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

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

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 ١٤ يوليو إلى ٢٠ يوليو. 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

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 14th of July to 20th of July. 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	Severity
CVE-2024-39736	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 is vulnerable to HTTP header injection, caused by improper validation of input by the HOST headers. This could allow an attacker to conduct various attacks against the vulnerable system, including cross-site scripting, cache poisoning or session hijacking. IBM X-Force ID: 296003.	2024-07-15	9.8	Critical
CVE-2024-21181	oracle - multiple products	Vulnerability in the Oracle WebLogic Server product of Oracle Fusion Middleware (component: Core). Supported versions that are affected are 12.2.1.4.0 and 14.1.1.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via T3, IIOP to compromise Oracle WebLogic Server. Successful attacks of this vulnerability can result in takeover of Oracle WebLogic Server. 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).	2024-07-16	9.8	Critical
CVE-2024-5471	zohocorp - manageengine_ddi_central	Zohocorp ManageEngine DDI Central versions 4001 and prior were vulnerable to agent takeover vulnerability due to the hard-coded sensitive keys.	2024-07-17	9.8	Critical
CVE-2023-46801	apache - linkis	In Apache Linkis <= 1.5.0, data source management module, when adding Mysql data source, exists remote code execution vulnerability for java version < 1.8.0_241. The deserialization vulnerability exploited through jrmp can inject malicious files into the server and execute them. This attack requires the attacker to obtain an authorized account from Linkis before it can be carried out. We recommend that users upgrade the java version to >= 1.8.0_241. Or users upgrade Linkis to version 1.6.0.	2024-07-15	8.8	High
CVE-2023-49566	apache - linkis	In Apache Linkis <=1.5.0, due to the lack of effective filtering of parameters, an attacker configuring malicious db2 parameters in the DataSource Manager Module will result in jndi injection. Therefore, the parameters in the DB2 URL should be blacklisted. This attack requires the attacker to obtain an authorized account from Linkis before it can be carried out. Versions of Apache Linkis <=1.5.0	2024-07-15	8.8	High

		will be affected. We recommend users upgrade the version of Linkis to version 1.6.0.			
CVE-2024-3168	google - chrome	Use after free in DevTools in Google Chrome prior to 122.0.6261.57 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)	2024-07-16	8.8	High
CVE-2024-3169	google - chrome	Use after free in V8 in Google Chrome prior to 121.0.6167.139 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-07-16	8.8	High
CVE-2024-3170	google - chrome	Use after free in WebRTC in Google Chrome prior to 121.0.6167.85 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-07-16	8.8	High
CVE-2024-3171	google - chrome	Use after free in Accessibility in Google Chrome prior to 122.0.6261.57 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via specific UI gestures. (Chromium security severity: Medium)	2024-07-16	8.8	High
CVE-2024-3172	google - chrome	Insufficient data validation in DevTools in Google Chrome prior to 121.0.6167.85 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)	2024-07-16	8.8	High
CVE-2024-3173	google - chrome	Insufficient data validation in Updater in Google Chrome prior to 120.0.6099.62 allowed a remote attacker to perform OS-level privilege escalation via a malicious file. (Chromium security severity: High)	2024-07-16	8.8	High
CVE-2024-3174	google - chrome	Inappropriate implementation in V8 in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to potentially exploit object corruption via a crafted HTML page. (Chromium security severity: High)	2024-07-16	8.8	High
CVE-2024-3176	google - chrome	Out of bounds write in SwiftShader in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to perform an out of bounds memory write via a crafted HTML page. (Chromium security severity: High)	2024-07-16	8.8	High
CVE-2024-39877	apache - airflow	Apache Airflow 2.4.0, and versions before 2.9.3, has a vulnerability that allows authenticated DAG authors to craft a doc_md parameter in a way that could execute arbitrary code in the scheduler context, which should be forbidden according to the Airflow Security model. Users should upgrade to version 2.9.3 or later which has removed the vulnerability.	2024-07-17	8.8	High
CVE-2024-27311	zohocorp - manageengine_ddi_central	Zohocorp ManageEngine DDI Central versions 4001 and prior were vulnerable to directory traversal vulnerability which allows the user to upload new files to the server folder.	2024-07-17	8.8	High
CVE-2022-48834	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: usb: usbtmc: Fix bug in pipe direction for control transfers The syzbot fuzzer reported a minor bug in the usbtmc driver: usb 5-1: BOGUS control dir, pipe 80001e80 doesn't match bRequestType 0 WARNING: CPU: 0 PID: 3813 at drivers/usb/core/urb.c:412 usb_submit_urb+0x13a5/0x1970 drivers/usb/core/urb.c:410 Modules linked in: CPU: 0 PID: 3813 Comm: syz-executor122 Not tainted 5.17.0-rc5-syzkaller-00306-g2293be58d6a1 #0 ... Call Trace: <TASK> usb_start_wait_urb+0x113/0x530 drivers/usb/core/message.c:58 usb_internal_control_msg drivers/usb/core/message.c:102 [inline] usb_control_msg+0x2a5/0x4b0 drivers/usb/core/message.c:153 usbtmc_ioctl_request drivers/usb/class/usbtmc.c:1947 [inline] The problem is that usbtmc_ioctl_request() uses usb_rcvctrlpipe() for all of its transfers, whether they are in or out. It's easy to fix.	2024-07-16	7.8	High
CVE-2022-48837	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: rndis: prevent integer overflow in rndis_set_response() If "BufOffset" is very large the "BufOffset + 8" operation can have an integer overflow.	2024-07-16	7.8	High
CVE-2022-48847	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2024-07-16	7.8	High

		watch_queue: Fix filter limit check
		
		In watch_queue_set_filter(), there are a couple of places where we check
		that the filter type value does not exceed what the type_filter bitmap
		can hold. One place calculates the number of bits by:
		
		if (tf[i].type >= sizeof(wfilter->type_filter) * 8)
		
		which is fine, but the second does:
		
		if (tf[i].type >= sizeof(wfilter->type_filter) * BITS_PER_LONG)
		
		which is not. This can lead to a couple of out-of-bounds writes due to
		a too-large type:
		
		(1) __set_bit() on wfilter->type_filter
		(2) Writing more elements in wfilter->filters[] than we allocated.
		
		Fix this by just using the proper WATCH_TYPE__NR instead, which is the
		number of types we actually know about.
		
		The bug may cause an oops looking something like:
		
		BUG: KASAN: slab-out-of-bounds in watch_queue_set_filter+0x659/0x740
		Write of size 4 at addr ffff88800d2c66bc by task watch_queue_oob/611
		...
		Call Trace:
		<TASK>
		dump_stack_lvl+0x45/0x59
		print_address_description.constprop.0+0x1f/0x150
		...
		kasan_report.cold+0x7f/0x11b
		...
		watch_queue_set_filter+0x659/0x740
		...
		__x64_sys_ioctl+0x127/0x190
		do_syscall_64+0x43/0x90
		entry_SYSCALL_64_after_hwframe+0x44/0xae
		
		Allocated by task 611:
		kasan_save_stack+0x1e/0x40
		__kasan_kmalloc+0x81/0xa0
		watch_queue_set_filter+0x23a/0x740
		__x64_sys_ioctl+0x127/0x190
		do_syscall_64+0x43/0x90
		entry_SYSCALL_64_after_hwframe+0x44/0xae
		
		The buggy address belongs to the object at ffff88800d2c66a0
		which belongs to the cache kmalloc-32 of size 32
		The buggy address is located 28 bytes inside of
		32-byte region [ffff88800d2c66a0, ffff88800d2c66c0)			

CVE-2022-48848	linux - linux_kernel	In the Linux kernel, the following vulnerability has been resolved: tracing/osnoise: Do not unregister events twice Nicolas reported that using: # trace-cmd record -e all -M 10 -p osnoise --poll Resulted in the following kernel warning: ------------[cut here]------------ WARNING: CPU: 0 PID: 1217 at kernel/tracepoint.c:404 tracepoint_probe_unregister+0x280/0x370 [...] CPU: 0 PID: 1217 Comm: trace-cmd Not tainted 5.17.0-rc6-next-20220307-nico+ #19 RIP: 0010:tracepoint_probe_unregister+0x280/0x370 [...] CR2: 00007ff919b29497 CR3: 0000000109da4005 CR4: 0000000000170ef0 Call Trace: <TASK> osnoise_workload_stop+0x36/0x90	2024-07-16	7.8	High

		tracing_set_tracer+0x108/0x260 tracing_set_trace_write+0x94/0xd0 ? __check_object_size.part.0+0x10a/0x150 ? selinux_file_permission+0x104/0x150 vfs_write+0xb5/0x290 ksys_write+0x5f/0xe0 do_syscall_64+0x3b/0x90 entry_SYSCALL_64_after_hwframe+0x44/0xae RIP: 0033:0x7ff919a18127 [...] ---[end trace 0000000000000000]--- The warning complains about an attempt to unregister an unregistered tracepoint. This happens on trace-cmd because it first stops tracing, and then switches the tracer to nop. Which is equivalent to: # cd /sys/kernel/tracing/ # echo osnoise > current_tracer # echo 0 > tracing_on # echo nop > current_tracer The osnoise tracer stops the workload when no trace instance is actually collecting data. This can be caused both by disabling tracing or disabling the tracer itself. To avoid unregistering events twice, use the existing trace_osnoise_callback_enabled variable to check if the events (and the workload) are actually active before trying to deactivate them.			
CVE-2022-48851	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: staging: gdm724x: fix use after free in gdm_lte_rx() The netif_rx_ni() function frees the skb so we can't dereference it to save the skb->len.	2024-07-16	7.8	High
CVE-2022-48854	linux - linux_kernel	In the Linux kernel, the following vulnerability has been resolved: net: arc_emac: Fix use after free in arc_mdio_probe() If bus->state is equal to MDIOBUS_ALLOCATED, mdiobus_free(bus) will free the "bus". But bus->name is still used in the next line, which will lead to a use after free. We can fix it by putting the name in a local variable and make the bus->name point to the rodata section "name",then use the name in the error message without referring to bus to avoid the uaf.	2024-07-16	7.8	High
CVE-2024-39731	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 uses weaker than expected cryptographic algorithms that could allow an attacker to decrypt highly sensitive information. IBM X-Force ID: 295970.	2024-07-15	7.5	High
CVE-2024-21175	oracle - multiple products	Vulnerability in the Oracle WebLogic Server product of Oracle Fusion Middleware (component: Core). Supported versions that are affected are 12.2.1.4.0 and 14.1.1.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle WebLogic Server. Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critical data or all Oracle WebLogic Server accessible data. CVSS 3.1 Base Score 7.5 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:H/A:N).	2024-07-16	7.5	High
CVE-2024-21182	oracle - multiple products	Vulnerability in the Oracle WebLogic Server product of Oracle Fusion Middleware (component: Core). Supported versions that are affected are 12.2.1.4.0 and 14.1.1.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via T3, IIOP to compromise Oracle WebLogic Server. Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle WebLogic Server accessible data. CVSS 3.1 Base Score 7.5 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N).	2024-07-16	7.5	High
CVE-2024-21183	oracle - multiple products	Vulnerability in the Oracle WebLogic Server product of Oracle Fusion Middleware (component: Core). Supported versions that are affected are 12.2.1.4.0 and 14.1.1.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via T3, IIOP to compromise Oracle WebLogic Server. Successful	2024-07-16	7.5	High

		attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Oracle WebLogic Server accessible data. CVSS 3.1 Base Score 7.5 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N).			
CVE-2024-32007	apache - multiple products	An improper input validation of the p2c parameter in the Apache CXF JOSE code before 4.0.5, 3.6.4 and 3.5.9 allows an attacker to perform a denial of service attack by specifying a large value for this parameter in a token.	2024-07-19	7.5	High
CVE-2024-21184	oracle - database_server	Vulnerability in the Oracle Database RDBMS Security component of Oracle Database Server. Supported versions that are affected are 19.3-19.23. Easily exploitable vulnerability allows high privileged attacker having Execute on SYS.XS_DIAG privilege with network access via Oracle Net to compromise Oracle Database RDBMS Security. Successful attacks of this vulnerability can result in takeover of Oracle Database RDBMS Security. CVSS 3.1 Base Score 7.2 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:H/I:H/A:H).	2024-07-16	7.2	High
CVE-2022-48855	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: sctp: fix kernel-infoleak for SCTP sockets syzbot reported a kernel infoleak [1] of 4 bytes. After analysis, it turned out r->idiag_expires is not initialized if inet_sctp_diag_fill() calls inet_diag_msg_common_fill() Make sure to clear idiag_timer/idiag_retrans/idiag_expires and let inet_diag_msg_sctpasoc_fill() fill them again if needed. [1] BUG: KMSAN: kernel-infoleak in instrument_copy_to_user include/linux/instrumented.h:121 [inline] BUG: KMSAN: kernel-infoleak in copyout lib/iov_iter.c:154 [inline] BUG: KMSAN: kernel-infoleak in _copy_to_iter+0x6ef/0x25a0 lib/iov_iter.c:668 instrument_copy_to_user include/linux/instrumented.h:121 [inline] copyout lib/iov_iter.c:154 [inline] _copy_to_iter+0x6ef/0x25a0 lib/iov_iter.c:668 copy_to_iter include/linux/uio.h:162 [inline] simple_copy_to_iter+0xf3/0x140 net/core/datagram.c:519 __skb_datagram_iter+0x2d5/0x11b0 net/core/datagram.c:425 skb_copy_datagram_iter+0xdc/0x270 net/core/datagram.c:533 skb_copy_datagram_msg include/linux/skbuff.h:3696 [inline] netlink_recvmsg+0x669/0x1c80 net/netlink/af_netlink.c:1977 sock_recvmsg_nosec net/socket.c:948 [inline] sock_recvmsg net/socket.c:966 [inline] __sys_recvfrom+0x795/0xa10 net/socket.c:2097 __do_sys_recvfrom net/socket.c:2115 [inline] __se_sys_recvfrom net/socket.c:2111 [inline] __x64_sys_recvfrom+0x19d/0x210 net/socket.c:2111 do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_64+0x54/0xd0 arch/x86/entry/common.c:82 entry_SYSCALL_64_after_hwframe+0x44/0xae Uninit was created at: slab_post_alloc_hook mm/slab.h:737 [inline] slab_alloc_node mm/slub.c:3247 [inline] __kmalloc_node_track_caller+0xe0c/0x1510 mm/slub.c:4975 kmalloc_reserve net/core/skbuff.c:354 [inline] __alloc_skb+0x545/0xf90 net/core/skbuff.c:426 alloc_skb include/linux/skbuff.h:1158 [inline] netlink_dump+0x3e5/0x16c0 net/netlink/af_netlink.c:2248 __netlink_dump_start+0xcf8/0xe90 net/netlink/af_netlink.c:2373 netlink_dump_start include/linux/netlink.h:254 [inline] inet_diag_handler_cmd+0x2e7/0x400 net/ipv4/inet_diag.c:1341 sock_diag_rcv_msg+0x24a/0x620 netlink_rcv_skb+0x40c/0x7e0 net/netlink/af_netlink.c:2494 sock_diag_rcv+0x63/0x80 net/core/sock_diag.c:277 netlink_unicast_kernel net/netlink/af_netlink.c:1317 [inline] netlink_unicast+0x1093/0x1360 net/netlink/af_netlink.c:1343 netlink_sendmsg+0x14d9/0x1720 net/netlink/af_netlink.c:1919 sock_sendmsg_nosec net/socket.c:705 [inline] sock_sendmsg net/socket.c:725 [inline] sock_write_iter+0x594/0x690 net/socket.c:1061 do_iter_readv_writev+0xa7f/0xc70 do_iter_write+0x52c/0x1500 fs/read_write.c:851 vfs_writev fs/read_write.c:924 [inline]	2024-07-16	7.1	High

		do_writev+0x645/0xe00 fs/read_write.c:967 __do_sys_writev fs/read_write.c:1040 [inline] __se_sys_writev fs/read_write.c:1037 [inline] __x64_sys_writev+0xe5/0x120 fs/read_write.c:1037 do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_64+0x54/0xd0 arch/x86/entry/common.c:82 entry_SYSCALL_64_after_hwframe+0x44/0xae Bytes 68-71 of 2508 are uninitialized Memory access of size 2508 starts at ffff888114f9b000 Data copied to user address 00007f7fe09ff2e0 CPU: 1 PID: 3478 Comm: syz-executor306 Not tainted 5.17.0-rc4-syzkaller #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/01/2011			
CVE-2022-48866	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: HID: hid-thrustmaster: fix OOB read in thrustmaster_interrupts Syzbot reported an slab-out-of-bounds Read in thrustmaster_probe() bug. The root case is in missing validation check of actual number of endpoints. Code should not blindly access usb_host_interface::endpoint array, since it may contain less endpoints than code expects. Fix it by adding missing validaion check and print an error if number of endpoints do not match expected number	2024-07-16	7.1	High
CVE-2022-48858	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Fix a race on command flush flow Fix a refcount use after free warning due to a race on command entry. Such race occurs when one of the commands releases its last refcount and frees its index and entry while another process running command flush flow takes refcount to this command entry. The process which handles commands flush may see this command as needed to be flushed if the other process released its refcount but didn't release the index yet. Fix it by adding the needed spin lock. It fixes the following warning trace: refcount_t: addition on 0; use-after-free. WARNING: CPU: 11 PID: 540311 at lib/refcount.c:25 refcount_warn_saturate+0x80/0xe0 ... RIP: 0010:refcount_warn_saturate+0x80/0xe0 ... Call Trace: <TASK> mlx5_cmd_trigger_completions+0x293/0x340 [mlx5_core] mlx5_cmd_flush+0x3a/0xf0 [mlx5_core] enter_error_state+0x44/0x80 [mlx5_core] mlx5_fw_fatal_reporter_err_work+0x37/0xe0 [mlx5_core] process_one_work+0x1be/0x390 worker_thread+0x4d/0x3d0 ? rescuer_thread+0x350/0x350 kthread+0x141/0x160 ? set_kthread_struct+0x40/0x40 ret_from_fork+0x1f/0x30 </TASK>	2024-07-16	7	High
CVE-2023-41916	apache - linkis	In Apache Linkis =1.4.0, due to the lack of effective filtering of parameters, an attacker configuring malicious Mysql JDBC parameters in the DataSource Manager Module will trigger arbitrary file reading. Therefore, the parameters in the Mysql JDBC URL should be blacklisted. This attack requires the attacker to obtain an authorized account from Linkis before it can be carried out. Versions of Apache Linkis = 1.4.0 will be affected. We recommend users upgrade the version of Linkis to version 1.5.0.	2024-07-15	6.5	Medium

CVE-2024-21177	oracle - multiple products	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported versions that are affected are 8.0.37 and prior and 8.4.0 and prior. Easily exploitable vulnerability allows low privileged attacker with network access via multiple protocols to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of MySQL Server. CVSS 3.1 Base Score 6.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:N/I:N/A:H).	2024-07-16	6.5	Medium
CVE-2024-2884	google - chrome	Out of bounds read in V8 in Google Chrome prior to 121.0.6167.139 allowed a remote attacker to potentially perform out of bounds memory access via a crafted HTML page. (Chromium security severity: Medium)	2024-07-16	6.5	Medium
CVE-2024-5500	google - chrome	Inappropriate implementation in Sign-In in Google Chrome prior to 1.3.36.351 allowed a remote attacker to bypass navigation restrictions via a crafted HTML page. (Chromium security severity: Medium)	2024-07-16	6.5	Medium
CVE-2024-3175	google - chrome	Insufficient data validation in Extensions in Google Chrome prior to 120.0.6099.62 allowed a remote attacker to perform privilege escalation via a crafted Chrome Extension. (Chromium security severity: Low)	2024-07-16	6.3	Medium
CVE-2024-21178	oracle - multiple products	Vulnerability in the PeopleSoft Enterprise PeopleTools product of Oracle PeopleSoft (component: Portal). Supported versions that are affected are 8.59, 8.60 and 8.61. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise PeopleSoft Enterprise PeopleTools. Successful attacks require human interaction from a person other than the attacker and while the vulnerability is in PeopleSoft Enterprise PeopleTools, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of PeopleSoft Enterprise PeopleTools accessible data as well as unauthorized read access to a subset of PeopleSoft Enterprise PeopleTools accessible data. CVSS 3.1 Base Score 6.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:C/C:L/I:L/A:N).	2024-07-16	6.1	Medium
CVE-2024-21188	oracle - multiple products	Vulnerability in the Oracle Financial Services Revenue Management and Billing product of Oracle Financial Services Applications (component: Chatbot). Supported versions that are affected are 6.0.0.0.0 and 6.1.0.0.0. Easily exploitable vulnerability allows unauthenticated attacker with network access via HTTP to compromise Oracle Financial Services Revenue Management and Billing. Successful attacks require human interaction from a person other than the attacker and while the vulnerability is in Oracle Financial Services Revenue Management and Billing, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle Financial Services Revenue Management and Billing accessible data as well as unauthorized read access to a subset of Oracle Financial Services Revenue Management and Billing accessible data. CVSS 3.1 Base Score 6.1 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:C/C:L/I:L/A:N).	2024-07-16	6.1	Medium
CVE-2022-48835	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: scsi: mpt3sas: Page fault in reply q processing A page fault was encountered in mpt3sas on a LUN reset error path: [145.763216] mpt3sas_cm1: Task abort tm failed: handle(0x0002),timeout(30) tr_method(0x0) smid(3) msix_index(0) [145.778932] scsi 1:0:0:0: task abort: FAILED scmd(0x0000000024ba29a2) [145.817307] scsi 1:0:0:0: attempting device reset! scmd(0x0000000024ba29a2) [145.827253] scsi 1:0:0:0: [sg1] tag#2 CDB: Receive Diagnostic 1c 01 01 ff fc 00 [145.837617] scsi target1:0:0: handle(0x0002), sas_address(0x500605b0000272b9), phy(0) [145.848598] scsi target1:0:0: enclosure logical id(0x500605b0000272b8), slot(0) [149.858378] mpt3sas_cm1: Poll ReplyDescriptor queues for completion of smid(0), task_type(0x05), handle(0x0002) [149.875202] BUG: unable to handle page fault for address: 00000007fffc445d [149.885617] #PF: supervisor read access in kernel mode	2024-07-16	5.5	Medium

		[149.894346] #PF: error_code(0x0000) - not-present page [149.903123] PGD 0 P4D 0 [149.909387] Oops: 0000 [#1] PREEMPT SMP NOPTI [149.917417] CPU: 24 PID: 3512 Comm: scsi_eh_1 Kdump: loaded Tainted: G S O 5.10.89-altav-1 #1 [149.934327] Hardware name: DDN 200NVX2 /200NVX2-MB , BIOS ATHG2.2.02.01 09/10/2021 [149.951871] RIP: 0010:_base_process_reply_queue+0x4b/0x900 [mpt3sas] [149.961889] Code: 0f 84 22 02 00 00 8d 48 01 49 89 fd 48 8d 57 38 f0 0f b1 4f 38 0f 85 d8 01 00 00 49 8b 45 10 45 31 e4 41 8b 55 0c 48 8d 1c d0 <0f> b6 03 83 e0 0f 3c 0f 0f 85 a2 00 00 00 e9 e6 01 00 00 0f b7 ee [149.991952] RSP: 0018:ffffc9000f1ebcb8 EFLAGS: 00010246 [150.000937] RAX: 0000000000000055 RBX: 00000007fffc445d RCX: 000000002548f071 [150.011841] RDX: 00000000ffff8881 RSI: 0000000000000001 RDI: ffff888125ed50d8 [150.022670] RBP: 0000000000000000 R08: 0000000000000000 R09: c0000000ffff7fff [150.033445] R10: ffffc9000f1ebb68 R11: ffffc9000f1ebb60 R12: 0000000000000000 [150.044204] R13: ffff888125ed50d8 R14: 0000000000000080 R15: 34cdc00034cdea80 [150.054963] FS: 0000000000000000(0000) GS:ffff88dfaf200000(0000) knlGS:0000000000000000 [150.066715] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [150.076078] CR2: 00000007fffc445d CR3: 000000012448a006 CR4: 0000000000770ee0 [150.086887] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 [150.097670] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [150.108323] PKRU: 55555554 [150.114690] Call Trace: [150.120497] ? printk+0x48/0x4a [150.127049] mpt3sas_scsih_issue_tm.cold.114+0x2e/0x2b3 [mpt3sas] [150.136453] mpt3sas_scsih_issue_locked_tm+0x86/0xb0 [mpt3sas] [150.145759] scsih_dev_reset+0xea/0x300 [mpt3sas] [150.153891] scsi_eh_ready_devs+0x541/0x9e0 [scsi_mod] [150.162206] ? __scsi_host_match+0x20/0x20 [scsi_mod] [150.170406] ? scsi_try_target_reset+0x90/0x90 [scsi_mod] [150.178925] ? blk_mq_tagset_busy_iter+0x45/0x60 [150.186638] ? scsi_try_target_reset+0x90/0x90 [scsi_mod] [150.195087] scsi_error_handler+0x3a5/0x4a0 [scsi_mod] [150.203206] ? __schedule+0x1e9/0x610 [150.209783] ? scsi_eh_get_sense+0x210/0x210 [scsi_mod] [150.217924] kthread+0x12e/0x150 [150.224041] ? kthread_worker_fn+0x130/0x130 [150.231206] ret_from_fork+0x1f/0x30 This is caused by mpt3sas_base_sync_reply_irqs() using an invalid reply_q pointer outside of the list_for_each_entry() loop. At the end of the full list traversal the pointer is invalid. Move the _base_process_reply_queue() call inside of the loop.			
CVE-2022-48836	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Input: aiptek - properly check endpoint type Syzbot reported warning in usb_submit_urb() which is caused by wrong endpoint type. There was a check for the number of endpoints, but not for the type of endpoint. Fix it by replacing old desc.bNumEndpoints check with usb_find_common_endpoints() helper for finding endpoints Fail log: usb 5-1: BOGUS urb xfer, pipe 1 != type 3 WARNING: CPU: 2 PID: 48 at drivers/usb/core/urb.c:502 usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb.c:502 Modules linked in:	2024-07-16	5.5	Medium

		CPU: 2 PID: 48 Comm: kworker/2:2 Not tainted 5.17.0-rc6-syzkaller-00226-g07ebd38a0da2 #0 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.14.0-2 04/01/2014 Workqueue: usb_hub_wq hub_event ... Call Trace: <TASK> aiptek_open+0xd5/0x130 drivers/input/tablet/aiptek.c:830 input_open_device+0x1bb/0x320 drivers/input/input.c:629 kbd_connect+0xfe/0x160 drivers/tty/vt/keyboard.c:1593			

CVE-2022-48838	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		usb: gadget: Fix use-after-free bug by not setting udc->dev.driver
		
		The syzbot fuzzer found a use-after-free bug:
		
		BUG: KASAN: use-after-free in dev_uevent+0x712/0x780 drivers/base/core.c:2320
		Read of size 8 at addr ffff88802b934098 by task udevd/3689
		
		CPU: 2 PID: 3689 Comm: udevd Not tainted 5.17.0-rc4-syzkaller-00229-g4f12b742eb2b #0
		Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.14.0-2 04/01/2014
		Call Trace:
		<TASK>
		__dump_stack lib/dump_stack.c:88 [inline]
		dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106
		print_address_description.constprop.0.cold+0x8d/0x303 mm/kasan/report.c:255
		__kasan_report mm/kasan/report.c:442 [inline]
		kasan_report.cold+0x83/0xdf mm/kasan/report.c:459
		dev_uevent+0x712/0x780 drivers/base/core.c:2320
		uevent_show+0x1b8/0x380 drivers/base/core.c:2391
		dev_attr_show+0x4b/0x90 drivers/base/core.c:2094
		
		Although the bug manifested in the driver core, the real cause was a
		race with the gadget core. dev_uevent() does:
		
		if (dev->driver)
		add_uevent_var(env, "DRIVER=%s", dev->driver->name);
		
		and between the test and the dereference of dev->driver, the gadget
		core sets dev->driver to NULL.
		
		The race wouldn't occur if the gadget core registered its devices on a real bus, using the standard synchronization techniques of the driver core. However, it's not necessary to make such a large change
		in order to fix this bug; all we need to do is make sure that udc->dev.driver is always NULL.
		
		In fact, there is no reason for udc->dev.driver ever to be set to anything, let alone to the value it currently gets: the address of the gadget's driver. After all, a gadget driver only knows how to manage
		a gadget, not how to manage a UDC.
		
		This patch simply removes the statements in the gadget core that touch
		udc->dev.driver.
			2024-07-16	5.5	Medium

CVE-2022-48839	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net/packet: fix slab-out-of-bounds access in packet_recvmsg() syzbot found that when an AF_PACKET socket is using PACKET_COPY_THRESH and mmap operations, tpacket_rcv() is queueing skbs with garbage in skb->cb[], triggering a too big copy [1] Presumably, users of af_packet using mmap() already gets correct metadata from the mapped buffer, we can simply make sure to clear 12 bytes that might be copied to user space later. BUG: KASAN: stack-out-of-bounds in memcpy include/linux/fortify-string.h:225 [inline]	2024-07-16	5.5	Medium

		BUG: KASAN: stack-out-of-bounds in packet_recvmsg+0x56c/0x1150 net/packet/af_packet.c:3489 Write of size 165 at addr ffffc9000385fb78 by task syz-executor233/3631 CPU: 0 PID: 3631 Comm: syz-executor233 Not tainted 5.17.0-rc7-syzkaller-02396-g0b3660695e80 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/01/2011 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106 print_address_description.constprop.0.cold+0xf/0x336 mm/kasan/report.c:255 __kasan_report mm/kasan/report.c:442 [inline] kasan_report.cold+0x83/0xdf mm/kasan/report.c:459 check_region_inline mm/kasan/generic.c:183 [inline] kasan_check_range+0x13d/0x180 mm/kasan/generic.c:189 memcpy+0x39/0x60 mm/kasan/shadow.c:66 memcpy include/linux/fortify-string.h:225 [inline] packet_recvmsg+0x56c/0x1150 net/packet/af_packet.c:3489 sock_recvmsg_nosec net/socket.c:948 [inline] sock_recvmsg net/socket.c:966 [inline] sock_recvmsg net/socket.c:962 [inline] ____sys_recvmsg+0x2c4/0x600 net/socket.c:2632 ___sys_recvmsg+0x127/0x200 net/socket.c:2674 __sys_recvmsg+0xe2/0x1a0 net/socket.c:2704 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x44/0xae RIP: 0033:0x7fdfd5954c29 Code: 28 00 00 00 75 05 48 83 c4 28 c3 e8 41 15 00 00 90 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d8 64 89 01 48 RSP: 002b:00007ffcf8e71e48 EFLAGS: 00000246 ORIG_RAX: 000000000000002f RAX: ffffffffffffffda RBX: 0000000000000003 RCX: 00007fdfd5954c29 RDX: 0000000000000000 RSI: 0000000020000500 RDI: 0000000000000005 RBP: 0000000000000000 R08: 000000000000000d R09: 000000000000000d R10: 0000000000000000 R11: 0000000000000246 R12: 00007ffcf8e71e60 R13: 00000000000f4240 R14: 000000000000c1ff R15: 00007ffcf8e71e54 </TASK> addr ffffc9000385fb78 is located in stack of task syz-executor233/3631 at offset 32 in frame: ____sys_recvmsg+0x0/0x600 include/linux/uio.h:246 this frame has 1 object: [32, 160) 'addr' Memory state around the buggy address: ffffc9000385fa80: 00 04 f3 f3 f3 f3 f3 00 00 00 00 00 00 00 00 00 ffffc9000385fb00: 00 00 00 00 00 00 00 00 00 00 00 00 f1 f1 f1 f1 00 >ffffc9000385fb80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 f3 ^ ffffc9000385fc00: f3 f3 f3 00 00 00 00 00 00 00 00 00 00 00 00 f1 ffffc9000385fc80: f1 f1 f1 00 f2 f2 f2 00 f2 f2 f2 00 00 00 00 00 ========== ==========			
CVE-2022-48840	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: iavf: Fix hang during reboot/shutdown Recent commit 974578017fc1 ("iavf: Add waiting so the port is initialized in remove") adds a wait-loop at the beginning of iavf_remove() to ensure that port initialization is finished prior unregistering net device. This causes a regression in reboot/shutdown scenario because in this case callback iavf_shutdown() is called and this callback detaches the device, makes it down if it is running and sets its state to __IAVF_REMOVE. Later shutdown callback of associated PF driver (e.g. ice_shutdown)	2024-07-16	5.5	Medium

		is called. That callback calls among other things sriov_disable() that calls indirectly iavf_remove() (see stack trace below). As the adapter state is already __IAVF_REMOVE then the mentioned loop is end-less and shutdown process hangs. The patch fixes this by checking adapter's state at the beginning of iavf_remove() and skips the rest of the function if the adapter is already in remove state (shutdown is in progress). Reproducer: 1. Create VF on PF driven by ice or i40e driver 2. Ensure that the VF is bound to iavf driver 3. Reboot [52625.981294] sysrq: SysRq : Show Blocked State [52625.988377] task:reboot state:D stack: 0 pid:17359 ppid: 1 f2 [52625.996732] Call Trace: [52625.999187] __schedule+0x2d1/0x830 [52626.007400] schedule+0x35/0xa0 [52626.010545] schedule_hrtimeout_range_clock+0x83/0x100 [52626.020046] usleep_range+0x5b/0x80 [52626.023540] iavf_remove+0x63/0x5b0 [iavf] [52626.027645] pci_device_remove+0x3b/0xc0 [52626.031572] device_release_driver_internal+0x103/0x1f0 [52626.036805] pci_stop_bus_device+0x72/0xa0 [52626.040904] pci_stop_and_remove_bus_device+0xe/0x20 [52626.045870] pci_iov_remove_virtfn+0xba/0x120 [52626.050232] sriov_disable+0x2f/0xe0 [52626.053813] ice_free_vfs+0x7c/0x340 [ice] [52626.057946] ice_remove+0x220/0x240 [ice] [52626.061967] ice_shutdown+0x16/0x50 [ice] [52626.065987] pci_device_shutdown+0x34/0x60 [52626.070086] device_shutdown+0x165/0x1c5 [52626.074011] kernel_restart+0xe/0x30 [52626.077593] __do_sys_reboot+0x1d2/0x210 [52626.093815] do_syscall_64+0x5b/0x1a0 [52626.097483] entry_SYSCALL_64_after_hwframe+0x65/0xca			
CVE-2022-48841	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ice: fix NULL pointer dereference in ice_update_vsi_tx_ring_stats() It is possible to do NULL pointer dereference in routine that updates Tx ring stats. Currently only stats and bytes are updated when ring pointer is valid, but later on ring is accessed to propagate gathered Tx stats onto VSI stats. Change the existing logic to move to next ring when ring is NULL.	2024-07-16	5.5	Medium
CVE-2022-48843	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/vrr: Set VRR capable prop only if it is attached to connector VRR capable property is not attached by default to the connector It is attached only if VRR is supported. So if the driver tries to call drm core set prop function without it being attached that causes NULL dereference.	2024-07-16	5.5	Medium
CVE-2022-48844	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_core: Fix leaking sent_cmd skb sent_cmd memory is not freed before freeing hci_dev causing it to leak it contents.	2024-07-16	5.5	Medium
CVE-2022-48845	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: MIPS: smp: fill in sibling and core maps earlier After enabling CONFIG_SCHED_CORE (landed during 5.14 cycle), 2-core 2-thread-per-core interAptiv (CPS-driven) started emitting the following: [0.025698] CPU1 revision is: 0001a120 (MIPS interAptiv (multi)) [0.048183] ------------[cut here]------------ [0.048187] WARNING: CPU: 1 PID: 0 at kernel/sched/core.c:6025 sched_core_cpu_starting+0x198/0x240 [0.048220] Modules linked in:	2024-07-16	5.5	Medium

		[0.048233] CPU: 1 PID: 0 Comm: swapper/1 Not tainted 5.17.0-rc3+ #35 b7b319f24073fd9a3c2aa7ad15fb7993eec0b26f [0.048247] Stack : 817f0000 00000004 327804c8 810eb050 00000000 00000004 00000000 c314fdd1 [0.048278] 830cbd64 819c0000 81800000 817f0000 83070bf4 00000001 830cbd08 00000000 [0.048307] 00000000 00000000 815fcbc4 00000000 00000000 00000000 00000000 00000000 [0.048334] 00000000 00000000 00000000 00000000 817f0000 00000000 00000000 817f6f34 [0.048361] 817f0000 818a3c00 817f0000 00000004 00000000 00000000 4dc33260 0018c933 [0.048389] ... [0.048396] Call Trace: [0.048399] [<8105a7bc>] show_stack+0x3c/0x140 [0.048424] [<8131c2a0>] dump_stack_lvl+0x60/0x80 [0.048440] [<8108b5c0>] __warn+0xc0/0xf4 [0.048454] [<8108b658>] warn_slowpath_fmt+0x64/0x10c [0.048467] [<810bd418>] sched_core_cpu_starting+0x198/0x240 [0.048483] [<810c6514>] sched_cpu_starting+0x14/0x80 [0.048497] [<8108c0f8>] cpuhp_invoke_callback_range+0x78/0x140 [0.048510] [<8108d914>] notify_cpu_starting+0x94/0x140 [0.048523] [<8106593c>] start_secondary+0xbc/0x280 [0.048539] [0.048543] ---[end trace 0000000000000000]--- [0.048636] Synchronize counters for CPU 1: done. ...for each but CPU 0/boot. Basic debug printks right before the mentioned line say: [0.048170] CPU: 1, smt_mask: So smt_mask, which is sibling mask obviously, is empty when entering the function. This is critical, as sched_core_cpu_starting() calculates core-scheduling parameters only once per CPU start, and it's crucial to have all the parameters filled in at that moment (at least it uses cpu_smt_mask() which in fact is `&cpu_sibling_map[cpu]` on MIPS). A bit of debugging led me to that set_cpu_sibling_map() performing the actual map calculation, was being invocated after notify_cpu_start(), and exactly the latter function starts CPU HP callback round (sched_core_cpu_starting() is basically a CPU HP callback). While the flow is same on ARM64 (maps after the notifier, although before calling set_cpu_online()), x86 started calculating sibling maps earlier than starting the CPU HP callbacks in Linux 4.14 (see [0] for the reference). Neither me nor my brief tests couldn't find any potential caveats in calculating the maps right after performing delay calibration, but the WARN splat is now gone. The very same debug prints now yield exactly what I expected from them: [0.048433] CPU: 1, smt_mask: 0-1 [0] https://git.kernel.org/pub/scm/linux/kernel/git/mips/linux.git/commit/?id=76ce7cfe35ef			
CVE-2022-48846	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: block: release rq qos structures for queue without disk blkcg_init_queue() may add rq qos structures to request queue, previously blk_cleanup_queue() calls rq_qos_exit() to release them, but commit 8e141f9eb803 ("block: drain file system I/O on del_gendisk") moves rq_qos_exit() into del_gendisk(), so memory leak is caused because queues may not have disk, such as un-present scsi luns, nvme admin queue, ...	2024-07-16	5.5	Medium

		Fixes the issue by adding rq_qos_exit() to blk_cleanup_queue() back. BTW, v5.18 won't need this patch any more since we move blkcg_init_queue()/blkcg_exit_queue() into disk allocation/release handler, and patches have been in for-5.18/block.			
CVE-2022-48849	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: bypass tiling flag check in virtual display case (v2) vkms leverages common amdgpu framebuffer creation, and also as it does not support FB modifier, there is no need to check tiling flags when initing framebuffer when virtual display is enabled. This can fix below calltrace: amdgpu 0000:00:08.0: GFX9+ requires FB check based on format modifier WARNING: CPU: 0 PID: 1023 at drivers/gpu/drm/amd/amdgpu/amdgpu_display.c:1150 amdgpu_display_framebuffer_init+0x8e7/0xb40 [amdgpu] v2: check adev->enable_virtual_display instead as vkms can be enabled in bare metal as well.	2024-07-16	5.5	Medium
CVE-2022-48850	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net-sysfs: add check for netdevice being present to speed_show When bringing down the netdevice or system shutdown, a panic can be triggered while accessing the sysfs path because the device is already removed. [755.549084] mlx5_core 0000:12:00.1: Shutdown was called [756.404455] mlx5_core 0000:12:00.0: Shutdown was called ... [757.937260] BUG: unable to handle kernel NULL pointer dereference at (null) [758.031397] IP: [<ffffffff8ee11acb>] dma_pool_alloc+0x1ab/0x280 crash> bt ... PID: 12649 TASK: ffff8924108f2100 CPU: 1 COMMAND: "amsd" ... #9 [ffff89240e1a38b0] page_fault at ffffffff8f38c778 [exception RIP: dma_pool_alloc+0x1ab] RIP: ffffffff8ee11acb RSP: ffff89240e1a3968 RFLAGS: 00010046 RAX: 0000000000000246 RBX: ffff89243d874100 RCX: 0000000000001000 RDX: 0000000000000000 RSI: 0000000000000246 RDI: ffff89243d874090 RBP: ffff89240e1a39c0 R8: 000000000001f080 R9: ffff8905ffc03c00 R10: ffffffffc04680d4 R11: ffffffff8edde9fd R12: 00000000000080d0 R13: ffff89243d874090 R14: ffff89243d874080 R15: 0000000000000000 ORIG_RAX: ffffffffffffffff CS: 0010 SS: 0018 #10 [ffff89240e1a39c8] mlx5_alloc_cmd_msg at ffffffffc04680f3 [mlx5_core] #11 [ffff89240e1a3a18] cmd_exec at ffffffffc046ad62 [mlx5_core] #12 [ffff89240e1a3ab8] mlx5_cmd_exec at ffffffffc046b4fb [mlx5_core] #13 [ffff89240e1a3ae8] mlx5_core_access_reg at ffffffffc0475434 [mlx5_core] #14 [ffff89240e1a3b40] mlx5e_get_fec_caps at ffffffffc04a7348 [mlx5_core] #15 [ffff89240e1a3bb0] get_fec_supported_advertised at ffffffffc04992bf [mlx5_core] #16 [ffff89240e1a3c08] mlx5e_get_link_ksettings at ffffffffc049ab36 [mlx5_core] #17 [ffff89240e1a3ce8] __ethtool_get_link_ksettings at ffffffff8f25db46	2024-07-16	5.5	Medium

		#18 [ffff89240e1a3d48] speed_show at ffffffff8f277208 #19 [ffff89240e1a3dd8] dev_attr_show at ffffffff8f0b70e3 #20 [ffff89240e1a3df8] sysfs_kf_seq_show at ffffffff8eedbedf #21 [ffff89240e1a3e18] kernfs_seq_show at ffffffff8eeda596 #22 [ffff89240e1a3e28] seq_read at ffffffff8ee76d10 #23 [ffff89240e1a3e98] kernfs_fop_read at ffffffff8eedaef5 #24 [ffff89240e1a3ed8] vfs_read at ffffffff8ee4e3ff #25 [ffff89240e1a3f08] sys_read at ffffffff8ee4f27f #26 [ffff89240e1a3f50] system_call_fastpath at ffffffff8f395f92 crash> net_device.state ffff89443b0c0000 state = 0x5 (__LINK_STATE_START| __LINK_STATE_NOCARRIER) To prevent this scenario, we also make sure that the netdevice is present.			
CVE-2022-48853	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: swiotlb: fix info leak with DMA_FROM_DEVICE The problem I'm addressing was discovered by the LTP test covering cve-2018-1000204. A short description of what happens follows: 1) The test case issues a command code 00 (TEST UNIT READY) via the SG_IO interface with: dxfer_len == 524288, dxdfer_dir == SG_DXFER_FROM_DEV and a corresponding dxferp. The peculiar thing about this is that TUR is not reading from the device. 2) In sg_start_req() the invocation of blk_rq_map_user() effectively bounces the user-space buffer. As if the device was to transfer into it. Since commit a45b599ad808 ("scsi: sg: allocate with __GFP_ZERO in sg_build_indirect()") we make sure this first bounce buffer is allocated with GFP_ZERO. 3) For the rest of the story we keep ignoring that we have a TUR, so the device won't touch the buffer we prepare as if the we had a DMA_FROM_DEVICE type of situation. My setup uses a virtio-scsi device and the buffer allocated by SG is mapped by the function virtqueue_add_split() which uses DMA_FROM_DEVICE for the "in" sgs (here scatter-gather and not scsi generics). This mapping involves bouncing via the swiotlb (we need swiotlb to do virtio in protected guest like s390 Secure Execution, or AMD SEV). 4) When the SCSI TUR is done, we first copy back the content of the second (that is swiotlb) bounce buffer (which most likely contains some previous IO data), to the first bounce buffer, which contains all zeros. Then we copy back the content of the first bounce buffer to the user-space buffer. 5) The test case detects that the buffer, which it zero-initialized, ain't all zeros and fails. One can argue that this is an swiotlb problem, because without swiotlb we leak all zeros, and the swiotlb should be transparent in a sense that it does not affect the outcome (if all other participants are well behaved). Copying the content of the original buffer into the swiotlb buffer is the only way I can think of to make swiotlb transparent in such scenarios. So let's do just that if in doubt, but allow the driver to tell us that the whole mapped buffer is going to be overwritten, in which case we can preserve the old behavior and avoid the performance impact of the extra bounce.	2024-07-16	5.5	Medium
CVE-2022-48856	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: gianfar: ethtool: Fix refcount leak in gfar_get_ts_info	2024-07-16	5.5	Medium

		The of_find_compatible_node() function returns a node pointer with refcount incremented, We should use of_node_put() on it when done Add the missing of_node_put() to release the refcount.			
CVE-2022-48857	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: NFC: port100: fix use-after-free in port100_send_complete Syzbot reported UAF in port100_send_complete(). The root case is in missing usb_kill_urb() calls on error handling path of ->probe function. port100_send_complete() accesses devm allocated memory which will be freed on probe failure. We should kill this urbs before returning an error from probe function to prevent reported use-after-free Fail log: BUG: KASAN: use-after-free in port100_send_complete+0x16e/0x1a0 drivers/nfc/port100.c:935 Read of size 1 at addr ffff88801bb59540 by task ksoftirqd/2/26 ... Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106 print_address_description.constprop.0.cold+0x8d/0x303 mm/kasan/report.c:255 __kasan_report mm/kasan/report.c:442 [inline] kasan_report.cold+0x83/0xdf mm/kasan/report.c:459 port100_send_complete+0x16e/0x1a0 drivers/nfc/port100.c:935 __usb_hcd_giveback_urb+0x2b0/0x5c0 drivers/usb/core/hcd.c:1670 ... Allocated by task 1255: kasan_save_stack+0x1e/0x40 mm/kasan/common.c:38 kasan_set_track mm/kasan/common.c:45 [inline] set_alloc_info mm/kasan/common.c:436 [inline] ____kasan_kmalloc mm/kasan/common.c:515 [inline] ____kasan_kmalloc mm/kasan/common.c:474 [inline] __kasan_kmalloc+0xa6/0xd0 mm/kasan/common.c:524 alloc_dr drivers/base/devres.c:116 [inline] devm_kmalloc+0x96/0x1d0 drivers/base/devres.c:823 devm_kzalloc include/linux/device.h:209 [inline] port100_probe+0x8a/0x1320 drivers/nfc/port100.c:1502 Freed by task 1255: kasan_save_stack+0x1e/0x40 mm/kasan/common.c:38 kasan_set_track+0x21/0x30 mm/kasan/common.c:45 kasan_set_free_info+0x20/0x30 mm/kasan/generic.c:370 ____kasan_slab_free mm/kasan/common.c:366 [inline] ____kasan_slab_free+0xff/0x140 mm/kasan/common.c:328 kasan_slab_free include/linux/kasan.h:236 [inline] __cache_free mm/slab.c:3437 [inline] kfree+0xf8/0x2b0 mm/slab.c:3794 release_nodes+0x112/0x1a0 drivers/base/devres.c:501 devres_release_all+0x114/0x190 drivers/base/devres.c:530 really_probe+0x626/0xcc0 drivers/base/dd.c:670	2024-07-16	5.5	Medium
CVE-2022-48859	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: net: marvell: prestera: Add missing of_node_put() in prestera_switch_set_base_mac_addr This node pointer is returned by of_find_compatible_node() with refcount incremented. Calling of_node_put() to aovid the refcount leak.	2024-07-16	5.5	Medium
CVE-2022-48860	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ethernet: Fix error handling in xemaclite_of_probe This node pointer is returned by of_parse_phandle() with refcount incremented in this function. Calling of_node_put() to avoid the refcount leak. As the remove function do.	2024-07-16	5.5	Medium

CVE-2022-48861	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: vdpa: fix use-after-free on vp_vdpa_remove When vp_vdpa driver is unbind, vp_vdpa is freed in vdpa_unregister_device and then vp_vdpa->mdev.pci_dev is dereferenced in vp_modern_remove, triggering use-after-free. Call Trace of unbinding driver free vp_vdpa : do_syscall_64 vfs_write kernfs_fop_write_iter device_release_driver_internal pci_device_remove vp_vdpa_remove vdpa_unregister_device kobject_release device_release kfree Call Trace of dereference vp_vdpa->mdev.pci_dev: vp_modern_remove pci_release_selected_regions pci_release_region pci_resource_len pci_resource_end (dev)->resource[(bar)].end	2024-07-16	5.5	Medium
CVE-2022-48862	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: vhost: fix hung thread due to erroneous iotlb entries In vhost_iotlb_add_range_ctx(), range size can overflow to 0 when start is 0 and last is ULONG_MAX. One instance where it can happen is when userspace sends an IOTLB message with iova=size=uaddr=0 (vhost_process_iotlb_msg). So, an entry with size = 0, start = 0, last = ULONG_MAX ends up in the iotlb. Next time a packet is sent, iotlb_access_ok() loops indefinitely due to that erroneous entry. Call Trace: <TASK> iotlb_access_ok+0x21b/0x3e0 drivers/vhost/vhost.c:1340 vq_meta_prefetch+0xbc/0x280 drivers/vhost/vhost.c:1366 vhost_transport_do_send_pkt+0xe0/0xfd0 drivers/vhost/vsock.c:104 vhost_worker+0x23d/0x3d0 drivers/vhost/vhost.c:372 kthread+0x2e9/0x3a0 kernel/kthread.c:377 ret_from_fork+0x1f/0x30 arch/x86/entry/entry_64.S:295 </TASK> Reported by syzbot at: https://syzkaller.appspot.com/bug?extid=0abd373e2e50d704db87 To fix this, do two things: 1. Return -EINVAL in vhost_chr_write_iter() when userspace asks to map a range with size 0. 2. Fix vhost_iotlb_add_range_ctx() to handle the range [0, ULONG_MAX] by splitting it into two entries.	2024-07-16	5.5	Medium
CVE-2022-48863	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mISDN: Fix memory leak in dsp_pipeline_build() dsp_pipeline_build() allocates dup pointer by kstrdup(cfg), but then it updates dup variable by strsep(&dup, "|"). As a result when it calls kfree(dup), the dup variable contains NULL. Found by Linux Driver Verification project (linuxtesting.org) with SVACE.	2024-07-16	5.5	Medium
CVE-2022-48864	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2024-07-16	5.5	Medium

		vdpa/mlx5: add validation for VIRTIO_NET_CTRL_MQ_VQ_PAIRS_SET command When control vq receives a VIRTIO_NET_CTRL_MQ_VQ_PAIRS_SET command request from the driver, presently there is no validation against the number of queue pairs to configure, or even if multiqueue had been negotiated or not is unverified. This may lead to kernel panic due to uninitialized resource for the queues were there any bogus request sent down by untrusted driver. Tie up the loose ends there.			

CVE-2022-48865	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		tipc: fix kernel panic when enabling bearer
		
		When enabling a bearer on a node, a kernel panic is observed:
		
		[4.498085] RIP: 0010:tipc_mon_prep+0x4e/0x130 [tipc]
		...
		[4.520030] Call Trace:
		[4.520689] <IRQ>
		[4.521236] tipc_link_build_proto_msg+0x375/0x750 [tipc]
		[4.522654] tipc_link_build_state_msg+0x48/0xc0 [tipc]
		[4.524034] __tipc_node_link_up+0xd7/0x290 [tipc]
		[4.525292] tipc_rcv+0x5da/0x730 [tipc]
		[4.526346] ? __netif_receive_skb_core+0xb7/0xfc0
		[4.527601] tipc_l2_rcv_msg+0x5e/0x90 [tipc]
		[4.528737] __netif_receive_skb_list_core+0x20b/0x260
		[4.530068] netif_receive_skb_list_internal+0x1bf/0x2e0
		[4.531450] ? dev_gro_receive+0x4c2/0x680
		[4.532512] napi_complete_done+0x6f/0x180
		[4.533570] virtnet_poll+0x29c/0x42e [virtio_net]
		...
		
		The node in question is receiving activate messages in another thread after changing bearer status to allow message sending/ receiving in current thread:
		
		thread 1 | thread 2
		-------- | --------
		|
		tipc_enable_bearer() |
		test_and_set_bit_lock() |
		tipc_bearer_xmit_skb() |
		| tipc_l2_rcv_msg()
		| tipc_rcv()
		| __tipc_node_link_up()
		| tipc_link_build_state_msg()
		| tipc_link_build_proto_msg()
		| tipc_mon_prep()
		| {
		| ...
		| // null-pointer dereference
		| u16 gen = mon->dom_gen;
		| ...
		| }
		// Not being executed yet |
		tipc_mon_create() |
		{ |
		... |
		// allocate |
		mon = kzalloc(); |
		... |
		} |
		
		Monitoring pointer in thread 2 is dereferenced before monitoring data
		is allocated in thread 1. This causes kernel panic.
		
		This commit fixes it by allocating the monitoring data before enabling
		the bearer to receive messages.
			2024-07-16	5.5	Medium

CVE-2024-41009	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix overrunning reservations in ringbuf	2024-07-17	5.5	Medium

		The BPF ring buffer internally is implemented as a power-of-2 sized circular buffer, with two logical and ever-increasing counters: consumer_pos is the consumer counter to show which logical position the consumer consumed the data, and producer_pos which is the producer counter denoting the amount of data reserved by all producers. Each time a record is reserved, the producer that "owns" the record will successfully advance producer counter. In user space each time a record is read, the consumer of the data advanced the consumer counter once it finished processing. Both counters are stored in separate pages so that from user space, the producer counter is read-only and the consumer counter is read-write. One aspect that simplifies and thus speeds up the implementation of both producers and consumers is how the data area is mapped twice contiguously back-to-back in the virtual memory, allowing to not take any special measures for samples that have to wrap around at the end of the circular buffer data area, because the next page after the last data page would be first data page again, and thus the sample will still appear completely contiguous in virtual memory. Each record has a struct bpf_ringbuf_hdr { u32 len; u32 pg_off; } header for book-keeping the length and offset, and is inaccessible to the BPF program. Helpers like bpf_ringbuf_reserve() return `(void *)hdr + BPF_RINGBUF_HDR_SZ` for the BPF program to use. Bing-Jhong and Muhammad reported that it is however possible to make a second allocated memory chunk overlapping with the first chunk and as a result, the BPF program is now able to edit first chunk's header. For example, consider the creation of a BPF_MAP_TYPE_RINGBUF map with size of 0x4000. Next, the consumer_pos is modified to 0x3000 /before/ a call to bpf_ringbuf_reserve() is made. This will allocate a chunk A, which is in [0x0,0x3008], and the BPF program is able to edit [0x8,0x3008]. Now, lets allocate a chunk B with size 0x3000. This will succeed because consumer_pos was edited ahead of time to pass the `new_prod_pos - cons_pos > rb->mask` check. Chunk B will be in range [0x3008,0x6010], and the BPF program is able to edit [0x3010,0x6010]. Due to the ring buffer memory layout mentioned earlier, the ranges [0x0,0x4000] and [0x4000,0x8000] point to the same data pages. This means that chunk B at [0x4000,0x4008] is chunk A's header. bpf_ringbuf_submit() / bpf_ringbuf_discard() use the header's pg_off to then locate the bpf_ringbuf itself via bpf_ringbuf_restore_from_rec(). Once chunk B modified chunk A's header, then bpf_ringbuf_commit() refers to the wrong page and could cause a crash. Fix it by calculating the oldest pending_pos and check whether the range			

		from the oldest outstanding record to the newest would span beyond the ring buffer size. If that is the case, then reject the request. We've tested with the ring buffer benchmark in BPF selftests (./benchs/run_bench_ringbufs.sh) before/after the fix and while it seems a bit slower on some benchmarks, it is still not significantly enough to matter.			

CVE-2024-41010	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:
		
		bpf: Fix too early release of tcx_entry
		
		Pedro Pinto and later independently also Hyunwoo Kim and Wongi Lee reported
		an issue that the tcx_entry can be released too early leading to a use
		after free (UAF) when an active old-style ingress or clsact qdisc with a
		shared tc block is later replaced by another ingress or clsact instance.
		
		Essentially, the sequence to trigger the UAF (one example) can be as follows:
		
		1. A network namespace is created
		2. An ingress qdisc is created. This allocates a tcx_entry, and
		&tcx_entry->miniq is stored in the qdisc's miniqp->p_miniq. At the
		same time, a tcf block with index 1 is created.
		3. chain0 is attached to the tcf block. chain0 must be connected to
		the block linked to the ingress qdisc to later reach the function
		tcf_chain0_head_change_cb_del() which triggers the UAF.
		4. Create and graft a clsact qdisc. This causes the ingress qdisc
		created in step 1 to be removed, thus freeing the previously linked
		tcx_entry:
		
		rtnetlink_rcv_msg()
		=> tc_modify_qdisc()
		=> qdisc_create()
		=> clsact_init() [a]
		=> qdisc_graft()
		=> qdisc_destroy()
		=> __qdisc_destroy()
		=> ingress_destroy() [b]
		=> tcx_entry_free()
		=> kfree_rcu() // tcx_entry freed
		
		5. Finally, the network namespace is closed. This registers the
		cleanup_net worker, and during the process of releasing the
		remaining clsact qdisc, it accesses the tcx_entry that was
		already freed in step 4, causing the UAF to occur:
		
		cleanup_net()
		=> ops_exit_list()
		=> default_device_exit_batch()
		=> unregister_netdevice_many()
		=> unregister_netdevice_many_notify()
		=> dev_shutdown()
		=> qdisc_put()
		=> clsact_destroy() [c]
		=> tcf_block_put_ext()
		=> tcf_chain0_head_change_cb_del()
		=> tcf_chain_head_change_item()
		=> clsact_chain_head_change()
		=> mini_qdisc_pair_swap() // UAF
		
		There are also other variants, the gist is to add an ingress (or clsact)
		qdisc with a specific shared block, then to replace that qdisc, waiting
		for the tcx_entry kfree_rcu() to be executed and subsequently accessing
		the current active qdisc's miniq one way or another.
		
		The correct fix is to turn the miniq_active boolean into a counter. What
			2024-07-17	5.5	Medium

		can be observed, at step 2 above, the counter transitions from 0->1, at step [a] from 1->2 (in order for the miniq object to remain active during the replacement), then in [b] from 2->1 and finally [c] 1->0 with the eventual release. The reference counter in general ranges from [0,2] and it does not need to be atomic since all access to the counter is protected by the rtnl mutex. With this in place, there is no longer a UAF happening and the tcx_entry is freed at the correct time.			
CVE-2024-39728	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 is vulnerable to stored cross-site scripting. This vulnerability allows users to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session. IBM X-Force ID: 295967.	2024-07-15	5.4	Medium
CVE-2024-39735	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 is vulnerable to cross-site scripting. This vulnerability allows an authenticated user to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials disclosure within a trusted session. IBM X-Force ID: 296002.	2024-07-15	5.4	Medium
CVE-2024-39863	apache - airflow	Apache Airflow versions before 2.9.3 have a vulnerability that allows an authenticated attacker to inject a malicious link when installing a provider. Users are recommended to upgrade to version 2.9.3, which fixes this issue.	2024-07-17	5.4	Medium
CVE-2024-39737	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 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. IBM X-Force ID: 296004.	2024-07-15	5.3	Medium
CVE-2024-39740	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 displays version information in HTTP requests that could allow an attacker to gather information for future attacks against the system. IBM X-Force ID: 296009.	2024-07-15	5.3	Medium
CVE-2024-39741	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 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. IBM X-Force ID: 296010.	2024-07-15	5.3	Medium
CVE-2024-21176	oracle - mysql_server	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Thread Pooling). Supported versions that are affected are 8.4.0 and prior. Difficult to exploit vulnerability allows low privileged attacker with network access via multiple protocols to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of MySQL Server. CVSS 3.1 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:L/UI:N/S:U/C:N/I:N/A:H).	2024-07-16	5.3	Medium
CVE-2024-6535	redhat - service_interconnect	A flaw was found in Skupper. When Skupper is initialized with the console-enabled and with console-auth set to Openshift, it configures the openshift oauth-proxy with a static cookie-secret. In certain circumstances, this may allow an attacker to bypass authentication to the Skupper console via a specially-crafted cookie.	2024-07-17	5.3	Medium
CVE-2024-21179	oracle - multiple products	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are affected are 8.0.37 and prior and 8.4.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multiple protocols to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of MySQL Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/A:H).	2024-07-16	4.9	Medium
CVE-2024-21185	oracle - multiple products	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are affected are 8.0.38, 8.4.1 and 9.0.0. Easily exploitable vulnerability allows high privileged attacker with network access via multiple protocols to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (complete DOS) of MySQL Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/A:H).	2024-07-16	4.9	Medium
CVE-2022-48842	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved:	2024-07-16	4.7	Medium

ice: Fix race condition during interface enslave

Commit 5dbbbd01cbba83 ("ice: Avoid RTNL lock when re-creating auxiliary device") changes a process of re-creation of aux device so ice_plug_aux_dev() is called from ice_service_task() context. This unfortunately opens a race window that can result in dead-lock
when interface has left LAG and immediately enters LAG again.

Reproducer:
```
#!/bin/sh

ip link add lag0 type bond mode 1 miimon 100
ip link set lag0

for n in {1..10}; do
	echo Cycle: $n
	ip link set ens7f0 master lag0
	sleep 1
	ip link set ens7f0 nomaster
done
```

This results in:
[20976.208697] Workqueue: ice ice_service_task [ice]
[20976.213422] Call Trace:
[20976.215871] __schedule+0x2d1/0x830
[20976.219364] schedule+0x35/0xa0
[20976.222510] schedule_preempt_disabled+0xa/0x10
[20976.227043] __mutex_lock.isra.7+0x310/0x420
[20976.235071] enum_all_gids_of_dev_cb+0x1c/0x100 [ib_core]
[20976.251215] ib_enum_roce_netdev+0xa4/0xe0 [ib_core]
[20976.256192] ib_cache_setup_one+0x33/0xa0 [ib_core]
[20976.261079] ib_register_device+0x40d/0x580 [ib_core]
[20976.266139] irdma_ib_register_device+0x129/0x250 [irdma]
[20976.281409] irdma_probe+0x2c1/0x360 [irdma]
[20976.285691] auxiliary_bus_probe+0x45/0x70
[20976.289790] really_probe+0x1f2/0x480
[20976.298509] driver_probe_device+0x49/0xc0
[20976.302609] bus_for_each_drv+0x79/0xc0
[20976.306448] __device_attach+0xdc/0x160
[20976.310286] bus_probe_device+0x9d/0xb0
[20976.314128] device_add+0x43c/0x890
[20976.321287] __auxiliary_device_add+0x43/0x60
[20976.325644] ice_plug_aux_dev+0xb2/0x100 [ice]
[20976.330109] ice_service_task+0xd0c/0xed0 [ice]
[20976.342591] process_one_work+0x1a7/0x360
[20976.350536] worker_thread+0x30/0x390
[20976.358128] kthread+0x10a/0x120
[20976.365547] ret_from_fork+0x1f/0x40
...
[20976.438030] task:ip   state:D stack: 0 pid:213658 ppid:213627 flags:0x00004084
[20976.446469] Call Trace:
[20976.448921] __schedule+0x2d1/0x830
[20976.452414] schedule+0x35/0xa0
[20976.455559] schedule_preempt_disabled+0xa/0x10
[20976.460090] __mutex_lock.isra.7+0x310/0x420
[20976.464364] device_del+0x36/0x3c0
[20976.467772] ice_unplug_aux_dev+0x1a/0x40 [ice]
[20976.472313] ice_lag_event_handler+0x2a2/0x520 [ice]
[20976.477288] notifier_call_chain+0x47/0x70
[20976.481386] __netdev_upper_dev_link+0x18b/0x280
[20976.489845] bond_enslave+0xe05/0x1790 [bonding]
[20976.494475] do_setlink+0x336/0xf50
[20976.502517] __rtnl_newlink+0x529/0x8b0
[20976.543441] rtnl_newlink+0x43/0x60
[20976.546934] rtnetlink_rcv_msg+0x2b1/0x360
[20976.559238] netlink_rcv_skb+0x4c/0x120
[20976.563079] netlink_unicast+0x196/0x230
[20976.567005] netlink_sendmsg+0x204/0x3d0
[20976.570930] sock_sendmsg+0x4c/0x50
[20976.574423] ____sys_sendmsg+0x1eb/0x250
[20976.586807] ___sys_sendmsg+0x7c/0xc0
[20976.606353] __sys_sendmsg+0x57/0xa0
[20976.609930] do_syscall_64+0x5b/0x1a0
[20976.613598] entry_SYSCALL_64_after_hwframe+0x65/0xca

		1. Command 'ip link ... set nomaster' causes that ice_plug_aux_dev() is called from ice_service_task() context, aux device is created and associated device->lock is taken. 2. Command 'ip link ... set master...' calls ice's notifier under RTNL lock and that notifier calls ice_unplug_aux_dev(). That function tries to take aux device->lock but this is already taken by ice_plug_aux_dev() in step 1 3. Later ice_plug_aux_dev() tries to take RTNL lock but this is already taken in step 2 4. Dead-lock The patch fixes this issue by following changes: - Bit ICE_FLAG_PLUG_AUX_DEV is kept to be set during ice_plug_aux_dev() call in ice_service_task() - The bit is checked in ice_clear_rdma_cap() and only if it is not set then ice_unplug_aux_dev() is called. If it is set (in other words plugging of aux device was requested and ice_plug_aux_dev() is potentially running) then the function only clears the ---truncated---			
CVE-2023-52291	apache - streampark	In streampark, the project module integrates Maven's compilation capabilities. The input parameter validation is not strict, allowing attackers to insert commands for remote command execution, The prerequisite for a successful attack is that the user needs to log in to the streampark system and have system-level permissions. Generally, only users of that system have the authorization to log in, and users would not manually input a dangerous operation command. Therefore, the risk level of this vulnerability is very low. Background: In the "Project" module, the maven build args “<” operator causes command injection. e.g : “< (curl http://xxx.com)” will be executed as a command injection, Mitigation: all users should upgrade to 2.1.4, The "<" operator will blocked。	2024-07-17	4.7	Medium
CVE-2024-29737	apache - streampark	In streampark, the project module integrates Maven's compilation capabilities. The input parameter validation is not strict, allowing attackers to insert commands for remote command execution, The prerequisite for a successful attack is that the user needs to log in to the streampark system and have system-level permissions. Generally, only users of that system have the authorization to log in, and users would not manually input a dangerous operation command. Therefore, the risk level of this vulnerability is very low. Mitigation: all users should upgrade to 2.1.4 Background info: Log in to Streampark using the default username (e.g. test1, test2, test3) and the default password (streampark). Navigate to the Project module, then add a new project. Enter the git repository address of the project and input `touch /tmp/success_2.1.2` as the "Build Argument". Note that there is no verification and interception of the special character "`". As a result, you will find that this injection command will be successfully executed after executing the build.  In the latest version, the special symbol ` is intercepted.	2024-07-17	4.7	Medium
CVE-2024-39739	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 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. IBM X-Force ID: 296008.	2024-07-15	4.3	Medium
CVE-2024-39729	ibm - multiple products	IBM Datacap Navigator 9.1.5, 9.1.6, 9.1.7, 9.1.8, and 9.1.9 could allow an authenticated user to obtain sensitive information from source code that could be used in further attacks against the system. IBM X-Force ID: 295968.	2024-07-15	4.3	Medium
CVE-2024-31979	apache - streampipes	Server-Side Request Forgery (SSRF) vulnerability in Apache StreamPipes during installation process of pipeline elements.	2024-07-17	4.3	Medium

		Previously, StreamPipes allowed users to configure custom endpoints from which to install additional pipeline elements. These endpoints were not properly validated, allowing an attacker to get StreamPipes to send an HTTP GET request to an arbitrary address. This issue affects Apache StreamPipes: through 0.93.0. Users are recommended to upgrade to version 0.95.0, which fixes the issue.			
CVE-2024-21180	oracle - multiple products	Vulnerability in the PeopleSoft Enterprise PeopleTools product of Oracle PeopleSoft (component: OpenSearch Dashboards). Supported versions that are affected are 8.59, 8.60 and 8.61. Easily exploitable vulnerability allows low privileged attacker with network access via HTTP to compromise PeopleSoft Enterprise PeopleTools. Successful attacks require human interaction from a person other than the attacker and while the vulnerability is in PeopleSoft Enterprise PeopleTools, attacks may significantly impact additional products (scope change). Successful attacks of this vulnerability can result in unauthorized read access to a subset of PeopleSoft Enterprise PeopleTools accessible data. CVSS 3.1 Base Score 4.1 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:R/S:C/C:L/I:N/A:N).	2024-07-16	4.1	Medium
CVE-2024-30471	apache - streampipes	Time-of-check Time-of-use (TOCTOU) Race Condition vulnerability in Apache StreamPipes in user self-registration. This allows an attacker to potentially request the creation of multiple accounts with the same email address until the email address is registered, creating many identical users and corrupting StreamPipe's user management. This issue affects Apache StreamPipes: through 0.93.0. Users are recommended to upgrade to version 0.95.0, which fixes the issue.	2024-07-17	3.7	Low
CVE-2024-41007	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: tcp: avoid too many retransmit packets If a TCP socket is using TCP_USER_TIMEOUT, and the other peer retracted its window to zero, tcp_retransmit_timer() can retransmit a packet every two jiffies (2 ms for HZ=1000), for about 4 minutes after TCP_USER_TIMEOUT has 'expired'. The fix is to make sure tcp_rtx_probe0_timed_out() takes icsk->icsk_user_timeout into account. Before blamed commit, the socket would not timeout after icsk->icsk_user_timeout, but would use standard exponential backoff for the retransmits. Also worth noting that before commit e89688e3e978 ("net: tcp: fix unexcepted socket die when snd_wnd is 0"), the issue would last 2 minutes instead of 4.	2024-07-15	3.3	Low
CVE-2022-48852	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: hdmi: Unregister codec device on unbind On bind we will register the HDMI codec device but we don't unregister it on unbind, leading to a device leakage. Unregister our device at unbind.	2024-07-16	3.3	Low

Where NCA provides the vulnerability information as published by NIST's NVD. In addition, it is the entity's or individual's responsibility to ensure the implementation of appropriate recommendations.

وحيث تقدم الهيئة تفاصيل الثغرات كما تم نشرها من قبل NIST's NVD. وإذ تبقى مسؤولية الجهة أو الشخص قائمة للتأكد من تطبيق التوصيات المناسبة.