat_get_dentry_set` to ensure refcount consistency. This patch relocate two checks to avoid possible leaks.	N/A	More Details
CVE-2024-27480	givanz VvvebJs 1.7.2 is vulnerable to Insecure File Upload.	N/A	More Details
CVE-2025-68350	In the Linux kernel, the following vulnerability has been resolved: exfat: fix divide-by-zero in exfat_allocate_bitmap The variable max_ra_count can be 0 in exfat_allocate_bitmap(), which causes a divide-by-zero error in the subsequent modulo operation (i % max_ra_count), leading to a system crash. When max_ra_count is 0, it means that readahead is not used. This patch load the bitmap without readahead.	N/A	More Details
CVE-2025-68349	In the Linux kernel, the following vulnerability has been resolved: NFSv4/pNFS: Clear NFS_INO_LAYOUTCOMMIT in pnfs_mark_layout_stateid_invalid Fixes a crash when layout is null during this call stack: write_inode -> nfs4_write_inode -> pnfs_layoutcommit_inode pnfs_set_layoutcommit relies on the lseg refcount to keep the layout around. Need to clear NFS_INO_LAYOUTCOMMIT otherwise we might attempt to reference a null layout.	N/A	More Details
CVE-2025-68348	In the Linux kernel, the following vulnerability has been resolved: block: fix memory leak in __blkdev_issue_zero_pages Move the fatal signal check before bio_alloc() to prevent a memory leak when BLKDEV_ZERO_KILLABLE is set and a fatal signal is pending. Previously, the bio was allocated before checking for a fatal signal. If a signal was pending, the code would break out of the loop without freeing or chaining the just-allocated bio, causing a memory leak. This matches the pattern already used in __blkdev_issue_write_zeroes() where the signal check precedes the allocation.	N/A	More Details
CVE-2023-54120	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: Fix race condition in hidp_session_thread There is a potential race condition in hidp_session_thread that may lead to use-after-free. For instance, the timer is active while hidp_del_timer is called in hidp_session_thread(). After hidp_session_put, then 'session' will be freed, causing kernel panic when hidp_idle_timeout is running. The solution is to use del_timer_sync instead of del_timer. Here is the call trace: ? hidp_session_probe+0x780/0x780 call_timer_fn+0x2d/0x1e0 __run_timers.part.0+0x569/0x940 hidp_session_probe+0x780/0x780 call_timer_fn+0x1e0/0x1e0 ktime_get+0x5c/0xf0 lapic_next_deadline+0x2c/0x40 clockevents_program_event+0x205/0x320 run_timer_softirq+0xa9/0x1b0 __do_softirq+0x1b9/0x641 __irq_exit_rcu+0xdc/0x190 irq_exit_rcu+0xe/0x20 sysvec_apic_timer_interrupt+0xa1/0xc0	N/A	More Details