الهيئة الوطنية للأمن السيبراني
National Cybersecurity Authority

تمت مشاركة هذه المعلومة بإشارة مشاركة ***أبيض*** حيث يسمح بتبادلها أو نشرها من خلال القنوات العامة.

في ضوء دور الهيئة الوطنية للأمن السيبراني للمساعدة في حماية الفضاء السيبراني الوطني، تود الهيئة مشاركتكم النشرة الأسبوعية للثغرات المسجلة من قبل the National Institute of Standards and Technology (NIST) National Vulnerability Database (NVD) للأسبوع من 24 مايو إلى 30 مايو. علماً أنه يتم تصنيف هذه الثغرات باستخدام معيار Common Vulnerability Scoring System (CVSS) حيث يتم تصنيف الثغرات بناء على التالي:

- عالي جدًا: النتيجة الأساسية لـCVSS 9.0-10.0
- عالي: النتيجة الأساسية لـCVSS 7.0-8.9
- متوسط: النتيجة الأساسية لـCVSS 4.0-6.9
- منخفض: النتيجة الأساسية لـ CVSS 0.0-3.9

Please note that this notification/advisory has been tagged as TLP ***WHITE*** where information can be shared or published on any public forums.

As part of NCA duties to help securing the cyberspace and protecting national interests, NCA provides the weekly summary of published vulnerabilities by the National Institute of Standards and Technology (NIST) National Vulnerability Database (NVD) for the week from 24th of May to 30th of May. Vulnerabilities are scored using the Common Vulnerability Scoring System (CVSS) standard as per the following severity:

- **Critical: CVSS base score of 9.0-10.0**
- **High: CVSS base score of 7.0-8.9**
- **Medium: CVSS base score 4.0-6.9**
- **Low: CVSS base score 0.0-3.9**

CVE ID & Source	Vendor - Product	Description	Publish Date	CVSS Score
CVE-2026-46840	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Backend-as-a-Service). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. While the vulnerability is in Oracle REST Data Services, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle REST Data Services. CVSS 3.1 Base Score 10.0 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	10
CVE-2026-7374	red hat - multiple products	A flaw was found in KubeVirt's virt-handler component. This vulnerability allows an authenticated OpenShift user with edit permissions in a single namespace to exploit improper symlink validation when connecting to virtual machine console sockets. By replacing the console socket with a symlink to the host's container runtime (CRI-O) socket, an attacker can hijack virt-handler's privileged connection. This enables the attacker to access any Unix socket on the host, potentially leading to full control of the node and the entire cluster.	2026-05-26	9.9
CVE-2026-46775	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle REST Data Services. While the vulnerability is in Oracle REST Data Services, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle REST Data Services. CVSS 3.1 Base Score 9.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	9.9
CVE-2026-46822	oracle - iassets	Vulnerability in the Oracle iAssets product of Oracle E-Business Suite (component: Internal Operations). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle iAssets. While the vulnerability is in Oracle iAssets, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle iAssets. CVSS 3.1 Base Score 9.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	9.9
CVE-2026-46824	oracle - universal_work_queue	Vulnerability in the Oracle Universal Work Queue product of Oracle E-Business Suite (component: Work Provider Site Level Administration). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Universal Work Queue. While the vulnerability is in Oracle Universal Work Queue, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle Universal Work Queue. CVSS 3.1 Base Score 9.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	9.9
CVE-2026-46839	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle REST Data Services. While the vulnerability is in Oracle REST Data Services, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Oracle REST Data Services. CVSS 3.1 Base Score 9.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	9.9
CVE-2026-8633	ibm - multiple products	IBM Web Server Plug-ins for WebSphere Application Server and WebSphere Liberty 8.5, 9.0 IBM WebSphere Application Server and WebSphere Application Server Liberty are vulnerable to remote code execution in the Web Server Plug-ins, through a specially crafted request.	2026-05-26	9.8

CVE-2026-9170	ibm - multiple products	IBM HTTP Server 8.5, and 9.0	2026-05-26	9.8
CVE-2026-3660	ibm - multiple products	IBM Engineering Lifecycle Management 7.0.3, 7.1.0, and 7.2.0 could allow an unauthenticated remote attacker to update server property files that would allow them to gain unauthorized access to the application.	2026-05-26	9.8
CVE-2025-12686	synology - beestation_os	Buffer copy without checking size of input ('Classic Buffer Overflow') vulnerability in AdminCenter in Synology BeeStation OS before 1.3.2-65648 allows remote attackers to execute arbitrary code via unspecified vectors.	2026-05-27	9.8
CVE-2026-45898	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/iwcm: Fix workqueue list corruption by removing work_list The commit e1168f0 ("RDMA/iwcm: Simplify cm_event_handler()") changed the work submission logic to unconditionally call queue_work() with the expectation that queue_work() would have no effect if work was already pending. The problem is that a free list of struct iwcm_work is used (for which struct work_struct is embedded), so each call to queue_work() is basically unique and therefore does indeed queue the work. This causes a problem in the work handler which walks the work_list until it's empty to process entries. This means that a single run of the work handler could process item N+1 and release it back to the free list while the actual workqueue entry is still queued. It could then get reused (INIT_WORK...) and lead to list corruption in the workqueue logic. Fix this by just removing the work_list. The workqueue already does this for us. This fixes the following error that was observed when stress testing with ucmatose on an Intel E830 in iWARP mode: [151.465780] list_del corruption. next->prev should be ffff9f0915c69c08, but was ffff9f0a1116be08. (next=ffff9f0a15b11c08) [151.466639] ------------[cut here]------------ [151.466986] kernel BUG at lib/list_debug.c:67! [151.467349] Oops: invalid opcode: 0000 [#1] SMP NOPTI [151.467753] CPU: 14 UID: 0 PID: 2306 Comm: kworker/u64:18 Not tainted 6.19.0-rc4+ #1 PREEMPT(voluntary) [151.468466] Hardware name: QEMU Ubuntu 24.04 PC (i440FX + PIIX, 1996), BIOS 1.16.3-debian-1.16.3-2 04/01/2014 [151.469192] Workqueue: 0x0 (iw_cm_wq) [151.469478] RIP: 0010:__list_del_entry_valid_or_report+0xf0/0x100 [151.469942] Code: c7 58 5f 4c b2 e8 10 50 aa ff 0f 0b 48 89 ef e8 36 57 cb ff 48 8b 55 08 48 89 e9 48 89 de 48 c7 c7 a8 5f 4c b2 e8 f0 4f aa ff <0f> 0b 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 40 00 90 90 90 90 90 90 [151.471323] RSP: 0000:ffffb15644e7bd68 EFLAGS: 00010046 [151.471712] RAX: 0000000000000006d RBX: ffff9f0915c69c08 RCX: 0000000000000027 [151.472243] RDX: 0000000000000000 RSI: 0000000000000000 RDI: ffff9f0a37d9c600 [151.472768] RBP: ffff9f0a15b11c08 R08: 0000000000000000 R09: c0000000ffff7fff [151.473294] R10: 0000000000000001 R11: ffffb15644e7bba8 R12: ffff9f092339ee68 [151.473817] R13: ffff9f0900059c28 R14: ffff9f092339ee78 R15: 0000000000000000 [151.474344] FS: 0000000000000000(0000) GS:ffff9f0a847b5000(0000) knlGS:0000000000000000 [151.474934] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [151.475362] CR2: 0000559e233a9088 CR3: 000000020296b004 CR4: 0000000000770ef0 [151.475895] PKRU: 55555554 [151.476118] Call Trace: [151.476331] <TASK> [151.476497] move_linked_works+0x49/0xa0 [151.476792] __pwq_activate_work.isra.46+0x2f/0xa0 [151.477151] pwq_dec_nr_in_flight+0x1e0/0x2f0 [151.477479] process_scheduled_works+0x1c8/0x410 [151.477823] worker_thread+0x125/0x260 [151.478108] ? __pfx_worker_thread+0x10/0x10 [151.478430] kthread+0xfe/0x240 [151.478671] ? __pfx_kthread+0x10/0x10 [151.478955] ? __pfx_kthread+0x10/0x10 [151.479240] ret_from_fork+0x208/0x270 [151.479523] ? __pfx_kthread+0x10/0x10 [151.479806] ret_from_fork_asm+0x1a/0x30 [151.480103] </TASK>	2026-05-27	9.8
CVE-2026-45972	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix potential UAF and double free in smb2_open_file() Zero out @err_iov and @err_buftype before retrying SMB2_open() to prevent an UAF bug if @data != NULL, otherwise a double free.	2026-05-27	9.8
CVE-2026-45988	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2026-05-27	9.8

		rxrpc: Fix re-decryption of RESPONSE packets If a RESPONSE packet gets a temporary failure during processing, it may end up in a partially decrypted state - and then get requeued for a retry. Fix this by just discarding the packet; we will send another CHALLENGE packet and thereby elicit a further response. Similarly, discard an incoming CHALLENGE packet if we get an error whilst generating a RESPONSE; the server will send another CHALLENGE.		
CVE-2026-46039	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: rxgk: Fix potential integer overflow in length check Fix potential integer overflow in rxgk_extract_token() when checking the length of the ticket. Rather than rounding up the value to be tested (which might overflow), round down the size of the available data.	2026-05-27	9.8
CVE-2026-8175	ibm - multiple products	IBM Aspera High-Speed Transfer Endpoint 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Server 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Endpoint are affected by a buffer overflow in the asperahttpd component. This vulnerability could be exploited to cause a denial of service and potentially lead to authentication bypass or remote code execution.	2026-05-27	9.8
CVE-2026-46115	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: block: add pgmap check to biovec_phys_mergeable biovec_phys_mergeable() is used by the request merge, DMA mapping, and integrity merge paths to decide if two physically contiguous bvec segments can be coalesced into one. It currently has no check for whether the segments belong to different dev_pagemaps. When zone device memory is registered in multiple chunks, each chunk gets its own dev_pagemap. A single bio can legitimately contain bvecs from different pgmaps -- iov_iter_extract_bvecs() breaks at pgmap boundaries but the outer loop in bio_iov_iter_get_pages() continues filling the same bio. If such bvecs are physically contiguous, biovec_phys_mergeable() will coalesce them, making it impossible to recover the correct pgmap for the merged segment via page_pgmap(). Add a zone_device_pages_have_same_pgmap() check to prevent merging bvec segments that span different pgmaps.	2026-05-28	9.8
CVE-2026-46135	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: nvmet-tcp: fix race between ICReq handling and queue teardown nvmet_tcp_handle_icreq() updates queue->state after sending an Initialization Connection Response (ICResp), but it does so without serializing against target-side queue teardown. If an NVMe/TCP host sends an Initialization Connection Request (ICReq) and immediately closes the connection, target-side teardown may start in softirq context before io_work drains the already buffered ICReq. In that case, nvmet_tcp_schedule_release_queue() sets queue->state to NVMET_TCP_Q_DISCONNECTING and drops the queue reference under state_lock. If io_work later processes that ICReq, nvmet_tcp_handle_icreq() can still overwrite the state back to NVMET_TCP_Q_LIVE. That defeats the DISCONNECTING-state guard in nvmet_tcp_schedule_release_queue() and allows a later socket state change to re-enter teardown and issue a second kref_put() on an already released queue. The ICResp send failure path has the same problem. If teardown has already moved the queue to DISCONNECTING, a send error can still overwrite the state with NVMET_TCP_Q_FAILED, again reopening the window for a second teardown path to drop the queue reference. Fix this by serializing both post-send state transitions with state_lock and bailing out if teardown has already started. Use -ESHUTDOWN as an internal sentinel for that bail-out path rather than propagating it as a transport error like -ECONNRESET. Keep nvmet_tcp_socket_error() setting rcv_state to NVMET_TCP_RECV_ERR before honoring that sentinel so receive-side parsing stays quiesced until the existing release path completes.	2026-05-28	9.8
CVE-2026-46137	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mptcp: pm: ADD_ADDR rtx: fix potential data-race This mptcp_pm_add_timer() helper is executed as a timer callback in softirq context. To avoid any data races, the socket lock needs to be held with bh_lock_sock().	2026-05-28	9.8

		If the socket is in use, retry again soon after, similar to what is done with the keepalive timer.		
CVE-2026-46195	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: smb: client: validate dacloffset before building DACL pointers parse_sec_desc(), build_sec_desc(), and the chown path in id_mode_to_cifs_acl() all add the server-supplied dacloffset to pntsd before proving a DACL header fits inside the returned security descriptor. On 32-bit builds a malicious server can return dacloffset near U32_MAX, wrap the derived DACL pointer below end_of_acl, and then slip past the later pointer-based bounds checks. build_sec_desc() and id_mode_to_cifs_acl() can then dereference DACL fields from the wrapped pointer in the chmod/chown rewrite paths. Validate dacloffset numerically before building any DACL pointer and reuse the same helper at the three DACL entry points.	2026-05-28	9.8
CVE-2026-34311	oracle - multiple products	Vulnerability in the Oracle Hospitality OPERA 5 Property Services product of Oracle Hospitality Applications (component: Opera). Supported versions that are affected are 5.6.19.24, 5.6.22, 5.6.25.19, 5.6.27.6 and 5.6.28. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Hospitality OPERA 5 Property Services. Successful attacks of this vulnerability can result in takeover of Oracle Hospitality OPERA 5 Property Services. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	2026-05-28	9.8
CVE-2026-46817	oracle - e-business_suite	Vulnerability in the Oracle Payments product of Oracle E-Business Suite (component: File Transmission). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Payments. Successful attacks of this vulnerability can result in takeover of Oracle Payments. CVSS 3.1 Base Score 9.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H).	2026-05-28	9.8
CVE-2026-9872	google - chrome	Out of bounds write in GPU in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	9.6
CVE-2026-9874	google - chrome	Use after free in Dawn in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	9.6
CVE-2026-9875	google - chrome	Out of bounds read in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	9.6
CVE-2026-9876	google - chrome	Use after free in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	9.6
CVE-2026-9886	google - chrome	Use after free in Base in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	9.6
CVE-2026-9918	google - chrome	Inappropriate implementation in Tint in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	9.6
CVE-2026-9967	google - chrome	Out of bounds write in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	9.6
CVE-2026-9739	google - MCP Toolbox for Databases	Vulnerable to DNS rebinding attacks when using SSE (http://b/499408790). During the beta phase, we implemented `allowed-origins` and `allowed-hosts` flags to align with MCP security guidelines. However, the hardcoded `Access-Control-Allow-Origin: *` header in the SSE initialization handler was inadvertently retained. This vulnerability specifically impacts users connecting via Toolbox using SSE under specification v2024-11-05.	2026-05-27	9.4
CVE-2026-32998	veeam - Service Provider Console	This vulnerability in Veeam Service Provider Console allows for remote code execution.	2026-05-28	9.4
CVE-2026-46043	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Validate pad and ICRC before payload_size() in rxe_rcv rxe_rcv() currently checks only that the incoming packet is at least header_size(pkt) bytes long before payload_size() is used. However, payload_size() subtracts both the attacker-controlled BTH pad field and RXE_ICRC_SIZE from pkt->paylen: payload_size = pkt->paylen - offset[RXE_PAYLOAD] - bth_pad(pkt) - RXE_ICRC_SIZE This means a short packet can still make payload_size() underflow even if it includes enough bytes for the fixed headers. Simply requiring header_size(pkt) + RXE_ICRC_SIZE is not sufficient either, because a packet with a forged non-zero BTH pad can still leave payload_size() negative and pass an underflowed value to later receive-path users. Fix this by validating pkt->paylen against the full minimum length	2026-05-27	9.1

		required by payload_size(): header_size(pkt) + bth_pad(pkt) + RXE_ICRC_SIZE.		
CVE-2026-7876	ibm - aspera_high-speed_transfer_server_for_cloud_pak_for_integration	IBM Aspera HSTS for CP4I 1.5.1 through 1.5.19	2026-05-27	9.1
CVE-2026-46119	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: libceph: Fix slab-out-of-bounds access in auth message processing If a (potentially corrupted) message of type CEPH_MSG_AUTH_REPLY contains a positive value in its result field, it is treated as an error code by ceph_handle_auth_reply() and returned to handle_auth_reply(). Thereafter, an attempt is made to send the preallocated message of type CEPH_MSG_AUTH, where the returned value is interpreted as the size of the front segment to send. If the result value in the message is greater than the size of the memory buffer allocated for the front segment, an out-of-bounds access occurs, and the content of the memory region beyond this buffer is sent out. This patch fixes the issue by treating only negative values in the result field as errors. Positive values are therefore treated as success in the same way as a zero value. Additionally, a BUG_ON is added to __send_prepared_auth_request() comparing the len parameter to front_alloc_len to prevent sending the message if it exceeds the bounds of the allocation and to make it easier to catch any logic flaws leading to this.	2026-05-28	9.1
CVE-2026-46155	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: smb/client: fix out-of-bounds read in smb2_compound_op() If a server sends a truncated response but a large OutputBufferLength, and terminates the EA list early, check_wsl_eas() returns success without validating that the entire OutputBufferLength fits within iov_len. Then smb2_compound_op() does: memcpy(idata->wsl.eas, data[0], size[0]); Where size[0] is OutputBufferLength. If iov_len is smaller than size[0], memcpy can read beyond the end of the rsp_iov allocation and leak adjacent kernel heap memory.	2026-05-28	9.1
CVE-2026-46185	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: smb/client: fix out-of-bounds read in symlink_data() Since smb2_check_message() returns success without length validation for the symlink error response, in symlink_data() it is possible for iov->iov_len to be smaller than sizeof(struct smb2_err_rsp). If the buffer only contains the base SMB2 header (64 bytes), accessing err->ErrorContextCount (at offset 66) or err->ByteCount later in symlink_data() will cause an out-of-bounds read.	2026-05-28	9.1
CVE-2026-46819	oracle - e-business_suite	Vulnerability in the Oracle Internet Procurement Connector product of Oracle E-Business Suite (component: Internal Operations). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Internet Procurement Connector. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Internet Procurement Connector accessible data as well as unauthorized access to critical data or complete access to all Oracle Internet Procurement Connector accessible data. CVSS 3.1 Base Score 9.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N).	2026-05-28	9.1
CVE-2026-4480	redhat - multiple products	A flaw was found in the Samba printing subsystem. Samba passes the client-controlled job description string to the command configured with the "print command" setting via the "%J" substitution character without escaping shell meta characters. A remote attacker could exploit this vulnerability by sending a specially crafted print job description that contains unescaped shell characters. This could lead to remote code execution on the affected system.	2026-05-26	9
CVE-2026-4408	red hat - multiple products	A flaw was found in Samba. A remote attacker can exploit a misconfiguration in Samba file servers and classic domain controllers that use the "check password script" feature. If this script is configured with the %u substitution character, the client-controlled username is passed without proper escaping of shell meta-characters. This vulnerability allows an attacker to achieve remote command execution on the affected system. This issue primarily affects non-standard configurations where the "check password script" is used with %u and the samba-dcerpcd service is started as a system service.	2026-05-28	9
CVE-2026-46833	oracle - database_server	Vulnerability in the Net Service component of Oracle Database Server. Supported versions that are affected are 23.4.0-23.26.2. Difficult to exploit vulnerability allows unauthenticated attacker with network access via TLS to compromise Net Service. While the vulnerability is in Net Service, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in takeover of Net Service. CVSS 3.1 Base Score 9.0 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:H/A:H).	2026-05-28	9

CVE-2026-9881	google - chrome	Use after free in Bluetooth in Google Chrome on Mac prior to 148.0.7778.216 allowed an attacker who convinced a user to install a malicious extension to potentially perform a sandbox escape via a crafted Chrome Extension. (Chromium security severity: Critical)	2026-05-28	9
CVE-2026-9891	google - chrome	Use after free in Extensions in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted Chrome Extension. (Chromium security severity: Critical)	2026-05-28	9

CVE-2026-45945	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		iommu/vt-d: Fix race condition during PASID entry replacement
		
		The Intel VT-d PASID table entry is 512 bits (64 bytes). When replacing an active PASID entry (e.g., during domain replacement), the current implementation calculates a new entry on the stack and copies it to the table using a single structure assignment.
		
		struct pasid_entry *pte, new_pte;
		
		pte = intel_pasid_get_entry(dev, pasid);
		pasid_pte_config_first_level(iommu, &new_pte, ...);
		*pte = new_pte;
		
		Because the hardware may fetch the 512-bit PASID entry in multiple 128-bit chunks, updating the entire entry while it is active (Present bit set) risks a "torn" read. In this scenario, the IOMMU hardware could observe an inconsistent state — partially new data and partially old data — leading to unpredictable behavior or spurious faults.
		
		Fix this by removing the unsafe "replace" helpers and following the "clear-then-update" flow, which ensures the Present bit is cleared and the required invalidation handshake is completed before the new configuration is applied.
			2026-05-27	8.8

CVE-2026-46056	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_event: fix potential UAF in SSP passkey handlers hci_conn lookup and field access must be covered by hdev lock in hci_user_passkey_notify_evt() and hci_keypress_notify_evt(), otherwise the connection can be freed concurrently. Extend the hci_dev_lock critical section to cover all conn usage in both handlers. Keep the existing keypress notification behavior unchanged by routing the early exits through a common unlock path.	2026-05-27	8.8
CVE-2026-5065	ibm - controller	IBM Controller 11.0.1, 11.1.0, 11.1.1, and 11.1.2 contains hard-coded credentials, such as a password or cryptographic key, which it uses for its own inbound authentication, outbound communication to external components, or encryption of internal data.	2026-05-27	8.8
CVE-2026-8179	ibm - multiple products	IBM Aspera High-Speed Transfer Endpoint 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Server 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Endpoint are affected by a buffer overflow in the asperahttpd component. This vulnerability could allow an authenticated user to execute arbitrary code on the system.	2026-05-27	8.8
CVE-2026-48920	jenkins - multiple products	Jenkins Email Extension Plugin 1933.v45cec755423f and earlier allows inlining images as `base64` in email content by setting the `data-inline` attribute, without restrictions on the image URLs that can be inlined, allowing attackers able to control the email content to specify `file:` URLs for images to read arbitrary files from the Jenkins controller filesystem.	2026-05-27	8.8
CVE-2026-46414	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. In 3.0.1-4-ge2626659, Microsoft UFO's WebSocket control plane trusts client-supplied identity and role fields in task messages. A client connection can register as a normal device, but later send a TASK message claiming client_type="constellation" and target_id=<victim-device-id>. The server trusts the role and target values from the wire message rather than enforcing the role registered for that WebSocket connection. As a result, any authenticated WebSocket client with the shared server token can spoof the higher-privilege constellation role and dispatch attacker-controlled tasks to another connected device. The same client registry also allows duplicate client_id registration, overwriting an existing live client's stored websocket, role, and task protocol. This is an authenticated WebSocket role/identity spoofing issue leading to peer task hijacking.	2026-05-27	8.8
CVE-2026-8915	samsung - escargot	Out-of-bounds write vulnerability in Samsung Open Source Escargot allows Overflow Buffers. This issue affects Escargot: 36f5fb58366a67b713c02f6fd985e924fcc09e31.	2026-05-28	8.8
CVE-2026-46113	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Fix shadow paging use-after-free due to unexpected GFN The shadow MMU computes GFNs for direct shadow pages using sp->gfn plus the SPTE index. This assumption breaks for shadow paging if the guest page tables are modified between VM entries (similar to commit aad885e77496, "KVM: x86/mmu: Drop/zap existing present SPTE even when creating an MMIO SPTE", 2026-03-27). The flow is as follows: - a PDE is installed for a 2MB mapping, and a page in that area is accessed. KVM creates a kvm_mmu_page consisting of 512 4KB pages;	2026-05-28	8.8

		the kvm_mmu_page is marked by FNAME(fetch) as direct-mapped because the guest's mapping is a huge page (and thus contiguous). - the PDE mapping is changed from outside the guest. - the guest accesses another page in the same 2MB area. KVM installs a new leaf SPTE and rmap entry; the SPTE uses the "correct" GFN (i.e. based on the new mapping, as changed in the previous step) but that GFN is outside of the [sp->gfn, sp->gfn + 511] range; therefore the rmap entry cannot be found and removed when the kvm_mmu_page is zapped. - the memslot that covers the first 2MB mapping is deleted, and the kvm_mmu_page for the now-invalid GPA is zapped. However, rmap_remove() only looks at the [sp->gfn, sp->gfn + 511] range established in step 1, and fails to find the rmap entry that was recorded by step 3. - any operation that causes an rmap walk for the same page accessed by step 3 then walks a stale rmap and dereferences a freed kvm_mmu_page. This includes dirty logging or MMU notifier invalidations (e.g., from MADV_DONTNEED). The underlying issue is that KVM's walking of shadow PTEs assumes that if a SPTE is present when KVM wants to install a non-leaf SPTE, then the existing kvm_mmu_page must be for the correct gfn. Because the only way for the gfn to be wrong is if KVM messed up and failed to zap a SPTE... which shouldn't happen, but *actually* only happens in response to a guest write. That bug dates back literally forever, as even the first version of KVM assumes that the GFN matches and walks into the "wrong" shadow page. However, that was only an imprecision until 2032a93d66fa ("KVM: MMU: Don't allocate gfns page for direct mmu pages") came along. Fix it by checking for a target gfn mismatch and zapping the existing SPTE. That way the old SP and rmap entries are gone, KVM installs the rmap in the right location, and everyone is happy.		
CVE-2026-46125	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: remove station if connection prep fails If connection preparation fails for MLO connections, then the interface is completely reset to non-MLD. In this case, we must not keep the station since it's related to the link of the vif being removed. Delete an existing station. Any "new_sta" is already being removed, so that doesn't need changes. This fixes a use-after-free/double-free in debugfs if that's enabled, because a vif going from MLD (and to MLD, but that's not relevant here) recreates its entire debugfs.	2026-05-28	8.8
CVE-2026-46152	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: drop stray 'static' from fast-RX rx_result ieee80211_invoke_fast_rx() is documented as safe for parallel RX, but its per-invocation rx_result is declared static. Concurrent callers then share one instance and can overwrite each other's result between ieee80211_rx_mesh_data() and the switch on res. That can make a packet that was queued or consumed by ieee80211_rx_mesh_data() fall through into ieee80211_rx_8023(), or make a packet that should continue return as queued. Make res an automatic variable so each invocation keeps its own result.	2026-05-28	8.8
CVE-2026-46166	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: use safe list iteration in radar detect work The call to ieee80211_dfs_cac_cancel can cause the iterated chanctx to be freed and removed from the list. Guard against this to avoid a slab-use-after-free error.	2026-05-28	8.8
CVE-2026-46174	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: x86/CPU/AMD: Prevent improper isolation of shared resources in Zen2's op cache Make sure resources are not improperly shared in the op cache and cause instruction corruption this way.	2026-05-28	8.8
CVE-2026-46198	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: batman-adv: fix integer overflow on buff_pos	2026-05-28	8.8

		Fixing an integer overflow present in batadv_iv_ogm_send_to_if. The size check is done using the int type in batadv_iv_ogm_aggr_packet whereas the buff_pos variable uses the s16 type. This could lead to an out-of-bound read.		
CVE-2026-46212	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: batman-adv: bla: prevent use-after-free when deleting claims When batadv_bla_del_backbone_claims() removes all claims for a backbone, it does this by dropping the link entry in the hash list. This list entry itself was one of the references which need to be dropped at the same time via batadv_claim_put(). But the batadv_claim_put() must not be done before the last access to the claim object in this function. Otherwise the claim might be freed already by the batadv_claim_release() function before the list entry was dropped.	2026-05-28	8.8
CVE-2026-46238	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: batman-adv: stop caching unowned originator pointers in BAT IV BAT IV keeps the last-hop neighbor address in each neigh_node, but some paths also cache an originator pointer derived from a temporary lookup. That pointer is not owned by the neigh_node and may no longer refer to a live originator entry after purge handling runs. Stop storing the auxiliary originator pointer in the BAT IV neighbor state. When BAT IV needs the neighbor originator data, resolve it from the stored neighbor address and drop the reference again after use. [sven: avoid bonding logic for outgoing OGM]	2026-05-28	8.8
CVE-2026-46826	oracle - e-business_suite	Vulnerability in the Oracle Payroll product of Oracle E-Business Suite (component: Internal Operations). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle Payroll. Successful attacks of this vulnerability can result in takeover of Oracle Payroll. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	2026-05-28	8.8
CVE-2026-46827	oracle - e-business_suite	Vulnerability in the Oracle Payroll product of Oracle E-Business Suite (component: Self Service Manager). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Payroll. Successful attacks of this vulnerability can result in takeover of Oracle Payroll. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	2026-05-28	8.8
CVE-2026-46837	oracle - e-business_suite	Vulnerability in the Oracle Flow Manufacturing product of Oracle E-Business Suite (component: Security). Supported versions that are affected are 12.2.9-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via SQL to compromise Oracle Flow Manufacturing. Successful attacks of this vulnerability can result in takeover of Oracle Flow Manufacturing. CVSS 3.1 Base Score 8.8 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H).	2026-05-28	8.8
CVE-2026-10002	google - chrome	Use after free in PDFium in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted PDF file. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-10007	google - chrome	Use after free in SVG in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-10013	google - chrome	Use after free in WebCodecs in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-10015	google - chrome	Integer overflow in WTF in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-10016	google - chrome	Use after free in DOM in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-10019	google - chrome	Integer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: Medium)	2026-05-28	8.8
CVE-2026-10021	google - chrome	Insufficient validation of untrusted input in USB in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: Medium)	2026-05-28	8.8
CVE-2026-9873	google - chrome	Use after free in Network in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.8
CVE-2026-9878	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.8
CVE-2026-9879	google - chrome	Out of bounds write in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.8
CVE-2026-9883	google - chrome	Use after free in Base in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.8

CVE-2026-9884	google - chrome	Use after free in Browser in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.8
CVE-2026-9887	google - chrome	Use after free in Proxy in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted PAC script. (Chromium security severity: Critical)	2026-05-28	8.8
CVE-2026-9896	google - chrome	Out of bounds write in V8 in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9897	google - chrome	Use after free in DOM in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9910	google - chrome	Out of bounds memory access in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9923	google - chrome	Use after free in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9927	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9928	google - chrome	Out of bounds read in ANGLE in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9938	google - chrome	Inappropriate implementation in V8 in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9939	google - chrome	Heap buffer overflow in WebCodecs in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9940	google - chrome	Heap buffer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9941	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9945	google - chrome	Use after free in Media in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9947	google - chrome	Use after free in XML in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9952	google - chrome	Use after free in WebAudio in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9957	google - chrome	Use after free in PDF in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted PDF file. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9958	google - chrome	Use after free in PDFium in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted PDF file. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9961	google - chrome	Use after free in SurfaceCapture in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9962	google - chrome	Use after free in WebRTC in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9965	google - chrome	Out of bounds write in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9968	google - chrome	Integer overflow in V8 in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9969	google - chrome	Insufficient validation of untrusted input in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9973	google - chrome	Out of bounds write in V8 in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9976	google - chrome	Inappropriate implementation in USB in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9978	google - chrome	Use after free in Glic in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9983	google - chrome	Type Confusion in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9984	google - chrome	Use after free in UI in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9992	google - chrome	Use after free in Network in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-9995	google - chrome	Use after free in WebXR in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8

CVE-2026-9999	google - chrome	Inappropriate implementation in ANGLE in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.8
CVE-2026-8697	tp-link - archer_c64_firmware	Due to improper enforcement of authentication rate-limiting on a debug SSH service in Archer C64 v1, the SSH service allows unlimited authentication attempts and uses the same credentials as the web interface. This enables an attacker to brute-force valid credentials via SSH. Successful exploitation could allow an attacker with adjacent network access to obtain administrative credentials through unrestricted authentication attempts and subsequently gain full administrative access to the device, impacting system confidentiality, integrity, and availability.	2026-05-28	8.7
CVE-2025-30028	synology - active_backup_for_business	A vulnerability in Active Backup for Business allows unauthorized remote attackers to read arbitrary files.	2026-05-27	8.6
CVE-2026-32997	veeam - Backup and Replication	A vulnerability allowing an authenticated user with the Backup Administrator role to write arbitrary files on Linux-based Veeam Backup & Replication server.	2026-05-28	8.6
CVE-2026-5509	tp-link - archer_be450_firmware	An authenticated command injection vulnerability exists in the Archer BE450 v1 and BE7200 v1 router that allows an administrator to execute arbitrary system commands through the web management interface. After successfully authenticating to the admin interface, an attacker can leverage the browser's developer console by supplying a crafted input that is passed to backend system commands without adequate sanitization. Successful exploitation enables execution of arbitrary commands with elevated privileges on the device, which may allow the attacker to start unauthorized services, modify system configuration, or otherwise fully compromise the router's operating environment.	2026-05-27	8.5
CVE-2025-48977	apache - ignite	Relative Path Traversal vulnerability in Apache Ignite REST API. Authenticated REST API users can read any file on the server with "cmd=log" command and a log path crafted in a certain way. This issue affects Apache Ignite: from 2.0.0 through 2.17.0. Users are recommended to upgrade to version 2.18.0, which fixes the issue.	2026-05-28	8.5
CVE-2026-46820	oracle - financials_common_modules	Vulnerability in the Oracle Financials Common Modules product of Oracle E-Business Suite (component: Common Components). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Financials Common Modules. While the vulnerability is in Oracle Financials Common Modules, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Financials Common Modules accessible data as well as unauthorized update, insert or delete access to some of Oracle Financials Common Modules accessible data. CVSS 3.1 Base Score 8.5 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:L/A:N).	2026-05-28	8.5
CVE-2026-7365	ibm - multiple products	IBM Operations Analytics - Log Analysis and IBM SmartCloud Analytics - Log Analysis uses default passwords default passwords from the manufacturing process for use during the installation process, which could allow an attacker to bypass authentication.	2026-05-27	8.4
CVE-2026-10000	google - chrome	Use after free in Passwords in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-10001	google - chrome	Use after free in PerformanceManager in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-10012	google - chrome	Use after free in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-10014	google - chrome	Use after free in WebMIDI in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-10017	google - chrome	Out of bounds read in Headless in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Medium)	2026-05-28	8.3
CVE-2026-10020	google - chrome	Insufficient validation of untrusted input in Skia in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Medium)	2026-05-28	8.3
CVE-2026-9877	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9880	google - chrome	Insufficient validation of untrusted input in WebGL in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9885	google - chrome	Insufficient validation of untrusted input in UI in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9888	google - chrome	Use after free in WebView in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9889	google - chrome	Out of bounds read and write in Dawn in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3

CVE-2026-9890	google - chrome	Use after free in XR in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9892	google - chrome	Inappropriate implementation in Skia in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9893	google - chrome	Use after free in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	8.3
CVE-2026-9894	google - chrome	Use after free in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9895	google - chrome	Out of bounds read in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9898	google - chrome	Insufficient validation of untrusted input in GPU in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9899	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9900	google - chrome	Out of bounds write in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9902	google - chrome	Use after free in Accessibility in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9904	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9905	google - chrome	Use after free in Accessibility in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9906	google - chrome	Out of bounds write in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9914	google - chrome	Insufficient validation of untrusted input in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9915	google - chrome	Heap buffer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9916	google - chrome	Out of bounds write in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9924	google - chrome	Heap buffer overflow in ANGLE in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9925	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9926	google - chrome	Heap buffer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9931	google - chrome	Use after free in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9932	google - chrome	Use after free in ANGLE in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9936	google - chrome	Use after free in GFX in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9937	google - chrome	Use after free in UI in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9946	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9948	google - chrome	Use after free in Views in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9949	google - chrome	Use after free in Core in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9951	google - chrome	Use after free in UI in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9966	google - chrome	Integer overflow in XML in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3

CVE-2026-9970	google - chrome	Use after free in WebGL in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9972	google - chrome	Uninitialized Use in Gamepad in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9974	google - chrome	Out of bounds write in GPU in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9975	google - chrome	Out of bounds read and write in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9977	google - chrome	Insufficient validation of untrusted input in WebShare in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9982	google - chrome	Insufficient validation of untrusted input in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9988	google - chrome	Use after free in WebRTC in Google Chrome on Linux prior to 148.0.7778.216 allowed a remote attacker to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9993	google - chrome	Use after free in Views in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted PDF file. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9994	google - chrome	Use after free in Core in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9997	google - chrome	Use after free in Input in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-9998	google - chrome	Integer overflow in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)	2026-05-28	8.3
CVE-2026-42013	red hat - multiple products	A flaw was found in gnutls. When validating certificates, an oversized Subject Alternative Name (SAN) could cause the validation process to incorrectly fall back to checking the Common Name (CN) field. This could allow a remote attacker to bypass proper certificate validation, potentially leading to spoofing or man-in-the-middle attacks.	2026-05-26	8.2
CVE-2026-5260	red hat - multiple products	A flaw was found in libgnutls. A remote attacker, by sending an extremely short premaster secret during an RSA key exchange to a server using an RSA key backed by a PKCS#11 token, could trigger a short heap overread. This memory corruption vulnerability could lead to information disclosure.	2026-05-26	8.2
CVE-2026-45843	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: slip: bound decode() reads against the compressed packet length slhc_uncompress() parses a VJ-compressed TCP header by advancing a pointer through the packet via decode() and pull16(). Neither helper bounds-checks against isize, and decode() masks its return with & 0xffff so it can never return the -1 that callers test for -- those error paths are dead code. A short compressed frame whose change byte requests optional fields lets decode() read past the end of the packet. The over-read bytes are folded into the cached cstate and reflected into subsequent reconstructed packets. Make decode() and pull16() take the packet end pointer and return -1 when exhausted. Add a bounds check before the TCP-checksum read. The existing == -1 tests now do what they were always meant to.	2026-05-27	8.2
CVE-2026-46037	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ipv4: icmp: validate reply type before using icmp_pointers Extended echo replies use ICMP_EXT_ECHOREPLY as the outbound reply type. That value is outside the range covered by icmp_pointers[], which only describes the traditional ICMP types up to NR_ICMP_TYPES. Avoid consulting icmp_pointers[] for reply types outside that range, and use array_index_nospec() for the remaining in-range lookup. Normal ICMP replies keep their existing behavior unchanged.	2026-05-27	8.2
CVE-2026-45361	apache - apache-airflow-providers-google	Apache Airflow providers-google's `ComputeEngineSSHHook` disables SSH host-key verification by default, exposing SSH traffic between an Airflow worker and a Compute Engine VM to in-path network attackers who can intercept or modify the session. Users are advised to upgrade to `apache-airflow-providers-google` 22.0.0 or later.	2026-05-25	8.1
CVE-2026-8855	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to remote code execution and denial of service in configurations with TLS mutual authentication (client authentication).	2026-05-26	8.1
CVE-2025-13392	synology - multiple products	Improper check for unusual or exceptional conditions vulnerability in SSO in Synology DiskStation Manager (DSM) before 7.2.2-72806-5 and 7.3.1-86003-1 (7.2.1-69057 is not affected) allows remote attackers to bypass authentication with prior knowledge of the distinguished name (DN).	2026-05-27	8.1

CVE-2026-46010	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix error handling in rxgk_extract_token() Fix a missing bit of error handling in rxgk_extract_token(): in the event that rxgk_decrypt_skb() returns -ENOMEM, it should just return that rather than continuing on (for anything else, it generates an abort).	2026-05-27	8.1
CVE-2026-46099	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: ipv6: fix NOREF dst use in seg6 and rpl lwtunnels seg6_input_core() and rpl_input() call ip6_route_input() which sets a NOREF dst on the skb, then pass it to dst_cache_set_ip6() invoking dst_hold() unconditionally. On PREEMPT_RT, ksoftirqd is preemptible and a higher-priority task can release the underlying pcpu_rt between the lookup and the caching through a concurrent FIB lookup on a shared nexthop. Simplified race sequence: ksoftirqd/X higher-prio task (same CPU X) ---------- ---------- seg6_input_core(,skb)/rpl_input(skb) dst_cache_get() -> miss ip6_route_input(skb) -> ip6_pol_route(,skb,flags) [RT6_LOOKUP_F_DST_NOREF in flags] -> FIB lookup resolves fib6_nh [nhid=N route] -> rt6_make_pcpu_route() [creates pcpu_rt, refcount=1] pcpu_rt->sernum = fib6_sernum [fib6_sernum=W] -> cmpxchg(fib6_nh.rt6i_pcpu, NULL, pcpu_rt) [slot was empty, store succeeds] -> skb_dst_set_noref(skb, dst) [dst is pcpu_rt, refcount still 1] rt_genid_bump_ipv6() -> bumps fib6_sernum [fib6_sernum from W to Z] ip6_route_output() -> ip6_pol_route() -> FIB lookup resolves fib6_nh [nhid=N] -> rt6_get_pcpu_route() pcpu_rt->sernum != fib6_sernum [W <> Z, stale] -> prev = xchg(rt6i_pcpu, NULL) -> dst_release(prev) [prev is pcpu_rt, refcount 1->0, dead] dst = skb_dst(skb) [dst is the dead pcpu_rt] dst_cache_set_ip6(dst) -> dst_hold() on dead dst -> WARN / use-after-free For the race to occur, ksoftirqd must be preemptible (PREEMPT_RT without PREEMPT_RT_NEEDS_BH_LOCK) and a concurrent task must be able to release the pcpu_rt. Shared nexthop objects provide such a path, as two routes pointing to the same nhid share the same fib6_nh and its rt6i_pcpu entry. Fix seg6_input_core() and rpl_input() by calling skb_dst_force() after ip6_route_input() to force the NOREF dst into a refcounted one before caching. The output path is not affected as ip6_route_output() already returns a refcounted dst.	2026-05-27	8.1
CVE-2026-46402	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. In 3.0.1-4-ge2626659, Microsoft UFO uses the user-controlled task_name value directly when constructing session log paths. An authenticated client can supply path traversal sequences in task_name and cause UFO to create log directories and log files outside the intended logs/ directory.	2026-05-27	8.1
CVE-2026-46138	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_event: Fix OOB read and infinite loop in hci_le_create_big_complete_evt	2026-05-28	8.1

		hci_le_create_big_complete_evt() iterates over BT_BOUND connections for a BIG handle using a while loop, accessing ev->bis_handle[i++] on each iteration. However, there is no check that i stays within ev->num_bis before the array access. When a controller sends a LE_Create_BIG_Complete event with fewer bis_handle entries than there are BT_BOUND connections for that BIG, or with num_bis=0, the loop reads beyond the valid bis_handle[] flex array into adjacent heap memory. Since the out-of-bounds values typically exceed HCI_CONN_HANDLE_MAX (0x0EFF), hci_conn_set_handle() rejects them and the connection remains in BT_BOUND state. The same connection is then found again by hci_conn_hash_lookup_big_state(), creating an infinite loop with hci_dev_lock held. Fix this by terminating the BIG if in case not all BIS could be setup properly.		
CVE-2026-46232	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: HID: playstation: Clamp num_touch_reports A device would never lie about the number of touch reports would it? If it does the loop in dualshock4_parse_report will read off the end of the touch_reports array, up to about 2 KiB for the maximum number of 256 loop iteraions. The data that is read is emitted via evdev if the DS4_TOUCH_POINT_INACTIVE bit happens to be set. Protect against this by clamping the num_touch_reports value provided by the device to the maximum size of the touch_reports array.	2026-05-28	8.1
CVE-2026-35277	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle REST Data Services accessible data as well as unauthorized access to critical data or complete access to all Oracle REST Data Services accessible data. CVSS 3.1 Base Score 8.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:N).	2026-05-28	8.1
CVE-2026-46828	oracle - e-business_suite	Vulnerability in the Oracle Payroll product of Oracle E-Business Suite (component: Internal Operations). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Payroll. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Payroll accessible data as well as unauthorized access to critical data or complete access to all Oracle Payroll accessible data. CVSS 3.1 Base Score 8.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:N).	2026-05-28	8.1
CVE-2026-9964	google - chrome	Use after free in Bluetooth in Google Chrome on Mac prior to 148.0.7778.216 allowed an attacker who convinced a user to install a malicious extension to execute arbitrary code via a crafted Chrome Extension. (Chromium security severity: High)	2026-05-28	8.1
CVE-2026-8834	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 contains a buffer overflow vulnerability. A privileged user, authenticated to the Administration Server, could exploit this vulnerability to execute remote code or cause a denial of service.	2026-05-26	8
CVE-2026-3012	red hat - multiple products	A flaw was found in Samba's certificate auto-enrollment Group Policy handling. When certificate auto-enrollment is enabled, Samba may retrieve a CA certificate over an unencrypted HTTP connection and install it into the local trust store without proper verification. An attacker with the ability to intercept or redirect network traffic could exploit this behavior to supply a malicious certificate authority certificate, potentially allowing interception or spoofing of trusted communications.	2026-05-27	8
CVE-2026-46076	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: KVM: nSVM: Raise #UD if unhandled VMMCALL isn't intercepted by L1 Explicitly synthesize a #UD for VMMCALL if L2 is active, L1 does NOT want to intercept VMMCALL, nested_svm_l2_tlb_flush_enabled() is true, and the hypercall is something other than one of the supported Hyper-V hypercalls. When all of the above conditions are met, KVM will intercept VMMCALL but never forward it to L1, i.e. will let L2 make hypercalls as if it were L1. The TLFS says a whole lot of nothing about this scenario, so go with the architectural behavior, which says that VMMCALL #UDs if it's not intercepted. Opportunistically do a 2-for-1 stub trade by stub-ifying the new API instead of the helpers it uses. The last remaining "single" stub will soon be dropped as well. [sean: rewrite changelog and comment, tag for stable, remove defunct stubs]	2026-05-27	7.9
CVE-2026-35266	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Difficult to exploit vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks require human interaction from a person other than the attacker and while the vulnerability is in Oracle REST Data Services, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data	2026-05-28	7.9

		or all Oracle REST Data Services accessible data as well as unauthorized access to critical data or complete access to all Oracle REST Data Services accessible data and unauthorized ability to cause a partial denial of service (partial DOS) of Oracle REST Data Services. CVSS 3.1 Base Score 7.9 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:L/UI:R/S:C/C:H/I:H/A:L).		
CVE-2025-43306	apple - multiple products	A logic issue was addressed with improved checks. This issue is fixed in macOS Sequoia 15.7, macOS Sonoma 14.8, macOS Tahoe 26. A malicious app may be able to gain root privileges.	2026-05-26	7.8
CVE-2023-52945	synology - beedrive	Uncontrolled search path element vulnerability in OpenSSL DLL component in Synology BeeDrive for desktop before 1.3.2-13814 allows local users to execute arbitrary code via unspecified vectors.	2026-05-27	7.8
CVE-2026-3623	ibm - netezza_performance_server_replication_services	IBM Netezza Performance Server Replication Services 3.0.2.0 through 3.0.5.0 allows an attacker with low-privileged access to escalate their privileges to root. By exploiting this flaw, the attacker can execute root-level commands, obtain a root shell, and change the root user's password. Successful exploitation also enables modification or removal of system-wide files and the installation of persistent backdoors. This results in full system compromise with complete loss of confidentiality, integrity, and availability.	2026-05-27	7.8
CVE-2026-45852	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix double free in rxe_srq_from_init In rxe_srq_from_init(), the queue pointer 'q' is assigned to 'srq->rq.queue' before copying the SRQ number to user space. If copy_to_user() fails, the function calls rxe_queue_cleanup() to free the queue, but leaves the now-invalid pointer in 'srq->rq.queue'. The caller of rxe_srq_from_init() (rxe_create_srq) eventually calls rxe_srq_cleanup() upon receiving the error, which triggers a second rxe_queue_cleanup() on the same memory, leading to a double free. The call trace looks like this: kmem_cache_free+0x.../0x... rxe_queue_cleanup+0x1a/0x30 [rdma_rxe] rxe_srq_cleanup+0x42/0x60 [rdma_rxe] rxe_elem_release+0x31/0x70 [rdma_rxe] rxe_create_srq+0x12b/0x1a0 [rdma_rxe] ib_create_srq_user+0x9a/0x150 [ib_core] Fix this by moving 'srq->rq.queue = q' after copy_to_user.	2026-05-27	7.8
CVE-2026-45861	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: gfs2: Fix slab-use-after-free in qd_put Commit a475c5dd16e5 ("gfs2: Free quota data objects synchronously") started freeing quota data objects during filesystem shutdown instead of putting them back onto the LRU list, but it failed to remove these objects from the LRU list, causing LRU list corruption. This caused use-after-free when the shrinker (gfs2_qd_shrink_scan) tried to access already-freed objects on the LRU list. Fix this by removing qd objects from the LRU list before freeing them in qd_put(). Initial fix from Deepanshu Kartikey <kartikey406@gmail.com>.	2026-05-27	7.8
CVE-2026-45862	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Flush cache for PASID table before using it When writing the address of a freshly allocated zero-initialized PASID table to a PASID directory entry, do that after the CPU cache flush for this PASID table, not before it, to avoid the time window when this PASID table may be already used by non-coherent IOMMU hardware while its contents in RAM is still some random old data, not zero-initialized.	2026-05-27	7.8
CVE-2026-45878	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: Fix watch_id bounds checking in debug address watch v2 The address watch clear code receives watch_id as an unsigned value (u32), but some helper functions were using a signed int and checked bits by shifting with watch_id. If a very large watch_id is passed from userspace, it can be converted to a negative value. This can cause invalid shifts and may access memory outside the watch_points array. drm/amdkfd: Fix watch_id bounds checking in debug address watch v2 Fix this by checking that watch_id is within MAX_WATCH_ADDRESSES before using it. Also use BIT(watch_id) to test and clear bits safely.	2026-05-27	7.8

		This keeps the behavior unchanged for valid watch IDs and avoids
		undefined behavior for invalid ones.
		
		Fixes the below:
		drivers/gpu/drm/amd/amdgpu/../amdkfd/kfd_debug.c:448
		kfd_dbg_trap_clear_dev_address_watch() error: buffer overflow
		'pdd->watch_points' 4 <= u32max user_rl='0-3,2147483648-u32max' uncapped
		
		drivers/gpu/drm/amd/amdgpu/../amdkfd/kfd_debug.c
		433 int kfd_dbg_trap_clear_dev_address_watch(struct kfd_process_device *pdd,
		434 uint32_t watch_id)
		435 {
		436 int r;
		437
		438 if (!kfd_dbg_owns_dev_watch_id(pdd, watch_id))
		
		kfd_dbg_owns_dev_watch_id() doesn't check for negative values so if
		watch_id is larger than INT_MAX it leads to a buffer overflow.
		(Negative shifts are undefined).
		
		439 return -EINVAL;
		440
		441 if (!pdd->dev->kfd->shared_resources.enable_mes) {
		442 r = debug_lock_and_unmap(pdd->dev->dqm);
		443 if (r)
		444 return r;
		445 }
		446
		447 amdgpu_gfx_off_ctrl(pdd->dev->adev, false);
		--> 448 pdd->watch_points[watch_id] = pdd->dev->kfd2kgd->clear_address_watch(
		449 pdd->dev->adev,
		450 watch_id);
		
		v2: (as per, Jonathan Kim)
		- Add early watch_id >= MAX_WATCH_ADDRESSES validation in the set path to
		match the clear path.
		- Drop the redundant bounds check in kfd_dbg_owns_dev_watch_id().		

CVE-2026-45894	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Clear Present bit before tearing down PASID entry The Intel VT-d Scalable Mode PASID table entry consists of 512 bits (64 bytes). When tearing down an entry, the current implementation zeros the entire 64-byte structure immediately using multiple 64-bit writes. Since the IOMMU hardware may fetch these 64 bytes using multiple internal transactions (e.g., four 128-bit bursts), updating or zeroing the entire entry while it is active (P=1) risks a "torn" read. If a hardware fetch occurs simultaneously with the CPU zeroing the entry, the hardware could observe an inconsistent state, leading to unpredictable behavior or spurious faults. Follow the "Guidance to Software for Invalidations" in the VT-d spec (Section 6.5.3.3) by implementing the recommended ownership handshake: 1. Clear only the 'Present' (P) bit of the PASID entry. 2. Use a dma_wmb() to ensure the cleared bit is visible to hardware before proceeding. 3. Execute the required invalidation sequence (PASID cache, IOTLB, and Device-TLB flush) to ensure the hardware has released all cached references. 4. Only after the flushes are complete, zero out the remaining fields of the PASID entry. Also, add a dma_wmb() in pasid_set_present() to ensure that all other fields of the PASID entry are visible to the hardware before the Present bit is set.	2026-05-27	7.8
CVE-2026-45909	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: Drop __initconst from gates Since commit 8ceff24a754a ("clk: mediatek: clk-gate: Refactor mtk_clk_register_gate to use mtk_gate struct") the mtk_gate structs are no longer just used for initialization/registration, but also at runtime. So drop __initconst annotations.	2026-05-27	7.8
CVE-2026-45910	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix race condition in QP timer handlers I encontered the following warning:	2026-05-27	7.8

		WARNING: drivers/infiniband/sw/rxe/rxe_task.c:249 at rxe_sched_task+0x1c8/0x238 [rdma_rxe], CPU#0: swapper/0/0
		...
		libsha1 [last unloaded: ip6_udp_tunnel]
		CPU: 0 UID: 0 PID: 0 Comm: swapper/0 Tainted: G C 6.19.0-rc5-64k-v8+ #37 PREEMPT
		Tainted: [C]=CRAP
		Hardware name: Raspberry Pi 4 Model B Rev 1.2
		Call trace:
		rxe_sched_task+0x1c8/0x238 [rdma_rxe] (P)
		retransmit_timer+0x130/0x188 [rdma_rxe]
		call_timer_fn+0x68/0x4d0
		__run_timers+0x630/0x888
		...
		WARNING: drivers/infiniband/sw/rxe/rxe_task.c:38 at rxe_sched_task+0x1c0/0x238 [rdma_rxe], CPU#0: swapper/0/0
		...
		WARNING: drivers/infiniband/sw/rxe/rxe_task.c:111 at do_work+0x488/0x5c8 [rdma_rxe], CPU#3: kworker/u17:4/93400
		...
		refcount_t: underflow; use-after-free.
		WARNING: lib/refcount.c:28 at refcount_warn_saturate+0x138/0x1a0, CPU#3: kworker/u17:4/93400
		
		The issue is caused by a race condition between retransmit_timer() and rxe_destroy_qp, leading to the Queue Pair's (QP) reference count dropping to zero during timer handler execution.
		
		It seems this warning is harmless because rxe_qp_do_cleanup() will flush all pending timers and requests.
		
		Example of flow causing the issue:
		
		CPU0 CPU1
		retransmit_timer() {
		spin_lock_irqsave
		rxe_destroy_qp()
		__rxe_cleanup()
		__rxe_put() // qp->ref_count decrease to 0
		rxe_qp_do_cleanup() {
		if (qp->valid) {
		rxe_sched_task() {
		WARN_ON(rxe_read(task->qp) <= 0);
		}
		}
		spin_unlock_irqrestore
		}
		spin_lock_irqsave
		qp->valid = 0
		spin_unlock_irqrestore
		}
		
		Ensure the QP's reference count is maintained and its validity is checked within the timer callbacks by adding calls to rxe_get(qp) and corresponding rxe_put(qp) after use.		

CVE-2026-45929	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ovpn: fix possible use-after-free in ovpn_net_xmit When building the skb_list in ovpn_net_xmit, skb_share_check will free the original skb if it is shared. The current implementation continues to use the stale skb pointer for subsequent operations: - peer lookup, - skb_dst_drop (even though all segments produced by skb_gso_segment will have a dst attached), - ovpn_peer_stats_increment_tx. Fix this by moving the peer lookup and skb_dst_drop before segmentation so that the original skb is still valid when used. Return early if all segments fail skb_share_check and the list ends up empty. Also switch ovpn_peer_stats_increment_tx to use skb_list.next; the next patch fixes the stats logic.	2026-05-27	7.8
CVE-2026-45931	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: accel/amdxdna: Hold mm structure across iommu_sva_unbind_device() Some tests trigger a crash in iommu_sva_unbind_device() due to accessing iommu_mm after the associated mm structure has been freed. Fix this by taking an explicit reference to the mm structure	2026-05-27	7.8

		after successfully binding the device, and releasing it only after the device is unbound. This ensures the mm remains valid for the entire SVA bind/unbind lifetime.		
CVE-2026-45933	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bpf: Preserve id of register in sync_linked_regs() sync_linked_regs() copies the id of known_reg to reg when propagating bounds of known_reg to reg using the off of known_reg, but when known_reg was linked to reg like: known_reg = reg ; both known_reg and reg get same id known_reg += 4 ; known_reg gets off = 4, and its id gets BPF_ADD_CONST now when a call to sync_linked_regs() happens, let's say with the following: if known_reg >= 10 goto pc+2 known_reg's new bounds are propagated to reg but now reg gets BPF_ADD_CONST from the copy. This means if another link to reg is created like: another_reg = reg ; another_reg should get the id of reg but assign_scalar_id_before_mov() sees BPF_ADD_CONST on reg and assigns a new id to it. As reg has a new id now, known_reg's link to reg is broken. If we find new bounds for known_reg, they will not be propagated to reg. This can be seen in the selftest added in the next commit: 0: (85) call bpf_get_prandom_u32#7 ; R0=scalar() 1: (57) r0 &= 255 ; R0=scalar(smin=smin32=0,smax=umax=smax32=umax32=255,var_off=(0x0; 0xff)) 2: (bf) r1 = r0 ; R0=scalar(id=1,smin=smin32=0,smax=umax=smax32=umax32=255,var_off=(0x0; 0xff)) R1=scalar(id=1,smin=smin32=0,smax=umax=smax32=umax32=255,var_off=(0x0; 0xff)) 3: (07) r1 += 4 ; R1=scalar(id=1+4,smin=umin=smin32=umin32=4,smax=umax=smax32=umax32=259,var_off=(0x0; 0x1ff)) 4: (a5) if r1 < 0xa goto pc+4 ; R1=scalar(id=1+4,smin=umin=smin32=umin32=10,smax=umax=smax32=umax32=259,var_off=(0x0; 0x1ff)) 5: (bf) r2 = r0 ; R0=scalar(id=2,smin=umin=smin32=umin32=6,smax=umax=smax32=umax32=255) R2=scalar(id=2,smin=umin=smin32=umin32=6,smax=umax=smax32=umax32=255) 6: (a5) if r1 < 0xe goto pc+2 ; R1=scalar(id=1+4,smin=umin=smin32=umin32=14,smax=umax=smax32=umax32=259,var_off=(0x0; 0x1ff)) 7: (35) if r0 >= 0xa goto pc+1 ; R0=scalar(id=2,smin=umin=smin32=umin32=6,smax=umax=smax32=umax32=9,var_off=(0x0; 0xf)) 8: (37) r0 /= 0 div by zero When 4 is verified, r1's bounds are propagated to r0 but r0 also gets BPF_ADD_CONST (bug). When 5 is verified, r0 gets a new id (2) and its link with r1 is broken. After 6 we know r1 has bounds [14, 259] and therefore r0 should have bounds [10, 255], therefore the branch at 7 is always taken. But because r0's id was changed to 2, r1's new bounds are not propagated to r0. The verifier still thinks r0 has bounds [6, 255] before 7 and execution can reach div by zero. Fix this by preserving id in sync_linked_regs() like off and subreg_def.	2026-05-27	7.8
CVE-2026-45935	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Fix slab-out-of-bounds read in DeleteIndexEntryRoot In the 'DeleteIndexEntryRoot' case of the 'do_action' function, the entry size ('esize') is retrieved from the log record without adequate bounds checking. Specifically, the code calculates the end of the entry ('e2') using: e2 = Add2Ptr(e1, esize); It then calculates the size for memmove using 'PtrOffset(e2, ...)', which subtracts the end pointer from the buffer limit. If 'esize' is maliciously large, 'e2' exceeds the used buffer size. This results in	2026-05-27	7.8

		a negative offset which, when cast to size_t for memmove, interprets as a massive unsigned integer, leading to a heap buffer overflow. This commit adds a check to ensure that the entry size ('esize') strictly fits within the remaining used space of the index header before performing memory operations.		
CVE-2026-45942	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ext4: fix e4b bitmap inconsistency reports A bitmap inconsistency issue was observed during stress tests under mixed huge-page workloads. Ext4 reported multiple e4b bitmap check failures like: ext4_mb_complex_scan_group:2508: group 350, 8179 free clusters as per group info. But got 8192 blocks Analysis and experimentation confirmed that the issue is caused by a race condition between page migration and bitmap modification. Although this timing window is extremely narrow, it is still hit in practice: folio_lock ext4_mb_load_buddy __migrate_folio check ref count folio_mc_copy __filemap_get_folio folio_try_get(folio) ..... mb_mark_used ext4_mb_unload_buddy __folio_migrate_mapping folio_ref_freeze folio_unlock The root cause of this issue is that the fast path of load_buddy only increments the folio's reference count, which is insufficient to prevent concurrent folio migration. We observed that the folio migration process acquires the folio lock. Therefore, we can determine whether to take the fast path in load_buddy by checking the lock status. If the folio is locked, we opt for the slow path (which acquires the lock) to close this concurrency window. Additionally, this change addresses the following issues: When the DOUBLE_CHECK macro is enabled to inspect bitmap-related issues, the following error may be triggered: corruption in group 324 at byte 784(6272): f in copy != ff on disk/prealloc Analysis reveals that this is a false positive. There is a specific race window where the bitmap and the group descriptor become momentarily inconsistent, leading to this error report: ext4_mb_load_buddy ext4_mb_load_buddy __filemap_get_folio(create|lock) folio_lock ext4_mb_init_cache folio_mark_uptodate __filemap_get_folio(no lock) ..... mb_mark_used mb_mark_used_double mb_cmp_bitmaps mb_set_bits(e4b->bd_bitmap) folio_unlock The original logic assumed that since mb_cmp_bitmaps is called when the bitmap is newly loaded from disk, the folio lock would be sufficient to prevent concurrent access. However, this overlooks a specific race condition: if another process attempts to load buddy and finds the folio is already in an uptodate state, it will immediately begin using it without holding folio lock.	2026-05-27	7.8
CVE-2026-45951	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix a potential use-after-free of BTF object Refcounting in the check_pseudo_btf_id() function is incorrect: the __check_pseudo_btf_id() function might get called with a zero refcounted btf. Fix this, and patch related code accordingly.	2026-05-27	7.8

		v3: rephrase a comment (AI) v2: fix a refcount leak introduced in v1 (AI)		
CVE-2026-45956	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/exynos: vidi: use priv->vidi_dev for ctx lookup in vidi_connection_ioctl() vidi_connection_ioctl() retrieves the driver_data from drm_dev->dev to obtain a struct vidi_context pointer. However, drm_dev->dev is the exynos-drm master device, and the driver_data contained therein is not the vidi component device, but a completely different device. This can lead to various bugs, ranging from null pointer dereferences and garbage value accesses to, in unlucky cases, out-of-bounds errors, use-after-free errors, and more. To resolve this issue, we need to store/delete the vidi device pointer in exynos_drm_private->vidi_dev during bind/unbind, and then read this exynos_drm_private->vidi_dev within ioctl() to obtain the correct struct vidi_context pointer.	2026-05-27	7.8
CVE-2026-45959	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Fix a crash due to incorrect cleanup usage of kfree Annotating a local pointer variable, which will be assigned with the kmalloc-family functions, with the `__cleanup(kfree)` attribute will make the address of the local variable, rather than the address returned by kmalloc, passed to kfree directly and lead to a crash due to invalid deallocation of stack address. According to other places in the repo, the correct usage should be `__free(kfree)`. The code coincidentally compiled because the parameter type `void *` of kfree is compatible with the desired type `struct { ... } **`.	2026-05-27	7.8
CVE-2026-45970	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bonding: alb: fix UAF in rlb_arp_recv during bond up/down The ALB RX path may access rx_hashtbl concurrently with bond teardown. During rapid bond up/down cycles, rlb_deinitialize() frees rx_hashtbl while RX handlers are still running, leading to a null pointer dereference detected by KASAN. However, the root cause is that rlb_arp_recv() can still be accessed after setting recv_probe to NULL, which is actually a use-after-free (UAF) issue. That is the reason for using the referenced commit in the Fixes tag. [214.174138] Oops: general protection fault, probably for non-canonical address 0xdffffc000000001d: 0000 [#1] SMP KASAN PTI [214.186478] KASAN: null-ptr-deref in range [0x00000000000000e8-0x00000000000000ef] [214.194933] CPU: 30 UID: 0 PID: 2375 Comm: ping Kdump: loaded Not tainted 6.19.0-rc8+ #2 PREEMPT(voluntary) [214.205907] Hardware name: Dell Inc. PowerEdge R730/0WCJNT, BIOS 2.14.0 01/14/2022 [214.214357] RIP: 0010:rlb_arp_recv+0x505/0xab0 [bonding] [214.220320] Code: 0f 85 2b 05 00 00 48 b8 00 00 00 00 00 fc ff df 40 0f b6 ed 48 c1 e5 06 49 03 ad 78 01 00 00 48 8d 7d 28 48 89 fa 48 c1 ea 03 <0f> b6 04 02 84 c0 74 06 0f 8e 12 05 00 00 80 7d 28 00 0f 84 8c 00 [214.241280] RSP: 0018:ffffc900073d8870 EFLAGS: 00010206 [214.247116] RAX: dfffffc0000000000 RBX: ffff888168556822 RCX: ffff88816855681e [214.255082] RDX: 000000000000001d RSI: dfffffc0000000000 RDI: 00000000000000e8 [214.263048] RBP: 00000000000000c0 R08: 0000000000000002 R09: ffffed11192021c8 [214.271013] R10: ffff8888c9010e43 R11: 0000000000000001 R12: 1ffff92000e7b119 [214.278978] R13: ffff8888c9010e00 R14: ffff888168556822 R15: ffff888168556810 [214.286943] FS: 00007f85d2d9cb80(0000) GS:ffff88886ccb3000(0000) knlGS:0000000000000000 [214.295966] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [214.302380] CR2: 00007f0d047b5e34 CR3: 00000008a1c2e002 CR4: 00000000001726f0 [214.310347] Call Trace: [214.313070] <IRQ> [214.315318] ? __pfx_rlb_arp_recv+0x10/0x10 [bonding] [214.320975] bond_handle_frame+0x166/0xb60 [bonding] [214.326537] ? __pfx_bond_handle_frame+0x10/0x10 [bonding] [214.332680] __netif_receive_skb_core.constprop.0+0x576/0x2710 [214.339199] ? __pfx_arp_process+0x10/0x10 [214.343775] ? sched_balance_find_src_group+0x98/0x630 [214.349513] ? __pfx___netif_receive_skb_core.constprop.0+0x10/0x10 [214.356513] ? arp_rcv+0x307/0x690 [214.360311] ? __pfx_arp_rcv+0x10/0x10 [214.364499] ? __lock_acquire+0x58c/0xbd0 [214.368975] __netif_receive_skb_one_core+0xae/0x1b0 [214.374518] ? __pfx___netif_receive_skb_one_core+0x10/0x10 [214.380743] ? lock_acquire+0x10b/0x140 [214.385026] process_backlog+0x3f1/0x13a0	2026-05-27	7.8

		[214.389502] ? process_backlog+0x3aa/0x13a0 [214.394174] __napi_poll.constprop.0+0x9f/0x370 [214.399233] net_rx_action+0x8c1/0xe60 [214.403423] ? __pfx_net_rx_action+0x10/0x10 [214.408193] ? lock_acquire.part.0+0xbd/0x260 [214.413058] ? sched_clock_cpu+0x6c/0x540 [214.417540] ? mark_held_locks+0x40/0x70 [214.421920] handle_softirqs+0x1fd/0x860 [214.426302] ? __pfx_handle_softirqs+0x10/0x10 [214.431264] ? __neigh_event_send+0x2d6/0xf50 [214.436131] do_softirq+0xb1/0xf0 [214.439830] </IRQ> The issue is reproducible by repeatedly running ip link set bond0 up/down while receiving ARP messages, where rlb_arp_recv() can race with rlb_deinitialize() and dereference a freed rx_hashtbl entry. Fix this by setting recv_probe to NULL and then calling synchronize_net() to wait for any concurrent RX processing to finish. This ensures that no RX handler can access rx_hashtbl after it is freed in bond_alb_deinitialize().		
CVE-2026-45980	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: accel/amdxdna: Stop job scheduling across aie2_release_resource() Running jobs on a hardware context while it is in the process of releasing resources can lead to use-after-free and crashes. Fix this by stopping job scheduling before calling aie2_release_resource() and restarting it after the release completes. Additionally, aie2_sched_job_run() now checks whether the hardware context is still active.	2026-05-27	7.8
CVE-2026-45984	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: gfs2: Fix use-after-free in iomap inline data write path The inline data buffer head (dibh) is being released prematurely in gfs2_iomap_begin() via release_metapath() while iomap->inline_data still points to dibh->b_data. This causes a use-after-free when iomap_write_end_inline() later attempts to write to the inline data area. The bug sequence: 1. gfs2_iomap_begin() calls gfs2_meta_inode_buffer() to read inode metadata into dibh 2. Sets iomap->inline_data = dibh->b_data + sizeof(struct gfs2_dinode) 3. Calls release_metapath() which calls brelse(dibh), dropping refcount to 0 4. kswapd reclaims the page (~39ms later in the syzbot report) 5. iomap_write_end_inline() tries to memcpy() to iomap->inline_data 6. KASAN detects use-after-free write to freed memory Fix by storing dibh in iomap->private and incrementing its refcount with get_bh() in gfs2_iomap_begin(). The buffer is then properly released in gfs2_iomap_end() after the inline write completes, ensuring the page stays alive for the entire iomap operation. Note: A C reproducer is not available for this issue. The fix is based on analysis of the KASAN report and code review showing the buffer head is freed before use. [agruenba: Take buffer head reference in gfs2_iomap_begin() to avoid leaks in gfs2_iomap_get() and gfs2_iomap_alloc().]	2026-05-27	7.8
CVE-2026-45991	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: udf: fix partition descriptor append bookkeeping Mounting a crafted UDF image with repeated partition descriptors can trigger a heap out-of-bounds write in part_descs_loc[]. handle_partition_descriptor() deduplicates entries by partition number, but appended slots never record partnum. As a result duplicate Partition Descriptors are appended repeatedly and num_part_descs keeps growing. Once the table is full, the growth path still sizes the allocation from partnum even though inserts are indexed by num_part_descs. If partnum is already aligned to PART_DESC_ALLOC_STEP, ALIGN(partnum, step) can keep the old capacity and the next append writes past the end of the table.	2026-05-27	7.8

		Store partnum in the appended slot and size growth from the next append count so deduplication and capacity tracking follow the same model.		
CVE-2026-46006	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau: fix u32 overflow in pushbuf reloc bounds check nouveau_gem_pushbuf_reloc_apply() validates each relocation with if (r->reloc_bo_offset + 4 > nvbo->bo.base.size) but reloc_bo_offset is __u32 (uapi/drm/nouveau_drm.h) and the integer literal 4 promotes to unsigned int, so the addition is performed in 32 bits and wraps before the comparison against the size_t bo size. Cast to u64 so the addition happens in 64-bit arithmetic. [Add Fixes: tag. - Danilo]	2026-05-27	7.8

CVE-2026-46011	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		media: mtk-jpeg: fix use-after-free in release path due to uncancelled work
		
		The mtk_jpeg_release() function frees the context structure (ctx) without
		first cancelling any pending or running work in ctx->jpeg_work. This
		creates a race window where the workqueue callback may still be accessing
		the context memory after it has been freed.
		
		Race condition:
		
		CPU 0 (release) CPU 1 (workqueue)
		---------- ----------
		close()
		mtk_jpeg_release()
		mtk_jpegenc_worker()
		ctx = work->data
		// accessing ctx
		
		kfree(ctx) // freed!
		access ctx // UAF!
		
		The work is queued via queue_work() during JPEG encode/decode operations
		(via mtk_jpeg_device_run). If the device is closed while work is pending
		or running, the work handler will access freed memory.
		
		Fix this by calling cancel_work_sync() BEFORE acquiring the mutex. This
		ordering is critical: if cancel_work_sync() is called after mutex_lock(),
		and the work handler also tries to acquire the same mutex, it would cause
		a deadlock.
		
		Note: The open error path does NOT need cancel_work_sync() because
		INIT_WORK() only initializes the work structure - it does not schedule
		it. Work is only scheduled later during ioctl operations.
			2026-05-27	7.8

CVE-2026-46015	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: tcp: call sk_data_ready() after listener migration When inet_csk_listen_stop() migrates an established child socket from a closing listener to another socket in the same SO_REUSEPORT group, the target listener gets a new accept-queue entry via inet_csk_reqsk_queue_add(), but that path never notifies the target listener's waiters. A nonblocking accept() still works because it checks the queue directly, but poll()/epoll_wait() waiters and blocking accept() callers can also remain asleep indefinitely. Call READ_ONCE(nsk->sk_data_ready)(nsk) after a successful migration in inet_csk_listen_stop(). However, after inet_csk_reqsk_queue_add() succeeds, the ref acquired in reuseport_migrate_sock() is effectively transferred to nreq->rsk_listener. Another CPU can then dequeue nreq via accept() or listener shutdown, hit reqsk_put(), and drop that listener ref. Since listeners are SOCK_RCU_FREE, wrap the post-queue_add() dereferences of nsk in rcu_read_lock()/rcu_read_unlock(), which also covers the existing sock_net(nsk) access in that path. The reqsk_timer_handler() path does not need the same changes for two reasons: half-open requests become readable only after the final ACK, where tcp_child_process() already wakes the listener; and once nreq is visible via inet_ehash_insert(), the success path no longer touches nsk directly.	2026-05-27	7.8

CVE-2026-46036	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: vfio/cdx: Serialize VFIO_DEVICE_SET_IRQS with a per-device mutex vfio_cdx_set_msi_trigger() reads vdev->config_msi and operates on the vdev->cdx_irqs array based on its value, but provides no serialization against concurrent VFIO_DEVICE_SET_IRQS ioctls. Two callers can race such that one observes config_msi as set while another clears it and frees cdx_irqs via vfio_cdx_msi_disable(), resulting in a use-after-free of the cdx_irqs array. Add a cdx_irqs_lock mutex to struct vfio_cdx_device and acquire it in vfio_cdx_set_msi_trigger(), which is the single chokepoint through which all updates to config_msi, cdx_irqs, and msi_count flow, covering both the ioctl path and the close-device cleanup path. This keeps the test of config_msi atomic with the subsequent enable, disable, or trigger operations. Drop the pre-call !cdx_irqs test from vfio_cdx_irqs_cleanup() as part of this change: the optimization it provided is redundant with the !config_msi early-return inside vfio_cdx_msi_disable(), and leaving the test in place would be an unsynchronized read of state the new lock is meant to protect.	2026-05-27	7.8
CVE-2026-46053	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: rds: fix MR cleanup on copy error __rds_rdma_map() hands sg/pages ownership to the transport after get_mr() succeeds. If copying the generated cookie back to user space fails after that point, the error path must not free those resources again before dropping the MR reference. Remove the duplicate unpin/free from the put_user() failure branch so that MR teardown is handled only through the existing final cleanup path.	2026-05-27	7.8
CVE-2026-46058	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: media: amphion: Fix race between m2m job_abort and device_run Fix kernel panic caused by race condition where v4l2_m2m_ctx_release() frees m2m_ctx while v4l2_m2m_try_run() is about to call device_run with the same context. Race sequence: v4l2_m2m_try_run(): v4l2_m2m_ctx_release(): lock/unlock v4l2_m2m_cancel_job() job_abort() v4l2_m2m_job_finish() kfree(m2m_ctx) <- frees ctx device_run() <- use-after-free crash at 0x538 Crash trace: Unable to handle kernel read from unreadable memory at virtual address 0000000000000538 v4l2_m2m_try_run+0x78/0x138 v4l2_m2m_device_run_work+0x14/0x20 The amphion vpu driver does not rely on the m2m framework's device_run callback to perform encode/decode operations. Fix the race by preventing m2m framework job scheduling entirely: - Add job_ready callback returning 0 (no jobs ready for m2m framework) - Remove job_abort callback to avoid the race condition	2026-05-27	7.8
CVE-2026-46062	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ntfs3: fix integer overflow in run_unpack() volume boundary check The volume boundary check `lcn + len > sbi->used.bitmap.nbits` uses raw addition which can wrap around for large lcn and len values, bypassing the validation. Use check_add_overflow() as is already done for the adjacent prev_lcn + dlcn and vcn64 + len checks added by commit 3ac37e100385 ("ntfs3: Fix integer overflow in run_unpack()"). Found by fuzzing with a source-patched harness (LibAFL + QEMU).	2026-05-27	7.8
CVE-2026-46065	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: fbdev: defio: Disconnect deferred I/O from the lifetime of struct fb_info Hold state of deferred I/O in struct fb_deferred_io_state. Allocate an instance as part of initializing deferred I/O and remove it only after	2026-05-27	7.8

		the final mapping has been closed. If the fb_info and the contained deferred I/O meanwhile goes away, clear struct fb_deferred_io_state.info to invalidate the mapping. Any access will then result in a SIGBUS signal. Fixes a long-standing problem, where a device hot-unplug happens while user space still has an active mapping of the graphics memory. The hot-unplug frees the instance of struct fb_info. Accessing the memory will operate on undefined state.		
CVE-2026-46081	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: crypto: acomp - fix wrong pointer stored by acomp_save_req() acomp_save_req() stores &req->chain in req->base.data. When acomp_reqchain_done() is invoked on asynchronous completion, it receives &req->chain as the data argument but casts it directly to struct acomp_req. Since data points to the chain member, all subsequent field accesses are at a wrong offset, resulting in memory corruption. The issue occurs when an asynchronous hardware implementation, such as the QAT driver, completes a request that uses the DMA virtual address interface (e.g. acomp_request_set_src_dma()). This combination causes crypto_acomp_compress() to enter the acomp_do_req_chain() path, which sets acomp_reqchain_done() as the completion callback via acomp_save_req(). With KASAN enabled, this manifests as a general protection fault in acomp_reqchain_done(): general protection fault, probably for non-canonical address 0xe000040000000000 KASAN: probably user-memory-access in range [0x0000400000000000-0x0000400000000007] RIP: 0010:acomp_reqchain_done+0x15b/0x4e0 Call Trace: <IRQ> qat_comp_alg_callback+0x5d/0xa0 [intel_qat] adf_ring_response_handler+0x376/0x8b0 [intel_qat] adf_response_handler+0x60/0x170 [intel_qat] tasklet_action_common+0x223/0x820 handle_softirqs+0x1ab/0x640 </IRQ> Fix this by storing the request itself in req->base.data instead of &req->chain, so that acomp_reqchain_done() receives the correct pointer. Simplify acomp_restore_req() accordingly to access req->chain directly.	2026-05-27	7.8
CVE-2026-46090	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ALSA: aloop: Fix peer runtime UAF during format-change stop loopback_check_format() may stop the capture side when playback starts with parameters that no longer match a running capture stream. Commit 826af7fa62e3 ("ALSA: aloop: Fix racy access at PCM trigger") moved the peer lookup under cable->lock, but the actual snd_pcm_stop() still runs after dropping that lock. A concurrent close can clear the capture entry from cable->streams[] and detach or free its runtime while the playback trigger path still holds a stale peer substream pointer. Keep a per-cable count of in-flight peer stops before dropping cable->lock, and make free_cable() wait for those stops before detaching the runtime. This preserves the existing behavior while making the peer runtime lifetime explicit.	2026-05-27	7.8
CVE-2026-46093	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mm/vmalloc: take vmap_purge_lock in shrinker decay_va_pool_node() can be invoked concurrently from two paths: __purge_vmap_area_lazy() when pools are being purged, and the shrinker via vmap_node_shrink_scan(). However, decay_va_pool_node() is not safe to run concurrently, and the shrinker path currently lacks serialization, leading to races and possible leaks. Protect decay_va_pool_node() by taking vmap_purge_lock in the shrinker path to ensure serialization with purge users.	2026-05-27	7.8
CVE-2026-46100	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: fs: afs: revert mmap_prepare() change	2026-05-27	7.8

		Partially reverts commit 9d5403b1036c ("fs: convert most other generic_file_*mmap() users to .mmap_prepare()"). This is because the .mmap invocation establishes a refcount, but .mmap_prepare is called at a point where a merge or an allocation failure might happen after the call, which would leak the refcount increment. Functionality is being added to permit the use of .mmap_prepare in this case, but in the interim, we need to fix this.		
CVE-2026-45322	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. Microsoft UFO tagged releases up to and including v3.0.0 contain an OS command injection vulnerability in the shell action replay path. In affected releases, ShellReceiver.run_shell() passes a command string from action parameters directly to subprocess.Popen() with shell=True and executable=powershell.exe. The same shell-execution behavior is also reachable through ShellReceiver.execute_command(). The shell receiver is invoked by action classes such as RunShellCommand.execute() and ExecuteCommand.execute(), which forward stored action parameters to the shell receiver. Because UFO stores planned and executed actions in per-session JSON records, an attacker who can write or modify a session/action JSON file can plant a shell action. When the session is resumed or replayed, UFO executes the attacker's command as the UFO process user.	2026-05-27	7.8
CVE-2026-46105	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: scsi: mpt3sas: Limit NVMe request size to 2 MiB The HBA firmware reports NVMe MDTS values based on the underlying drive capability. However, because the driver allocates a fixed 4K buffer for the PRP list, accommodating at most 512 entries, the driver supports a maximum I/O transfer size of 2 MiB. Limit max_hw_sectors to the smaller of the reported MDTS and the 2 MiB driver limit to prevent issuing oversized I/O that may lead to a kernel oops.	2026-05-28	7.8
CVE-2026-46107	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: dm-thin: fix metadata refcount underflow There's a bug in dm-thin in the function rebalance_children. If the internal btree node has one entry, the code tries to copy all btree entries from the node's child to the node itself and then decrement the child's reference count. If the child node is shared (it has reference count > 1), we won't free it, so there would be two pointers to each of the grandchildren nodes. But the reference counts of the grandchildren is not increased, thus the reference count doesn't match the number of pointers that point to the grandchildren. This results in "device mapper: space map common: unable to decrement block" errors. Fix this bug by incrementing reference counts on the grandchildren if the btree node is shared.	2026-05-28	7.8
CVE-2026-46111	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_conn: fix potential UAF in create_big_sync Add hci_conn_valid() check in create_big_sync() to detect stale connections before proceeding with BIG creation. Handle the resulting -ECANCELED in create_big_complete() and re-validate the connection under hci_dev_lock() before dereferencing, matching the pattern used by create_le_conn_complete() and create_pa_complete(). Keep the hci_conn object alive across the async boundary by taking a reference via hci_conn_get() when queueing create_big_sync(), and dropping it in the completion callback. The refcount and the lock are complementary: the refcount keeps the object allocated, while hci_dev_lock() serializes hci_conn_hash_del()'s list_del_rcu() on hdev->conn_hash, as required by hci_conn_del(). hci_conn_put() is called outside hci_dev_unlock() so the final put (which resolves to kfree() via bt_link_release) does not run under hdev->lock, though the release path would be safe either way. Without this, create_big_complete() would unconditionally dereference the conn pointer on error, causing a use-after-free via hci_connect_cfm() and hci_conn_del().	2026-05-28	7.8
CVE-2026-46112	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/hns: Fix unlocked call to hns_roce_qp_remove() Sashiko points out that hns_roce_qp_remove() requires the caller to hold locks. The error flow in hns_roce_create_qp_common() doesn't hold those	2026-05-28	7.8

		locks for the error unwind so it risks corrupting memory. Grab the same locks the other two callers use.		

CVE-2026-46116	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		xfrm: defensively unhash xfrm_state lists in __xfrm_state_delete
		
		KASAN reproduces a slab-use-after-free in __xfrm_state_delete()'s
		hlist_del_rcu calls under syzkaller load on linux-6.12.y stable
		(reproduced on 6.12.47, also reachable via the same code path on
		torvalds/master and on the ipsec tree). Nine unique signatures cluster
		in the xfrm_state lifecycle, the load-bearing one being:
		
		BUG: KASAN: slab-use-after-free in __hlist_del include/linux/list.h:990 [inline]
		BUG: KASAN: slab-use-after-free in hlist_del_rcu include/linux/rculist.h:516 [inline]
		BUG: KASAN: slab-use-after-free in __xfrm_state_delete net/xfrm/xfrm_state.c
		Write of size 8 at addr ffff8881198bcb70 by task kworker/u8:9/435
		
		Workqueue: netns cleanup_net
		Call Trace:
		__hlist_del / hlist_del_rcu
		__xfrm_state_delete
		xfrm_state_delete
		xfrm_state_flush
		xfrm_state_fini
		ops_exit_list
		cleanup_net
		
		The other observed signatures hit the same slab object from
		__xfrm_state_lookup, xfrm_alloc_spi, __xfrm_state_insert and an OOB
		write variant of __xfrm_state_delete, all on the byseq/byspi
		hash chains.
		
		__xfrm_state_delete() guards its byseq and byspi unhashes with
		value-based predicates:
		
		if (x->km.seq)
		hlist_del_rcu(&x->byseq);
		if (x->id.spi)
		hlist_del_rcu(&x->byspi);
		
		while everywhere else in the file (e.g. state_cache, state_cache_input)
		the safer hlist_unhashed() check is used. xfrm_alloc_spi() sets
		x->id.spi = newspi inside xfrm_state_lock and then immediately inserts
		into byspi, but a path that observes x->id.spi != 0 outside of
		xfrm_state_lock can still skip-or-hit the byspi unhash inconsistently
		with whether x is actually on the list. The same holds for x->km.seq
		versus byseq, and the bydst/bysrc unhashes have no predicate at all,
		so a second __xfrm_state_delete() on the same object writes through
		LIST_POISON pprev.
		
		The defensive change here:
		
		- Use hlist_del_init_rcu() instead of hlist_del_rcu() on bydst,
		bysrc, byseq and byspi so a second deletion is a no-op rather
		than a write through LIST_POISON pprev. The byseq/byspi nodes
		are already initialised in xfrm_state_alloc().
		- Test hlist_unhashed() rather than the value predicate for
		byseq/byspi, so the unhash decision tracks list state rather than
		mutable scalar fields.
		
		Empirical verification: applied this patch on top of v6.12.47, rebuilt,
		and re-ran the same syzkaller harness for 1h16m on a previously-crashy
		configuration that produced ~100 hits each of slab-use-after-free
		Read in xfrm_alloc_spi / Read in __xfrm_state_lookup / Write in
		__xfrm_state_delete. After the patch, 7.1M execs across 32 VMs at
		~1550 exec/sec produced zero xfrm_state UAF/OOB hits. /proc/slabinfo
		confirms the xfrm_state slab is actively allocated and freed during
		the run (~143 KiB resident), so the fuzzer is still exercising those
		code paths -- they just no longer crash.
		
		Reproduction:
		
		- Linux 6.12.47 x86_64 + KASAN_GENERIC + KASAN_INLINE + KCOV
		- syzkaller @ 746545b8b1e4c3a128db8652b340d3df90ce61db
		- 32 QEMU/KVM VMs x 2 vCPU on AWS c5.metal bare metal
		- 9 unique signatures collected in ~9h, all within xfrm_state
		lifecycle
			2026-05-28	7.8

CVE-2026-46117	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2026-05-28	7.8

		RDMA/mana: Remove user triggerable WARN_ON() in mana_ib_create_qp_rss() Sashiko points out that the user can specify WQs sharing the same CQ as a part of the uAPI and this will trigger the WARN_ON() then go on to corrupt the kernel. Just reject it outright and fail the QP creation.		
CVE-2026-46120	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ip6_gre: Use cached t->net in ip6erspan_changelink(). After commit 5e72ce3e3980 ("net: ipv6: Use link netns in newlink() of rtnl_link_ops"), ip6erspan_newlink() correctly resolves the per-netns ip6gre hash via link_net. ip6erspan_changelink() was not converted in that series and still uses dev_net(dev), which diverges from the device's creation netns after IFLA_NET_NS_FD migration. This re-inserts the tunnel into the wrong per-netns hash. The original netns keeps a stale entry. When that netns is later destroyed, ip6gre_exit_rtnl_net() walks the stale entry, producing a slab-use-after-free reported by KASAN, followed by a kernel BUG at net/core/dev.c (LIST_POISON1) in unregister_netdevice_many_notify(). Reachable from an unprivileged user namespace (unshare --user --map-root-user --net). ip6gre_changelink() earlier in the same file already uses the cached t->net; only ip6erspan_changelink() has the wrong shape.	2026-05-28	7.8
CVE-2026-46129	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix double free in create_space_info() error path When kobject_init_and_add() fails, the call chain is: create_space_info() -> btrfs_sysfs_add_space_info_type() -> kobject_init_and_add() -> failure -> kobject_put(&space_info->kobj) -> space_info_release() -> kfree(space_info) Then control returns to create_space_info(): btrfs_sysfs_add_space_info_type() returns error -> goto out_free -> kfree(space_info) This causes a double free. Keep the direct kfree(space_info) for the earlier failure path, but after btrfs_sysfs_add_space_info_type() has called kobject_put(), let the kobject release callback handle the cleanup.	2026-05-28	7.8
CVE-2026-46145	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/mana: Validate rx_hash_key_len Sashiko points out that rx_hash_key_len comes from a uAPI structure and is blindly passed to memcpy, allowing the userspace to trash kernel memory. Bounds check it so the memcpy cannot overflow.	2026-05-28	7.8
CVE-2026-46157	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ALSA: pcm: oss: Fix data race at accessing runtime.oss.trigger Currently the runtime.oss.trigger field may be accessed concurrently without protection, which may lead to the data race. And, in this case, it may lead to more severe problem because it's a bit field; as writing the data, it may overwrite other bit fields as well, which confuses the operation completely, as spotted by fuzzing. Fix it by covering runtime.oss.trigger bit fled also with the existing params_lock mutex in both snd_pcm_oss_get_trigger() and snd_pcm_oss_poll().	2026-05-28	7.8
CVE-2026-46173	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: exit: prevent preemption of oopsing TASK_DEAD task When an already-exiting task oopses, make_task_dead() currently calls do_task_dead() with preemption enabled. That is forbidden: do_task_dead() calls __schedule(), which has a comment saying "WARNING:	2026-05-28	7.8

		must be called with preemption disabled!". If an oopsing task is preempted in do_task_dead(), between becoming TASK_DEAD and entering the scheduler explicitly, bad things happen: finish_task_switch() assumes that once the scheduler has switched away from a TASK_DEAD task, the task can never run again and its stack is no longer needed; but that assumption apparently doesn't hold if the dead task was preempted (the SM_PREEMPT case). This means that the scheduler ends up repeatedly dropping references on the dead task's stack, which can lead to use-after-free or double-free of the entire task stack; in other words, two tasks can end up running on the same stack, resulting in various kinds of memory corruption. (This does not just affect "recursively oopsing" tasks; it is enough to oops once during task exit, for example in a file_operations::release handler)		
CVE-2026-46176	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx5: Fix error path fall-through in mlx5_ib_dev_res_srq_init() mlx5_ib_dev_res_srq_init() allocates two SRQs, s0 and s1. When ib_create_srq() fails for s1, the error branch destroys s0 but falls through and unconditionally assigns the freed s0 and the ERR_PTR s1 to devr->s0 and devr->s1. This leads to several problems: the lock-free fast path checks "if (devr->s1) return 0;" and treats the ERR_PTR as already initialised; users in mlx5_ib_create_qp() dereference the freed SRQ or ERR_PTR via to_msrq(devr->s0)->msrq.srqn; and mlx5_ib_dev_res_cleanup() dereferences the ERR_PTR and double-frees s0 on teardown. Fix by adding the same `goto unlock` in the s1 failure path.	2026-05-28	7.8
CVE-2026-46178	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx4: Fix resource leak on error in mlx4_ib_create_srq() Sashiko points out that mlx4_srq_alloc() was not undone during error unwind, add the missing call to mlx4_srq_free().	2026-05-28	7.8
CVE-2026-46181	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx4: Fix mis-use of RCU in mlx4_srq_event() Sashiko points out the radix_tree itself is RCU safe, but nothing ever frees the mlx4_srq struct with RCU, and it isn't even accessed within the RCU critical section. It also will crash if an event is delivered before the srq object is finished initializing. Use the spinlock since it isn't easy to make RCU work, use refcount_inc_not_zero() to protect against partially initialized objects, and order the refcount_set() to be after the srq is fully initialized.	2026-05-28	7.8
CVE-2026-46197	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: validate SVM ioctl nattr against buffer size Validate nattr field against the buffer size, preventing out-of-bounds buffer access via user-controlled attribute count. (cherry picked from commit 5eca8bfdfa456c3304ca77523718fe24254c172f)	2026-05-28	7.8
CVE-2026-46201	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Fix dma-buf attachment leak in xe_gem_prime_import() When xe_dma_buf_init_obj() fails, the attachment from dma_buf_dynamic_attach() is not detached. Add dma_buf_detach() before returning the error. Note: we cannot use goto out_err here because xe_dma_buf_init_obj() already frees bo on failure, and out_err would double-free it. (cherry picked from commit a828eb185aac41800df8eae4b60501ccc0dbbe51)	2026-05-28	7.8
CVE-2026-46205	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: staging: media: atomisp: Disallow all private IOCTLs Disallow all private IOCTLs. These aren't quite as safe as one could assume of IOCTL handlers; disable them for now. Instead of removing the code, return in the beginning of the function if cmd is non-zero in order to keep static checkers happy.	2026-05-28	7.8
CVE-2026-46206	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2026-05-28	7.8

		batman-adv: reject new tp_meter sessions during teardown Prevent tp_meter from starting new sender or receiver sessions after mesh_state has left BATADV_MESH_ACTIVE.		
CVE-2026-46208	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: batman-adv: stop tp_meter sessions during mesh teardown TP meter sessions remain linked on bat_priv->tp_list after the netlink request has already finished. When the mesh interface is removed, batadv_mesh_free() currently tears down the mesh without first draining these sessions. A running sender thread or a late incoming tp_meter packet can then keep processing against a mesh instance which is already shutting down. Synchronize tp_meter with the mesh lifetime by stopping all active sessions from batadv_mesh_free() and waiting for sender threads to exit before teardown continues.	2026-05-28	7.8

CVE-2026-46209	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		drm/gem: Fix inconsistent plane dimension calculation in drm_gem_fb_init_with_funcs()
		
		drm_gem_fb_init_with_funcs() computes sub-sampled plane dimensions using plain integer division:
		
		unsigned int width = mode_cmd->width / (i ? info->hsub : 1);
		unsigned int height = mode_cmd->height / (i ? info->vsub : 1);
		
		However, the ioctl-level framebuffer_check() in drm_framebuffer.c uses drm_format_info_plane_width/height() which round up dimensions via DIV_ROUND_UP(). This inconsistency corrupts the subsequent GEM object size check for certain pixel format and dimension combinations.
		
		For example, with NV12 (vsub=2) and a 1-pixel-tall framebuffer the GEM size validation path sees height=0 instead of height=1. The expression (height - 1) then wraps to UINT_MAX as an unsigned int, causing min_size to overflow and wrap back to a small value. A tiny GEM object therefore passes the size guard, yet when the GPU accesses the chroma plane it will read or write memory beyond the object's bounds.
		
		Fix by replacing the open-coded divisions with drm_format_info_plane_width() and drm_format_info_plane_height(), which use DIV_ROUND_UP() and match the calculation already used in framebuffer_check().
			2026-05-28	7.8

CVE-2026-46210	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: media: iris: fix use-after-free of fmt_src during MBPF check During concurrency testing, multiple instances can run in parallel, and each instance uses its own inst->lock while the core->lock protects the list of active instances. The race happens because these locks cover different scopes, inst->lock protects only the internals of a single instance, while the Macro Blocks Per Frame (MBPF) checker walks the core list under core->lock and reads fields like fmt_src->width and fmt_src->height. At the same time, iris_close() may free fmt_src and fmt_dst under inst->lock while the instance is still present in the core list. This allows a situation where the MBPF checker, still iterating through the core list, reaches an instance whose fmt_src was already freed by another thread and ends up dereferencing a dangling pointer, resulting in a use-after-free. This happens because the MBPF checker assumes that any instance in the core list is fully valid, but the freeing of fmt_src and fmt_dst without removing the instance from the core list is not correct. The correct ordering is to defer freeing fmt_src and fmt_dst until after the instance has been removed from the core list and all teardown under the core lock has completed, ensuring that no dangling pointers are ever exposed during MBPF checks.	2026-05-28	7.8
CVE-2026-46215	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm: Set old handle to NULL before prime swap in change_handle There was a potential race condition in change_handle. The ioctl briefly had a single object with two idr entries; a concurrent gem_close could delete the object and remove one of the handles while leaving the other one dangling, which could subsequently be dereferenced for a use-after-free. To fix this, do the same dance that gem_close itself does. (f6cd7daecff5 drm: Release driver references to handle before making it available again)	2026-05-28	7.8

		First idr_replace the old handle to NULL. Later, if the prime operations are successful, actually close it. create_tail required a similar dance to avoid a similar problem. (bd46cece51a3 drm/gem: Fix race in drm_gem_handle_create_tail()) It idr_allocs the new handle with NULL, then swaps in the correct object later to avoid races. We don't need to do that here, since the only operations that could race are drm_prime, and change_handle holds the prime lock for the entire duration. v2: cleanups of error paths		
CVE-2026-46227	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: sctp: revalidate list cursor after sctp_sendmsg_to_asoc() in SCTP_SENDALL The SCTP_SENDALL path in sctp_sendmsg() iterates ep->asocs with list_for_each_entry_safe(), which caches the next entry in @tmp before the loop body runs. The body calls sctp_sendmsg_to_asoc(), which may drop the socket lock inside sctp_wait_for_sndbuf(). While the lock is dropped, another thread can SCTP_SOCKOPT_PEELOFF the association cached in @tmp, migrating it to a new endpoint via sctp_sock_migrate() (list_del_init() + list_add_tail() to newep->asocs), and optionally close the new socket which frees the association via kfree_rcu(). The cached @tmp can also be freed by a network ABORT for that association, processed in softirq while the lock is dropped. sctp_wait_for_sndbuf() revalidates @asoc (the current entry) on re-lock via the "sk != asoc->base.sk" and "asoc->base.dead" checks, but nothing revalidates @tmp. After a successful return, the iterator advances to the stale @tmp, yielding either a use-after-free (if the peeled socket was closed) or a list-walk onto the new endpoint's list head (type confusion of &newep->asocs as a struct sctp_association *). Both are reachable from CapEff=0; the type-confusion path gives controlled indirect call via the outqueue.sched->init_sid pointer. Fix by re-deriving @tmp from @asoc after sctp_sendmsg_to_asoc() returns. @asoc is known to still be on ep->asocs at that point: the only callers that list_del an association from ep->asocs are sctp_association_free() (which sets asoc->base.dead) and sctp_assoc_migrate() (which changes asoc->base.sk), and sctp_wait_for_sndbuf() checks both under the lock before any successful return; a tripped check propagates as err < 0 and the loop bails before the re-derive. The SCTP_ABORT path in sctp_sendmsg_check_sflags() returns 0 and the loop hits 'continue' before sctp_sendmsg_to_asoc() is ever called, so the @tmp cached by list_for_each_entry_safe() still covers the lock-held free that ba59fb027307 ("sctp: walk the list of asoc safely") was added for.	2026-05-28	7.8
CVE-2026-46240	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: media: iris: Fix use-after-free in iris_release_internal_buffers() The recent change in commit 1dabf00ee206 ("media: iris: gen1: Destroy internal buffers after FW releases") introduced a regression where session_release_buf() may free the buffer. The caller, iris_release_internal_buffers(), continued to access `buffer` after the call, leading to a potential use-after-free. Fix this by setting BUF_ATTR_PENDING_RELEASE before calling session_release_buf(), and reverting the flag if the call fails. This ensures no dereference occurs after potential freeing.	2026-05-28	7.8
CVE-2026-9987	google - chrome	Insufficient validation of untrusted input in WebAppInstalls in Google Chrome on Android prior to 148.0.7778.216 allowed a local attacker to execute arbitrary code via a malicious file. (Chromium security severity: High)	2026-05-28	7.8
CVE-2026-8856	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to denial of service in configurations where an attacker has write access to parts of the server configuration.	2026-05-26	7.7
CVE-2026-9804	red hat - multiple products	A flaw was found in KubeVirt's virt-exportserver component. An attacker with specific namespace-level access can exploit a path traversal vulnerability in the VMExport directory endpoint. By placing a symbolic link (symlink) within an exported filesystem Persistent Volume Claim (PVC) that points outside its designated mount root, the attacker can read arbitrary files from the exporter pod's filesystem. This leads to information disclosure, potentially exposing sensitive data.	2026-05-28	7.7
CVE-2026-46123	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: virtio_bt: clamp rx length before skb_put virtbt_rx_work() calls skb_put(skb, len) where len comes directly	2026-05-28	7.7

		from virtqueue_get_buf() with no validation against the buffer we posted to the device. The RX skb is allocated in virtbt_add_inbuf() and exposed to virtio as exactly 1000 bytes via sg_init_one(). Checking len against skb_tailroom(skb) is not sufficient because alloc_skb() can leave more tailroom than the 1000 bytes actually handed to the device. A malicious or buggy backend can therefore report used.len between 1001 and skb_tailroom(skb), causing skb_put() to include uninitialized kernel heap bytes that were never written by the device. The same path also accepts len == 0, in which case skb_put(skb, 0) leaves the skb empty but virtbt_rx_handle() still reads the pkt_type byte from skb->data, consuming uninitialized memory. Define VIRTBT_RX_BUF_SIZE once and reuse it in alloc_skb() and sg_init_one(), and gate virtbt_rx_work() on that same constant so the bound checked matches the buffer actually exposed to the device. Reject used.len == 0 in the same gate so an empty completion can no longer reach virtbt_rx_handle(). Use bt_dev_err_ratelimited() because the length value comes from an untrusted backend that can otherwise flood the kernel log. Same class of bug as commit c04db81cd028 ("net/9p: Fix buffer overflow in USB transport layer"), which hardened the USB 9p transport against unchecked device-reported length.		
CVE-2026-46821	oracle - financials_common_modules	Vulnerability in the Oracle Financials Common Modules product of Oracle E-Business Suite (component: Common Components). Supported versions that are affected are 12.2.3-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise Oracle Financials Common Modules. While the vulnerability is in Oracle Financials Common Modules, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Financials Common Modules accessible data. CVSS 3.1 Base Score 7.7 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:N/A:N).	2026-05-28	7.7
CVE-2026-46823	oracle - public_sector_financials	Vulnerability in the Oracle Public Sector Financials (International) product of Oracle E-Business Suite (component: Authorization). Supported versions that are affected are 12.2.6-12.2.15. Easily exploitable vulnerability allows low privileged attacker with network access via HTTPS to compromise Oracle Public Sector Financials (International). While the vulnerability is in Oracle Public Sector Financials (International), attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle Public Sector Financials (International) accessible data. CVSS 3.1 Base Score 7.7 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:N/A:N).	2026-05-28	7.7
CVE-2026-42965	red hat - multiple products	A flaw was found in the OpenShift Router. A user with EndpointSlice write access can exploit this vulnerability by creating a Service backed by an FQDN (Fully Qualified Domain Name) EndpointSlice that resolves to a cloud metadata endpoint. This allows the router to proxy requests to the cloud metadata endpoint, leading to the disclosure of instance credentials and other sensitive metadata. This bypasses previous security measures for validating IP addresses.	2026-05-29	7.7
CVE-2026-48829	gnu - GNU SASL	In GNU SASL before 2.2.3, DIGEST-MD5 has a NULL pointer dereference affecting both clients and servers, via a known token with no accompanying = character. This occurs in lib/digest-md5/getsubopt.c.	2026-05-24	7.5
CVE-2026-8850	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to denial of service via the optional module mod_ibm_upload.	2026-05-26	7.5
CVE-2026-8620	ibm - multiple products	IBM Web Server Plug-ins for WebSphere Application Server and WebSphere Liberty 8.5, 9.0 IBM WebSphere Application Server and WebSphere Application Server Liberty are vulnerable to HTTP request smuggling in the Web Server Plug-ins through a specially crafted request.	2026-05-26	7.5
CVE-2026-8854	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to denial of service via the optional module mod_mem_cache.	2026-05-26	7.5
CVE-2025-14713	synology - c2_identity_edge_server	An Exposed Dangerous Method or Function vulnerability in Synology C2 Identity Edge Server package in DSM before 1.76.0-0307 allows remote attackers to obtain user credentials from the edge server.	2026-05-27	7.5
CVE-2026-3366	ibm - multiple products	IBM InfoSphere Optim Test Data Fabrication 1.0.0, 1.0.0.1, 1.0.0.2, 1.0.2, 1.0.2.2, 1.0.2.3, 1.0.2.4, 1.0.2.5, 1.0.2.6, 1.0.2.7 could allow a remote attacker to traverse directories on the system. An attacker could send a specially crafted URL request containing "dot dot" sequences (/../) to view arbitrary files on the system	2026-05-27	7.5
CVE-2026-45859	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: netfilter: nfnetlink_queue: do shared-unconfirmed check before segmentation Ulrich reports a regression with nfqueue: If an application did not set the 'F_GSO' capability flag and a gso packet with an unconfirmed nf_conn entry is received all packets are now dropped instead of queued, because the check happens after skb_gso_segment(). In that case, we did have exclusive ownership of the skb and its associated conntrack entry. The elevated use count is due to skb_clone happening via skb_gso_segment(). Move the check so that its peformed vs. the aggregated packet.	2026-05-27	7.5

		Then, annotate the individual segments except the first one so we can do a 2nd check at reinject time. For the normal case, where userspace does in-order reinjects, this avoids packet drops: first reinjected segment continues traversal and confirms entry, remaining segments observe the confirmed entry. While at it, simplify nf_ct_drop_unconfirmed(): We only care about unconfirmed entries with a refcnt > 1, there is no need to special-case dying entries. This only happens with UDP. With TCP, the only unconfirmed packet will be the TCP SYN, those aren't aggregated by GRO. Next patch adds a udpgro test case to cover this scenario.		
CVE-2026-45860	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_conncount: increase the connection clean up limit to 64 After the optimization to only perform one GC per jiffy, a new problem was introduced. If more than 8 new connections are tracked per jiffy the list won't be cleaned up fast enough possibly reaching the limit wrongly. In order to prevent this issue, only skip the GC if it was already triggered during the same jiffy and the increment is lower than the clean up limit. In addition, increase the clean up limit to 64 connections to avoid triggering GC too often and do more effective GCs. This has been tested using a HTTP server and several performance tools while having nft_connlimit/xt_connlimit or OVS limit configured. Output of slowhttptest + OVS limit at 52000 connections: slow HTTP test status on 340th second: initializing: 0 pending: 432 connected: 51998 error: 0 closed: 0 service available: YES	2026-05-27	7.5
CVE-2026-45944	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Clear Present bit before tearing down context entry When tearing down a context entry, the current implementation zeros the entire 128-bit entry using multiple 64-bit writes. This creates a window where the hardware can fetch a "torn" entry — where some fields are already zeroed while the 'Present' bit is still set — leading to unpredictable behavior or spurious faults. While x86 provides strong write ordering, the compiler may reorder writes to the two 64-bit halves of the context entry. Even without compiler reordering, the hardware fetch is not guaranteed to be atomic with respect to multiple CPU writes. Align with the "Guidance to Software for Invalidations" in the VT-d spec (Section 6.5.3.3) by implementing the recommended ownership handshake: 1. Clear only the 'Present' (P) bit of the context entry first to signal the transition of ownership from hardware to software. 2. Use dma_wmb() to ensure the cleared bit is visible to the IOMMU. 3. Perform the required cache and context-cache invalidation to ensure hardware no longer has cached references to the entry. 4. Fully zero out the entry only after the invalidation is complete. Also, add a dma_wmb() to context_set_present() to ensure the entry is fully initialized before the 'Present' bit becomes visible.	2026-05-27	7.5
CVE-2026-46024	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: libceph: Prevent potential null-ptr-deref in ceph_handle_auth_reply() If a message of type CEPH_MSG_AUTH_REPLY contains a zero value for both protocol and result, this is currently not treated as an error. In case of ac->negotiating == true and ac->protocol > 0, this leads to setting ac->protocol = 0 and ac->ops = NULL. Thereafter, the check for ac->protocol != protocol returns false, and init_protocol() is not called. Subsequently, ac->ops->handle_reply() is called, which leads to	2026-05-27	7.5

		a null pointer dereference, because ac->ops is still NULL. This patch changes the check for ac->protocol != protocol to !ac->protocol, as this also includes the case when the protocol was set to zero in the message. This causes the message to be treated as containing a bad auth protocol.		
CVE-2026-46027	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net/smc: avoid early lgr access in smc_clc_wait_msg A CLC decline can be received while the handshake is still in an early stage, before the connection has been associated with a link group. The decline handling in smc_clc_wait_msg() updates link-group level sync state for first-contact declines, but that state only exists after link group setup has completed. Guard the link-group update accordingly and keep the per-socket peer diagnosis handling unchanged. This preserves the existing sync_err handling for established link-group contexts and avoids touching link-group state before it is available.	2026-05-27	7.5
CVE-2026-46031	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: ks8851: Reinstate disabling of BHs around IRQ handler If the driver executes ks8851_irq() AND a TX packet has been sent, then the driver enables TX queue via netif_wake_queue() which schedules TX softirq to queue packets for this device. If CONFIG_PREEMPT_RT=y is set AND a packet has also been received by the MAC, then ks8851_rx_pkts() calls netdev_alloc_skb_ip_align() to allocate SKBs for the received packets. If netdev_alloc_skb_ip_align() is called with BH enabled, then local_bh_enable() at the end of netdev_alloc_skb_ip_align() will trigger the pending softirq processing, which may ultimately call the .xmit callback ks8851_start_xmit_par(). The ks8851_start_xmit_par() will try to lock struct ks8851_net_par .lock spinlock, which is already locked by ks8851_irq() from which ks8851_start_xmit_par() was called. This leads to a deadlock, which is reported by the kernel, including a trace listed below. If CONFIG_PREEMPT_RT is not set, then since commit 0913ec336a6c0 ("net: ks8851: Fix deadlock with the SPI chip variant") the deadlock can also be triggered without received packet in the RX FIFO. The pending softirqs will be processed on return from spin_unlock_bh(&ks->statelock) in ks8851_irq(), which triggers the deadlock as well. Fix the problem by disabling BH around critical sections, including the IRQ handler, thus preventing the net_tx_action() softirq from triggering during these critical sections. The net_tx_action() softirq is triggered once BH are re-enabled and at the end of the IRQ handler, once all the other IRQ handler actions have been completed. __schedule from schedule_rtlock+0x1c/0x34 schedule_rtlock from rtlock_slowlock_locked+0x548/0x904 rtlock_slowlock_locked from rt_spin_lock+0x60/0x9c rt_spin_lock from ks8851_start_xmit_par+0x74/0x1a8 ks8851_start_xmit_par from netdev_start_xmit+0x20/0x44 netdev_start_xmit from dev_hard_start_xmit+0xd0/0x188 dev_hard_start_xmit from sch_direct_xmit+0xb8/0x25c sch_direct_xmit from __qdisc_run+0x1f8/0x4ec __qdisc_run from qdisc_run+0x1c/0x28 qdisc_run from net_tx_action+0x1f0/0x268 net_tx_action from handle_softirqs+0x1a4/0x270 handle_softirqs from __local_bh_enable_ip+0xcc/0xe0 __local_bh_enable_ip from __alloc_skb+0xd8/0x128 __alloc_skb from __netdev_alloc_skb+0x3c/0x19c __netdev_alloc_skb from ks8851_irq+0x388/0x4d4 ks8851_irq from irq_thread_fn+0x24/0x64 irq_thread_fn from irq_thread+0x178/0x28c irq_thread from kthread+0x12c/0x138 kthread from ret_from_fork+0x14/0x28	2026-05-27	7.5
CVE-2026-46052	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ceph: only d_add() negative dentries when they are unhashed Ceph can call d_add(dentry, NULL) on a negative dentry that is already present in the primary dcache hash. In the current VFS that is not safe. d_add() goes through __d_add() to __d_rehash(), which unconditionally reinserts dentry->d_hash into	2026-05-27	7.5

		the hlist_bl bucket. If the dentry is already hashed, reinserting the same node can corrupt the bucket, including creating a self-loop. Once that happens, __d_lookup() can spin forever in the hlist_bl walk, typically looping only on the d_name.hash mismatch check and eventually triggering RCU stall reports like this one: rcu: INFO: rcu_sched self-detected stall on CPU rcu: 87-....: (2100 ticks this GP) idle=3a4c/1/0x4000000000000000 softirq=25003319/25003319 fqs=829 rcu: (t=2101 jiffies g=79058445 q=698988 ncpus=192) CPU: 87 UID: 2952868916 PID: 3933303 Comm: php-cgi8.3 Not tainted 6.18.17-i1-amd #950 NONE Hardware name: Dell Inc. PowerEdge R7615/0G9DHV, BIOS 1.6.6 09/22/2023 RIP: 0010:__d_lookup+0x46/0xb0 Code: c1 e8 07 48 8d 04 c2 48 8b 00 49 89 fc 49 89 f5 48 89 c3 48 83 e3 fe 48 83 f8 01 77 0f eb 2d 0f 1f 44 00 00 48 8b 1b 48 85 db <74> 20 39 6b 18 75 f3 48 8d 7b 78 e8 ba 85 d0 00 4c 39 63 10 74 1f RSP: 0018:ff745a70c8253898 EFLAGS: 00000282 RAX: ff26e470054cb208 RBX: ff26e470054cb208 RCX: 000000006e958966 RDX: ff26e48267340000 RSI: ff745a70c82539b0 RDI: ff26e458f74655c0 RBP: 000000006e958966 R08: 0000000000000180 R09: 9cd08d909b919a89 R10: ff26e458f74655c0 R11: 0000000000000000 R12: ff26e458f74655c0 R13: ff745a70c82539b0 R14: d0d0d0d0d0d0d0d0 R15: 2f2f2f2f2f2f2f2f FS: 00007f5770896980(0000) GS:ff26e482c5d88000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007f5764de50c0 CR3: 000000a72abb5001 CR4: 0000000000771ef0 PKRU: 55555554 Call Trace: <TASK> lookup_fast+0x9f/0x100 walk_component+0x1f/0x150 link_path_walk+0x20e/0x3d0 path_lookupat+0x68/0x180 filename_lookup+0xdc/0x1e0 vfs_statx+0x6c/0x140 vfs_fstatat+0x67/0xa0 __do_sys_newfstatat+0x24/0x60 do_syscall_64+0x6a/0x230 entry_SYSCALL_64_after_hwframe+0x76/0x7e This is reachable with reused cached negative dentries. A Ceph lookup or atomic_open can be handed a negative dentry that is already hashed, and fs/ceph/dir.c then hits one of two paths that incorrectly assume "negative" also means "unhashed": - ceph_finish_lookup(): MDS reply is -ENOENT with no trace -> d_add(dentry, NULL) - ceph_lookup(): local ENOENT fast path for a complete directory with shared caps -> d_add(dentry, NULL) Both paths can therefore re-add an already-hashed negative dentry. Ceph already uses the correct pattern elsewhere: ceph_fill_trace() only calls d_add(dn, NULL) for a negative null-dentry reply when d_unhashed(dn) is true. Fix both fs/ceph/dir.c sites the same way: only call d_add() for a negative dentry when it is actually unhashed. If the negative dentry is already hashed, leave it in place and reuse it as-is. This preserves the existing behavior for unhashed dentries while avoiding d_hash list corruption for reused hashed negatives.		
CVE-2026-46085	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix rxkad crypto unalignment handling Fix handling of a packet with a misaligned crypto length. Also handle non-ENOMEM errors from decryption by aborting. Further, remove the WARN_ON_ONCE() so that it can't be remotely triggered (a trace line can still be emitted).	2026-05-27	7.5
CVE-2026-46102	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: strparser: fix skb_head leak in strp_abort_strp() When the stream parser is aborted, for example after a message assembly timeout, it can still hold a reference to a partially assembled message in strp->skb_head. That skb is not released in strp_abort_strp(), which leaks the partially	2026-05-27	7.5

		assembled message and can be triggered repeatedly to exhaust memory. Fix this by freeing strp->skb_head and resetting the parser state in the abort path. Leave strp_stop() unchanged so final cleanup still happens in strp_done() after the work and timer have been synchronized.		
CVE-2026-8180	ibm - multiple products	IBM Aspera High-Speed Transfer Endpoint 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Server 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Endpoint are affected by a potential denial of service in the asperahttpd component. An unauthenticated user can cause the asperahttpd service to crash.	2026-05-27	7.5
CVE-2026-48921	jenkins - pipeline\	Jenkins Pipeline: Groovy Libraries Plugin 797.v90ea_a_9b_e45a_0 and earlier does not prohibit symbolic links in shared libraries, allowing attackers able to control the content of a library used by a Pipeline job to read arbitrary files on the Jenkins controller filesystem.	2026-05-27	7.5
CVE-2026-48922	jenkins - credentials_binding	Jenkins Credentials Binding Plugin 720.v3f6decef43ea_ and earlier does not properly sanitize file names for file and zip file credentials, allowing attackers able to provide credentials to a job to write files to arbitrary locations on the node filesystem, which can lead to remote code execution if Jenkins is configured to allow a low-privileged user to configure file or zip file credentials used for a job running on the built-in node.	2026-05-27	7.5
CVE-2026-46110	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: stmmac: Prevent NULL deref when RX memory exhausted The CPU receives frames from the MAC through conventional DMA: the CPU allocates buffers for the MAC, then the MAC fills them and returns ownership to the CPU. For each hardware RX queue, the CPU and MAC coordinate through a shared ring array of DMA descriptors: one descriptor per DMA buffer. Each descriptor includes the buffer's physical address and a status flag ("OWN") indicating which side owns the buffer: OWN=0 for CPU, OWN=1 for MAC. The CPU is only allowed to set the flag and the MAC is only allowed to clear it, and both must move through the ring in sequence: thus the ring is used for both "submissions" and "completions." In the stmmac driver, stmmac_rx() bookmarks its position in the ring with the `cur_rx` index. The main receive loop in that function checks for rx_descs[cur_rx].own=0, gives the corresponding buffer to the network stack (NULLing the pointer), and increments `cur_rx` modulo the ring size. After the loop exits, stmmac_rx_refill(), which bookmarks its position with `dirty_rx`, allocates fresh buffers and rearms the descriptors (setting OWN=1). If it fails any allocation, it simply stops early (leaving OWN=0) and will retry where it left off when next called. This means descriptors have a three-stage lifecycle (terms my own): - `empty` (OWN=1, buffer valid) - `full` (OWN=0, buffer valid and populated) - `dirty` (OWN=0, buffer NULL) But because stmmac_rx() only checks OWN, it confuses `full`/`dirty`. In the past (see 'Fixes:'), there was a bug where the loop could cycle `cur_rx` all the way back to the first descriptor it dirtied, resulting in a NULL dereference when mistaken for `full`. The aforementioned commit resolved that *specific* failure by capping the loop's iteration limit at `dma_rx_size - 1`, but this is only a partial fix: if the previous stmmac_rx_refill() didn't complete, then there are leftover `dirty` descriptors that the loop might encounter without needing to cycle fully around. The current code therefore panics (see 'Closes:') when stmmac_rx_refill() is memory-starved long enough for `cur_rx` to catch up to `dirty_rx`. Fix this by explicitly checking, before advancing `cur_rx`, if the next entry is dirty; exit the loop if so. This prevents processing of the final, used descriptor until stmmac_rx_refill() succeeds, but fully prevents the `cur_rx == dirty_rx` ambiguity as the previous bugfix intended: so remove the clamp as well. Since stmmac_rx_zc() is a copy-paste-and-tweak of stmmac_rx() and the code structure is identical, any fix to stmmac_rx() will also need a corresponding fix for stmmac_rx_zc(). Therefore, apply the same check there. In stmmac_rx() (not stmmac_rx_zc()), a related bug remains: after the MAC sets OWN=0 on the final descriptor, it will be unable to send any further DMA-complete IRQs until it's given more `empty` descriptors. Currently, the driver simply *hopes* that the next stmmac_rx_refill() succeeds, risking an indefinite stall of the receive process if not. But this is not a regression, so it can be addressed in a future change.	2026-05-28	7.5
CVE-2026-46114	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Reject non-8-byte ATOMIC_WRITE payloads atomic_write_reply() at drivers/infiniband/sw/rxe/rxe_resp.c unconditionally dereferences 8 bytes at payload_addr(pkt):	2026-05-28	7.5

		value = *(u64 *)payload_addr(pkt); check_rkey() previously accepted an ATOMIC_WRITE request with pktlen == resid == 0 because the length validation only compared pktlen against resid. A remote initiator that sets the RETH length to 0 therefore reaches atomic_write_reply() with a zero-byte logical payload, and the responder reads sizeof(u64) bytes from past the logical end of the packet into skb->head tailroom, then writes those 8 bytes into the attacker's MR via rxe_mr_do_atomic_write(). That is a remote disclosure of 4 bytes of kernel tailroom per probe (the other 4 bytes are the packet's own trailing ICRC). IBA oA19-28 defines ATOMIC_WRITE as exactly 8 bytes. Anything else is protocol-invalid. Hoist a strict length check into check_rkey() so the responder never reaches the unchecked dereference, and keep the existing WRITE-family length logic for the normal RDMA WRITE path. Reproduced on mainline with an unmodified rxe driver: a sustained zero-length ATOMIC_WRITE probe repeatedly leaks adjacent skb head-buffer bytes into the attacker's MR, including recognisable kernel strings and partial kernel-direct-map pointer words. With this patch applied the responder rejects the PDU and the MR stays all-zero.		
CVE-2026-46124	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: isofs: validate block number from NFS file handle in isofs_export_iget isofs_fh_to_dentry() and isofs_fh_to_parent() pass an attacker-controlled block number (ifid->block or ifid->parent_block) from the NFS file handle to isofs_export_iget(), which only rejects block == 0 before calling isofs_iget() and ultimately sb_bread(). A crafted file handle with fh_len sufficient to pass the check added by commit 0405d4b63d08 ("isofs: Prevent the use of too small fid") can still drive the server to read any in-range block on the backing device as if it were an iso_directory_record. That earlier fix was assigned CVE-2025-37780. sb_bread() on an out-of-range block returns NULL cleanly via the EIO path, so there is no memory-safety violation. For in-range reads of adjacent-partition data on the same block device, the unrelated bytes end up in iso_inode_info fields that reach the NFS client as dentry metadata. The deployment surface (isofs exported over NFS from loop-mounted images) is narrow and requires an authenticated NFS peer, but the malformed-file-handle class is reportable as hardening next to the existing CVE-2025-37780 fix. Reject block >= ISOFS_SB(sb)->s_nzones in isofs_export_iget() so the check covers both isofs_fh_to_dentry() and isofs_fh_to_parent() call sites with a single line.	2026-05-28	7.5
CVE-2026-46133	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Reject unknown opcodes before ICRC processing Even after applying commit 7244491dab34 ("RDMA/rxe: Validate pad and ICRC before payload_size() in rxe_rcv"), a single unauthenticated UDP packet can still trigger panic. That patch handled payload_size() underflow only for valid opcodes with short packets, not for packets carrying an unknown opcode. The unknown-opcode OOB read described below predates that commit and reaches back to the initial Soft RoCE driver. The check added there reads pkt->paylen < header_size(pkt) + bth_pad(pkt) + RXE_ICRC_SIZE where header_size(pkt) expands to rxe_opcode[pkt->opcode].length. The rxe_opcode[] array has 256 entries but is only populated for defined IB opcodes; any other entry (for example opcode 0xff) is zero-initialized, so length == 0 and the check degenerates to pkt->paylen < 0 + bth_pad(pkt) + RXE_ICRC_SIZE which does not constrain pkt->paylen enough. rxe_icrc_hdr() then computes rxe_opcode[pkt->opcode].length - RXE_BTH_BYTES which underflows when length == 0 and passes a huge value to rxe_crc32(), causing an out-of-bounds read of the skb payload. Reproduced on v7.0-rc7 with that fix applied, QEMU/KVM with CONFIG_RDMA_RXE=y and CONFIG_KASAN=y, after rdma link add rxe0 type rxe netdev eth0	2026-05-28	7.5

		A single 48-byte UDP packet to port 4791 with BTH opcode=0xff and QPN=IB_MULTICAST_QPN triggers: BUG: KASAN: slab-out-of-bounds in crc32_le+0x115/0x170 Read of size 1 at addr ... The buggy address is located 0 bytes to the right of allocated 704-byte region Call Trace: crc32_le+0x115/0x170 rxe_icrc_hdr.isra.0+0x226/0x300 rxe_icrc_check+0x13f/0x3a0 rxe_rcv+0x6e1/0x16e0 rxe_udp_encap_recv+0x20a/0x320 udp_queue_rcv_one_skb+0x7ed/0x12c0 Subsequent packets with the same shape fault on unmapped memory and panic the kernel. The trigger requires only module load and "rdma link add"; no QP, no connection, and no authentication. Fix this by rejecting packets whose opcode has no rxe_opcode[] entry, detected via the zero mask or zero length, before any length arithmetic runs.		
CVE-2026-46177	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ipmi: Add limits to event and receive message requests The driver would just fetch events and receive messages until the BMC said it was done. To avoid issues with BMCs that never say they are done, add a limit of 10 fetches at a time. In addition, an si interface has an attn state it can return from the hardware which is supposed to cause a flag fetch to see if the driver needs to fetch events or message or a few other things. If the attn bit gets stuck, it's a similar problem. So allow messages in between flag fetches so the driver itself doesn't get stuck. This is a more general fix than the previous fix for the specific bad BMC, but should fix the more general issue of a BMC that won't stop saying it has data. This has been there from the beginning of the driver. It's not a bug per-se, but it is accounting for bugs in BMCs.	2026-05-28	7.5
CVE-2026-46829	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Mongoapi). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of Oracle REST Data Services. CVSS 3.1 Base Score 7.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:H).	2026-05-28	7.5
CVE-2026-46834	oracle - database_server	Vulnerability in the Net Service component of Oracle Database Server. Supported versions that are affected are 23.4.0-23.26.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via TLS to compromise Net Service. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of Net Service. CVSS 3.1 Base Score 7.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:H).	2026-05-28	7.5
CVE-2026-46835	oracle - database_server	Vulnerability in the Net Service component of Oracle Database Server. Supported versions that are affected are 23.4.0-23.26.2. Easily exploitable vulnerability allows unauthenticated attacker with network access via TLS to compromise Net Service. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of Net Service. CVSS 3.1 Base Score 7.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:H).	2026-05-28	7.5
CVE-2026-10003	google - chrome	Use after free in Views in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-10005	google - chrome	Use after free in WebAppInstalls in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-10006	google - chrome	Race in WebAudio in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-10009	google - chrome	Integer overflow in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-10022	google - chrome	Type Confusion in V8 in Google Chrome prior to 148.0.7778.216 allowed an attacker who convinced a user to install a malicious extension to execute arbitrary code inside a sandbox via a crafted Chrome Extension. (Chromium security severity: Medium)	2026-05-28	7.5
CVE-2026-9901	google - chrome	Use after free in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5

CVE-2026-9909	google - chrome	Integer overflow in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9922	google - chrome	Use after free in GPU in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9933	google - chrome	Use after free in Input in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9934	google - chrome	Use after free in Aura in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9954	google - chrome	Use after free in TabStrip in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9956	google - chrome	Use after free in iOS in Google Chrome on iOS prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to execute arbitrary code via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9960	google - chrome	Integer overflow in PDFium in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to execute arbitrary code inside a sandbox via a crafted font file. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9963	google - chrome	Uninitialized Use in iOS in Google Chrome on iOS prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to execute arbitrary code inside a sandbox via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-9990	google - chrome	Use after free in WebAppInstalls in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2026-05-28	7.5
CVE-2026-46818	oracle - e-business_suite	Vulnerability in the Oracle Payments product of Oracle E-Business Suite (component: File Transmission). Supported versions that are affected are 12.2.3-12.2.15. Difficult to exploit vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle Payments. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle Payments accessible data as well as unauthorized access to critical data or complete access to all Oracle Payments accessible data. CVSS 3.1 Base Score 7.4 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:N).	2026-05-28	7.4
CVE-2026-46579	red hat - multiple products	A flaw was found in the OpenShift Router. When a Route has `insecureEdgeTerminationPolicy` set to Allow, the HTTP frontend does not remove `X-SSL-Client-*` headers from incoming requests. This allows an unauthenticated attacker to send plain HTTP requests with crafted `X-SSL-Client-*` headers. As a result, backends relying on these headers for mutual TLS (Transport Layer Security) authentication can be bypassed, enabling the attacker to impersonate client certificate identities.	2026-05-29	7.4
CVE-2026-10062	trendnet - tew-432brp_firmware	A vulnerability was determined in TRENDnet TEW-432BRP 3.10B20. Affected by this vulnerability is the function formSetRoute of the file /goform/formSetRoute. This manipulation of the argument ip/mask/gateway causes stack-based buffer overflow. The attack is possible to be carried out remotely. The exploit has been publicly disclosed and may be utilized. The vendor explains: "This product has been EOL for 15 years (since 2009). As the item has been EOL for such a long time, we are not able to replicate or fix any vulnerabilities." This vulnerability only affects products that are no longer supported by the maintainer.	2026-05-29	7.4
CVE-2026-10063	trendnet - tew-432brp_firmware	A vulnerability was identified in TRENDnet TEW-432BRP 3.10B20. Affected by this issue is the function formWPS of the file /goform/formWPS. Such manipulation of the argument peerPin leads to stack-based buffer overflow. The attack may be performed from remote. The exploit is publicly available and might be used. The vendor explains: "This product has been EOL for 15 years (since 2009). As the item has been EOL for such a long time, we are not able to replicate or fix any vulnerabilities." This vulnerability only affects products that are no longer supported by the maintainer.	2026-05-29	7.4
CVE-2026-8835	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to invalid pointer dereference. A privileged user, authenticated to the Administration Server, could exploit this vulnerability to expose sensitive information or cause a denial of service.	2026-05-26	7.3
CVE-2026-45932	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix tcx/netkit detach permissions when prog fd isn't given This commit fixes a security issue where BPF_PROG_DETACH on tcx or netkit devices could be executed by any user when no program fd was provided, bypassing permission checks. The fix adds a capability check for CAP_NET_ADMIN or CAP_SYS_ADMIN in this case.	2026-05-27	7.3
CVE-2026-32996	veeam - Backup and Replication	This vulnerability in Veeam Agent for Microsoft Windows allows for Local Privilege Escalation.	2026-05-28	7.3
CVE-2026-9795	redhat - build_of_keycloak	A flaw was found in Keycloak's Fine-Grained Admin Permissions (FGAPv2) feature. An administrator with limited client management permissions can exploit this vulnerability to assign any realm role, including highly privileged roles, to a client's scope mapping. This bypasses intended security controls, allowing the injected role to be projected into a user's authentication token when they access the modified client. This could lead to unauthorized privilege escalation within the Keycloak realm.	2026-05-28	7.3
CVE-2026-34126	tp-link - tapo_l535e_firmware	TP-Link has identified a vulnerability in Tapo L535E v1.0 and v3.0, Tapo P300 v1.0, and Tapo D100C v1.0, where Bluetooth communication during the initial setup phase is transmitted in cleartext without encryption. Bluetooth is only used during initialization. An attacker within the Bluetooth range could exploit this behavior using Bluetooth sniffing or man-in-the-middle techniques, which may allow eavesdropping on Bluetooth communication,	2026-05-28	7.3

		manipulate transmitted setup data and potentially gain unauthorized control of the device during initialization. An attacker within the Bluetooth range could exploit this behavior using Bluetooth sniffing or man-in-the-middle techniques, which may allow eavesdropping on Bluetooth communication, manipulate transmitted setup data and potentially gain unauthorized control of the device during initialization. D100C is the chime delivered with your Tapo camera, and it is delivered with the following Tapo products: D130, D210, D235, D225, TD21, TDB21 and TD25		
CVE-2026-42782	apache - multiple products	Improper Isolation or Compartmentalization vulnerability in Apache Syncope. An administrator with adequate entitlements for Implementations can create a malicious Groovy class containing untrusted code reaching a non-sandboxed execution path via the class static initializer. This issue affects Apache Syncope: 3.0 through 3.0.16, 4.0 through 4.0.5, 4.1.0. Users are recommended to upgrade to version 4.0.6 / 4.1.1, which fix this issue by forcing even the static initializer in Groovy code to run in a sandbox.	2026-05-25	7.2
CVE-2026-4051	ibm - multiple products	IBM Engineering Lifecycle Management 7.0.3, 7.1.0, and 7.2.0 could allow an attacker with administrative privileges to execute remote code due to exposed method that is not properly restricted.	2026-05-26	7.2
CVE-2024-56462	ibm - multiple products	IBM QRadar 7.5.0 through 7.5.0 UP15 Interim Fix 002 could allow a privileged user to upload a malicious backup archive that could be restored and used to gain access to the underlying operating system.	2026-05-27	7.2
CVE-2026-3603	ibm - multiple products	IBM Engineering Lifecycle Management 7.0.3 Interim Fix 001 through Interim Fix 021, 7.1.0 Interim Fix 001 through Interim Fix 009, and 7.2.0 and 7.2.0 Interim Fix 001 is vulnerable to an XML external entity injection (XXE) attack when processing XML data. An authenticated attacker could exploit this vulnerability to expose sensitive information or consume memory resources.	2026-05-26	7.1
CVE-2026-42012	red hat - multiple products	A flaw was found in gnutls. A remote attacker could exploit this vulnerability by presenting a specially crafted certificate that contains Uniform Resource Identifier (URI) or Service (SRV) Subject Alternative Names (SANs). This could cause the certificate validation process to incorrectly fall back to checking DNS hostnames against the Common Name (CN), potentially allowing the attacker to spoof legitimate services or intercept sensitive information.	2026-05-26	7.1
CVE-2026-1718	ibm - multiple products	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to a denial of service with a specially crafted query when autonomous transactions are enabled.	2026-05-27	7.1
CVE-2026-1933	redhat - multiple products	A flaw was found in Samba's handling of NTFS-style reparse points on shares configured with read only = yes. Due to missing SMB-layer access checks, authenticated users with underlying filesystem write permissions may create or delete reparse point metadata through SMB operations even on read-only exports. This could allow modification of SMB-visible file behavior, including converting files into symbolic links or other reparse point types.	2026-05-27	7.1
CVE-2026-45856	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: RDMA/uverbs: Validate wqe_size before using it in ib_uverbs_post_send ib_uverbs_post_send() uses cmd.wqe_size from userspace without any validation before passing it to kmalloc() and using the allocated buffer as struct ib_uverbs_send_wr. If a user provides a small wqe_size value (e.g., 1), kmalloc() will succeed, but subsequent accesses to user_wr->opcode, user_wr->num_sge, and other fields will read beyond the allocated buffer, resulting in an out-of-bounds read from kernel heap memory. This could potentially leak sensitive kernel information to userspace. Additionally, providing an excessively large wqe_size can trigger a WARNING in the memory allocation path, as reported by syzkaller. This is inconsistent with ib_uverbs_unmarshall_recv() which properly validates that wqe_size >= sizeof(struct ib_uverbs_recv_wr) before proceeding.	2026-05-27	7.1

		Add the same validation for ib_uverbs_post_send() to ensure wqe_size is at least sizeof(struct ib_uverbs_send_wr).		
CVE-2026-45955	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: md/md-llbitmap: fix percpu_ref not resurrected on suspend timeout When llbitmap_suspend_timeout() times out waiting for percpu_ref to become zero, it returns -ETIMEDOUT without resurrecting the percpu_ref. The caller (md_llbitmap_daemon_fn) then continues to the next page without calling llbitmap_resume(), leaving the percpu_ref in a killed state permanently. Fix this by resurrecting the percpu_ref before returning the error, ensuring the page control structure remains usable for subsequent operations.	2026-05-27	7.1
CVE-2026-45958	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/exynos: vidi: fix to avoid directly dereferencing user pointer In vidi_connection_ioctl(), vidi->edid(user pointer) is directly dereferenced in the kernel. This allows arbitrary kernel memory access from the user space, so instead of directly accessing the user pointer in the kernel, we should modify it to copy edid to kernel memory using copy_from_user() and use it.	2026-05-27	7.1

CVE-2026-45999	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		erofs: fix unsigned underflow in z_erofs_lz4_handle_overlap()
		
		Some crafted images can have illegal (!partial_decoding && m_llen < m_plen) extents, and the LZ4 inplace decompression path can be wrongly hit, but it cannot handle (outpages < inpages) properly: "outpages - inpages" wraps to a large value and the subsequent rq->out[] access reads past the decompressed_pages array.
		
		However, such crafted cases can correctly result in a corruption report in the normal LZ4 non-inplace path.
		
		Let's add an additional check to fix this for backporting.
		
		Reproducible image (base64-encoded gzipped blob):
		
		H4sIAJGR12kCA+3SPUoDQRgG4MkmkkZk8QRbRFIIi9hbpEjrHQI5ghfwCN5BLCzTGtLbBI+g
		dilSJo1CnIm7GEXFxhT6PDDwfrs73/ywIQD/1ePD4r7Ou6ETsrq4mu7XcWfj++Pb58nJU/9i
		PNtbjhan04/9GtX4qVYc814WDqt6FaX5s+ZwXXeq52lndT6IuVvlblytLMvh4Gzwaf90nsvz
		2DF/21+20T/ldgp5s1jXRaN4t/8izsy/OUB6e/Qa79r+JwAAAAAAAL52vQVuGQAAAP6+my1w
		ywAAAAAAADwu14ATsEYtgBQAAA=
		
		$ mount -t erofs -o cache_strategy=disabled foo.erofs /mnt
		$ dd if=/mnt/data of=/dev/null bs=4096 count=1
			2026-05-27	7.1

CVE-2026-46054	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: selinux: fix overlayfs mmap() and mprotect() access checks The existing SELinux security model for overlayfs is to allow access if the current task is able to access the top level file (the "user" file) and the mounter's credentials are sufficient to access the lower level file (the "backing" file). Unfortunately, the current code does not properly enforce these access controls for both mmap() and mprotect() operations on overlayfs filesystems. This patch makes use of the newly created security_mmap_backing_file() LSM hook to provide the missing backing file enforcement for mmap() operations, and leverages the backing file API and new LSM blob to provide the necessary information to properly enforce the mprotect() access controls.	2026-05-27	7.1
CVE-2026-46055	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: apparmor: Fix string overrun due to missing termination When booting Ubuntu 26.04 with Linux 7.0-rc4 on an ARM64 Qualcomm Snapdragon X1 we see a string buffer overrun: BUG: KASAN: slab-out-of-bounds in aa_dfa_match (security/apparmor/match.c:535) Read of size 1 at addr ffff0008901cc000 by task snap-update-ns/2120 CPU: 5 UID: 60578 PID: 2120 Comm: snap-update-ns Not tainted 7.0.0-rc4+ #22 PREEMPTLAZY Hardware name: LENOVO 83ED/LNVNB161216, BIOS NHCN60WW 09/11/2025 Call trace: show_stack (arch/arm64/kernel/stacktrace.c:501) (C) dump_stack_lvl (lib/dump_stack.c:122) print_report (mm/kasan/report.c:379 mm/kasan/report.c:482) kasan_report (mm/kasan/report.c:597)	2026-05-27	7.1

		__asan_report_load1_noabort (mm/kasan/report_generic.c:378) aa_dfa_match (security/apparmor/match.c:535) match_mnt_path_str (security/apparmor/mount.c:244 security/apparmor/mount.c:336) match_mnt (security/apparmor/mount.c:371) aa_bind_mount (security/apparmor/mount.c:447 (discriminator 4)) apparmor_sb_mount (security/apparmor/lsm.c:719 (discriminator 1)) security_sb_mount (security/security.c:1062 (discriminator 31)) path_mount (fs/namespace.c:4101) __arm64_sys_mount (fs/namespace.c:4172 fs/namespace.c:4361 fs/namespace.c:4338 fs/namespace.c:4338) invoke_syscall.constprop.0 (arch/arm64/kernel/syscall.c:35 arch/arm64/kernel/syscall.c:49) el0_svc_common.constprop.0 (./include/linux/thread_info.h:142 (discriminator 2) arch/arm64/kernel/syscall.c:140 (discriminator 2)) do_el0_svc (arch/arm64/kernel/syscall.c:152) el0_svc (arch/arm64/kernel/entry-common.c:80 arch/arm64/kernel/entry-common.c:725) el0t_64_sync_handler (arch/arm64/kernel/entry-common.c:744) el0t_64_sync (arch/arm64/kernel/entry.S:596) Allocated by task 2120: kasan_save_stack (mm/kasan/common.c:58) kasan_save_track (./arch/arm64/include/asm/current.h:19 mm/kasan/common.c:70 mm/kasan/common.c:79) kasan_save_alloc_info (mm/kasan/generic.c:571) __kasan_kmalloc (mm/kasan/common.c:419) __kmalloc_noprof (./include/linux/kasan.h:263 mm/slub.c:5260 mm/slub.c:5272) aa_get_buffer (security/apparmor/lsm.c:2201) aa_bind_mount (security/apparmor/mount.c:442) apparmor_sb_mount (security/apparmor/lsm.c:719 (discriminator 1)) security_sb_mount (security/security.c:1062 (discriminator 31)) path_mount (fs/namespace.c:4101) __arm64_sys_mount (fs/namespace.c:4172 fs/namespace.c:4361 fs/namespace.c:4338 fs/namespace.c:4338) invoke_syscall.constprop.0 (arch/arm64/kernel/syscall.c:35 arch/arm64/kernel/syscall.c:49) el0_svc_common.constprop.0 (./include/linux/thread_info.h:142 (discriminator 2) arch/arm64/kernel/syscall.c:140 (discriminator 2)) do_el0_svc (arch/arm64/kernel/syscall.c:152) el0_svc (arch/arm64/kernel/entry-common.c:80 arch/arm64/kernel/entry-common.c:725) el0t_64_sync_handler (arch/arm64/kernel/entry-common.c:744) el0t_64_sync (arch/arm64/kernel/entry.S:596) The buggy address belongs to the object at ffff0008901ca000 which belongs to the cache kmalloc-rnd-06-8k of size 8192 The buggy address is located 0 bytes to the right of allocated 8192-byte region [ffff0008901ca000, ffff0008901cc000) The buggy address belongs to the physical page: page: refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x9101c8 head: order:3 mapcount:0 entire_mapcount:0 nr_pages_mapped:-1 pincount:0 flags: 0x8000000000000040(head|zone=2) page_type: f5(slab) raw: 8000000000000040 ffff000800016c40 fffffdffe2d14e10 ffff000800015c70 raw: 0000000000000000 0000000800010001 00000000f5000000 0000000000000000 head: 8000000000000040 ffff000800016c40 fffffdffe2d14e10 ffff000800015c70 head: 0000000000000000 0000000800010001 00000000f5000000 0000000000000000 head: 8000000000000003 fffffdffe2407201 ffffffffffffffff 00000000ffffffff head: ffffffffffffffff 0000000000000000 00000000ffffffff 0000000000000008 page dumped because: kasan: bad access detected Memory state around the buggy address: ffff0008901cbf00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ffff0008 ---truncated---		
CVE-2026-46070	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: md/raid5: validate payload size before accessing journal metadata r5c_recovery_analyze_meta_block() and r5l_recovery_verify_data_checksum_for_mb() iterate over payloads in a journal metadata block using on-disk payload size fields without validating them against the remaining space in the metadata block. A corrupted journal contains payload sizes extending beyond the PAGE_SIZE boundary can cause out-of-bounds reads when accessing payload fields or computing offsets. Add bounds validation for each payload type to ensure the full payload fits within meta_size before processing.	2026-05-27	7.1
CVE-2026-46078	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: erofs: fix the out-of-bounds nameoff handling for trailing dirents Currently we already have boundary-checks for nameoffs, but the trailing dirents are special since the namelens are calculated with strnlen() with unchecked nameoffs.	2026-05-27	7.1

		If a crafted EROFS has a trailing dirent with nameoff >= maxsize,maxsize - nameoff can underflow, causing strnlen() to read past the directory block. nameoff0 should also be verified to be a multiple of `sizeof(struct erofs_dirent)` as well [1]. [1] https://sashiko.dev/#/patchset/20260416063511.3173774-1-hsiangkao%40linux.alibaba.com		
CVE-2026-46149	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: scsi: target: configfs: Bound snprintf() return in tg_pt_gp_members_show() target_tg_pt_gp_members_show() formats LUN paths with snprintf() into a 256-byte stack buffer, then will memcpy() cur_len bytes from that buffer. snprintf() returns the length the output would have had, which can exceed the buffer size when the fabric WWN is long because iSCSI IQN names can be up to 223 bytes. The check at the memcpy() site only guards the destination page write, not the source read, so memcpy() will read past the stack buffer and copy adjacent stack contents to the sysfs reader, which when CONFIG_FORTIFY_SOURCE is enabled, fortify_panic() will be triggered. Commit 27e06650a5ea ("scsi: target: target_core_configfs: Add length check to avoid buffer overflow") added the same bound to the target_lu_gp_members_show() but the tg_pt_gp variant was missed so resolve that here.	2026-05-28	7.1
CVE-2026-46150	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: fanotify: fix false positive on permission events fsnotify_get_mark_safe() may return false for a mark on an unrelated group, which results in bypassing the permission check. Fix by skipping over detached marks that are not in the current group.	2026-05-28	7.1
CVE-2026-46175	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix fsck inconsistency caused by FGGC of node block During FGGC node block migration, fsck may incorrectly treat the migrated node block as fsync-written data. The reproduction scenario: root@vm:/mnt/f2fs# seq 1 2048 | xargs -n 1 ./test_sync // write inline inode and sync root@vm:/mnt/f2fs# rm -f 1 root@vm:/mnt/f2fs# sync root@vm:/mnt/f2fs# f2fs_io gc_range // move data block in sync mode and not write CP SPO, "fsck --dry-run" find inode has already checkpointed but still with DENT_BIT_SHIFT set The root cause is that GC does not clear the dentry mark and fsync mark during node block migration, leading fsck to misinterpret them as user-issued fsync writes. In BGGC mode, node block migration is handled by f2fs_sync_node_pages(), which guarantees the dentry and fsync marks are cleared before writing. This patch move the set/clear of the fsync|dentry marks into __write_node_folio to make the logic clearer, and ensures the fsync|dentry mark is cleared in FGGC.	2026-05-28	7.1
CVE-2026-46190	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mtd: spi-nor: debugfs: fix out-of-bounds read in spi_nor_params_show() Sashiko noticed an out-of-bounds read [1]. In spi_nor_params_show(), the snor_f_names array is passed to spi_nor_print_flags() using sizeof(snor_f_names). Since snor_f_names is an array of pointers, sizeof() returns the total number of bytes occupied by the pointers (element_count * sizeof(void *)) rather than the element count itself. On 64-bit systems, this makes the passed length 8x larger than intended. Inside spi_nor_print_flags(), the 'names_len' argument is used to bounds-check the 'names' array access. An out-of-bounds read occurs if a flag bit is set that exceeds the array's actual element count but is within the inflated byte-size count. Correct this by using ARRAY_SIZE() to pass the actual number of string pointers in the array.	2026-05-28	7.1
CVE-2026-46199	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn4: Prevent OOB reads when parsing dec msg Check bounds against the end of the BO whenever we access the msg.	2026-05-28	7.1
CVE-2026-46204	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn4: Prevent OOB reads when parsing IB Rewrite the IB parsing to use amdgpu_ib_get_value() which handles the bounds checks.	2026-05-28	7.1
CVE-2026-46218	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Add bounds checking to ib_{get,set}_value The uvd/vce/vcn code accesses the IB at predefined offsets without checking that the IB is large enough. Check the bounds here. The caller is responsible for making sure it can handle arbitrary return values. Also make the idx a uint32_t to prevent overflows causing the condition to fail.	2026-05-28	7.1
CVE-2026-46230	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn3: Prevent OOB reads when parsing dec msg Check bounds against the end of the BO whenever we access the msg.	2026-05-28	7.1

CVE-2026-46237	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu/vcn3: Avoid overflow on msg bound check As pointed out by SDL, the previous condition may be vulnerable to overflow. (cherry picked from commit db00257ac9e4a51eb2515aaea161a019f7125e10)	2026-05-28	7.1
CVE-2025-46284	apple - macos	A race condition was addressed with additional validation. This issue is fixed in macOS Sequoia 15.7, macOS Tahoe 26. An app may be able to gain root privileges.	2026-05-26	7
CVE-2026-46029	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mm/slab: return NULL early from kmalloc_nolock() in NMI on UP On UP kernels (!CONFIG_SMP), spin_trylock() is a no-op that unconditionally succeeds even when the lock is already held. As a result, kmalloc_nolock() called from NMI context can re-enter the slab allocator and acquire n->list_lock that the interrupted context is already holding, corrupting slab state. With CONFIG_DEBUG_SPINLOCK on UP, the following BUG is triggered with the slub_kunit test module: BUG: spinlock trylock failure on UP on CPU#0, kunit_try_catch/243 [...] Call Trace: <NMI> dump_stack_lvl+0x3f/0x60 do_raw_spin_trylock+0x41/0x50 _raw_spin_trylock+0x24/0x50 get_from_partial_node+0x120/0x4d0 ___slab_alloc+0x8a/0x4c0 kmalloc_nolock_noprof+0x164/0x310 [...] </NMI> Fix this by returning NULL early when invoked from NMI on a UP kernel.	2026-05-27	7
CVE-2026-44604	red hat - multiple products	A command injection vulnerability was discovered in the `rpmuncompress` utility of RPM. When extracting certain archive formats (ZIP, 7z, GEM) to a specified destination directory, the tool inserts the archive's top-level folder name into a shell command without properly sanitizing it. A specially crafted archive containing shell metacharacters in its folder name can execute arbitrary commands as the user running the extraction.	2026-05-28	7
CVE-2026-46154	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: sched_ext: Read scx_root under scx_cgroup_ops_rwsem in cgroup setters scx_group_set_{weight,idle,bandwidth}() cache scx_root before acquiring scx_cgroup_ops_rwsem, so the pointer can be stale by the time the op runs. If the loaded scheduler is disabled and freed (via RCU work) and another is enabled between the naked load and the rwsem acquire, the reader sees scx_cgroup_enabled=true (the new scheduler's) but dereferences the freed one - UAF on SCX_HAS_OP(sch, ...) / SCX_CALL_OP(sch, ...). scx_cgroup_enabled is toggled only under scx_cgroup_ops_rwsem write (scx_cgroup_{init,exit}), so reading scx_root inside the rwsem read section correlates @sch with the enabled snapshot.	2026-05-28	7
CVE-2026-46164	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix double free in create_space_info_sub_group() error path When kobject_init_and_add() fails, the call chain is: create_space_info_sub_group() -> btrfs_sysfs_add_space_info_type() -> kobject_init_and_add() -> failure -> kobject_put(&sub_group->kobj) -> space_info_release() -> kfree(sub_group) Then control returns to create_space_info_sub_group(), where: btrfs_sysfs_add_space_info_type() returns error -> kfree(sub_group) Thus, sub_group is freed twice. Keep parent->sub_group[index] = NULL for the failure path, but after btrfs_sysfs_add_space_info_type() has called kobject_put(), let the kobject release callback handle the cleanup.	2026-05-28	7
CVE-2024-11399	synology - beedrive	Files or directories accessible to external parties vulnerability in redis-server component in Synology BeeDrive for desktop before 1.3.2-13814 allows local users to conduct denial-of-service attacks via unspecified vectors.	2026-05-27	6.8
CVE-2026-9704	redhat - build_of_keycloak	A flaw was found in Keycloak. An authenticated user with low privileges can exploit this vulnerability by sending an oversized subject_token JSON Web Token (JWT) to the TokenEndpoint. When the token exceeds a 4000-character limit, it is silently dropped, causing the system to fall back to client credentials. This allows the user to gain the permissions of the client's service account, leading to privilege escalation.	2026-05-27	6.8
CVE-2026-9802	redhat - build_of_keycloak	A flaw was found in Keycloak. When revokeRefreshToken=true is enabled and persistent session storage is in use, a server restart can reset internal timing mechanisms. This allows a remote attacker, who has previously captured a user's refresh token, to replay that token even after it has	2026-05-28	6.8

		been revoked. Successful exploitation grants the attacker unauthorized access to the victim's account, potentially leading to information disclosure or privilege escalation.		
CVE-2026-48916	jenkins - multiple products	Jenkins LDAP Plugin 807.v7d7de30930cf and earlier follows LDAP referrals.	2026-05-27	6.6
CVE-2026-48917	jenkins - multiple products	Jenkins LDAP Plugin 807.v7d7de30930cf and earlier deserializes data from LDAP referrals without validation.	2026-05-27	6.6
CVE-2026-48918	jenkins - active_directory	Jenkins Active Directory Plugin 2.41 and earlier follows LDAP referrals by default.	2026-05-27	6.6
CVE-2026-48919	jenkins - active_directory	Jenkins Active Directory Plugin 2.41 and earlier deserializes data from LDAP referrals without validation.	2026-05-27	6.6
CVE-2026-41863	vmware - spring_ai	Spring AI's support for Anthropic's Skills API used LLM-influenced filenames unsanitized in Path.resolve before writing files to disk. This could allow a malicious user to write files outside the intended target directory, including restricted directories. Affected versions: Spring AI: 1.1.0 through 1.1.x	2026-05-25	6.5
CVE-2026-4795	zyxel - multiple products	A missing authorization vulnerability in Zyxel GS1200-5v3 firmware versions through 1.00(ACPS.2)C0, GS1200-8v3 firmware versions through 1.00(ACPT.2)C0, GS1200-5HPv3 firmware versions through 1.00(ACPU.2)C0, GS1200-8HPv3 firmware versions through 1.00(ACPV.2)C0, and GS1200-10v3 firmware versions through 1.00(ACPW.2)C0 could allow a LAN-based, unauthenticated attacker to read the system configuration from a log file via a crafted HTTP request.	2026-05-26	6.5
CVE-2026-40564	apache - flink_kubernetes_operator	Files or Directories Accessible to External Parties, Server-Side Request Forgery (SSRF) vulnerability in Apache Flink Kubernetes Operator. The FlinkSessionJob jarURI is currently not validated so that it points to user-owned files or addresses. This lets a user with CR create permissions read files from the operator pod's filesystem and pull content from any backing store reachable through Flink's pluggable filesystem layer and access them through the submitted Flink job. Furthermore for fetching from http/https addresses there is currently no allowlist on the URI scheme, no host check, no IP-range restriction, and no protection against pointing the URI at internal or link-local addresses.This issue affects Apache Flink Kubernetes Operator: from 1.3.0 before 1.15.0. Users are recommended to upgrade to version 1.15.0, which fixes the issue.	2026-05-26	6.5
CVE-2026-2340	redhat - multiple products	A flaw was found in Samba’s vfs_worm module. The module is intended to provide write-once, read-many (WORM) protections by preventing modification of files after a configurable grace period. Due to insufficient validation during rename operations, an authenticated user with write access to a share could overwrite a protected file by renaming a newly created file over the existing WORM-protected file.	2026-05-27	6.5
CVE-2026-3676	ibm - multiple products	IBM Cloud APM, Base Private 8.1.4 and IBM Cloud APM, Advanced Private 8.1.4 IBM Db2 for Linux, UNIX and Windows (includes DB2 Connect Server) could allow an authenticated user to cause a denial of service due to improper neutralization of special elements in the data query logic of the Fenced environment.	2026-05-27	6.5
CVE-2026-6052	ibm - multiple products	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to running out of memory when executing certain queries with MDC tables.	2026-05-27	6.5
CVE-2026-6936	ibm - i	IBM i 7.6, 7.5, 7.4, and 7.3 s vulnerable to a denial-of-service attack due to uncontrolled recursion in the Integrated Language Environment (ILE) compiler. An authenticated attacker could exploit this vulnerability by compiling specially crafted source code containing a specific combination of statements.	2026-05-27	6.5
CVE-2026-6938	ibm - db2	IBM Db2 12.1.0 through 12.1.4 is vulnerable to authorization bypass when uploading to a remote object storage path with a special query.	2026-05-27	6.5
CVE-2026-8405	ibm - multiple products	IBM Guardium Data Protection 12.2.1, and 12.2.2 's add-on feature of Guardium Data Protection named "Long Term Retention" (LTR) can expose sensitive credentials in debug mode.	2026-05-27	6.5
CVE-2026-9035	ibm - multiple products	IBM Aspera High-Speed Transfer Endpoint 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Server 3.7.4 through 4.4.7 Fix Pack 1 and IBM Aspera High-Speed Transfer Endpoint are affected by a potential arbitrary file read in the asperahttpd component. An authenticated user may be able to take advantage of this vulnerability to access files in the server’s local storage that they should not have access to.	2026-05-27	6.5
CVE-2026-9792	redhat - build_of_keycloak	A flaw was found in Keycloak's Client Policies, specifically within the `org.keycloak.protocol.oidc` component. When certain condition providers (client-type, client-roles, client-attributes, client-scopes) are used to enforce security restrictions, the `reject-ropc-grant` executor is silently bypassed. This allows an unauthenticated remote attacker to obtain tokens via a Resource Owner Password Credentials (ROPC) grant, even when a policy is explicitly configured to block it. This bypass can lead to unauthorized access and information disclosure.	2026-05-28	6.5
CVE-2026-9796	redhat - build_of_keycloak	A flaw was found in Keycloak. An authenticated administrator with the `manage-clients` role can exploit a Time-of-check to time-of-use (TOCTOU) vulnerability in the name-based admin role checks. This allows the attacker to escalate their privileges to `realm-admin` for all users within the realm, granting them extensive control over the system. The composite role relationship persists even after the attacker's own permissions are revoked and across system reboots.	2026-05-28	6.5
CVE-2026-10004	google - chrome	Insufficient validation of untrusted input in Passwords in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to perform UI spoofing via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-10008	google - chrome	Uninitialized Use in GPU in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-10018	google - chrome	Integer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: Medium)	2026-05-28	6.5
CVE-2026-9882	google - chrome	Integer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: Critical)	2026-05-28	6.5
CVE-2026-9908	google - chrome	Out of bounds read in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5

CVE-2026-9912	google - chrome	Inappropriate implementation in GPU in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-9917	google - chrome	Uninitialized Use in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-9953	google - chrome	Out of bounds read in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-9981	google - chrome	Inappropriate implementation in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2026-9996	google - chrome	Out of bounds read in WebRTC in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	6.5
CVE-2025-36126	ibm - multiple products	IBM Cognos Analytics 11.2.0, 12.0, and 12.1.0 and IBM Cognos Transformer 12.0, 11.2.4, and 12.1.0 is vulnerable to stored cross-site scripting (XSS) in Cognos Adminstration. This vulnerability allows a privileged user to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session.	2026-05-26	6.4
CVE-2026-46416	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. In 3.0.1-4-ge2626659, Microsoft UFO creates one shared UFOWebSocketHandler instance and reuses it for multiple authenticated WebSocket connections. The handler stores per-connection protocol objects in mutable instance fields. Each new WebSocket connection overwrites those fields. Later, message handlers send responses through the shared fields instead of through protocol objects bound to the originating connection. As a result, the most recently connected authenticated client can receive protocol responses that belong to another authenticated client.	2026-05-27	6.3
CVE-2026-9989	google - chrome	Inappropriate implementation in Media in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to bypass same origin policy via a crafted video file. (Chromium security severity: High)	2026-05-28	6.3
CVE-2026-10101	red hat - Multicluster Engine for Kubernetes	ACM/MCE assisted-service writes raw referenced pull-secret contents into `InfraEnv.status.conditions[].message` when pull-secret validation fails. A namespace principal with the stock `view` ClusterRole cannot directly read Secrets, but can read `InfraEnv` objects and recover the referenced Secret's `.dockerconfigjson` data from status. This bypasses the Kubernetes/OpenShift RBAC separation between read-only namespace viewers and Secret readers. In the reproduced proof, the same ServiceAccount was denied `get` and `list` on Secrets, but recovered synthetic pull-secret `username`, `password`, `email`, and base64 `auth` fields through `InfraEnv.status`.	2026-05-29	6.3
CVE-2026-8852	ibm - multiple products	IBM HTTP Server 8.5, and 9.0 is vulnerable to denial of service via the optional module mod_fastcgi module.	2026-05-26	6.2
CVE-2026-2237	synology - storage_manager	A use of get request method with sensitive query strings vulnerability in volume encryption of Synology Storage Manager package before 1.0.1-1100 allows local users on Windows to obtain sensitive information.	2026-05-27	6.2
CVE-2026-45249	apache - echarts	A cross-site scripting (XSS) vulnerability exists in Apache ECharts in the Lines series tooltip rendering logic.This issue affects Apache ECharts: from before 6.1.0. In versions prior to 6.1.0, if both Lines series and tooltip are used, and no user-specified tooltip.formatter is provided, and series.data[i].name is specified, raw HTML string series.data[i].name can be rendered through innerHTML sink into tooltip content. Although tooltip is allowed to accept user-provided raw HTML via a custom tooltip.formatter, the built-in tooltip formatters conventionally perform HTML escaping automatically. This case breaks that convention and may unexpectedly lead to script execution when tooltips are displayed. Users are recommended to upgrade to version 6.1.0 if using the Lines series in this way, which fixes the issue.	2026-05-25	6.1
CVE-2025-13593	synology - activeprotect_agent	Origin validation error vulnerability in Synology ActiveProtect Agent before 1.1.0-0439 allows local users to write arbitrary files with restricted content and conduct denial-of-service during installation.	2026-05-27	6.1
CVE-2025-66592	synology - active_backup_for_business_agent	An origin validation error vulnerability in Synology Active Backup for Business Agent before 3.1.0-4967 allows local users to write arbitrary files with restricted content and conduct denial-of-service during installation.	2026-05-27	6.1
CVE-2025-66593	synology - assistant	An origin validation error vulnerability in Synology Assistant before 7.0.6-50085 allows local users to write arbitrary files with restricted content and conduct denial-of-service during installation.	2026-05-27	6.1
CVE-2026-43827	apache - multiple products	Default configurations of Apache Shiro have a session fixation vulnerability. This issue affects Apache Shiro from 1.0 to 2.1.0, and 3.0.0-alpha-1. Users are recommended to upgrade to version 2.1.1, or 3.0.0-alpha-2 or later, which fixes the issue. In the affected versions, when a session already exists, it is not invalidated upon successful login, nor is a new session being generated with a new ID.	2026-05-25	5.9
CVE-2026-43828	apache - multiple products	Default configurations of Apache Shiro send sensitive cookies in HTTPS session without 'Secure' attribute. This issue affects Apache Shiro from 1.0 to 2.1.0, and 3.0.0-alpha-1.Users are recommended to upgrade to version 2.1.1, or 3.0.0-alpha-2 or later, which fixes the issue. In the affected versions, Shiro-native session manager, as well as Remember-Me manager sends JSESSIONID and rememberMe cookies without 'secure' attribute by default.	2026-05-25	5.9
CVE-2025-10466	synology - safe_access	Improper neutralization of input during web page generation ('Cross-site Scripting') vulnerability in Safe Access in Synology Safe Access before 1.3.1-0329 allows remote authenticated users with administrator privileges to read or write specific files containing non-sensitive information or conduct limited denial-of-service in SRM.	2026-05-27	5.9
CVE-2024-40684	ibm - multiple products	IBM Operations Analytics - Log Analysis 1.3.5.0, 1.3.5.1, 1.3.5.2, 1.3.5.3, 1.3.6.0, 1.3.6.1, 1.3.7.0, 1.3.7.1, 1.3.7.2, and 1.3.8.0, 1.3.8.1, 1.3.8.2, 1.3.8.3, 1.3.8.4 IBM SmartCloud Analytics - Log Analysis does not require that users should have strong passwords by default, which makes it easier for attackers to compromise user accounts.	2026-05-27	5.9
CVE-2026-46538	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. In 3.0.1-4-ge2626659, Microsoft UFO's constellation client tracks pending task responses by session_id	2026-05-27	5.9

		only and does not verify that a TASK_END message came from the device that originally received the task. When the constellation sends a task to a target device, it records a pending Future under a session key. The pending task record stores the expected device ID, but the completion path ignores that binding. If another authenticated peer device sends a forged TASK_END with the same session_id, the constellation accepts the response and completes the victim device's pending Future with attacker-controlled result data. This is an authenticated cross-device task-result injection issue.		
CVE-2026-9793	redhat - build_of_keycloak	A flaw was found in Keycloak. When a JSON Web Encryption (JWE) encrypted request object is submitted, Keycloak may incorrectly process unsigned claims if the decrypted content is raw JSON, bypassing the configured signature policy. This allows a remote attacker to submit unauthorized claims, leading to a compromise of data integrity within the OpenID Connect (OIDC) authorization flow. While a redirect URI allowlist acts as a compensating control, this vulnerability violates OIDC Core and Financial-grade API (FAPI) signing requirements.	2026-05-28	5.9
CVE-2026-8174	zohocorp - Zoho Mail wordpress plugin	Zohocorp Zoho Mail wordpress plugin is vulnerable to Cross-Site request forgery (CSRF). This issue affects Zoho Mail wordpress plugin versions before 1.6.2.	2026-05-26	5.7
CVE-2025-13755	ibm - multiple products	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 for Linux, UNIX and Windows (includes DB2 Connect Server) stores potentially sensitive information in log files that could be read by a local user.	2026-05-26	5.5
CVE-2025-43289	apple - multiple products	A logic issue was addressed with improved validation. This issue is fixed in macOS Sequoia 15.7, macOS Sonoma 14.8, macOS Tahoe 26. A malicious app may be able to access sensitive user data.	2026-05-26	5.5
CVE-2025-43290	apple - multiple products	A permissions issue was addressed with additional restrictions. This issue is fixed in macOS Sequoia 15.7, macOS Sonoma 14.8, macOS Tahoe 26. An app may be able to modify protected parts of the file system.	2026-05-26	5.5
CVE-2025-43451	apple - macos	A permissions issue was addressed by removing the vulnerable code. This issue is fixed in macOS Tahoe 26. An app may be able to access sensitive user data.	2026-05-26	5.5
CVE-2025-46280	apple - macos	An out-of-bounds read was addressed with improved bounds checking. This issue is fixed in macOS Tahoe 26. An app may be able to cause unexpected system termination.	2026-05-26	5.5
CVE-2025-46307	apple - macos	A logic issue was addressed with improved restrictions. This issue is fixed in macOS Tahoe 26. An app may be able to access sensitive user data.	2026-05-26	5.5
CVE-2026-9605	gnu - libredwg	A flaw has been found in GNU libredwg up to 0.13.4.8160. This issue affects the function bit_read_RC of the file bits.c of the component Dwgbmp Utility. This manipulation causes heap-based buffer overflow. The attack is possible to be carried out remotely. The exploit has been published and may be used. Patch name: 8f03865f37f5d4ffd616fef802acc980be54d300. Applying a patch is the recommended action to fix this issue.	2026-05-27	5.5
CVE-2026-5515	ibm - app_connect_enterprise	IBM App Connect Enterprise 13.0.1.0 through 13.0.7.0 stores potentially sensitive information in log files that could be read by a local user.	2026-05-27	5.5
CVE-2026-6051	ibm - multiple products	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to a denial of service when executing a specially crafted query with a small statement heap.	2026-05-27	5.5
CVE-2026-6053	ibm - multiple products	IBM Db2 11.5.0 through 11.5.9, and 12.1.0 through 12.1.4 is vulnerable to a denial of service when a specially crafted query is run with range partitioned tables.	2026-05-27	5.5
CVE-2026-48927	jenkins - buildgraph-view	Jenkins buildgraph-view Plugin 1.8 and earlier does not escape the build URL, resulting in a stored cross-site scripting (XSS) vulnerability exploitable by attackers able to configure jobs or views.	2026-05-27	5.5
CVE-2026-9078	mozilla - firefox	Firefox for iOS displayed specially crafted right-to-left (RTL) and internationalized domain names (IDNs) incorrectly in link preview UI surfaces. A crafted RTL hostname could visually reorder portions of the displayed domain, causing attacker-controlled sites to appear as trusted origins. This vulnerability was fixed in Firefox for iOS 151.1.	2026-05-25	5.4
CVE-2025-14290	ibm - multiple products	IBM webMethods Integration (on prem) -Integration Server 10.15 through IS_10.15_Core_Fix2611.1 to IS_11.1_Core_Fix10 IBM webMethods Integration is vulnerable to server-side request forgery (SSRF). This may allow an authenticated attacker to send unauthorized requests from the system, potentially leading to network enumeration or facilitating other attacks.	2026-05-26	5.4
CVE-2025-36145	ibm - watsonx.data	IBM watsonx.data 2.2 through 2.3.1 IBM Lakehouse does not properly restrict inbound and outbound connections which could allow an attacker to transfer or modify files without restrictions.	2026-05-26	5.4
CVE-2025-36148	ibm - financial_transaction_manager_for_multiplatform	IBM Financial Transaction Manager for SWIFT Services for Multiplatforms 3.2.4.0 through 3.2.4.15 IBM Financial Transaction Manager SWIFT is vulnerable to cross-site scripting. This vulnerability allows an unauthenticated attacker to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session.	2026-05-26	5.4
CVE-2025-13167	synology - contacts	Improper neutralization of input during web page generation ('Cross-site Scripting') vulnerability in contact functionality in Synology Contacts before 1.0.10-20659 allows remote authenticated users to read or write specific files containing non-sensitive information via unspecified vectors.	2026-05-27	5.4
CVE-2025-3633	ibm - multiple products	IBM Cognos Analytics 11.2.0, 11.2.4, 12.0, and 12.1.0 and IBM Cognos Transformer 11.2.4, 12.0, and 12.1.0 are vulnerable to cross-site scripting (XSS). This vulnerability allows a remote attacker to inject arbitrary JavaScript code into the web user interface, which may alter the intended functionality and could lead to the disclosure of credentials within a trusted session.	2026-05-27	5.4
CVE-2026-9971	google - chrome	Inappropriate implementation in iOS in Google Chrome on iOS prior to 148.0.7778.216 allowed a remote attacker who convinced a user to engage in specific UI gestures to inject arbitrary scripts or HTML (UXSS) via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5.4
CVE-2026-46745	apache - apache-airflow-providers-fab	Apache Airflow FAB Auth Manager contains an LDAP filter injection vulnerability (CWE-90) that allows unauthenticated attackers to exfiltrate directory data or bypass authentication. Upgrade to apache-airflow-providers-fab 3.6.4 or later. If immediate upgrade is not possible, disable LDAP authentication until the provider can be updated.	2026-05-25	5.3
CVE-2025-36221	ibm - multiple products	IBM Cloud Pak for Data System - Cyclops 11.3.0.2 through Interim Fix 002 IBM Cloud Pak for Data System uses default passwords default passwords from the manufacturing process for use during the installation process, which could allow an attacker to bypass authentication.	2026-05-26	5.3
CVE-2026-42015	red hat - multiple products	A flaw was found in gnutls. An off-by-one error exists in the PKCS#12 bag element bounds check. This vulnerability allows an remote attacker to write past the internal array of a PKCS#12 bag when appending to a bag that already contains 32 elements. This memory corruption could lead to a denial of service (DoS) or potentially other unspecified impacts.	2026-05-26	5.3

CVE-2024-28765	ibm - multiple products	IBM SDI 7.2.0.0 through 7.2.0.14 and IBM Security Directory Integrator 10.0.0.0 through 10.0.0.2 could allow a remote attacker to obtain sensitive information when a detailed technical error message is returned in the browser. This information could be used in further attacks against the system.	2026-05-27	5.3
CVE-2026-7254	ibm - openbmc	IBM OPENBMC FW1110.00 through FW1110.11 is vulnerable to denial of service attacks by unauthenticated network users.	2026-05-27	5.3
CVE-2026-46544	microsoft - UFO	Microsoft UFO open-source framework for intelligent automation across devices and platforms. In 3.0.1-4-ge2626659, Microsoft UFO accepts client-supplied session_id values in WebSocket task messages and reuses an existing in-memory session object if that session_id already exists. If a prior session has completed and remains in memory with populated results, a different authenticated client can send a new TASK message using the same session_id. The server re-enters the existing session object and sends the stale stored result to the new requester through the normal send_task_end() callback path. This is an authenticated cross-client stale result replay issue. The issue requires that the attacker knows or can predict a live or recently completed session_id.	2026-05-27	5.3
CVE-2026-9794	redhat - build_of_keycloak	A flaw was found in Keycloak. A remote, unauthenticated attacker can exploit this vulnerability by sending specially crafted SOAP requests to the SAML ECP (Security Assertion Markup Language Enhanced Client or Proxy) endpoint with varying client IDs. By observing distinct faultstrings in the responses, the attacker can determine the client's protocol type, leading to information disclosure.	2026-05-28	5.3
CVE-2026-9803	redhat - build_of_keycloak	A flaw was found in Keycloak's ClientRegistrationAuth component. A remote unauthenticated attacker can exploit this vulnerability by sending a specially crafted POST request with a malformed 'Authorization: Bearer' header to any client registration endpoint. This can lead to an ArrayIndexOutOfBoundsException, causing the server to return an HTTP 500 error and resulting in a Denial of Service (DoS) for the affected service.	2026-05-28	5.3
CVE-2026-46830	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Mongoapi). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle REST Data Services accessible data. CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N).	2026-05-28	5.3
CVE-2026-46841	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: General). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle REST Data Services accessible data. CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N).	2026-05-28	5.3
CVE-2026-46842	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle REST Data Services accessible data. CVSS 3.1 Base Score 5.3 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:L/A:N).	2026-05-28	5.3
CVE-2026-46843	oracle - rest_data_services	Vulnerability in Oracle REST Data Services (component: Core). Supported versions that are affected are 24.2.0-26.1.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTPS to compromise Oracle REST Data Services. Successful attacks of this vulnerability can result in unauthorized ability to cause a partial denial of service (partial DOS) of Oracle REST Data Services. CVSS 3.1 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L).	2026-05-28	5.3
CVE-2026-9985	google - chrome	Insufficient validation of untrusted input in Media in Google Chrome on ChromeOS prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to obtain potentially sensitive information from process memory via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5.3
CVE-2026-44598	apache - multiple products	With valid login credentials, URL Redirection to Untrusted Site ('Open Redirect'), Server-Side Request Forgery (SSRF) vulnerability in Apache Shiro. This issue affects Apache Shiro from 2.0-alpha to 2.1.0, and 3.0.0-alpha-1, only when using shiro-jakarta-ee integration module. Users are recommended to upgrade to version 2.1.1, or 3.0.0-alpha-2 or later, which fixes the issue by encrypting the cookie. After successful login, Jakarta EE integration module uses shiroSavedRequest cookie to redirect to a particular web page after login. This cookie was not validated, and can be forged to send a HTTP GET request from the server itself to an arbitrary URL from the cookie.	2026-05-25	5.1
CVE-2026-2607	ibm - multiple products	IBM MQ Operator SC2: v3.2.0 through 3.2.23CD: v3.3.0, v3.4.0, v3.4.1, v3.5.0, v3.5.1 - v3.5.3, v3.6.0 - v3.6.4, v3.7.0 - v3.7.2, v3.8.0, v3.8.1, v3.9.0, v3.9.1LTS: v2.0.0 - 2.0.29 and IBM supplied MQ Advanced container images SC2: 9.4.0.6 through r1, 9.4.0.6-r2, 9.4.0.7-r1, 9.4.0.10-r1, 9.4.0.10-r2, 9.4.0.11-r1, 9.4.0.11-r2, 9.4.0.11-r3, 9.4.0.12-r1, 9.4.0.15-r1 - 9.4.0.15-r4, 9.4.0.16-r1, 9.4.0.16-r2, 9.4.0.17-r1, 9.4.0.17-r2, 9.4.0.20-r1CD: 9.4.1.0-r1, 9.4.1.0-r2, 9.4.1.1-r1, 9.4.2.0-r1, 9.4.2.0-r2, 9.4.2.1-r1, 9.4.2.1-r2, 9.4.3.0-r1, 9.4.3.0-r2, 9.4.3.1-r1 - 9.4.3.1-r3, 9.4.4.0-r1 - 9.4.4.0-r4, 9.4.4.1-r1, 9.4.5.0-r1, 9.4.5.0-r2LTS: 9.3.0.0-r1, 9.3.0.0-r2, 9.3.0.0-r3, 9.3.0.1-r1, 9.3.0.1-r2, 9.3.0.1-r3, 9.3.0.1-r4, 9.3.0.3-r1, 9.3.0.4-r1, 9.3.0.4-r2, 9.3.0.5-r1, 9.3.0.5-r2, 9.3.0.5-r3, 9.3.0.6-r1, 9.3.0.10-r1, 9.3.0.10-r2, 9.3.0.11-r1,9.3.0.11-r2, 9.3.0.15-r1, 9.3.0.16-r1, 9.3.0.16-r2, 9.3.0.17-r1, 9.3.0.17-r2, 9.3.0.17-r3, 9.3.0.20-r1, 9.3.0.20-r2, 9.3.0.21-r1, 9.3.0.21-r2, 9.3.0.21-r3, 9.3.0.25-r1, 9.4.0.0-r1, 9.4.0.0-r2, 9.4.0.0-r3, 9.4.0.5-r1, 9.4.0.5-r2 IBM MQ stores potentially sensitive information in log files that could be read by a local user.	2026-05-27	5.1
CVE-2026-10010	google - chrome	Inappropriate implementation in Input in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5
CVE-2026-9903	google - chrome	Insufficient validation of untrusted input in Site Isolation in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted MHTML page. (Chromium security severity: High)	2026-05-28	5
CVE-2026-9942	google - chrome	Uninitialized Use in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5

CVE-2026-9979	google - chrome	Insufficient validation of untrusted input in Input in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5
CVE-2026-9980	google - chrome	Insufficient validation of untrusted input in Printing in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass site isolation via a crafted HTML page. (Chromium security severity: High)	2026-05-28	5
CVE-2026-42797	apache - multiple products	Exposure of Sensitive Information Through Data Queries vulnerability in Apache Syncope. An administrator with adequate entitlements for Derived Schemas can create a malicious JEXL expression which allows any administrator with sufficient entitlements for User read to access User-related security-sensitive information. This issue affects Apache Syncope: 3.0 through 3.0.16, 4.0 through 4.0.5, 4.1.0. Users are recommended to upgrade to version 4.0.6 / 4.1.1, which fix this issue by further restricting the JEXL expression definition.	2026-05-25	4.9
CVE-2024-47268	synology - surveillance_station	Missing authorization vulnerability in AddOns functionality in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to obtain sensitive information via unspecified vectors.	2026-05-27	4.9
CVE-2024-47269	synology - surveillance_station	Cleartext transmission of sensitive information vulnerability in Export Key functionality in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to obtain sensitive information via unspecified vectors.	2026-05-27	4.9
CVE-2024-47271	synology - surveillance_station	Insufficiently protected credentials vulnerability in IPSpeaker component in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to obtain sensitive information via unspecified vectors.	2026-05-27	4.9
CVE-2026-9801	redhat - build_of_keycloak	A flaw was found in Keycloak. A remote attacker with high privileges, such as a realm administrator configuring a malicious Lightweight Directory Access Protocol (LDAP) server or an attacker compromising an upstream LDAP server, could exploit this vulnerability. By sending a malformed LDAP password policy response during a password authentication request, the attacker can trigger an OutOfMemoryError. This causes the Keycloak Java Virtual Machine (JVM) to terminate, leading to a denial of service (DoS) for all realms on the affected node.	2026-05-28	4.9
CVE-2026-4410	ibm - multiple products	IBM WebSphere Application Server - Liberty 19.0.0.7 through 26.0.0.5 and IBM WebSphere Application Server 9.0, and 8.5 and WebSphere Application Server Liberty are vulnerable to a denial of service, caused by sending a specially-crafted request. A remote attacker could exploit this vulnerability to cause the server to consume memory resources.	2026-05-27	4.8
CVE-2026-6324	red hat - multiple products	A flaw was found in libsoup. A remote attacker could exploit an unsigned to signed conversion error in the `soup_body_input_stream_read_chunked()` function by sending a malicious HTTP request. This vulnerability occurs when libsoup operates behind a non-libsoup proxy server or as a proxy in front of a non-libsoup backend server. Successful exploitation can allow an attacker to bypass security controls, poison web caches, or gain unauthorized access.	2026-05-29	4.8
CVE-2026-5516	ibm - websphere_application_server	IBM WebSphere Application Server - Liberty 22.0.0.11 through 26.0.0.5 IBM WebSphere Application Server Liberty could allow a remote attacker to bypass security under limited conditions by exploiting a specific timing window.	2026-05-27	4.4
CVE-2025-36220	ibm - multiple products	IBM Cloud Pak for Data System - Cyclops 11.3.0.2 through Interim Fix 002 IBM Cloud Pak for Data System is vulnerable to SQL injection. A remote attacker could send specially crafted SQL statements, which could allow the attacker to view, add, modify, or delete information in the back-end database.	2026-05-26	4.3
CVE-2026-1248	ibm - multiple products	IBM Business Automation Workflow containers and traditional may leak information about its database structure in error messages.	2026-05-27	4.3
CVE-2026-48923	jenkins - appspider	Jenkins AppSpider Plugin 1.0.17 and earlier does not perform a permission check in a method implementing form validation, allowing attackers with Overall/Read permission to connect to an attacker-specified URL.	2026-05-27	4.3
CVE-2026-48924	jenkins - bitbucket_oauth	Jenkins Bitbucket OAuth Plugin 0.17 and earlier does not restrict the redirect URL after login, allowing attackers to perform phishing attacks.	2026-05-27	4.3
CVE-2026-48926	jenkins - multiple products	Jenkins Job Import Plugin 143.v044a_2e819b_27 and earlier does not perform a permission check in an HTTP endpoint, allowing attackers with Overall/Read permission to enumerate credentials IDs of credentials stored in Jenkins.	2026-05-27	4.3
CVE-2026-9674	jenkins - multijob	A cross-site request forgery (CSRF) vulnerability in Jenkins Multijob Plugin 662.vd2e0001f6b_b_d and earlier allows attackers to resume failed Multijob builds.	2026-05-27	4.3
CVE-2026-9791	redhat - build_of_keycloak	A flaw was found in Keycloak. An authenticated user with existing organization membership can exploit this flaw by accessing user-facing APIs, such as the account API or by requesting an OpenID Connect (OIDC) token with the 'organization' scope. This allows organization metadata to be disclosed in tokens, even after an administrator has explicitly disabled the Organizations feature, potentially leading to incorrect authorization decisions by resource servers.	2026-05-28	4.3
CVE-2026-9798	redhat - build_of_keycloak	A flaw was found in Keycloak, an open-source identity and access management solution. When a user account is temporarily locked due to repeated failed login attempts, an attacker with valid client credentials can exploit the Client-Initiated Backchannel Authentication (CIBA) flow to bypass this brute-force protection. This allows continued authentication attempts and token issuance even when the account should be locked, potentially enabling further unauthorized access attempts.	2026-05-28	4.3
CVE-2026-40914	apache - multiple products	A vulnerability exists in Apache Artemis whereby an application using the STOMP protocol with security credentials that grant either the consume or send permission on an address can augment the routing-type supported by that address even if said user doesn't have the createAddress permission for that particular address. A user could successfully send a message to an address or consume a message from a queue with a routing-type not supported by the corresponding address when that operation should actually be rejected on the basis that the user doesn't have permission to change the routing-type of the address. Even though the user was already granted permission to send and/or consume messages, they should not be able to augment the routing-type of the address without the createAddress permission. This issue affects Apache Artemis: from 2.50.0 through 2.53.0; Apache ActiveMQ Artemis: from 2.0.0 through 2.44.0. Users are recommended to upgrade to version 2.54.0, which fixes the issue.	2026-05-28	4.3

CVE-2026-10028	red hat - multiple products	A flaw was found in glib-networking. A remote attacker can exploit this vulnerability by presenting a specially crafted certificate chain to an application that uses glib-networking with the GnuTLS backend enabled and performs certificate verification. This crafted chain, which contains circular issuer relationships, can cause an infinite loop during certificate verification. The unbounded traversal consumes excessive CPU resources, leading to a denial of service for the affected process or worker.	2026-05-28	4.3
CVE-2026-9907	google - chrome	Out of bounds read in Dawn in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9911	google - chrome	Integer overflow in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to perform an out of bounds memory read via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9913	google - chrome	Inappropriate implementation in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to potentially perform out of bounds memory access via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9919	google - chrome	Out of bounds read in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9921	google - chrome	Uninitialized Use in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin information via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9929	google - chrome	Inappropriate implementation in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9930	google - chrome	Out of bounds write in Dawn in Google Chrome on Mac prior to 148.0.7778.216 allowed a remote attacker to perform an out of bounds memory write via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9935	google - chrome	Uninitialized Use in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9943	google - chrome	Out of bounds read in WebGL in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9955	google - chrome	Inappropriate implementation in iOS in Google Chrome on iOS prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.3
CVE-2026-9689	redhat - build_of_keycloak	A flaw was found in Keycloak, an open-source identity and access management solution. When a client application is configured to accept broad redirect Uniform Resource Identifiers (URIs), a remote attacker can manipulate the authentication process by crafting a special web address. If a user clicks this link, the client application might incorrectly prioritize attacker-controlled information over legitimate data. This vulnerability, known as HTTP parameter pollution, could allow an attacker to bypass security measures or gain unauthorized access to resources.	2026-05-27	4.2
CVE-2026-9986	google - chrome	Insufficient validation of untrusted input in OptimizationGuide in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to perform UI spoofing via a crafted HTML page. (Chromium security severity: High)	2026-05-28	4.2
CVE-2026-10052	red hat - multiple products	A flaw was found in the Quay config-tool's LDAP and SMTP validation functions. An attacker with config editor access can exploit these functions, which make outbound connections to user-supplied endpoints without proper IP or host filtering. This allows the attacker to perform internal network reconnaissance from the Quay pod's network position, potentially mapping the internal network infrastructure.	2026-05-29	4.1
CVE-2026-10011	google - chrome	Inappropriate implementation in Skia in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2026-9920	google - chrome	Uninitialized Use in GPU in Google Chrome on Android prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2026-9944	google - chrome	Uninitialized Use in ANGLE in Google Chrome prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2026-9950	google - chrome	Insufficient validation of untrusted input in iOS in Google Chrome on iOS prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to bypass same origin policy via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2026-9959	google - chrome	Race in WebRTC in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2026-9991	google - chrome	Inappropriate implementation in Media in Google Chrome on Windows prior to 148.0.7778.216 allowed a remote attacker who had compromised the renderer process to leak cross-origin data via a crafted HTML page. (Chromium security severity: High)	2026-05-28	3.1
CVE-2024-47267	synology - surveillance_station	Improper limitation of a pathname to a restricted directory ('Path Traversal') vulnerability in Archiving Pull functionality in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to limited file write via unspecified vectors.	2026-05-27	2.7
CVE-2024-47270	synology - surveillance_station	Improper preservation of permissions vulnerability in Archiving Push functionality in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to limited file write via unspecified vectors.	2026-05-27	2.7
CVE-2024-47272	synology - surveillance_station	Incorrect authorization vulnerability in IO Module functionality in Synology Surveillance Station before 9.2.2-11575 and 9.2.2-9575 allows remote authenticated users with administrator privileges to limited file write via unspecified vectors.	2026-05-27	2.7
CVE-2026-10078	red hat - multiple products	A flaw was found in the Quay config-tool's GitLab OAuth validator. This vulnerability causes sensitive credentials, specifically client_id and client_secret, to be transmitted as plaintext in URL query parameters during POST requests to the GitLab endpoint. This insecure transmission can lead	2026-05-29	2.7

		to the disclosure of these credentials in various system logs, such as server access logs, reverse proxy logs, and other monitoring systems. An attacker with access to these logs could potentially obtain these credentials, leading to unauthorized information disclosure.		
CVE-2026-10060	trendnet - tew-432brp_firmware	A vulnerability has been found in TRENDnet TEW-432BRP 3.10B20. This impacts the function formSetRoute of the file /goform/formSetRoute. The manipulation of the argument ip/mask/gateway leads to command injection. Remote exploitation of the attack is possible. The exploit has been disclosed to the public and may be used. The vendor explains: "This product has been EOL for 15 years (since 2009). As the item has been EOL for such a long time, we are not able to replicate or fix any vulnerabilities." This vulnerability only affects products that are no longer supported by the maintainer.	2026-05-29	2.1
CVE-2026-10061	trendnet - tew-432brp_firmware	A vulnerability was found in TRENDnet TEW-432BRP 3.10B20. Affected is the function formWPS of the file /goform/formWPS. The manipulation of the argument peerPin results in command injection. The attack can be executed remotely. The exploit has been made public and could be used. The vendor explains: "This product has been EOL for 15 years (since 2009). As the item has been EOL for such a long time, we are not able to replicate or fix any vulnerabilities." This vulnerability only affects products that are no longer supported by the maintainer.	2026-05-29	2.1
CVE-2026-10064	trendnet - tew-432brp_firmware	A security flaw has been discovered in TRENDnet TEW-432BRP 3.10B20. This affects the function formSetPortTr of the file /goform/formSetPortTr. Performing a manipulation of the argument special_name results in stack-based buffer overflow. It is possible to initiate the attack remotely. The exploit has been released to the public and may be used for attacks. The vendor explains: "This product has been EOL for 15 years (since 2009). As the item has been EOL for such a long time, we are not able to replicate or fix any vulnerabilities." This vulnerability only affects products that are no longer supported by the maintainer.	2026-05-29	2.1
CVE-2026-9500	gnu - LibreDWG	A vulnerability was found in GNU LibreDWG up to 0.14. The affected element is the function read_2004_compressed_section of the file src/decode.c of the component Dwgread Utility. Performing a manipulation results in heap-based buffer overflow. The attack is only possible with local access. 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.	2026-05-25	1.9
CVE-2026-9501	gnu - LibreDWG	A vulnerability was determined in GNU LibreDWG up to 0.14. The impacted element is the function decompress_R2004_section of the file src/decode.c of the component Dwgread Utility. Executing a manipulation can lead to reachable assertion. The attack is restricted to local execution. The exploit has been publicly disclosed and may be utilized. This patch is called e501cb9926c1e9a07a0d1cc997f3e69e9be801c9. A patch should be applied to remediate this issue.	2026-05-25	1.9
CVE-2026-9502	gnu - LibreDWG	A vulnerability was identified in GNU LibreDWG up to 0.14. This affects the function decompress_R2004_section of the file src/decode.c of the component Dwgread Utility. The manipulation leads to heap-based buffer overflow. The attack must be carried out locally. The exploit is publicly available and might be used. The identifier of the patch is e501cb9926c1e9a07a0d1cc997f3e69e9be801c9. To fix this issue, it is recommended to deploy a patch.	2026-05-25	1.9
CVE-2026-9503	gnu - LibreDWG	A security flaw has been discovered in GNU LibreDWG up to 0.14. This impacts the function dwg_next_entity of the file src/decode.c of the component DWG File Handler. The manipulation results in null pointer dereference. The attack must be initiated from a local position. The exploit has been released to the public and may be used for attacks. The patch is identified as 8f03865f37f5d4ffd616fef802acc980be54d300. Upgrading the affected component is advised.	2026-05-25	1.9
CVE-2026-9504	gnu - LibreDWG	A weakness has been identified in GNU LibreDWG up to 0.14. Affected is the function bit_convert_TU of the file programs/dwggrep.c of the component Dwggrep Utility. This manipulation causes out-of-bounds read. The attack needs to be launched locally. The exploit has been made available to the public and could be used for attacks. Patch name: be996bf2178a40e98720f18c2414815d244413db. Applying a patch is the recommended action to fix this issue.	2026-05-25	1.9
CVE-2026-9529	gnu - LibreDWG	A security flaw has been discovered in GNU LibreDWG up to 0.14. The affected element is the function match_BLOCK_HEADER of the file dwggrep.c of the component Dwggrep Utility. Performing a manipulation results in null pointer dereference. The attack requires a local approach. The exploit has been released to the public and may be used for attacks.	2026-05-26	1.9
CVE-2026-9530	gnu - LibreDWG	A weakness has been identified in GNU LibreDWG up to 0.14. The impacted element is the function read_2004_compressed_section of the file src/decode.c of the component Dwgbmp Utility. Executing a manipulation can lead to out-of-bounds read. The attack requires local access. The exploit has been made available to the public and could be used for attacks. This patch is called 8f03865f37f5d4ffd616fef802acc980be54d300. It is advisable to implement a patch to correct this issue.	2026-05-26	1.9
CVE-2026-48589	apache - multiple products	Apache Shiro’s Jakarta EE module used the HTTP Referer header in certain cases to issue redirect after a user login. In affected versions, insufficient validation of this client-controlled value could allow an attacker to influence the redirect target in applications using the Jakarta EE module. This issue affects Apache Shiro from 2.0-alpha to 2.2.0, and 3.0.0-alpha-1, only when using shiro-jakarta-ee integration module.	2026-05-25	0