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 ٣٠ يونيو إلى ٦ 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 30th of July to 6th of June. 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-2024-6376	mongodb - compass	MongoDB Compass may be susceptible to code injection due to insufficient sandbox protection settings with the usage of ejson shell parser in Compass' connection handling. This issue affects MongoDB Compass versions prior to version 1.42.2	2024-07-01	9.8	Critical
CVE-2024-38346	apache - multiple products	The CloudStack cluster service runs on unauthenticated port (default 9090) that can be misused to run arbitrary commands on targeted hypervisors and CloudStack management server hosts. Some of these commands were found to have command injection vulnerabilities that can result in arbitrary code execution via agents on the hosts that may run as a privileged user. An attacker that can reach the cluster service on the unauthenticated port (default 9090), can exploit this to perform remote code execution on CloudStack managed hosts and result in complete compromise of the confidentiality, integrity, and availability of CloudStack managed infrastructure. Users are recommended to restrict the network access to the cluster service port (default 9090) on a CloudStack management server host to only its peer CloudStack management server hosts. Users are recommended to upgrade to version 4.18.2.1, 4.19.0.2 or later, which addresses this issue.	2024-07-05	9.8	Critical
CVE-2024-39864	apache - multiple products	The CloudStack integration API service allows running its unauthenticated API server (usually on port 8096 when configured and enabled via integration.api.port global setting) for internal portal integrations and for testing purposes. By default, the integration API service port is disabled and is considered disabled when integration.api.port is set to 0 or negative. Due to an improper initialisation logic, the integration API service would listen on a random port when its port value is set to 0 (default value). An attacker that can access the CloudStack management network could scan and find the randomised integration API service port and exploit it to perform unauthorised administrative actions and perform remote code execution on CloudStack managed hosts and result in complete compromise of the confidentiality, integrity, and availability of CloudStack managed infrastructure. Users are recommended to restrict the network access on the CloudStack management server hosts to only essential ports. Users are recommended to upgrade to version 4.18.2.1, 4.19.0.2 or later, which addresses this issue.	2024-07-05	9.8	Critical
CVE-2024-20890	samsung - multiple products	Improper input validation in BLE prior to SMR Jul-2024 Release 1 allows adjacent attackers to trigger abnormal behavior.	2024-07-02	8.8	High
CVE-2024-34593	samsung - multiple products	Improper input validation in parsing and distributing RTCP packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote	2024-07-02	8.8	High

		attackers to execute arbitrary code with system privilege. User interaction is required for triggering this vulnerability.			
CVE-2024-6387	openbsd - multiple products	A security regression (CVE-2006-5051) was discovered in OpenSSH's server (sshd). There is a race condition which can lead to sshd to handle some signals in an unsafe manner. An unauthenticated, remote attacker may be able to trigger it by failing to authenticate within a set time period.	2024-07-01	8.1	High
CVE-2024-21461	qualcomm - 315_5g_iot_modem_firmware	Memory corruption while performing finish HMAC operation when context is freed by keymaster.	2024-07-01	7.8	High
CVE-2024-21465	qualcomm - 9205_lte_modem_firmware	Memory corruption while processing key blob passed by the user.	2024-07-01	7.8	High
CVE-2024-21469	qualcomm - 9205_lte_modem_firmware	Memory corruption when an invoke call and a TEE call are bound for the same trusted application.	2024-07-01	7.8	High
CVE-2024-21482	qualcomm - csr8811_firmware	Memory corruption during the secure boot process, when the `bootm` command is used, it bypasses the authentication of the kernel/rootfs image.	2024-07-01	7.8	High
CVE-2024-23368	qualcomm - apq8064au_firmware	Memory corruption when allocating and accessing an entry in an SMEM partition.	2024-07-01	7.8	High
CVE-2024-23372	qualcomm - fastconnect_6200_firmware	Memory corruption while invoking IOCTL call for GPU memory allocation and size param is greater than expected size.	2024-07-01	7.8	High
CVE-2024-23373	qualcomm - 315_5g_iot_modem_firmware	Memory corruption when IOMMU unmap operation fails, the DMA and anon buffers are getting released.	2024-07-01	7.8	High
CVE-2024-23380	qualcomm - fastconnect_6200_firmware	Memory corruption while handling user packets during VBO bind operation.	2024-07-01	7.8	High
CVE-2024-32853	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.7.0.2 contain an execution with unnecessary privileges vulnerability. A local low privileged attacker could potentially exploit this vulnerability, leading to escalation of privileges.	2024-07-02	7.8	High
CVE-2024-20888	samsung - multiple products	Improper access control in OneUIHome prior to SMR Jul-2024 Release 1 allows local attackers to launch privileged activities. User interaction is required for triggering this vulnerability.	2024-07-02	7.8	High
CVE-2024-20891	samsung - multiple products	Improper access control in launchFullscreenIntent of SystemUI prior to SMR Jul-2024 Release 1 allows local attackers to launch privileged activities.	2024-07-02	7.8	High
CVE-2024-20892	samsung - multiple products	Improper verification of signature in FilterProvider prior to SMR Jul-2024 Release 1 allows local attackers to execute privileged behaviors. User interaction is required for triggering this vulnerability.	2024-07-02	7.8	High
CVE-2024-20893	samsung - multiple products	Improper input validation in libmediaextractorservice.so prior to SMR Jul-2024 Release 1 allows local attackers to trigger memory corruption.	2024-07-02	7.8	High
CVE-2024-20901	samsung - multiple products	Improper input validation in copying data to buffer cache in libsaped prior to SMR Jul-2024 Release 1 allows local attackers to write out-of-bounds memory.	2024-07-02	7.8	High
CVE-2024-34585	samsung - multiple products	Improper access control in launchApp of SystemUI prior to SMR Jul-2024 Release 1 allows local attackers to launch privileged activities.	2024-07-02	7.8	High
CVE-2024-34595	samsung - multiple products	Improper access control in clickAdapterItem of SystemUI prior to SMR Jul-2024 Release 1 allows local attackers to launch privileged activities.	2024-07-02	7.8	High
CVE-2024-39479	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: drm/i915/hwmon: Get rid of devm When both hwmon and hwmon drvdata (on which hwmon depends) are device managed resources, the expectation, on device unbind, is that hwmon will be released before drvdata. However, in i915 there are two separate code paths, which both release either drvdata or hwmon and either can be released before the other. These code paths (for device unbind) are as follows (see also the bug referenced below): Call Trace: release_nodes+0x11/0x70 devres_release_group+0xb2/0x110 component_unbind_all+0x8d/0xa0 component_del+0xa5/0x140 intel_pxp_tee_component_fini+0x29/0x40 [i915]	2024-07-05	7.8	High

		intel_pxp_fini+0x33/0x80 [i915] i915_driver_remove+0x4c/0x120 [i915] i915_pci_remove+0x19/0x30 [i915] pci_device_remove+0x32/0xa0 device_release_driver_internal+0x19c/0x200 unbind_store+0x9c/0xb0 and Call Trace: release_nodes+0x11/0x70 devres_release_all+0x8a/0xc0 device_unbind_cleanup+0x9/0x70 device_release_driver_internal+0x1c1/0x200 unbind_store+0x9c/0xb0 This means that in i915, if use devm, we cannot gurantee that hwmon will always be released before drvdata. Which means that we have a uaf if hwmon sysfs is accessed when drvdata has been released but hwmon hasn't. The only way out of this seems to be do get rid of devm_ and release/free everything explicitly during device unbind. v2: Change commit message and other minor code changes v3: Cleanup from i915_hwmon_register on error (Armin Wolf) v4: Eliminate potential static analyzer warning (Rodrigo) Eliminate fetch_and_zero (Jani) v5: Restore previous logic for ddat_gt->hwmon_dev error return (Andi)			
CVE-2024-39480	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: kdb: Fix buffer overflow during tab-complete Currently, when the user attempts symbol completion with the Tab key, kdb will use strncpy() to insert the completed symbol into the command buffer. Unfortunately it passes the size of the source buffer rather than the destination to strncpy() with predictably horrible results. Most obviously if the command buffer is already full but cp, the cursor position, is in the middle of the buffer, then we will write past the end of the supplied buffer. Fix this by replacing the dubious strncpy() calls with memmove()/memcpy() calls plus explicit boundary checks to make sure we have enough space before we start moving characters around.	2024-07-05	7.8	High
CVE-2024-21457	qualcomm - ar8035_firmware	INformation disclosure while handling Multi-link IE in beacon frame.	2024-07-01	7.5	High
CVE-2024-21458	qualcomm - ar8035_firmware	Information disclosure while handling SA query action frame.	2024-07-01	7.5	High
CVE-2024-21466	qualcomm - fastconnect_7800_firmware	Information disclosure while parsing sub-IE length during new IE generation.	2024-07-01	7.5	High
CVE-2024-32852	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.7.0.0 contain use of a broken or risky cryptographic algorithm vulnerability. An unprivileged network malicious attacker could potentially exploit this vulnerability, leading to data leaks.	2024-07-02	7.5	High
CVE-2024-34596	samsung - smartthings	Improper authentication in SmartThings prior to version 1.8.17 allows remote attackers to bypass the expiration date for members set by the owner.	2024-07-02	7.5	High
CVE-2023-52340	linux - linux_kernel	The IPv6 implementation in the Linux kernel before 6.3 has a net/ipv6/route.c max_size threshold that can be consumed easily, e.g., leading to a denial of service (network is unreachable errors) when IPv6 packets are sent in a loop via a raw socket.	2024-07-05	7.5	High
CVE-2024-34587	samsung - multiple products	Improper input validation in parsing application information from RTCP packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to execute arbitrary code with system privilege. User interaction is required for triggering this vulnerability.	2024-07-02	6.8	Medium

CVE-2024-20399	cisco - multiple products	A vulnerability in the CLI of Cisco NX-OS Software could allow an authenticated, local attacker to execute arbitrary commands as root on the underlying operating system of an affected device. This vulnerability is due to insufficient validation of arguments that are passed to specific configuration CLI commands. An attacker could exploit this vulnerability by including crafted input as the argument of an affected configuration CLI command. A successful exploit could allow the attacker to execute arbitrary commands on the underlying operating system with the privileges of root. Note: To successfully exploit this vulnerability on a Cisco NX-OS device, an attacker must have Administrator credentials.	2024-07-01	6.7	Medium
CVE-2024-32854	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.8.0.0 contain an improper privilege management vulnerability. A local high privilege attacker could potentially exploit this vulnerability, leading to privilege escalation.	2024-07-02	6.7	Medium
CVE-2024-37126	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.8.0.0 contain an improper privilege management vulnerability. A local high privileged attacker could potentially exploit this vulnerability, leading to unauthorized gain of root-level access.	2024-07-02	6.7	Medium
CVE-2024-37132	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.8.0.0 contain an incorrect privilege assignment vulnerability. A high privileged attacker with local access could potentially exploit this vulnerability, leading to Denial of service and Elevation of privileges.	2024-07-02	6.7	Medium
CVE-2024-37133	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.8.0.0 contain an improper privilege management vulnerability. A local high privileged attacker could potentially exploit this vulnerability, leading to unauthorized gain of root-level access.	2024-07-02	6.7	Medium
CVE-2024-37134	dell - multiple products	Dell PowerScale OneFS versions 8.2.2.x through 9.8.0.0 contain an improper privilege management vulnerability. A local high privileged attacker could potentially exploit this vulnerability to gain root-level access.	2024-07-02	6.7	Medium
CVE-2024-21460	qualcomm - fastconnect_6900_firmware	Information disclosure when ASLR relocates the IMEM and Secure DDR portions as one chunk in virtual address space.	2024-07-01	6.5	Medium
CVE-2024-6375	mongodb - multiple products	A command for refining a collection shard key is missing an authorization check. This may cause the command to run directly on a shard, leading to either degradation of query performance, or to revealing chunk boundaries through timing side channels. This affects MongoDB Server v5.0 versions, prior to 5.0.22, MongoDB Server v6.0 versions, prior to 6.0.11 and MongoDB Server v7.0 versions prior to 7.0.3.	2024-07-01	6.5	Medium
CVE-2024-34588	samsung - multiple products	Improper input validation?in parsing RTCP SR packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to trigger temporary denial of service. User interaction is required for triggering this vulnerability.	2024-07-02	6.5	Medium
CVE-2024-34589	samsung - multiple products	Improper input validation in parsing RTCP RR packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to trigger temporary denial of service. User interaction is required for triggering this vulnerability.	2024-07-02	6.5	Medium
CVE-2024-21462	qualcomm - 315_5g_iot_modem_firmware	Transient DOS while loading the TA ELF file.	2024-07-01	5.5	Medium
CVE-2024-20895	samsung - multiple products	Improper access control in Dar service prior to SMR Jul-2024 Release 1 allows local attackers to bypass restriction for calling SDP features.	2024-07-02	5.5	Medium
CVE-2024-20896	samsung - multiple products	Use of implicit intent for sensitive communication in Configuration message prior to SMR Jul-2024 Release 1 allows local attackers to get sensitive information.	2024-07-02	5.5	Medium
CVE-2024-20897	samsung - multiple products	Use of implicit intent for sensitive communication in FCM function in IMS service prior to SMR Jul-2024 Release 1 allows local attackers to get sensitive information.	2024-07-02	5.5	Medium
CVE-2024-20898	samsung - multiple products	Use of implicit intent for sensitive communication in SoftphoneClient in IMS service prior to SMR Jul-2024 Release 1 allows local attackers to get sensitive information.	2024-07-02	5.5	Medium
CVE-2024-20899	samsung - multiple products	Use of implicit intent for sensitive communication in RCS function in IMS service prior to SMR Jul-2024 Release 1 allows local attackers to get sensitive information.	2024-07-02	5.5	Medium
CVE-2024-34594	samsung - multiple products	Exposure of sensitive information in proc file system prior to SMR Jul-2024 Release 1 allows local attackers to read kernel memory address.	2024-07-02	5.5	Medium
CVE-2024-39472	linux - linux_kernel	In the Linux kernel, the following vulnerability has been resolved: xfs: fix log recovery buffer allocation for the legacy h_size fixup	2024-07-05	5.5	Medium

		Commit a70f9fe52daa ("xfs: detect and handle invalid iclog size set by mkfs") added a fixup for incorrect h_size values used for the initial umount record in old xfsprogs versions. Later commit 0c771b99d6c9 ("xfs: clean up calculation of LR header blocks") cleaned up the log reover buffer calculation, but stoped using the fixed up h_size value to size the log recovery buffer, which can lead to an out of bounds access when the incorrect h_size does not come from the old mkfs tool, but a fuzzer. Fix this by open coding xlog_logrec_hblks and taking the fixed h_size into account for this calculation.			
CVE-2024-39473	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: ASoC: SOF: ipc4-topology: Fix input format query of process modules without base extension If a process module does not have base config extension then the same format applies to all of it's inputs and the process->base_config_ext is NULL, causing NULL dereference when specifically crafted topology and sequences used.	2024-07-05	5.5	Medium
CVE-2024-39474	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mm/vmalloc: fix vmalloc which may return null if called with __GFP_NOFAIL commit a421ef303008 ("mm: allow !GFP_KERNEL allocations for kvmalloc") includes support for __GFP_NOFAIL, but it presents a conflict with commit dd544141b9eb ("vmalloc: back off when the current task is OOM-killed"). A possible scenario is as follows: process-a __vmalloc_node_range(GFP_KERNEL | __GFP_NOFAIL) __vmalloc_area_node() vm_area_alloc_pages() --> oom-killer send SIGKILL to process-a if (fatal_signal_pending(current)) break; --> return NULL; To fix this, do not check fatal_signal_pending() in vm_area_alloc_pages() if __GFP_NOFAIL set. This issue occurred during OPLUS KASAN TEST. Below is part of the log -> oom-killer sends signal to process [65731.222840] [T1308] oom-kill:constraint=CONSTRAINT_NONE,nodemask=(null),cpuset=/,mems_allowed=0,global_oom,task_memcg=/apps/uid_10198,task=gs.intelligence,pid=32454,uid=10198 [65731.259685] [T32454] Call trace: [65731.259698] [T32454] dump_backtrace+0xf4/0x118 [65731.259734] [T32454] show_stack+0x18/0x24 [65731.259756] [T32454] dump_stack_lvl+0x60/0x7c [65731.259781] [T32454] dump_stack+0x18/0x38 [65731.259800] [T32454] mrdump_common_die+0x250/0x39c [mrdump] [65731.259936] [T32454] ipanic_die+0x20/0x34 [mrdump] [65731.260019] [T32454] atomic_notifier_call_chain+0xb4/0xfc [65731.260047] [T32454] notify_die+0x114/0x198 [65731.260073] [T32454] die+0xf4/0x5b4 [65731.260098] [T32454] die_kernel_fault+0x80/0x98 [65731.260124] [T32454] __do_kernel_fault+0x160/0x2a8 [65731.260146] [T32454] do_bad_area+0x68/0x148 [65731.260174] [T32454] do_mem_abort+0x151c/0x1b34 [65731.260204] [T32454] el1_abort+0x3c/0x5c [65731.260227] [T32454] el1h_64_sync_handler+0x54/0x90 [65731.260248] [T32454] el1h_64_sync+0x68/0x6c	2024-07-05	5.5	Medium

		[65731.260269] [T32454] z_erofs_decompress_queue+0x7f0/0x2258 --> be->decompressed_pages = kvcalloc(be->nr_pages, sizeof(struct page *), GFP_KERNEL | __GFP_NOFAIL); kernel panic by NULL pointer dereference. erofs assume kvmalloc with __GFP_NOFAIL never return NULL. [65731.260293] [T32454] z_erofs_runqueue+0xf30/0x104c [65731.260314] [T32454] z_erofs_readahead+0x4f0/0x968 [65731.260339] [T32454] read_pages+0x170/0xadc [65731.260364] [T32454] page_cache_ra_unbounded+0x874/0xf30 [65731.260388] [T32454] page_cache_ra_order+0x24c/0x714 [65731.260411] [T32454] filemap_fault+0xbf0/0x1a74 [65731.260437] [T32454] __do_fault+0xd0/0x33c [65731.260462] [T32454] handle_mm_fault+0xf74/0x3fe0 [65731.260486] [T32454] do_mem_abort+0x54c/0x1b34 [65731.260509] [T32454] el0_da+0x44/0x94 [65731.260531] [T32454] el0t_64_sync_handler+0x98/0xb4 [65731.260553] [T32454] el0t_64_sync+0x198/0x19c			
CVE-2024-39475	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: fbdev: savage: Handle err return when savagefb_check_var failed The commit 04e5eac8f3ab("fbdev: savage: Error out if pixclock equals zero") checks the value of pixclock to avoid divide-by-zero error. However the function savagefb_probe doesn't handle the error return of savagefb_check_var. When pixclock is 0, it will cause divide-by-zero error.	2024-07-05	5.5	Medium
CVE-2024-39476	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: md/raid5: fix deadlock that raid5d() wait for itself to clear MD_SB_CHANGE_PENDING Xiao reported that lvm2 test lvconvert-raid-takeover.sh can hang with small possibility, the root cause is exactly the same as commit bed9e27baf52 ("Revert "md/raid5: Wait for MD_SB_CHANGE_PENDING in raid5d"") However, Dan reported another hang after that, and junxiao investigated the problem and found out that this is caused by plugged bio can't issue from raid5d(). Current implementation in raid5d() has a weird dependence: 1) md_check_recovery() from raid5d() must hold 'reconfig_mutex' to clear MD_SB_CHANGE_PENDING; 2) raid5d() handles IO in a deadloop, until all IO are issued; 3) IO from raid5d() must wait for MD_SB_CHANGE_PENDING to be cleared; This behaviour is introduce before v2.6, and for consequence, if other context hold 'reconfig_mutex', and md_check_recovery() can't update super_block, then raid5d() will waste one cpu 100% by the deadloop, until 'reconfig_mutex' is released. Refer to the implementation from raid1 and raid10, fix this problem by skipping issue IO if MD_SB_CHANGE_PENDING is still set after md_check_recovery(), daemon thread will be woken up when 'reconfig_mutex' is released. Meanwhile, the hang problem will be fixed as well.	2024-07-05	5.5	Medium
CVE-2024-39477	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mm/hugetlb: do not call vma_add_reservation upon ENOMEM sysbot reported a splat [1] on __unmap_hugepage_range(). This is because vma_needs_reservation() can return -ENOMEM if allocate_file_region_entries() fails to allocate the file_region struct for the reservation.	2024-07-05	5.5	Medium

		Check for that and do not call vma_add_reservation() if that is the case, otherwise region_abort() and region_del() will see that we do not have any file_regions. If we detect that vma_needs_reservation() returned -ENOMEM, we clear the hugetlb_restore_reserve flag as if this reservation was still consumed, so free_huge_folio() will not increment the resv count. [1] https://lore.kernel.org/linux-mm/00000000000004096100617c58d54@google.com/T/#ma5983bc1ab18a54910da83416b3f89f3c7ee43aa			
CVE-2024-39478	linux - linux_kernel	In the Linux kernel, the following vulnerability has been resolved: crypto: starfive - Do not free stack buffer RSA text data uses variable length buffer allocated in software stack. Calling kfree on it causes undefined behaviour in subsequent operations.	2024-07-05	5.5	Medium
CVE-2024-39481	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: media: mc: Fix graph walk in media_pipeline_start The graph walk tries to follow all links, even if they are not between pads. This causes a crash with, e.g. a MEDIA_LNK_FL_ANCILLARY_LINK link. Fix this by allowing the walk to proceed only for MEDIA_LNK_FL_DATA_LINK links.	2024-07-05	5.5	Medium
CVE-2024-39482	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: bcache: fix variable length array abuse in btree_iter btree_iter is used in two ways: either allocated on the stack with a fixed size MAX_BSETS, or from a mempool with a dynamic size based on the specific cache set. Previously, the struct had a fixed-length array of size MAX_BSETS which was indexed out-of-bounds for the dynamically-sized iterators, which causes UBSAN to complain. This patch uses the same approach as in bcachefs's sort_iter and splits the iterator into a btree_iter with a flexible array member and a btree_iter_stack which embeds a btree_iter as well as a fixed-length data array.	2024-07-05	5.5	Medium
CVE-2024-39483	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: KVM: SVM: WARN on vNMI + NMI window iff NMIs are outright masked When requesting an NMI window, WARN on vNMI support being enabled if and only if NMIs are actually masked, i.e. if the vCPU is already handling an NMI. KVM's ABI for NMIs that arrive simultanesouly (from KVM's point of view) is to inject one NMI and pend the other. When using vNMI, KVM pends the second NMI simply by setting V_NMI_PENDING, and lets the CPU do the rest (hardware automatically sets V_NMI_BLOCKING when an NMI is injected). However, if KVM can't immediately inject an NMI, e.g. because the vCPU is in an STI shadow or is running with GIF=0, then KVM will request an NMI window and trigger the WARN (but still function correctly). Whether or not the GIF=0 case makes sense is debatable, as the intent of KVM's behavior is to provide functionality that is as close to real	2024-07-05	5.5	Medium

		hardware as possible. E.g. if two NMIs are sent in quick succession, the probability of both NMIs arriving in an STI shadow is infinitesimally low on real hardware, but significantly larger in a virtual environment, e.g. if the vCPU is preempted in the STI shadow. For GIF=0, the argument isn't as clear cut, because the window where two NMIs can collide is much larger in bare metal (though still small). That said, KVM should not have divergent behavior for the GIF=0 case based on whether or not vNMI support is enabled. And KVM has allowed simultaneous NMIs with GIF=0 for over a decade, since commit 7460fb4a3400 ("KVM: Fix simultaneous NMIs"). I.e. KVM's GIF=0 handling shouldn't be modified without a *really* good reason to do so, and if KVM's behavior were to be modified, it should be done irrespective of vNMI support.			
CVE-2024-39484	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: mmc: davinci: Don't strip remove function when driver is builtin Using __exit for the remove function results in the remove callback being discarded with CONFIG_MMC_DAVINCI=y. When such a device gets unbound (e.g. using sysfs or hotplug), the driver is just removed without the cleanup being performed. This results in resource leaks. Fix it by compiling in the remove callback unconditionally. This also fixes a W=1 modpost warning: WARNING: modpost: drivers/mmc/host/davinci_mmc: section mismatch in reference: davinci_mmcsd_driver+0x10 (section: .data) -> davinci_mmcsd_remove (section: .exit.text)	2024-07-05	5.5	Medium
CVE-2024-39485	linux - multiple products	In the Linux kernel, the following vulnerability has been resolved: media: v4l: async: Properly re-initialise notifier entry in unregister The notifier_entry of a notifier is not re-initialised after unregistering the notifier. This leads to dangling pointers being left there so use list_del_init() to return the notifier_entry an empty list.	2024-07-05	5.5	Medium
CVE-2024-34601	samsung - galaxystore	Improper verification of intent by broadcast receiver vulnerability in GalaxyStore prior to version 4.5.81.0 allows local attackers to launch unexported activities of GalaxyStore.	2024-07-02	5.3	Medium
CVE-2024-20889	samsung - multiple products	Improper authentication in BLE prior to SMR Jul-2024 Release 1 allows adjacent attackers to pair with devices.	2024-07-02	4.3	Medium
CVE-2024-20894	samsung - multiple products	Improper handling of exceptional conditions in Secure Folder prior to SMR Jul-2024 Release 1 allows physical attackers to bypass authentication under certain condition. User interaction is required for triggering this vulnerability.	2024-07-02	4.3	Medium
CVE-2024-34590	samsung - multiple products	Improper input validation혺in parsing an item type from RTCP SDES packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to trigger temporary denial of service. User interaction is required for triggering this vulnerability.	2024-07-02	4.3	Medium
CVE-2024-34591	samsung - multiple products	Improper input validation in parsing an item data from RTCP SDES packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to trigger temporary denial of service. User interaction is required for triggering this vulnerability.	2024-07-02	4.3	Medium
CVE-2024-34592	samsung - multiple products	Improper input validation in parsing RTCP SDES packet in librtp.so prior to SMR Jul-2024 Release 1 allows remote attackers to trigger temporary denial of service. User interaction is required for triggering this vulnerability.	2024-07-02	4.3	Medium
CVE-2024-20900	samsung - multiple products	Improper authentication in MTP application prior to SMR Jul-2024 Release 1 allows local attackers to enter MTP mode without proper authentication.	2024-07-02	3.3	Low
CVE-2024-34583	samsung - multiple products	Improper access control in system property prior to SMR Jul-2024 Release 1 allows local attackers to get device identifier.	2024-07-02	3.3	Low

CVE-2024-34586	samsung - multiple products	Improper access control in KnoxCustomManagerService prior to SMR Jul-2024 Release 1 allows local attackers to configure Knox privacy policy.	2024-07-02	3.3	Low
CVE-2024-34597	samsung - health	Improper input validation in Samsung Health prior to version 6.27.0.113 allows local attackers to write arbitrary document files to the sandbox of Samsung Health. User interaction is required for triggering this vulnerability.	2024-07-02	3.3	Low
CVE-2024-34599	samsung - tips	Improper input validation in Tips prior to version 6.2.9.4 in Android 14 allows local attacker to send broadcast with Tips' privilege.	2024-07-02	3.3	Low
CVE-2024-34600	samsung - flow	Improper verification of intent by broadcast receiver vulnerability in Samsung Flow prior to version 4.9.13.0 allows local attackers to copy image files to external storage.	2024-07-02	3.3	Low

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.

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.