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 ٤ أغسطس إلى ١٠ أغسطس. 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 4th of August to 10th of August. 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-2024-43111	Mozilla	Long pressing on a download link could potentially allow Javascript commands to be executed within the browser This vulnerability affects Firefox for iOS < 129.	2024-08-06	9.8	Critical
CVE-2024-36130	Ivanti	An insufficient authorization vulnerability in web component of EPMM prior to 12.1.0.1 allows an unauthorized attacker within the network to execute arbitrary commands on the underlying operating system of the appliance.	2024-08-07	9.8	Critical
CVE-2024-20450	Cisco	Multiple vulnerabilities in the web-based management interface of Cisco Small Business SPA300 Series IP Phones and Cisco Small Business SPA500 Series IP Phones could allow an unauthenticated, remote attacker to execute arbitrary commands on the underlying operating system with root privileges. These vulnerabilities exist because incoming HTTP packets are not properly checked for errors, which could result in a buffer overflow. An attacker could exploit this vulnerability by sending a crafted HTTP request to an affected device. A successful exploit could allow the attacker to overflow an internal buffer and execute arbitrary commands at the root privilege level.	2024-08-07	9.8	Critical
CVE-2024-20454	Cisco	Multiple vulnerabilities in the web-based management interface of Cisco Small Business SPA300 Series IP Phones and Cisco Small Business SPA500 Series IP Phones could allow an unauthenticated, remote attacker to execute arbitrary commands on the underlying operating system with root privileges. These vulnerabilities exist because incoming HTTP packets are not properly checked for errors, which could result in a buffer overflow. An attacker could exploit this vulnerability by sending a crafted HTTP request to an affected device. A successful exploit could allow the attacker to overflow an internal buffer and execute arbitrary commands at the root privilege level.	2024-08-07	9.8	Critical
CVE-2024-7519	Mozilla	Insufficient checks when processing graphics shared memory could have led to memory corruption. This could be leveraged by an attacker to perform a sandbox escape. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.	2024-08-06	9.6	Critical
CVE-2024-42037	Huawei	Vulnerability of uncaught exceptions in the Graphics module Impact: Successful exploitation of this vulnerability may affect service confidentiality.	2024-08-08	9.3	Critical
CVE-2024-7520	Mozilla	A type confusion bug in WebAssembly could be leveraged by an attacker to potentially achieve code execution. This vulnerability affects Firefox < 129, Firefox ESR < 128.1, and Thunderbird < 128.1.	2024-08-06	8.8	High
CVE-2024-7521	Mozilla	Incomplete WebAssembly exception handing could have led to a use-after-free. This vulnerability affects Firefox < 129, Firefox ESR <	2024-08-06	8.8	High

		115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.			
CVE-2024-7522	Mozilla	Editor code failed to check an attribute value. This could have led to an out-of-bounds read. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.	2024-08-06	8.8	High
CVE-2024-7527	Mozilla	Unexpected marking work at the start of sweeping could have led to a use-after-free. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.	2024-08-06	8.8	High
CVE-2024-7528	Mozilla	Incorrect garbage collection interaction in IndexedDB could have led to a use-after-free. This vulnerability affects Firefox < 129, Firefox ESR < 128.1, and Thunderbird < 128.1.	2024-08-06	8.8	High
CVE-2024-7530	Mozilla	Incorrect garbage collection interaction could have led to a use-after-free. This vulnerability affects Firefox < 129.	2024-08-06	8.8	High
CVE-2024-6988	Google	Use after free in Downloads in Google Chrome on iOS prior to 127.0.6533.72 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-6989	Google	Use after free in Loader in Google Chrome prior to 127.0.6533.72 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-6991	Google	Use after free in Dawn in Google Chrome prior to 127.0.6533.72 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-6994	Google	Heap buffer overflow in Layout in Google Chrome prior to 127.0.6533.72 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	8.8	High
CVE-2024-6997	Google	Use after free in Tabs in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	8.8	High
CVE-2024-6998	Google	Use after free in User Education in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	8.8	High
CVE-2024-7000	Google	Use after free in CSS in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	8.8	High
CVE-2024-7532	Google	Out of bounds memory access in ANGLE in Google Chrome prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Critical)	2024-08-06	8.8	High
CVE-2024-7533	Google	Use after free in Sharing in Google Chrome on iOS prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-7534	Google	Heap buffer overflow in Layout in Google Chrome prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-7535	Google	Inappropriate implementation in V8 in Google Chrome prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-7536	Google	Use after free in WebAudio in Google Chrome prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-7550	Google	Type Confusion in V8 in Google Chrome prior to 127.0.6533.99 allowed a remote attacker to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)	2024-08-06	8.8	High
CVE-2024-36131	Ivanti	An insecure deserialization vulnerability in web component of EPMM prior to 12.1.0.1 allows an authenticated remote attacker to execute arbitrary commands on the underlying operating system of the appliance.	2024-08-07	8.8	High
CVE-2024-42038	Huawei	Vulnerability of PIN enhancement failures in the screen lock module Impact: Successful exploitation of this vulnerability may affect service confidentiality, integrity, and availability.	2024-08-08	8.8	High
CVE-2024-42035	Huawei	Permission control vulnerability in the App Multiplier module Impact:Successful exploitation of this vulnerability may affect functionality and confidentiality.	2024-08-08	8.4	High
CVE-2024-7525	Mozilla	It was possible for a web extension with minimal permissions to create a `StreamFilter` which could be used to read and modify the response body of requests on any site. This vulnerability affects	2024-08-06	8.1	High

		Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.			
CVE-2024-36132	Ivanti	Insufficient verification of authentication controls in EPMM prior to 12.1.0.1 allows a remote attacker to bypass authentication and access sensitive resources.	2024-08-07	7.5	High
CVE-2024-20451	Cisco	Multiple vulnerabilities in the web-based management interface of Cisco Small Business SPA300 Series IP Phones and Cisco Small Business SPA500 Series IP Phones could allow an unauthenticated, remote attacker to cause an affected device to reload unexpectedly. These vulnerabilities exist because HTTP packets are not properly checked for errors. An attacker could exploit this vulnerability by sending a crafted HTTP packet to the remote interface of an affected device. A successful exploit could allow the attacker to cause a DoS condition on the device.	2024-08-07	7.5	High
CVE-2024-38202	Microsoft	Summary Microsoft was notified that an elevation of privilege vulnerability exists in Windows Update, potentially enabling an attacker with basic user privileges to reintroduce previously mitigated vulnerabilities or circumvent some features of Virtualization Based Security (VBS). However, an attacker attempting to exploit this vulnerability requires additional interaction by a privileged user to be successful. Microsoft is developing a security update to mitigate this threat, but it is not yet available. Guidance to help customers reduce the risks associated with this vulnerability and to protect their systems until the mitigation is available in a Windows security update is provided in the Recommended Actions section of this CVE. This CVE will be updated, and customers will be notified when the official mitigation is available in a Windows security update. We highly encourage customers to subscribe to Security Update Guide notifications to receive an alert when this update occurs. Details A security researcher informed Microsoft of an elevation of privilege vulnerability in Windows Update potentially enabling an attacker with basic user privileges to reintroduce previously mitigated vulnerabilities or circumvent some features of VBS. For exploitation to succeed, an attacker must trick or convince an Administrator or a user with delegated permissions into performing a system restore which inadvertently triggers the vulnerability. Microsoft is developing a security update that will mitigate this vulnerability, but it is not yet available. This CVE will be updated with new information and links to the security updates once available. We highly encourage customers subscribe to Security Update Guide notifications to be alerted of updates. See Microsoft Technical Security Notifications and Security Update Guide Notification System News: Create your profile now – Microsoft Security Response Center. Microsoft is not aware of any attempts to exploit this vulnerability. However, a public presentation regarding this vulnerability was hosted at BlackHat on August 7, 2024. The presentation was appropriately coordinated with Microsoft but may change the threat landscape. Customers concerned with these risks should reference the guidance provided in the Recommended Actions section to protect their systems. Recommended Actions The following recommendations do not mitigate the vulnerability but can be used to reduce the risk of exploitation until the security update...	2024-08-08	7.3	High
CVE-2024-42033	Huawei	Access control vulnerability in the security verification module mpact: Successful exploitation of this vulnerability will affect integrity and confidentiality.	2024-08-08	6.9	Medium
CVE-2024-35143	IBM	IBM Planning Analytics Local 2.0 and 2.1 connects to a MongoDB server. MongoDB, a document-oriented database system, is listening on the remote port, and it is configured to allow connections without password authentication. A remote attacker can gain unauthorized access to the database. IBM X-Force ID: 292420.	2024-08-04	6.7	Medium
CVE-2024-21302	Microsoft	Summary: Microsoft was notified that an elevation of privilege vulnerability exists in Windows based systems supporting Virtualization Based Security (VBS), including a subset of Azure Virtual Machine SKUS. This vulnerability enables an attacker with administrator privileges to replace current versions of Windows system files with outdated versions. By exploiting this vulnerability, an attacker could reintroduce previously mitigated vulnerabilities, circumvent some	2024-08-08	6.7	Medium

		features of VBS, and exfiltrate data protected by VBS. Microsoft is developing a security update to mitigate this threat, but it is not yet available. Guidance to help customers reduce the risks associated with this vulnerability and to protect their systems until the mitigation is available in a Windows security update is provided in the Recommended Actions section of this CVE. This CVE will be updated when the mitigation is available in a Windows security update. We highly encourage customers to subscribe to Security Update Guide notifications to receive an alert when this update occurs. Update: August 13, 2024 Microsoft has released the August 2024 security updates that include an opt-in revocation policy mitigation to address this vulnerability. Customers running affected versions of Windows are encouraged to review KB5042562: Guidance for blocking rollback of virtualization-based security related updates to assess if this opt-in policy meets the needs of their environment before implementing this mitigation. There are risks associated with this mitigation that should be understood prior to applying it to your systems. Detailed information about these risks is also available in KB5042562. Details: A security researcher informed Microsoft of an elevation of privilege vulnerability in Windows 10, Windows 11, Windows Server 2016, and higher based systems including Azure Virtual Machines (VM) that support VBS. For more information on Windows versions and VM SKUs supporting VBS, reference: Virtualization-based Security (VBS) | Microsoft Learn. The vulnerability enables an attacker with administrator privileges on the target system to replace current Windows system files with outdated versions. Successful exploitation provides an attacker with the ability to reintroduce previously mitigated vulnerabilities, circumvent VBS security features, and exfiltrate data protected by VBS. Microsoft is developing a security update that will revoke outdated, unpatched VBS system files to mitigate this...			
CVE-2024-42034	Huawei	LaunchAnywhere vulnerability in the account module. Impact: Successful exploitation of this vulnerability may affect service confidentiality.	2024-08-08	6.6	Medium
CVE-2024-28962	Dell	Dell Command | Update, Dell Update, and Alienware Update UWP, versions prior to 5.4, contain an Exposed Dangerous Method or Function vulnerability. An unauthenticated attacker with remote access could potentially exploit this vulnerability, leading to denial of service.	2024-08-06	6.5	Medium
CVE-2024-7526	Mozilla	ANGLE failed to initialize parameters which led to reading from uninitialized memory. This could be leveraged to leak sensitive data from memory. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.	2024-08-06	6.5	Medium
CVE-2024-7529	Mozilla	The date picker could partially obscure security prompts. This could be used by a malicious site to trick a user into granting permissions. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, Firefox ESR < 128.1, Thunderbird < 128.1, and Thunderbird < 115.14.	2024-08-06	6.5	Medium
CVE-2024-7531	Mozilla	Calling `PK11_Encrypt()` in NSS using CKM_CHACHA20 and the same buffer for input and output can result in plaintext on an Intel Sandy Bridge processor. In Firefox this only affects the QUIC header protection feature when the connection is using the ChaCha20-Poly1305 cipher suite. The most likely outcome is connection failure, but if the connection persists despite the high packet loss it could be possible for a network observer to identify packets as coming from the same source despite a network path change. This vulnerability affects Firefox < 129, Firefox ESR < 115.14, and Firefox ESR < 128.1.	2024-08-06	6.5	Medium
CVE-2024-38206	Microsoft	An authenticated attacker can bypass Server-Side Request Forgery (SSRF) protection in Microsoft Copilot Studio to leak sensitive information over a network.	2024-08-06	6.5	Medium
CVE-2024-34788	Ivanti	An improper authentication vulnerability in web component of EPMM prior to 12.1.0.1 allows a remote malicious user to access potentially sensitive information	2024-08-07	6.5	Medium
CVE-2024-7246	Google	It's possible for a gRPC client communicating with a HTTP/2 proxy to poison the HPACK table between the proxy and the backend such that other clients see failed requests. It's also possible to use this vulnerability to leak other clients HTTP header keys, but not values. This occurs because the error status for a misencoded header is not cleared between header reads, resulting in subsequent (incrementally indexed) added headers in the first request being poisoned until cleared from the HPACK table.	2024-08-06	6.3	Medium

		Please update to a fixed version of gRPC as soon as possible. This bug has been fixed in 1.58.3, 1.59.5, 1.60.2, 1.61.3, 1.62.3, 1.63.2, 1.64.3, 1.65.4.			
CVE-2024-42030	Huawei	Access permission verification vulnerability in the content sharing pop-up module Impact: Successful exploitation of this vulnerability may affect service confidentiality.	2024-08-08	6.2	Medium
CVE-2024-38166	Microsoft	An unauthenticated attacker can exploit improper neutralization of input during web page generation in Microsoft Dynamics 365 to spoof over a network by tricking a user to click on a link.	2024-08-06	6.1	Medium
CVE-2024-37403	Ivanti	Ivanti Docs@Work for Android, before 2.26.0 is affected by the 'Dirty Stream' vulnerability. The application fails to properly sanitize file names, resulting in a path traversal-affiliated vulnerability. This potentially enables other malicious apps on the device to read sensitive information stored in the app root.	2024-08-07	5.5	Medium
CVE-2024-42232	Linux	In the Linux kernel, the following vulnerability has been resolved: libceph: fix race between delayed_work() and ceph_monc_stop() The way the delayed work is handled in ceph_monc_stop() is prone to races with mon_fault() and possibly also finish_hunting(). Both of these can requeue the delayed work which wouldn't be canceled by any of the following code in case that happens after cancel_delayed_work_sync() runs -- __close_session() doesn't mess with the delayed work in order to avoid interfering with the hunting interval logic. This part was missed in commit b5d91704f53e ("libceph: behave in mon_fault() if cur_mon < 0") and use-after-free can still ensue on monc and objects that hang off of it, with monc->auth and monc->monmap being particularly susceptible to quickly being reused. To fix this: - clear monc->cur_mon and monc->hunting as part of closing the session in ceph_monc_stop() - bail from delayed_work() if monc->cur_mon is cleared, similar to how it's done in mon_fault() and finish_hunting() (based on monc->hunting) - call cancel_delayed_work_sync() after the session is closed	2024-08-07	5.5	Medium
CVE-2024-42234	Linux	In the Linux kernel, the following vulnerability has been resolved: mm: fix crashes from deferred split racing folio migration Even on 6.10-rc6, I've been seeing elusive "Bad page state"s (often on flags when freeing, yet the flags shown are not bad: PG_locked had been set and cleared??), and VM_BUG_ON_PAGE(page_ref_count(page) == 0)s from deferred_split_scan()'s folio_put(), and a variety of other BUG and WARN symptoms implying double free by deferred split and large folio migration. 6.7 commit 9bcef5973e31 ("mm: memcg: fix split queue list crash when large folio migration") was right to fix the memcg-dependent locking broken in 85ce2c517ade ("memcontrol: only transfer the memcg data for migration"), but missed a subtlety of deferred_split_scan(): it moves folios to its own local list to work on them without split_queue_lock, during which time folio->_deferred_list is not empty, but even the "right" lock does nothing to secure the folio and the list it is on. Fortunately, deferred_split_scan() is careful to use folio_try_get(): so folio_migrate_mapping() can avoid the race by folio_undo_large_rmappable()	2024-08-07	5.5	Medium

		while the old folio's reference count is temporarily frozen to 0 - adding such a freeze in the !mapping case too (originally, folio lock and unmapping and no swap cache left an anon folio unreachable, so no freezing was needed there: but the deferred split queue offers a way to reach it).			
CVE-2024-42235	Linux	In the Linux kernel, the following vulnerability has been resolved: s390/mm: Add NULL pointer check to crst_table_free() base_crst_free() crst_table_free() used to work with NULL pointers before the conversion to ptdescs. Since crst_table_free() can be called with a NULL pointer (error handling in crst_table_upgrade() add an explicit check. Also add the same check to base_crst_free() for consistency reasons. In real life this should not happen, since order two GFP_KERNEL allocations will not fail, unless FAIL_PAGE_ALLOC is enabled and used.	2024-08-07	5.5	Medium
CVE-2024-42236	Linux	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: configfs: Prevent OOB read/write in usb_string_copy() Userspace provided string 's' could trivially have the length zero. Left unchecked this will firstly result in an OOB read in the form `if (str[0 - 1] == '\n') followed closely by an OOB write in the form `str[0 - 1] = '\0'`. There is already a validating check to catch strings that are too long. Let's supply an additional check for invalid strings that are too short.	2024-08-07	5.5	Medium
CVE-2024-42237	Linux	In the Linux kernel, the following vulnerability has been resolved: firmware: cs_dsp: Validate payload length before processing block Move the payload length check in cs_dsp_load() and cs_dsp_coeff_load() to be done before the block is processed. The check that the length of a block payload does not exceed the number of remaining bytes in the firwmware file buffer was being done near the end of the loop iteration. However, some code before that check used the length field without validating it.	2024-08-07	5.5	Medium
CVE-2024-42238	Linux	In the Linux kernel, the following vulnerability has been resolved: firmware: cs_dsp: Return error if block header overflows file Return an error from cs_dsp_power_up() if a block header is longer than the amount of data left in the file. The previous code in cs_dsp_load() and cs_dsp_load_coeff() would loop while there was enough data left in the file for a valid region. This protected against overrunning the end of the file data, but it didn't abort the file processing with an error.	2024-08-07	5.5	Medium
CVE-2024-42239	Linux	In the Linux kernel, the following vulnerability has been resolved: bpf: Fail bpf_timer_cancel when callback is being cancelled Given a schedule: timer1 cb timer2 cb bpf_timer_cancel(timer2); bpf_timer_cancel(timer1); Both bpf_timer_cancel calls would wait for the other callback to finish executing, introducing a lockup.	2024-08-07	5.5	Medium

		Add an atomic_t count named 'cancelling' in bpf_hrtimer. This keeps
		track of all in-flight cancellation requests for a given BPF timer.
		Whenever cancelling a BPF timer, we must check if we have outstanding
		cancellation requests, and if so, we must fail the operation with an
		error (-EDEADLK) since cancellation is synchronous and waits for the
		callback to finish executing. This implies that we can enter a deadlock
		situation involving two or more timer callbacks executing in parallel
		and attempting to cancel one another.
		
		Note that we avoid incrementing the cancelling counter for the target
		timer (the one being cancelled) if bpf_timer_cancel is not invoked from
		a callback, to avoid spurious errors. The whole point of detecting
		cur->cancelling and returning -EDEADLK is to not enter a busy wait loop
		(which may or may not lead to a lockup). This does not apply in case the
		caller is in a non-callback context, the other side can continue to
		cancel as it sees fit without running into errors.
		
		Background on prior attempts:
		
		Earlier versions of this patch used a bool 'cancelling' bit and used the
		following pattern under timer->lock to publish cancellation status.
		
		lock(t->lock);
		t->cancelling = true;
		mb();
		if (cur->cancelling)
		return -EDEADLK;
		unlock(t->lock);
		hrtimer_cancel(t->timer);
		t->cancelling = false;
		
		The store outside the critical section could overwrite a parallel
		requests t->cancelling assignment to true, to ensure the parallely
		executing callback observes its cancellation status.
		
		It would be necessary to clear this cancelling bit once hrtimer_cancel
		is done, but lack of serialization introduced races. Another option was
		explored where bpf_timer_start would clear the bit when (re)starting the
		timer under timer->lock. This would ensure serialized access to the
		cancelling bit, but may allow it to be cleared before in-flight
		hrtimer_cancel has finished executing, such that lockups can occur
		again.
		
		Thus, we choose an atomic counter to keep track of all outstanding
		cancellation requests and use it to prevent lockups in case callbacks
		attempt to cancel each other while executing in parallel.			

CVE-2024-42240	Linux	In the Linux kernel, the following vulnerability has been resolved: x86/bhi: Avoid warning in #DB handler due to BHI mitigation When BHI mitigation is enabled, if SYSENTER is invoked with the TF flag set then entry_SYSENTER_compat() uses CLEAR_BRANCH_HISTORY and calls the clear_bhb_loop() before the TF flag is cleared. This causes the #DB handler (exc_debug_kernel()) to issue a warning because single-step is used outside the entry_SYSENTER_compat() function. To address this issue, entry_SYSENTER_compat() should use CLEAR_BRANCH_HISTORY after making sure the TF flag is cleared. The problem can be reproduced with the following sequence:	2024-08-07	5.5	Medium

		$ cat sysenter_step.c
		int main()
		{ asm("pushf; pop %ax; bts $8,%ax; push %ax; popf; sysenter"); }
		
		$ gcc -o sysenter_step sysenter_step.c
		
		$ ./sysenter_step
		Segmentation fault (core dumped)
		
		The program is expected to crash, and the #DB handler will issue a warning.
		
		Kernel log:
		
		WARNING: CPU: 27 PID: 7000 at arch/x86/kernel/traps.c:1009 exc_debug_kernel+0xd2/0x160
		...
		RIP: 0010:exc_debug_kernel+0xd2/0x160
		...
		Call Trace:
		<#DB>
		? show_regs+0x68/0x80
		? __warn+0x8c/0x140
		? exc_debug_kernel+0xd2/0x160
		? report_bug+0x175/0x1a0
		? handle_bug+0x44/0x90
		? exc_invalid_op+0x1c/0x70
		? asm_exc_invalid_op+0x1f/0x30
		? exc_debug_kernel+0xd2/0x160
		exc_debug+0x43/0x50
		asm_exc_debug+0x1e/0x40
		RIP: 0010:clear_bhb_loop+0x0/0xb0
		...
		</#DB>
		<TASK>
		? entry_SYSENTER_compat_after_hwframe+0x6e/0x8d
		</TASK>
		
		[bp: Massage commit message.]			

CVE-2024-42241	Linux	In the Linux kernel, the following vulnerability has been resolved: mm/shmem: disable PMD-sized page cache if needed For shmem files, it's possible that PMD-sized page cache can't be supported by xarray. For example, 512MB page cache on ARM64 when the base page size is 64KB can't be supported by xarray. It leads to errors as the following messages indicate when this sort of xarray entry is split. WARNING: CPU: 34 PID: 7578 at lib/xarray.c:1025 xas_split_alloc+0xf8/0x128 Modules linked in: binfmt_misc nft_fib_inet nft_fib_ipv4 nft_fib_ipv6 \ nft_fib nft_reject_inet nf_reject_ipv4 nf_reject_ipv6 nft_reject \ nft_ct nft_chain_nat nf_nat nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 \ ip_set rfkill nf_tables nfnetlink vfat fat virtio_balloon drm fuse xfs \ libcrc32c crct10dif_ce ghash_ce sha2_ce sha256_arm64 sha1_ce virtio_net \ net_failover virtio_console virtio_blk failover dimlib virtio_mmio CPU: 34 PID: 7578 Comm: test Kdump: loaded Tainted: G W 6.10.0-rc5-gavin+ #9 Hardware name: QEMU KVM Virtual Machine, BIOS edk2-20240524-1.el9 05/24/2024 pstate: 83400005 (Nzcv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=--) pc : xas_split_alloc+0xf8/0x128 lr : split_huge_page_to_list_to_order+0x1c4/0x720 sp : ffff8000882af5f0 x29: ffff8000882af5f0 x28: ffff8000882af650 x27: ffff8000882af768 x26: 0000000000000cc0 x25: 000000000000000d x24: ffff00010625b858 x23: ffff8000882af650 x22: fffffffdfc0900000 x21: 0000000000000000 x20: 0000000000000000 x19: fffffffdfc0900000 x18:	2024-08-07	5.5	Medium

		0000000000000000 x17: 0000000000000000 x16: 0000018000000000 x15: 52f8004000000000 x14: 0000e00000000000 x13: 0000000000002000 x12: 0000000000000020 x11: 52f8000000000000 x10: 52f8e1c0ffff6000 x9 : ffffbeb9619a681c x8 : 0000000000000003 x7 : 0000000000000000 x6 : ffff00010b02ddb0 x5 : ffffbeb96395e378 x4 : 0000000000000000 x3 : 0000000000000cc0 x2 : 000000000000000d x1 : 000000000000000c x0 : 0000000000000000 Call trace: xas_split_alloc+0xf8/0x128 split_huge_page_to_list_to_order+0x1c4/0x720 truncate_inode_partial_folio+0xdc/0x160 shmem_undo_range+0x2bc/0x6a8 shmem_fallocate+0x134/0x430 vfs_fallocate+0x124/0x2e8 ksys_fallocate+0x4c/0xa0 __arm64_sys_fallocate+0x24/0x38 invoke_syscall.constprop.0+0x7c/0xd8 do_el0_svc+0xb4/0xd0 el0_svc+0x44/0x1d8 el0t_64_sync_handler+0x134/0x150 el0t_64_sync+0x17c/0x180 Fix it by disabling PMD-sized page cache when HPAGE_PMD_ORDER is larger than MAX_PAGECACHE_ORDER. As Matthew Wilcox pointed, the page cache in a shmem file isn't represented by a multi-index entry and doesn't have this limitation when the xarry entry is split until commit 6b24ca4a1a8d ("mm: Use multi-index entries in the page cache").			

CVE-2024-42242	Linux	In the Linux kernel, the following vulnerability has been resolved:
		
		mmc: sdhci: Fix max_seg_size for 64KiB PAGE_SIZE
		
		blk_queue_max_segment_size() ensured:
		
		if (max_size < PAGE_SIZE)
		max_size = PAGE_SIZE;
		
		whereas:
		
		blk_validate_limits() makes it an error:
		
		if (WARN_ON_ONCE(lim->max_segment_size < PAGE_SIZE))
		return -EINVAL;
		
		The change from one to the other, exposed sdhci which was setting maximum
		segment size too low in some circumstances.
		
		Fix the maximum segment size when it is too low.
			2024-08-07	5.5	Medium

CVE-2024-42243	Linux	In the Linux kernel, the following vulnerability has been resolved: mm/filemap: make MAX_PAGECACHE_ORDER acceptable to xarray Patch series "mm/filemap: Limit page cache size to that supported by xarray", v2. Currently, xarray can't support arbitrary page cache size. More details can be found from the WARN_ON() statement in xas_split_alloc(). In our test whose code is attached below, we hit the WARN_ON() on ARM64 system where the base page size is 64KB and huge page size is 512MB. The issue was reported long time ago and some discussions on it can be found here [1]. [1] https://www.spinics.net/lists/linux-xfs/msg75404.html	2024-08-07	5.5	Medium

		In order to fix the issue, we need to adjust
		MAX_PAGECACHE_ORDER to one
		supported by xarray and avoid PMD-sized page cache if needed.
		The code
		changes are suggested by David Hildenbrand.
		
		PATCH[1] adjusts MAX_PAGECACHE_ORDER to that supported by
		xarray
		PATCH[2-3] avoids PMD-sized page cache in the synchronous
		readahead path
		PATCH[4] avoids PMD-sized page cache for shmem files if needed
		
		Test program
		==========
		# cat test.c
		#define _GNU_SOURCE
		#include <stdio.h>
		#include <stdlib.h>
		#include <unistd.h>
		#include <string.h>
		#include <fcntl.h>
		#include <errno.h>
		#include <sys/syscall.h>
		#include <sys/mman.h>
		
		#define TEST_XFS_FILENAME "/tmp/data"
		#define TEST_SHMEM_FILENAME "/dev/shm/data"
		#define TEST_MEM_SIZE 0x20000000
		
		int main(int argc, char **argv)
		{
		const char *filename;
		int fd = 0;
		void *buf = (void *)-1, *p;
		int pgsize = getpagesize();
		int ret;
		
		if (pgsize != 0x10000) {
		fprintf(stderr, "64KB base page size is required\n");
		return -EPERM;
		}
		
		system("echo force >
		/sys/kernel/mm/transparent_hugepage/shmem_enabled");
		system("rm -fr /tmp/data");
		system("rm -fr /dev/shm/data");
		system("echo 1 > /proc/sys/vm/drop_caches");
		
		/* Open xfs or shmem file */
		filename = TEST_XFS_FILENAME;
		if (argc > 1 && !strcmp(argv[1], "shmem"))
		filename = TEST_SHMEM_FILENAME;
		
		fd = open(filename, O_CREAT | O_RDWR | O_TRUNC);
		if (fd < 0) {
		fprintf(stderr, "Unable to open <%s>\n", filename);
		return -EIO;
		}
		
		/* Extend file size */
		ret = ftruncate(fd, TEST_MEM_SIZE);
		if (ret) {
		fprintf(stderr, "Error %d to ftruncate()\n", ret);
		goto cleanup;
		}
		
		/* Create VMA */
		buf = mmap(NULL, TEST_MEM_SIZE,
		PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
		if (buf == (void *)-1) {
		fprintf(stderr, "Unable to mmap <%s>\n", filename);
		goto cleanup;
		}
		
		fprintf(stdout, "mapped buffer at 0x%p\n", buf);
		ret = madvise(buf, TEST_MEM_SIZE, MADV_HUGEPAGE);
		if (ret) {
		fprintf(stderr, "Unable to madvise(MADV_HUGEPAGE)\n");
		goto cleanup;			

		}
		
		/* Populate VMA */
		ret = madvise(buf, TEST_MEM_SIZE, MADV_POPULATE_WRITE);
		if (ret) {
		fprintf(stderr, "Error %d to madvise(MADV_POPULATE_WRITE)\n",
		ret);
		goto cleanup;
		}
		
		/* Punch the file to enforce xarray split */
		ret = fallocate(fd, FALLOC_FL_KEEP_SIZE |
		FALLOC_FL_PUNCH_HOLE,
		TEST_MEM_SIZE - pgsize, pgsize);
		if (ret)
		fprintf(stderr, "Error %d to fallocate()\n", ret);
		
		cleanup:
		if (buf != (void *)-1)
		munmap(buf, TEST_MEM_SIZE);
		if (fd > 0)
		close(fd);
		
		return 0;
		}
		
		# gcc test.c -o test
		# cat /proc/1/smaps | grep KernelPageSize | head -n 1
		KernelPageSize: 64 kB
		# ./test shmem
		:
		----------[cut here]----------
		WARNING: CPU: 17 PID: 5253 at lib/xarray.c:1025
		xas_split_alloc+0xf8/0x128
		Modules linked in: nft_fib_inet nft_fib_ipv4 nft_fib_ipv6 nft_fib \
		nft_reject_inet nf_reject_ipv4 nf_reject_ipv6 nft_reject nft_ct \
		nft_chain_nat nf_nat nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4
		\
		ip_set nf_tables rfkill nfnetlink vfat fat virtio_balloon \
		drm fuse xfs libcrc32c crct10dif_ce ghash_ce sha2_ce
		sha256_arm64 \
		virtio_net sha1_ce net_failover failover virtio_console virtio_blk \
		dimlib virtio_mmio
		CPU: 17 PID: 5253 Comm: test Kdump: loaded Tainted: G W
		6.10.0-rc5-gavin+ #12
		Hardware name: QEMU KVM Virtual Machine, BIOS edk2-
		20240524-1.el9 05/24/2024
		pstate: 83400005 (Nzcv daif +PAN -UAO +TC
		---truncated---			

CVE-2024-42244	Linux	In the Linux kernel, the following vulnerability has been resolved: USB: serial: mos7840: fix crash on resume Since commit c49cfa917025 ("USB: serial: use generic method if no alternative is provided in usb serial layer"), USB serial core calls the generic resume implementation when the driver has not provided one. This can trigger a crash on resume with mos7840 since support for multiple read URBs was added back in 2011. Specifically, both port read URBs are now submitted on resume for open ports, but the context pointer of the second URB is left set to the core rather than mos7840 port structure. Fix this by implementing dedicated suspend and resume functions for mos7840. Tested with Delock 87414 USB 2.0 to 4x serial adapter. [johan: analyse crash and rewrite commit message; set busy flag on resume; drop bulk-in check; drop unnecessary usb_kill_urb()]	2024-08-07	5.5	Medium
CVE-2024-42245	Linux	In the Linux kernel, the following vulnerability has been resolved: Revert "sched/fair: Make sure to try to detach at least one movable task"	2024-08-07	5.5	Medium

		This reverts commit b0defa7ae03ecf91b8bfd10ede430cff12fcbd06. b0defa7ae03ec changed the load balancing logic to ignore env.max_loop if all tasks examined to that point were pinned. The goal of the patch was to make it more likely to be able to detach a task buried in a long list of pinned tasks. However, this has the unfortunate side effect of creating an O(n) iteration in detach_tasks(), as we now must fully iterate every task on a cpu if all or most are pinned. Since this load balance code is done with rq lock held, and often in softirq context, it is very easy to trigger hard lockups. We observed such hard lockups with a user who affined O(10k) threads to a single cpu. When I discussed this with Vincent he initially suggested that we keep the limit on the number of tasks to detach, but increase the number of tasks we can search. However, after some back and forth on the mailing list, he recommended we instead revert the original patch, as it seems likely no one was actually getting hit by the original issue.			

CVE-2024-42246	Linux	In the Linux kernel, the following vulnerability has been resolved:
		
		net, sunrpc: Remap EPERM in case of connection failure in xs_tcp_setup_socket
		
		When using a BPF program on kernel_connect(), the call can return -EPERM. This
		causes xs_tcp_setup_socket() to loop forever, filling up the syslog and causing
		the kernel to potentially freeze up.
		
		Neil suggested:
		
		  This will propagate -EPERM up into other layers which might not be ready
		  to handle it. It might be safer to map EPERM to an error we would be more
		  likely to expect from the network system - such as ECONNREFUSED or ENETDOWN.
		
		ECONNREFUSED as error seems reasonable. For programs setting a different error
		can be out of reach (see handling in 4fbac77d2d09) in particular on kernels
		which do not have f10d05966196 ("bpf: Make BPF_PROG_RUN_ARRAY return -err
		instead of allow boolean"), thus given that it is better to simply remap for
		consistent behavior. UDP does handle EPERM in xs_udp_send_request().
			2024-08-07	5.5	Medium

CVE-2024-42247	Linux	In the Linux kernel, the following vulnerability has been resolved: wireguard: allowedips: avoid unaligned 64-bit memory accesses On the parisc platform, the kernel issues kernel warnings because swap_endian() tries to load a 128-bit IPv6 address from an unaligned memory location: Kernel: unaligned access to 0x55f4688c in wg_allowedips_insert_v6+0x2c/0x80 [wireguard] (iir 0xf3010df) Kernel: unaligned access to 0x55f46884 in wg_allowedips_insert_v6+0x38/0x80 [wireguard] (iir 0xf2010dc) Avoid such unaligned memory accesses by instead using the get_unaligned_be64() helper macro. [Jason: replace src[8] in original patch with src+8]	2024-08-07	5.5	Medium
CVE-2024-42248	Linux	In the Linux kernel, the following vulnerability has been resolved: tty: serial: ma35d1: Add a NULL check for of_node	2024-08-07	5.5	Medium

		The pdev->dev.of_node can be NULL if the "serial" node is absent. Add a NULL check to return an error in such cases.			
CVE-2024-42250	Linux	In the Linux kernel, the following vulnerability has been resolved: cachefiles: add missing lock protection when polling Add missing lock protection in poll routine when iterating xarray, otherwise: Even with RCU read lock held, only the slot of the radix tree is ensured to be pinned there, while the data structure (e.g. struct cachefiles_req) stored in the slot has no such guarantee. The poll routine will iterate the radix tree and dereference cachefiles_req accordingly. Thus RCU read lock is not adequate in this case and spinlock is needed here.	2024-08-07	5.5	Medium
CVE-2024-20443	Cisco	A vulnerability in the web-based management interface of Cisco ISE could allow an authenticated, remote attacker to conduct an XSS attack against a user of the interface. This vulnerability is due to insufficient validation of user-supplied input by the web-based management interface of an affected system. An attacker could exploit this vulnerability by injecting malicious code into specific pages of the interface. A successful exploit could allow the attacker to execute arbitrary script code in the context of the affected interface or access sensitive, browser-based information. To exploit this vulnerability, the attacker must have at least a low-privileged account on an affected device.	2024-08-07	5.4	Medium
CVE-2024-20479	Cisco	A vulnerability in the web-based management interface of Cisco ISE could allow an authenticated, remote attacker to conduct an XSS attack against a user of the interface. This vulnerability is due to insufficient validation of user-supplied input by the web-based management interface of an affected system. An attacker could exploit this vulnerability by injecting malicious code into specific pages of the interface. A successful exploit could allow the attacker to execute arbitrary script code in the context of the affected interface or access sensitive, browser-based information. To exploit this vulnerability, the attacker must have Admin privileges on an affected device.	2024-08-07	4.8	Medium
CVE-2024-6995	Google	Inappropriate implementation in Fullscreen in Google Chrome on Android prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to spoof the contents of the Omnibox (URL bar) via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	4.7	Medium
CVE-2024-42032	Huawei	Access permission verification vulnerability in the Contacts module Impact: Successful exploitation of this vulnerability may affect service confidentiality.	2024-08-08	4.4	Medium
CVE-2024-39751	IBM	IBM InfoSphere Information Server 11.7 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: 297429	2024-08-06	4.3	Medium
CVE-2024-6999	Google	Inappropriate implementation in FedCM in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to perform UI spoofing via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	4.3	Medium
CVE-2024-7001	Google	Inappropriate implementation in HTML in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to perform UI spoofing via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	4.3	Medium
CVE-2024-7003	Google	Inappropriate implementation in FedCM in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to perform UI spoofing via a crafted HTML page. (Chromium security severity: Low)	2024-08-06	4.3	Medium
CVE-2024-7004	Google	Insufficient validation of untrusted input in Safe Browsing in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to bypass discretionary access control via a malicious file. (Chromium security severity: Low)	2024-08-06	4.3	Medium
CVE-2024-7005	Google	Insufficient validation of untrusted input in Safe Browsing in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to bypass discretionary access control via a malicious file. (Chromium security severity: Low)	2024-08-06	4.3	Medium

CVE-2023-7265	Huawei	Permission verification vulnerability in the lock screen module Impact: Successful exploitation of this vulnerability may affect availability	2024-08-08	4	Medium
CVE-2024-42233	Linux	In the Linux kernel, the following vulnerability has been resolved: filemap: replace pte_offset_map() with pte_offset_map_nolock() The vmf->ptl in filemap_fault_recheck_pte_none() is still set from handle_pte_fault(). But at the same time, we did a pte_unmap(vmf->pte). After a pte_unmap(vmf->pte) unmap and rcu_read_unlock(), the page table may be racily changed and vmf->ptl maybe fails to protect the actual page table. Fix this by replacing pte_offset_map() with pte_offset_map_nolock(). As David said, the PTL pointer might be stale so if we continue to use it infilemap_fault_recheck_pte_none(), it might trigger UAF. Also, if the PTL fails, the issue fixed by commit 58f327f2ce80 ("filemap: avoid unnecessary major faults in filemap_fault()") might reappear.	2024-08-07	3.3	Low
CVE-2024-42249	Linux	In the Linux kernel, the following vulnerability has been resolved: spi: don't unoptimize message in spi_async() Calling spi_maybe_unoptimize_message() in spi_async() is wrong because the message is likely to be in the queue and not transferred yet. This can corrupt the message while it is being used by the controller driver. spi_maybe_unoptimize_message() is already called in the correct place in spi_finalize_current_message() to balance the call to spi_maybe_optimize_message() in spi_async().	2024-08-07	3.3	Low
CVE-2024-6996	Google	Race in Frames in Google Chrome prior to 127.0.6533.72 allowed a remote attacker who convinced a user to engage in specific UI gestures to perform UI spoofing via a crafted HTML page. (Chromium security severity: Medium)	2024-08-06	3.1	Low
CVE-2024-42036	Huawei	Access permission verification vulnerability in the Notepad module Impact: Successful exploitation of this vulnerability may affect service confidentiality.	2024-08-08	2.5	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.