```
commit 0283cbd1473380d3ae46f653eac128aebb288423
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Fri Jul 22 10:21:59 2022 +0200

    Linux 5.18.13

    Link: https://lore.kernel.org/r/20220719114714.247441733@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Tested-by: Sudip Mukherjee <sudip.mukherjee@codethink.co.uk>
    Link: https://lore.kernel.org/r/20220721180421.921249751@linuxfoundation.org
    Link: https://lore.kernel.org/r/20220721182818.743726259@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6e3a6dddad55010a893b07e4e9f2ba30ff95b6d8
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Fri Jun 17 14:10:27 2022 +0200

    drm/aperture: Run fbdev removal before internal helpers

    commit bf43e4521ff3223a613f3a496991a22a4d78e04b upstream.

    Always run fbdev removal first to remove simpledrm via
    sysfb_disable(). This clears the internal state. The later call
    to drm_aperture_detach_drivers() then does nothing. Otherwise,
    with drm_aperture_detach_drivers() running first, the call to
    sysfb_disable() uses inconsistent state.

    Example backtrace show below:

    [   11.663422] ==================================================================
    [   11.663426] BUG: KASAN: use-after-free in device_del+0x79/0x5f0
    [   11.663435] Read of size 8 at addr ffff888108185050 by task systemd-udevd/311
    [   11.663440] CPU: 0 PID: 311 Comm: systemd-udevd Tainted: G            E     5
        .19.0-rc2-1-default+ #1689
    [   11.663445] Hardware name: HP ProLiant DL120 G7, BIOS J01 04/21/2011
    [   11.663447] Call Trace:
    [   11.663449]  <TASK>
    [   11.663451]  ? device_del+0x79/0x5f0
    [   11.663456]  dump_stack_lvl+0x5b/0x73
    [   11.663462]  print_address_description.constprop.0+0x1f/0x1b0
    [   11.663468]  ? device_del+0x79/0x5f0
    [   11.663471]  ? device_del+0x79/0x5f0
    [   11.663475]  print_report.cold+0x3c/0x21c
    [   11.663481]  ? lock_acquired+0x87/0x1e0
    [   11.663484]  ? lock_acquired+0x87/0x1e0
    [   11.663489]  ? device_del+0x79/0x5f0
    [   11.663492]  kasan_report+0xbf/0xf0
    [   11.663498]  ? device_del+0x79/0x5f0
    [   11.663503]  device_del+0x79/0x5f0
    [   11.663509]  ? device_remove_attrs+0x170/0x170
    [   11.663514]  ? lock_is_held_type+0xe8/0x140
    [   11.663523]  platform_device_del.part.0+0x19/0xe0
    [   11.663530]  platform_device_unregister+0x1c/0x30
    [   11.663535]  sysfb_disable+0x2d/0x70
    [   11.663540]  remove_conflicting_framebuffers+0x1c/0xf0
```

```
     [   11.663546]  remove_conflicting_pci_framebuffers+0x130/0x1a0
     [   11.663554]  drm_aperture_remove_conflicting_pci_framebuffers+0x86/0xb0
     [   11.663561]  ? mgag200_pci_remove+0x30/0x30 [mgag200]
     [   11.663578]  mgag200_pci_probe+0x2d/0x140 [mgag200]

     Reported-by: Zack Rusin <zackr@vmware.com>
     Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
     Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
     Reviewed-by: Zack Rusin <zackr@vmware.com>
     Fixes: ee7a69aa38d8 ("fbdev: Disable sysfb device registration when removing conflicting
 FBs")
     Cc: Javier Martinez Canillas <javierm@redhat.com>
     Cc: Daniel Vetter <daniel.vetter@ffwll.ch>
     Cc: Daniel Vetter <daniel@ffwll.ch>
     Cc: Sam Ravnborg <sam@ravnborg.org>
     Cc: Helge Deller <deller@gmx.de>
     Cc: Thomas Zimmermann <tzimmermann@suse.de>
     Cc: Alex Deucher <alexander.deucher@amd.com>
     Cc: Zhen Lei <thunder.leizhen@huawei.com>
     Cc: Changcheng Deng <deng.changcheng@zte.com.cn>
     Link: https://patchwork.freedesktop.org/patch/msgid/20220617121027.30273-1-
 tzimmermann@suse.de
     (cherry picked from commit fb84efa28a48e30b87fa1122e8aab8016c7347cd)
     Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 7642e42a64b09177060c4712fbbcc1e1882fb98e
 Author: Matthew Auld <matthew.auld@intel.com>
 Date:   Tue Jul 12 18:40:50 2022 +0100

     drm/i915/ttm: fix 32b build

     commit ced7866db39fc5c59ee05e154d4abc0977a17f6b upstream.

     Since segment_pages is no longer a compile time constant, it looks the
     DIV_ROUND_UP(node->size, segment_pages) breaks the 32b build. Simplest
     is just to use the ULL variant, but really we should need not need more
     than u32 for the page alignment (also we are limited by that due to the
     sg->length type), so also make it all u32.

     Reported-by: Ville Syrjälä <ville.syrjala@linux.intel.com>
     Fixes: aff1e0b09b54 ("drm/i915/ttm: fix sg_table construction")
     Signed-off-by: Matthew Auld <matthew.auld@intel.com>
     Cc: Nirmoy Das <nirmoy.das@linux.intel.com>
     Reviewed-by: Nirmoy Das <nirmoy.das@intel.com>
     Link: https://patchwork.freedesktop.org/patch/msgid/20220712174050.592550-1-
 matthew.auld@intel.com
     (cherry picked from commit 9306b2b2dfce6931241ef804783692cee526599c)
     Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 99c8a3c41e0791c39482f34ee8ca3f7014ee2e3f
 Author: Juergen Gross <jgross@suse.com>
 Date:   Fri Jul 8 15:14:56 2022 +0200

     x86/pat: Fix x86_has_pat_wp()

     commit 230ec83d4299b30c51a1c133b4f2a669972cc08a upstream.

     x86_has_pat_wp() is using a wrong test, as it relies on the normal
     PAT configuration used by the kernel. In case the PAT MSR has been
     setup by another entity (e.g. Xen hypervisor) it might return false
     even if the PAT configuration is allowing WP mappings. This due to the
     fact that when running as Xen PV guest the PAT MSR is setup by the
     hypervisor and cannot be changed by the guest. This results in the WP
```

```
    related entry to be at a different position when running as Xen PV
    guest compared to the bare metal or fully virtualized case.

    The correct way to test for WP support is:

    1. Get the PTE protection bits needed to select WP mode by reading
       __cachemode2pte_tbl[_PAGE_CACHE_MODE_WP] (depending on the PAT MSR
       setting this might return protection bits for a stronger mode, e.g.
       UC-)
    2. Translate those bits back into the real cache mode selected by those
       PTE bits by reading __pte2cachemode_tbl[__pte2cm_idx(prot)]
    3. Test for the cache mode to be _PAGE_CACHE_MODE_WP

    Fixes: f88a68facd9a ("x86/mm: Extend early_memremap() support with additional attrs")
    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Cc: <stable@vger.kernel.org> # 4.14
    Link: https://lore.kernel.org/r/20220503132207.17234-1-jgross@suse.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec02687a33e7ef731afbf77c7e60a7561e9cb441
Author: Pali Rohár <pali@kernel.org>
Date:   Tue Jun 28 12:09:22 2022 +0200

    serial: mvebu-uart: correctly report configured baudrate value

    commit 4f532c1e25319e42996ec18a1f473fd50c8e575d upstream.

    Functions tty_termios_encode_baud_rate() and uart_update_timeout() should
    be called with the baudrate value which was set to hardware. Linux then
    report exact values via ioctl(TCGETS2) to userspace.

    Change mvebu_uart_baud_rate_set() function to return baudrate value which
    was set to hardware and propagate this value to above mentioned functions.

    With this change userspace would see precise value in termios c_ospeed
    field.

    Fixes: 68a0db1d7da2 ("serial: mvebu-uart: add function to change baudrate")
    Cc: stable <stable@kernel.org>
    Reviewed-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Link: https://lore.kernel.org/r/20220628100922.10717-1-pali@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5df66302f03f87ae8953785a882d78e911f00c55
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Wed Jun 29 12:48:41 2022 +0300

    serial: 8250: Fix PM usage_count for console handover

    commit f9b11229b79c0fb2100b5bb4628a101b1d37fbf6 upstream.

    When console is enabled, univ8250_console_setup() calls
    serial8250_console_setup() before .dev is set to uart_port. Therefore,
    it will not call pm_runtime_get_sync(). Later, when the actual driver
    is going to take over univ8250_console_exit() is called. As .dev is
    already set, serial8250_console_exit() makes pm_runtime_put_sync() call
    with usage count being zero triggering PM usage count warning
    (extra debug for univ8250_console_setup(), univ8250_console_exit(), and
    serial8250_register_ports()):

    [    0.068987] univ8250_console_setup ttyS0 nodev
    [    0.499670] printk: console [ttyS0] enabled
```

```
     [    0.717955] printk: console [ttyS0] printing thread started
     [    1.960163] serial8250_register_ports assigned dev for ttyS0
     [    1.976830] printk: console [ttyS0] disabled
     [    1.976888] printk: console [ttyS0] printing thread stopped
     [    1.977073] univ8250_console_exit ttyS0 usage:0
     [    1.977075] serial8250 serial8250: Runtime PM usage count underflow!
     [    1.977429] dw-apb-uart.6: ttyS0 at MMIO 0x4010006000 (irq = 33, base_baud = 115200)
 is a 16550A
     [    1.977812] univ8250_console_setup ttyS0 usage:2
     [    1.978167] printk: console [ttyS0] printing thread started
     [    1.978203] printk: console [ttyS0] enabled

     To fix the issue, call pm_runtime_get_sync() in
     serial8250_register_ports() as soon as .dev is set for an uart_port
     if it has console enabled.

     This problem became apparent only recently because 82586a721595 ("PM:
     runtime: Avoid device usage count underflows") added the warning
     printout. I confirmed this problem also occurs with v5.18 (w/o the
     warning printout, obviously).

     Fixes: bedb404e91bb ("serial: 8250_port: Don't use power management for kernel console")
     Cc: stable <stable@kernel.org>
     Tested-by: Tony Lindgren <tony@atomide.com>
     Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
     Reviewed-by: Tony Lindgren <tony@atomide.com>
     Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
     Link: https://lore.kernel.org/r/b4f428e9-491f-daf2-2232-819928dc276e@linux.intel.com
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit b2f63fefefc5c1cba4e192949751a67bc166a671
 Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
 Date:   Tue Jun 14 10:56:37 2022 +0300

     serial: pl011: UPSTAT_AUTORTS requires .throttle/unthrottle

     commit 211565b100993c90b53bf40851eacaefc830cfe0 upstream.

     The driver must provide throttle and unthrottle in uart_ops when it
     sets UPSTAT_AUTORTS. Add them using existing stop_rx &
     enable_interrupts functions.

     Fixes: 2a76fa283098 (serial: pl011: Adopt generic flag to store auto RTS status)
     Cc: stable <stable@kernel.org>
     Cc: Lukas Wunner <lukas@wunner.de>
     Reported-by: Nuno Gonçalves <nunojpg@gmail.com>
     Tested-by: Nuno Gonçalves <nunojpg@gmail.com>
     Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
     Link: https://lore.kernel.org/r/20220614075637.8558-1-ilpo.jarvinen@linux.intel.com
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 484ec3ca92e47d7fb7bad005ddc66d8965485ad9
 Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
 Date:   Mon Jun 27 18:07:52 2022 +0300

     serial: stm32: Clear prev values before setting RTS delays

     commit 5c5f44e36217de5ead789ff25da71c31c2331c96 upstream.

     The code lacks clearing of previous DEAT/DEDT values. Thus, changing
     values on the fly results in garbage delays tending towards the maximum
     value as more and more bits are ORed together. (Leaving RS485 mode
     would have cleared the old values though).

     Fixes: 1bcda09d2910 ("serial: stm32: add support for RS485 hardware control mode")
```

```
    Cc: stable <stable@kernel.org>
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Link: https://lore.kernel.org/r/20220627150753.34510-1-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a762cee5d933fe4e2e1b773d60fc74fb8248d8c4
Author: Dorian Rudolph <mail@dorianrudolph.com>
Date:   Sat May 14 17:23:40 2022 +0200

    power: supply: core: Fix boundary conditions in interpolation

    commit 093d27bb6f2d1963f927ef59c9a2d37059175426 upstream.

    The functions power_supply_temp2resist_simple and power_supply_ocv2cap_simple
    handle boundary conditions incorrectly.
    The change was introduced in a4585ba2050f460f749bbaf2b67bd56c41e30283
    ("power: supply: core: Use library interpolation").
    There are two issues: First, the lines "high = i - 1" and "high = i" in ocv2cap
    have the wrong order compared to temp2resist. As a consequence, ocv2cap
    sets high=-1 if ocv>table[0].ocv, which causes an out-of-bounds read.
    Second, the logic of temp2resist is also not correct.
    Consider the case table[] = {{20, 100}, {10, 80}, {0, 60}}.
    For temp=5, we expect a resistance of 70% by interpolation.
    However, temp2resist sets high=low=2 and returns 60.

    Cc: stable@vger.kernel.org
    Signed-off-by: Dorian Rudolph <mail@dorianrudolph.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Fixes: a4585ba2050f ("power: supply: core: Use library interpolation")
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dda20f46124a80e3c832a3aa024befd998f847e1
Author: Yi Yang <yiyang13@huawei.com>
Date:   Tue Jun 28 16:35:15 2022 +0800

    serial: 8250: fix return error code in serial8250_request_std_resource()

    commit 6e690d54cfa802f939cefbd2fa2c91bd0b8bd1b6 upstream.

    If port->mapbase = NULL in serial8250_request_std_resource() , it need
    return a error code instead of 0. If uart_set_info() fail to request new
    regions by serial8250_request_std_resource() but the return value of
    serial8250_request_std_resource() is 0, The system incorrectly considers
    that the resource application is successful and does not attempt to
    restore the old setting. A null pointer reference is triggered when the
    port resource is later invoked.

    Signed-off-by: Yi Yang <yiyang13@huawei.com>
    Cc: stable <stable@kernel.org>
    Link: https://lore.kernel.org/r/20220628083515.64138-1-yiyang13@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 89c3c9dd7e0331d4e2529e91f273bc31e3c04f33
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Wed Jul 13 12:53:46 2022 -0500

    ACPI: CPPC: Fix enabling CPPC on AMD systems with shared memory

    commit fbd74d16890b9f5d08ea69b5282b123c894f8860 upstream.

    When commit 72f2ecb7ece7 ("ACPI: bus: Set CPPC _OSC bits for all
    and when CPPC_LIB is supported") was introduced, we found collateral
    damage that a number of AMD systems that supported CPPC but
    didn't advertise support in _OSC stopped having a functional
```

```
    amd-pstate driver. The _OSC was only enforced on Intel systems at that
    time.

    This was fixed for the MSR based designs by commit 8b356e536e69f
    ("ACPI: CPPC: Don't require _OSC if X86_FEATURE_CPPC is supported")
    but some shared memory based designs also support CPPC but haven't
    advertised support in the _OSC.  Add support for those designs as well by
    hardcoding the list of systems.

    Fixes: 72f2ecb7ece7 ("ACPI: bus: Set CPPC _OSC bits for all and when CPPC_LIB is
supported")
    Fixes: 8b356e536e69f ("ACPI: CPPC: Don't require _OSC if X86_FEATURE_CPPC is supported")
    Link: https://lore.kernel.org/all/3559249.JlDtxWtqDm@natalenko.name/
    Cc: 5.18+ <stable@vger.kernel.org> # 5.18+
    Reported-and-tested-by: Oleksandr Natalenko <oleksandr@natalenko.name>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ceb7a7e7ab0640f6f75d7eaa7611aabed7981c35
Author: Tony Krowiak <akrowiak@linux.ibm.com>
Date:   Wed Jul 6 17:43:29 2022 -0400

    s390/ap: fix error handling in __verify_queue_reservations()

    commit 2f23256c0ea20627c91ea2d468cda945f68c3395 upstream.

    The AP bus's __verify_queue_reservations function increments the ref count
    for the device driver passed in as a parameter, but fails to decrement it
    before returning control to the caller. This will prevents any subsequent
    removal of the module.

    Signed-off-by: Tony Krowiak <akrowiak@linux.ibm.com>
    Reported-by: Tony Krowiak <akrowiak@linux.ibm.com>
    Reviewed-by: Harald Freudenberger <freude@linux.ibm.com>
    Fixes: 4f8206b88286 ("s390/ap: driver callback to indicate resource in use")
    Link: https://lore.kernel.org/r/20220706222619.602094-1-akrowiak@linux.ibm.com
    Cc: stable@vger.kernel.org
    [agordeev@linux.ibm.com fixed description, added Fixes and Link]
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 14d2cc21ca622310babf373e3a8f0b40acfe8265
Author: Yangxi Xiang <xyangxi5@gmail.com>
Date:   Tue Jun 28 17:33:22 2022 +0800

    vt: fix memory overlapping when deleting chars in the buffer

    commit 39cdb68c64d84e71a4a717000b6e5de208ee60cc upstream.

    A memory overlapping copy occurs when deleting a long line. This memory
    overlapping copy can cause data corruption when scr_memcpyw is optimized
    to memcpy because memcpy does not ensure its behavior if the destination
    buffer overlaps with the source buffer. The line buffer is not always
    broken, because the memcpy utilizes the hardware acceleration, whose
    result is not deterministic.

    Fix this problem by using replacing the scr_memcpyw with scr_memmovew.

    Fixes: 81732c3b2fed ("tty vt: Fix line garbage in virtual console on command line
edition")
    Cc: stable <stable@kernel.org>
    Signed-off-by: Yangxi Xiang <xyangxi5@gmail.com>
    Link: https://lore.kernel.org/r/20220628093322.5688-1-xyangxi5@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit f6c04bde97baf7a0a62a7b4954b223db404f9757
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Wed Jul 13 23:11:01 2022 +0200

    ACPI: video: Fix acpi_video_handles_brightness_key_presses()

    [ Upstream commit 5ad26161a371e4aa2d2553286f0cac580987a493 ]

    Commit 3a0cf7ab8df3 ("ACPI: video: Change how we determine if brightness
    key-presses are handled") made acpi_video_handles_brightness_key_presses()
    report false when none of the ACPI Video Devices support backlight control.

    But it turns out that at least on a Dell Inspiron N4010 there is no ACPI
    backlight control, yet brightness hotkeys are still reported through
    the ACPI Video Bus; and since acpi_video_handles_brightness_key_presses()
    now returns false, brightness keypresses are now reported twice.

    To fix this rename the has_backlight flag to may_report_brightness_keys and
    also set it the first time a brightness key press event is received.

    Depending on the delivery of the other ACPI (WMI) event vs the ACPI Video
    Bus event this means that the first brightness key press might still get
    reported twice, but all further keypresses will be filtered as before.

    Note that this relies on other drivers reporting brightness key events
    calling acpi_video_handles_brightness_key_presses() when delivering
    the events (rather then once during driver probe). This is already
    required and documented in include/acpi/video.h:

    /*
     * Note: The value returned by acpi_video_handles_brightness_key_presses()
     * may change over time and should not be cached.
     */

    Fixes: 3a0cf7ab8df3 ("ACPI: video: Change how we determine if brightness key-presses are
handled")
    Link:
https://lore.kernel.org/regressions/CALF=6jEe5G8+r1Wo0vvz4GjNQQhdkLT5p8uCHn6ZXhg4nsOWow@mail
.gmail.com/
    Reported-and-tested-by: Ben Greening <bgreening@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Link: https://lore.kernel.org/r/20220713211101.85547-2-hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db738f76b03beaccfe9de35d838d9ec825226478
Author: Linus Torvalds <torvalds@linux-foundation.org>
Date:   Wed Jul 6 12:20:59 2022 -0700

    signal handling: don't use BUG_ON() for debugging

    [ Upstream commit a382f8fee42ca10c9bfce0d2352d4153f931f5dc ]

    These are indeed "should not happen" situations, but it turns out recent
    changes made the 'task_is_stopped_or_trace()' case trigger (fix for that
    exists, is pending more testing), and the BUG_ON() makes it
    unnecessarily hard to actually debug for no good reason.

    It's been that way for a long time, but let's make it clear: BUG_ON() is
    not good for debugging, and should never be used in situations where you
    could just say "this shouldn't happen, but we can continue".

    Use WARN_ON_ONCE() instead to make sure it gets logged, and then just
    continue running.  Instead of making the system basically unusable
```

```
    because you crashed the machine while potentially holding some very core
    locks (eg this function is commonly called while holding 'tasklist_lock'
    for writing).

    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 700364130bfa1689baa9d4385a92a7f122093d1e
Author: Jason A. Donenfeld <Jason@zx2c4.com>
Date:   Thu Jul 7 02:31:54 2022 +0200

    wireguard: selftests: always call kernel makefile

    [ Upstream commit 1a087eec257154e26a81a7a0a15380d7a2431765 ]

    These selftests are used for much more extensive changes than just the
    wireguard source files. So always call the kernel's build file, which
    will do something or nothing after checking the whole tree, per usual.

    Signed-off-by: Jason A. Donenfeld <Jason@zx2c4.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d2f02e532200add4dba5ff47d44c218bb9d214ad
Author: Jason A. Donenfeld <Jason@zx2c4.com>
Date:   Thu Jul 7 02:31:52 2022 +0200

    wireguard: selftests: set fake real time in init

    [ Upstream commit 829be057dbc1e71383b8d7de8edb31dcf07b4aa0 ]

    Not all platforms have an RTC, and rather than trying to force one into
    each, it's much easier to just set a fixed time. This is necessary
    because WireGuard's latest handshakes parameter is returned in wallclock
    time, and if the system time isn't set, and the system is really fast,
    then this returns 0, which trips the test.

    Turning this on requires setting CONFIG_COMPAT_32BIT_TIME=y, as musl
    doesn't support settimeofday without it.

    Signed-off-by: Jason A. Donenfeld <Jason@zx2c4.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3d8b35387e01cab217dc4691a6f770cbb6ed852c
Author: Keith Busch <kbusch@kernel.org>
Date:   Wed Jul 6 10:05:05 2022 -0700

    nvme: use struct group for generic command dwords

    [ Upstream commit 5c629dc9609dc43492a7bc8060cc6120875bf096 ]

    This will allow the trace event to know the full size of the data
    intended to be copied and silence read overflow checks.

    Reported-by: John Garry <john.garry@huawei.com>
    Suggested-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Keith Busch <kbusch@kernel.org>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97f9cba7686e395669ffb2808c4a7b140bfdb5bb
Author: Keith Busch <kbusch@kernel.org>
Date:   Tue Jul 5 10:21:02 2022 -0700
```

```
    nvme-pci: phison e16 has bogus namespace ids

    [ Upstream commit 73029c9b23cf1213e5f54c2b59efce08665199e7 ]

    Add the quirk.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216049
    Reported-by: Chris Egolf <cegolf@ugholf.net>
    Signed-off-by: Keith Busch <kbusch@kernel.org>
    Reviewed-by: Chaitanya Kulkarni <kch@nvidia.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7bd54d31155e3da12eb357f4a56d2612c942c4db
Author: Egor Vorontsov <sdoregor@sdore.me>
Date:   Mon Jun 27 13:00:35 2022 +0300

    ALSA: usb-audio: Add quirk for Fiero SC-01 (fw v1.0.0)

    [ Upstream commit 2307a0e1ca0b5c1337b37ac6302f96e017ebac3c ]

    The patch applies the same quirks used for SC-01 at firmware v1.1.0 to
    the ones running v1.0.0, with respect to hard-coded sample rates.

    I got two more units and successfully tested the patch series with both
    firmwares.

    The support is now complete (not accounting ASIO).

    Signed-off-by: Egor Vorontsov <sdoregor@sdore.me>
    Link: https://lore.kernel.org/r/20220627100041.2861494-2-sdoregor@sdore.me
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit da7fdaadc13506a902667fb98cfb584d70b10e0a
Author: Egor Vorontsov <sdoregor@sdore.me>
Date:   Mon Jun 27 13:00:34 2022 +0300

    ALSA: usb-audio: Add quirk for Fiero SC-01

    [ Upstream commit 4fb7c24f69c48fdc02ea7858dbd5a60ff08bf7e5 ]

    Fiero SC-01 is a USB sound card with two mono inputs and a single
    stereo output. The inputs are composed into a single stereo stream.

    The device uses a vendor-provided driver on Windows and does not work
    at all without it. The driver mostly provides ASIO functionality, but
    also alters the way the sound card is queried for sample rates and
    clocks.

    ALSA queries those failing with an EPIPE (same as Windows 10 does).
    Presumably, the vendor-provided driver does not query it at all, simply
    matching by VID:PID. Thus, I consider this a buggy firmware and adhere
    to a set of fixed endpoint quirks instead.

    The soundcard has an internal clock. Implicit feedback mode is required
    for the playback.

    I have updated my device to v1.1.0 from a Windows 10 VM using a vendor-
    provided binary prior to the development, hoping for it to just begin
    working. The device provides no obvious way to downgrade the firmware,
    and regardless, there's no binary available for v1.0.0 anyway.

    Thus, I will be getting another unit to extend the patch with support
    for that. Expected to be a simple copy-paste of the existing one,
```

```
    though.

    There were no previous reports of that device in context of Linux
    anywhere. Other issues have been reported though, but that's out of the
    scope.

    Signed-off-by: Egor Vorontsov <sdoregor@sdore.me>
    Link: https://lore.kernel.org/r/20220627100041.2861494-1-sdoregor@sdore.me
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b47d6dfaecb60c490a7c43d896e101e85490257
Author: John Veness <john-linux@pelago.org.uk>
Date:   Fri Jun 24 15:07:57 2022 +0100

    ALSA: usb-audio: Add quirks for MacroSilicon MS2100/MS2106 devices

    [ Upstream commit 6e2c9105e0b743c92a157389d40f00b81bdd09fe ]

    Treat the claimed 96kHz 1ch in the descriptors as 48kHz 2ch, so that
    the audio stream doesn't sound mono. Also fix initial stream
    alignment, so that left and right channels are in the correct order.

    Signed-off-by: John Veness <john-linux@pelago.org.uk>
    Link: https://lore.kernel.org/r/20220624140757.28758-1-john-linux@pelago.org.uk
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 449f51b5d6e10eeda2a59d1319334e5853e7cb3a
Author: Srinivas Neeli <srinivas.neeli@xilinx.com>
Date:   Thu Jun 9 13:54:32 2022 +0530

    Revert "can: xilinx_can: Limit CANFD brp to 2"

    [ Upstream commit c6da4590fe819dfe28a4f8037a8dc1e056542fb4 ]

    This reverts commit 05ca14fdb6fe65614e0652d03e44b02748d25af7.

    On early silicon engineering samples observed bit shrinking issue when
    we use brp as 1. Hence updated brp_min as 2. As in production silicon
    this issue is fixed, so reverting the patch.

    Link: https://lore.kernel.org/all/20220609082433.1191060-2-srinivas.neeli@xilinx.com
    Signed-off-by: Srinivas Neeli <srinivas.neeli@xilinx.com>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2764762d0d6a2fab76f3f68b5298e4ab634053b7
Author: Gabriel Fernandez <gabriel.fernandez@foss.st.com>
Date:   Fri Jun 24 11:27:13 2022 +0200

    ARM: dts: stm32: use the correct clock source for CEC on stm32mp151

    [ Upstream commit 78ece8cce1ba0c3f3e5a7c6c1b914b3794f04c44 ]

    The peripheral clock of CEC is not LSE but CEC.

    Signed-off-by: Gabriel Fernandez <gabriel.fernandez@foss.st.com>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bd77b8298f851916f4c9fd555dbc65298a4750b6
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun Jun 26 09:43:15 2022 +0200
```

```
    soc: ixp4xx/npe: Fix unused match warning

    [ Upstream commit 620f83b8326ce9706b1118334f0257ae028ce045 ]

    The kernel test robot found this inconsistency:

      drivers/soc/ixp4xx/ixp4xx-npe.c:737:34: warning:
      'ixp4xx_npe_of_match' defined but not used [-Wunused-const-variable=]
         737 | static const struct of_device_id ixp4xx_npe_of_match[] = {

    This is because the match is enclosed in the of_match_ptr()
    which compiles into NULL when OF is disabled and this
    is unnecessary.

    Fix it by dropping of_match_ptr() around the match.

    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20220626074315.61209-1-linus.walleij@linaro.org'
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2334bdfc2da469c9807767002a2831274b82c39a
Author: Juergen Gross <jgross@suse.com>
Date:   Thu Jun 30 09:14:40 2022 +0200

    x86: Clear .brk area at early boot

    [ Upstream commit 38fa5479b41376dc9d7f57e71c83514285a25ca0 ]

    The .brk section has the same properties as .bss: it is an alloc-only
    section and should be cleared before being used.

    Not doing so is especially a problem for Xen PV guests, as the
    hypervisor will validate page tables (check for writable page tables
    and hypervisor private bits) before accepting them to be used.

    Make sure .brk is initially zero by letting clear_bss() clear the brk
    area, too.

    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20220630071441.28576-3-jgross@suse.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4dc6fad52c5ca66ded0c5c00b2a25dc6c1e6a793
Author: Jacky Bai <ping.bai@nxp.com>
Date:   Mon Jun 13 11:18:54 2022 +0800

    pinctrl: imx: Add the zero base flag for imx93

    [ Upstream commit fbc24ebc65507feb9728dc38197f90486148dda0 ]

    On i.MX93, the pin mux reg offset is from 0x0,
    so need to add the 'ZERO_OFFSET_VALID' flag to make
    sure the pin at mux offset 0 can be found.

    Signed-off-by: Jacky Bai <ping.bai@nxp.com>
    Link: https://lore.kernel.org/r/20220613031854.1571357-1-ping.bai@nxp.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d6ef5aca6fe46c48f4d3b3c53d1fd9b5b86c07b4
Author: Stafford Horne <shorne@gmail.com>
Date:   Wed Jun 15 08:54:26 2022 +0900
```

```
    irqchip: or1k-pic: Undefine mask_ack for level triggered hardware

    [ Upstream commit 8520501346ed8d1c4a6dfa751cb57328a9c843f1 ]

    The mask_ack operation clears the interrupt by writing to the PICSR
    register.  This we don't want for level triggered interrupt because
    it does not actually clear the interrupt on the source hardware.

    This was causing issues in qemu with multi core setups where
    interrupts would continue to fire even though they had been cleared in
    PICSR.

    Just remove the mask_ack operation.

    Acked-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Stafford Horne <shorne@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7f97af49f6013c684192f63720d6fa81c974874f
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Thu Jun 23 11:51:19 2022 +0100

    ASoC: madera: Fix event generation for rate controls

    [ Upstream commit 980555e95f7cabdc9c80a07107622b097ba23703 ]

    madera_adsp_rate_put always returns zero regardless of if the control
    value was updated. This results in missing notifications to user-space
    of the control change. Update the handling to return 1 when the
    value is changed.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220623105120.1981154-5-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0303f460b85680f8d49ddfd9b7dbc5876e24f5c7
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Thu Jun 23 11:51:18 2022 +0100

    ASoC: madera: Fix event generation for OUT1 demux

    [ Upstream commit e3cabbef3db8269207a6b8808f510137669f8deb ]

    madera_out1_demux_put returns the value of
    snd_soc_dapm_mux_update_power, which returns a 1 if a path was found for
    the kcontrol. This is obviously different to the expected return a 1 if
    the control was updated value. This results in spurious notifications to
    user-space. Update the handling to only return a 1 when the value is
    changed.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220623105120.1981154-4-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2308fdb5191ad292d626502a9c99ca7da8bf3b90
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Thu Jun 23 11:51:17 2022 +0100

    ASoC: cs47l15: Fix event generation for low power mux control

    [ Upstream commit 7f103af4a10f375b9b346b4d0b730f6a66b8c451 ]

    cs47l15_in1_adc_put always returns zero regardless of if the control
```

```
    value was updated. This results in missing notifications to user-space
    of the control change. Update the handling to return 1 when the value is
    changed.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220623105120.1981154-3-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14937a06ab710601906fcc71fe73fcd5f20e376c
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Thu Jun 23 11:51:16 2022 +0100

    ASoC: cs35l41: Add ASP TX3/4 source to register patch

    [ Upstream commit 46b0d050c8c7df6dfb2c376aaa149bf2cfc5ca3e ]

    The mixer controls for ASP TX3/4 are set to values that are not included
    in their enumeration control. This will cause spurious event
    notifications when the controls are first changed, as the register value
    changes whilst the actual visible enumeration value does not. Use the
    register patch to set them to a known value, zero, which equates to zero
    fill, thereby avoiding the spurious notifications.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220623105120.1981154-2-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed5ddd50ffb5dbbb51ec4b525baf069991eebb4b
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Thu Jun 23 11:51:15 2022 +0100

    ASoC: dapm: Initialise kcontrol data for mux/demux controls

    [ Upstream commit 11d7a12f7f50baa5af9090b131c9b03af59503e7 ]

    DAPM keeps a copy of the current value of mux/demux controls,
    however this value is only initialised in the case of autodisable
    controls. This leads to false notification events when first
    modifying a DAPM kcontrol that has a non-zero default.

    Autodisable controls are left as they are, since they already
    initialise the value, and there would be more work required to
    support autodisable muxes where the first option isn't disabled
    and/or that isn't the default.

    Technically this issue could affect mixer/switch elements as well,
    although not on any of the devices I am currently running. There
    is also a little more work to do to address the issue there due to
    that side supporting stereo controls, so that has not been tackled
    in this patch.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220623105120.1981154-1-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7bb71133cae88d3003a3490b97864af76533072b
Author: Shuming Fan <shumingf@realtek.com>
Date:   Tue Jun 21 17:07:19 2022 +0800

    ASoC: rt711-sdca: fix kernel NULL pointer dereference when IO error

    [ Upstream commit 1df793d479bef546569fc2e409ff8bb3f0fb8e99 ]
```

```
    The initial settings will be written before the codec probe function.
    But, the rt711->component doesn't be assigned yet.
    If IO error happened during initial settings operations, it will cause the kernel panic.
    This patch changed component->dev to slave->dev to fix this issue.

    Signed-off-by: Shuming Fan <shumingf@realtek.com>
    Link: https://lore.kernel.org/r/20220621090719.30558-1-shumingf@realtek.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3f5f2a9e2bbccc6f4e3a0d2e588a02703cdfd7f
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Tue Jun 21 11:20:40 2022 +0100

    ASoC: cs35l41: Correct some control names

    [ Upstream commit c6a5f22f9b4fd5f21414be690ce34046d9712f05 ]

    Various boolean controls on cs35l41 are missing the required "Switch" in
    the name, add these.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220621102041.1713504-3-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 83431e500b2c5d5c95712053c1cea2f462a4366e
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Tue Jun 21 11:20:39 2022 +0100

    ASoC: wm5110: Fix DRE control

    [ Upstream commit 0bc0ae9a5938d512fd5d44f11c9c04892dcf4961 ]

    The DRE controls on wm5110 should return a value of 1 if the DRE state
    is actually changed, update to fix this.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220621102041.1713504-2-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89c2c0dd28250e1def6fdf37af311a02edb4d702
Author: Charles Keepax <ckeepax@opensource.cirrus.com>
Date:   Tue Jun 21 11:20:38 2022 +0100

    ASoC: wm_adsp: Fix event for preloader

    [ Upstream commit 9896c029f0df628c6cb108253d09b1d61f1d4a88 ]

    The preloader controls on ADSP should return a value of 1 if the
    preloader value was changed, update to correct this.

    Signed-off-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220621102041.1713504-1-ckeepax@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03f97b638dd9a8eb8c7e63b243bcef12c99264a9
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Jun 12 17:56:52 2022 +0200

    ASoC: Intel: bytcr_wm5102: Fix GPIO related probe-ordering problem
```

```
    [ Upstream commit 4e07479eab8a044cc9542414ccb4aeb8eb033bde ]

    The "wlf,spkvdd-ena" GPIO needed by the bytcr_wm5102 driver
    is made available through a gpio-lookup table.

    This gpio-lookup table is registered by drivers/mfd/arizona-spi.c, which
    may get probed after the bytcr_wm5102 driver.

    If the gpio-lookup table has not registered yet then the gpiod_get()
    will return -ENOENT. Treat -ENOENT as -EPROBE_DEFER to still keep
    things working in this case.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Acked-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20220612155652.107310-1-hdegoede@redhat.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3071f04cf125c1d4fb0698b524e1c030c1b0f17
Author: Mark Brown <broonie@kernel.org>
Date:   Fri Jun 3 14:25:26 2022 +0200

    ASoC: wcd938x: Fix event generation for some controls

    [ Upstream commit 10e7ff0047921e32b919ecee7be706dd33c107f8 ]

    Currently wcd938x_*_put() unconditionally report that the value of the
    control changed, resulting in spurious events being generated. Return 0 in
    that case instead as we should. There is still an issue in the compander
    control which is a bit more complex.

    Signed-off-by: Mark Brown <broonie@kernel.org>
    Reported-by: kernel test robot <lkp@intel.com>
    Link: https://lore.kernel.org/r/20220603122526.3914942-1-broonie@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d93ab0d2f6a82661d0bd9b79cc1a557f7a73353
Author: Mark Brown <broonie@kernel.org>
Date:   Fri Jun 3 14:46:09 2022 +0200

    ASoC: wcd9335: Fix spurious event generation

    [ Upstream commit a7786cbae4b2732815da98efa39df96746b5bd0d ]

    The slimbus mux put operation unconditionally reports a change in value
    which means that spurious events are generated. Fix this by exiting early
    in that case.

    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20220603124609.4024666-1-broonie@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e58a32e25ef12810a4684ce31f5d48645fcedec5
Author: Yassine Oudjana <y.oudjana@protonmail.com>
Date:   Mon Jun 6 19:22:26 2022 +0400

    ASoC: wcd9335: Remove RX channel from old list before adding it to a new one

    [ Upstream commit be6dd72edb216f20fc80e426ece9fe9b8aabf033 ]

    Currently in slim_rx_mux_put, an RX channel gets added to a new list
    even if it is already in one. This can mess up links and make either
    it, the new list head, or both, get linked to the wrong entries.
```

```
    This can cause an entry to link to itself which in turn ends up
    making list_for_each_entry in other functions loop infinitely.
    To avoid issues, always remove the RX channel from any list it's in
    before adding it to a new list.

    Signed-off-by: Yassine Oudjana <y.oudjana@protonmail.com>
    Link: https://lore.kernel.org/r/20220606152226.149164-1-y.oudjana@protonmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 79096060364c41b5627facc431c851f1367eb03e
Author: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
Date:   Thu Jun 9 11:59:49 2022 +0300

    ASoC: SOF: Intel: hda-loader: Clarify the cl_dsp_init() flow

    [ Upstream commit bbfef046c6613404c01aeb9e9928bebb78dd327a ]

    Update the comment for the cl_dsp_init() to clarify what is done by the
    function and use the chip->init_core_mask instead of BIT(0) when
    unstalling/running the init core.

    Complements: 2a68ff846164 ("ASoC: SOF: Intel: hda: Revisit IMR boot sequence")
    Signed-off-by: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Link: https://lore.kernel.org/r/20220609085949.29062-4-peter.ujfalusi@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71d199d622c4c9bdf864a9b1f26218efe85613e1
Author: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
Date:   Thu Jun 9 11:59:48 2022 +0300

    ASoC: SOF: Intel: hda-loader: Make sure that the fw load sequence is followed

    [ Upstream commit c31691e0d126ec5d60d2b6b03f699c11b613b219 ]

    The hda_dsp_enable_core() is powering up _and_ unstall the core in one
    call while the first step of the firmware loading  must not unstall the
    core.
    The core can be unstalled only after the set cpb_cfp and the configuration
    of the IPC register for the ROM_CONTROL message.

    Complements: 2a68ff846164 ("ASoC: SOF: Intel: hda: Revisit IMR boot sequence")
    Signed-off-by: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Link: https://lore.kernel.org/r/20220609085949.29062-3-peter.ujfalusi@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d275601a18cf0d98e57acf370221566477f95eb5
Author: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
Date:   Thu Jun 9 11:59:47 2022 +0300

    ASoC: SOF: Intel: hda-dsp: Expose hda_dsp_core_power_up()

    [ Upstream commit 08f8a93198e300dff9649bbae424cd805d313326 ]

    The hda_dsp_core_power_up() needs to be exposed so that it can be used in
    hda-loader.c to correct the boot flow.
    The first step must not unstall the core, it should only power up the
```

```
    core(s).

    Add sanity check for the core_mask while exposing it to be safe.

    Complements: 2a68ff846164 ("ASoC: SOF: Intel: hda: Revisit IMR boot sequence")
    Signed-off-by: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Link: https://lore.kernel.org/r/20220609085949.29062-2-peter.ujfalusi@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 11e35a6c6607216eeaf047ba85251f902f622a72
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:51 2022 -0500

    ASoC: codecs: rt700/rt711/rt711-sdca: initialize workqueues in probe

    [ Upstream commit ba98d7d8b60ba410aa03834f6aa48fd3b2e68478 ]

    The workqueues are initialized in the io_init functions, which isn't
    quite right. In some tests, this leads to warnings throw from
    __queue_delayed_work()

    WARN_ON_FUNCTION_MISMATCH(timer->function, delayed_work_timer_fn);

    Move all the initializations to the probe functions.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-7-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1d75b73ec6d6b705cca528b36d8315e43e8d7fa5
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:50 2022 -0500

    ASoC: rt7*-sdw: harden jack_detect_handler

    [ Upstream commit 0484271ab0ce50649329fa9dc23c50853c5b26a4 ]

    Realtek headset codec drivers typically check if the card is
    instantiated before proceeding with the jack detection.

    The rt700, rt711 and rt711-sdca are however missing a check on the
    card pointer, which can lead to NULL dereferences encountered in
    driver bind/unbind tests.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-6-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 517fa7405a48c702591887bdf378963b0e1e1328
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:49 2022 -0500

    ASoC: rt711: fix calibrate mutex initialization
```

```
    [ Upstream commit 08bb5dc6ce02374169213cea772b1c297eaf32d5 ]

    Follow the same flow as rt711-sdca and initialize all mutexes at probe
    time.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-5-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 09bca0ffc95c50369f1345d80ecfaca51864126f
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:48 2022 -0500

    ASoC: Intel: sof_sdw: handle errors on card registration

    [ Upstream commit fe154c4ff376bc31041c6441958a08243df09c99 ]

    If the card registration fails, typically because of deferred probes,
    the device properties added for headset codecs are not removed, which
    leads to kernel oopses in driver bind/unbind tests.

    We already clean-up the device properties when the card is removed,
    this code can be moved as a helper and called upon card registration
    errors.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-4-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f85a8e86015f2074a79d30b9425eb8b66eaacf60
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:47 2022 -0500

    ASoC: rt711-sdca-sdw: fix calibrate mutex initialization

    [ Upstream commit ed0a7fb29c9fd4f53eeb37d1fe2354df7a038047 ]

    In codec driver bind/unbind test, the following warning is thrown:

    DEBUG_LOCKS_WARN_ON(lock->magic != lock)
    ...
    [  699.182495]  rt711_sdca_jack_init+0x1b/0x1d0 [snd_soc_rt711_sdca]
    [  699.182498]  rt711_sdca_set_jack_detect+0x3b/0x90 [snd_soc_rt711_sdca]
    [  699.182500]  snd_soc_component_set_jack+0x24/0x50 [snd_soc_core]

    A quick check in the code shows that the 'calibrate_mutex' used by
    this driver are not initialized at probe time. Moving the
    initialization to the probe removes the issue.

    BugLink: https://github.com/thesofproject/linux/issues/3644
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-3-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
```

```
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3a2c0631302613f41fd8f7ce9882c633d66b9a5
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 6 15:37:46 2022 -0500

    ASoC: Realtek/Maxim SoundWire codecs: disable pm_runtime on remove

    [ Upstream commit ac63716da3070f8cb6baaba3a058a0c7f22aeb5b ]

    When binding/unbinding codec drivers, the following warnings are
    thrown:

    [ 107.266879] rt715-sdca sdw:3:025d:0714:01: Unbalanced pm_runtime_enable!
    [  306.879700] rt711-sdca sdw:0:025d:0711:01: Unbalanced pm_runtime_enable!

    Add a remove callback for all Realtek/Maxim SoundWire codecs and remove this
    warning.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20220606203752.144159-2-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e162a24f1dd06c0dcae71f2565c9f3da2827b98e
Author: Haowen Bai <baihaowen@meizu.com>
Date:   Thu Apr 21 10:26:59 2022 +0800

    pinctrl: aspeed: Fix potential NULL dereference in aspeed_pinmux_set_mux()

    [ Upstream commit 84a85d3fef2e75b1fe9fc2af6f5267122555a1ed ]

    pdesc could be null but still dereference pdesc->name and it will lead to
    a null pointer access. So we move a null check before dereference.

    Signed-off-by: Haowen Bai <baihaowen@meizu.com>
    Link: https://lore.kernel.org/r/1650508019-22554-1-git-send-email-baihaowen@meizu.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 55ad380852bcb7fe9b70935af0a7d2e8c51166af
Author: Mark Brown <broonie@kernel.org>
Date:   Sat Jun 4 11:52:46 2022 +0100

    ASoC: ops: Fix off by one in range control validation

    [ Upstream commit 5871321fb4558c55bf9567052b618ff0be6b975e ]

    We currently report that range controls accept a range of 0..(max-min) but
    accept writes in the range 0..(max-min+1). Remove that extra +1.

    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20220604105246.4055214-1-broonie@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f22ddc8a5278d7fb6369a0aeb0d8775a0aefaaee
Author: Jianglei Nie <niejianglei2021@163.com>
Date:   Wed Jun 29 15:55:50 2022 +0800

    net: sfp: fix memory leak in sfp_probe()
```

```
    [ Upstream commit 0a18d802d65cf662644fd1d369c86d84a5630652 ]

    sfp_probe() allocates a memory chunk from sfp with sfp_alloc(). When
    devm_add_action() fails, sfp is not freed, which leads to a memory leak.

    We should use devm_add_action_or_reset() instead of devm_add_action().

    Signed-off-by: Jianglei Nie <niejianglei2021@163.com>
    Reviewed-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Link: https://lore.kernel.org/r/20220629075550.2152003-1-niejianglei2021@163.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a8ba4bd0f542a2bdf28751896ebe515e7117d0dc
Author: Ruozhu Li <liruozhu@huawei.com>
Date:   Thu Jun 23 14:45:39 2022 +0800

    nvme: fix regression when disconnect a recovering ctrl

    [ Upstream commit f7f70f4aa09dc43d7455c060143e86a017c30548 ]

    We encountered a problem that the disconnect command hangs.
    After analyzing the log and stack, we found that the triggering
    process is as follows:
    CPU0                            CPU1
                                      nvme_rdma_error_recovery_work
                                        nvme_rdma_teardown_io_queues
    nvme_do_delete_ctrl                   nvme_stop_queues
      nvme_remove_namespaces
      --clear ctrl->namespaces
                                          nvme_start_queues
                                          --no ns in ctrl->namespaces
        nvme_ns_remove                    return(because ctrl is deleting)
          blk_freeze_queue
            blk_mq_freeze_queue_wait
            --wait for ns to unquiesce to clean infligt IO, hang forever

    This problem was not found in older kernels because we will flush
    err work in nvme_stop_ctrl before nvme_remove_namespaces.It does not
    seem to be modified for functional reasons, the patch can be revert
    to solve the problem.

    Revert commit 794a4cb3d2f7 ("nvme: remove the .stop_ctrl callout")

    Signed-off-by: Ruozhu Li <liruozhu@huawei.com>
    Reviewed-by: Sagi Grimberg <sagi@grimberg.me>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 569d1c493070eca57425c1db1875157a37c452b7
Author: Sagi Grimberg <sagi@grimberg.me>
Date:   Sun Jun 26 12:24:51 2022 +0300

    nvme-tcp: always fail a request when sending it failed

    [ Upstream commit 41d07df7de841bfbc32725ce21d933ad358f2844 ]

    queue stoppage and inflight requests cancellation is fully fenced from
    io_work and thus failing a request from this context. Hence we don't
    need to try to guess from the socket retcode if this failure is because
    the queue is about to be torn down or not.

    We are perfectly safe to just fail it, the request will not be cancelled
    later on.
```

```
    This solves possible very long shutdown delays when the users issues a
    'nvme disconnect-all'

    Reported-by: Daniel Wagner <dwagner@suse.de>
    Signed-off-by: Sagi Grimberg <sagi@grimberg.me>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e84a77bc836160b6b85bbc002465d5aaea307807
Author: Michael Walle <michael@walle.cc>
Date:   Mon Jun 27 19:06:43 2022 +0200

    NFC: nxp-nci: don't print header length mismatch on i2c error

    [ Upstream commit 9577fc5fdc8b07b891709af6453545db405e24ad ]

    Don't print a misleading header length mismatch error if the i2c call
    returns an error. Instead just return the error code without any error
    message.

    Signed-off-by: Michael Walle <michael@walle.cc>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b2957fc09fe1ac7f07f40dd50dd5f93e3f3a7a2
Author: Hangyu Hua <hbh25y@gmail.com>
Date:   Wed Jun 29 14:34:18 2022 +0800

    net: tipc: fix possible refcount leak in tipc_sk_create()

    [ Upstream commit 00aff3590fc0a73bddd3b743863c14e76fd35c0c ]

    Free sk in case tipc_sk_insert() fails.

    Signed-off-by: Hangyu Hua <hbh25y@gmail.com>
    Reviewed-by: Tung Nguyen <tung.q.nguyen@dektech.com.au>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07186778cf645cc79e6913a28dadf445cd3e2439
Author: Javier Martinez Canillas <javierm@redhat.com>
Date:   Tue Jun 7 20:23:36 2022 +0200

    fbdev: Disable sysfb device registration when removing conflicting FBs

    [ Upstream commit ee7a69aa38d87a3bbced7b8245c732c05ed0c6ec ]

    The platform devices registered by sysfb match with firmware-based DRM or
    fbdev drivers, that are used to have early graphics using a framebuffer
    provided by the system firmware.

    DRM or fbdev drivers later are probed and remove conflicting framebuffers,
    leading to these platform devices for generic drivers to be unregistered.

    But the current solution has a race, since the sysfb_init() function could
    be called after a DRM or fbdev driver is probed and request to unregister
    the devices for drivers with conflicting framebuffes.

    To prevent this, disable any future sysfb platform device registration by
    calling sysfb_disable(), if a driver requests to remove the conflicting
    framebuffers.

    Suggested-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
```

```
    Reviewed-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220607182338.344270-4-
javierm@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 38e549484675b6eaf20bbae4bad1eb33452ab95e
Author: Javier Martinez Canillas <javierm@redhat.com>
Date:   Tue Jun 7 20:23:35 2022 +0200

    firmware: sysfb: Add sysfb_disable() helper function

    [ Upstream commit bde376e9de3c0bc55eedc8956b0f114c05531595 ]

    This can be used by subsystems to unregister a platform device registered
    by sysfb and also to disable future platform device registration in sysfb.

    Suggested-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Reviewed-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220607182338.344270-3-
javierm@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f36471e568f95d5440ba89412e5086f57baa1af
Author: Javier Martinez Canillas <javierm@redhat.com>
Date:   Tue Jun 7 20:23:34 2022 +0200

    firmware: sysfb: Make sysfb_create_simplefb() return a pdev pointer

    [ Upstream commit 9e121040e54abef9ed5542e5fdfa87911cd96204 ]

    This function just returned 0 on success or an errno code on error, but it
    could be useful for sysfb_init() callers to have a pointer to the device.

    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Reviewed-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Reviewed-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220607182338.344270-2-
javierm@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05b92733fd6cc44622b7869358aca46e004f0e27
Author: Kai-Heng Feng <kai.heng.feng@canonical.com>
Date:   Tue Jun 28 20:37:26 2022 +0800

    platform/x86: hp-wmi: Ignore Sanitization Mode event

    [ Upstream commit 9ab762a84b8094540c18a170e5ddd6488632c456 ]

    After system resume the hp-wmi driver may complain:
    [ 702.620180] hp_wmi: Unknown event_id - 23 - 0x0

    According to HP it means 'Sanitization Mode' and it's harmless to just
    ignore the event.

    Cc: Jorge Lopez <jorge.lopez2@hp.com>
    Signed-off-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Link: https://lore.kernel.org/r/20220628123726.250062-1-kai.heng.feng@canonical.com
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2c27eb27e70f869e078d3179b3a12a527d33894f
Author: Mark Pearson <markpearson@lenovo.com>
Date:   Mon Jun 27 14:14:49 2022 -0400
```

```
    platform/x86: thinkpad_acpi: do not use PSC mode on Intel platforms

    [ Upstream commit bce6243f767f7da88aa4674d5d678f9f156eaba9 ]

    PSC platform profile mode is only supported on Linux for AMD platforms.

    Some older Intel platforms (e.g T490) are advertising it's capability
    as Windows uses it - but on Linux we should only be using MMC profile
    for Intel systems.

    Add a check to prevent it being enabled incorrectly.

    Signed-off-by: Mark Pearson <markpearson@lenovo.com>
    Link: https://lore.kernel.org/r/20220627181449.3537-1-markpearson@lenovo.com
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01453386fb9233ff274c60ff2ba0689e62cffb4a
Author: Mark Pearson <markpearson@lenovo.com>
Date:   Fri Jun 3 13:02:09 2022 -0400

    platform/x86: thinkpad-acpi: profile capabilities as integer

    [ Upstream commit 42504af775361ca2330a2bfde496a5ebc5655c86 ]

    Currently the active mode (PSC/MMC) is stored in an enum and queried
    throughout the driver.

    Other driver changes will enumerate additional submodes that are relevant
    to be tracked, so instead track PSC/MMC in a single integer variable.

    Co-developed-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Mark Pearson <markpearson@lenovo.com>
    Link: https://lore.kernel.org/r/20220603170212.164963-1-markpearson@lenovo.com
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4513018d0bd739097570d26a7760551cba3deb56
Author: Liang He <windhl@126.com>
Date:   Sat Jun 18 10:25:45 2022 +0800

    cpufreq: pmac32-cpufreq: Fix refcount leak bug

    [ Upstream commit ccd7567d4b6cf187fdfa55f003a9e461ee629e36 ]

    In pmac_cpufreq_init_MacRISC3(), we need to add corresponding
    of_node_put() for the three node pointers whose refcount have
    been incremented by of_find_node_by_name().

    Signed-off-by: Liang He <windhl@126.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 811ebff1ed15ff5860ed753b5d5de46a458a974f
Author: John Garry <john.garry@huawei.com>
Date:   Thu Jun 23 20:41:59 2022 +0800

    scsi: hisi_sas: Limit max hw sectors for v3 HW

    [ Upstream commit fce54ed027577517df1e74b7d54dc2b1bd536887 ]
```

```
    If the controller is behind an IOMMU then the IOMMU IOVA caching range can
    affect performance, as discussed in [0].

    Limit the max HW sectors to not exceed this limit. We need to hardcode the
    value until a proper DMA mapping API is available.

    [0] https://lore.kernel.org/linux-iommu/20210129092120.1482-1-
thunder.leizhen@huawei.com/

    Link: https://lore.kernel.org/r/1655988119-223714-1-git-send-email-john.garry@huawei.com
    Signed-off-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05dfa88e73c9fa9c495384534afbc0d2589b373e
Author: Florian Westphal <fw@strlen.de>
Date:   Tue Jun 21 18:26:03 2022 +0200

    netfilter: br_netfilter: do not skip all hooks with 0 priority

    [ Upstream commit c2577862eeb0be94f151f2f1fff662b028061b00 ]

    When br_netfilter module is loaded, skbs may be diverted to the
    ipv4/ipv6 hooks, just like as if we were routing.

    Unfortunately, bridge filter hooks with priority 0 may be skipped
    in this case.

    Example:
    1. an nftables bridge ruleset is loaded, with a prerouting
       hook that has priority 0.
    2. interface is added to the bridge.
    3. no tcp packet is ever seen by the bridge prerouting hook.
    4. flush the ruleset
    5. load the bridge ruleset again.
    6. tcp packets are processed as expected.

    After 1) the only registered hook is the bridge prerouting hook, but its
    not called yet because the bridge hasn't been brought up yet.

    After 2), hook order is:
       0 br_nf_pre_routing // br_netfilter internal hook
       0 chain bridge f prerouting // nftables bridge ruleset

    The packet is diverted to br_nf_pre_routing.
    If call-iptables is off, the nftables bridge ruleset is called as expected.

    But if its enabled, br_nf_hook_thresh() will skip it because it assumes
    that all 0-priority hooks had been called previously in bridge context.

    To avoid this, check for the br_nf_pre_routing hook itself, we need to
    resume directly after it, even if this hook has a priority of 0.

    Unfortunately, this still results in different packet flow.
    With this fix, the eval order after in 3) is:
    1. br_nf_pre_routing
    2. ip(6)tables (if enabled)
    3. nftables bridge

    but after 5 its the much saner:
    1. nftables bridge
    2. br_nf_pre_routing
    3. ip(6)tables (if enabled)

    Unfortunately I don't see a solution here:
```

```
    It would be possible to move br_nf_pre_routing to a higher priority
    so that it will be called later in the pipeline, but this also impacts
    ebtables evaluation order, and would still result in this very ordering
    problem for all nftables-bridge hooks with the same priority as the
    br_nf_pre_routing one.

    Searching back through the git history I don't think this has
    ever behaved in any other way, hence, no fixes-tag.

    Reported-by: Radim Hrazdil <rhrazdil@redhat.com>
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0016d5d46d7440729a3132f61a8da3bf7f84e2ba
Author: Florian Westphal <fw@strlen.de>
Date:   Wed Jun 22 16:43:57 2022 +0200

    netfilter: nf_tables: avoid skb access on nf_stolen

    [ Upstream commit e34b9ed96ce3b06c79bf884009b16961ca478f87 ]

    When verdict is NF_STOLEN, the skb might have been freed.

    When tracing is enabled, this can result in a use-after-free:
    1. access to skb->nf_trace
    2. access to skb->mark
    3. computation of trace id
    4. dump of packet payload

    To avoid 1, keep a cached copy of skb->nf_trace in the
    trace state struct.
    Refresh this copy whenever verdict is != STOLEN.

    Avoid 2 by skipping skb->mark access if verdict is STOLEN.

    3 is avoided by precomputing the trace id.

    Only dump the packet when verdict is not "STOLEN".

    Reported-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 387c23776010b6f6dbdf53ca267e34c092f23c58
Author: Stephan Gerhold <stephan.gerhold@kernkonzept.com>
Date:   Tue Jun 21 13:06:21 2022 +0200

    virtio_mmio: Restore guest page size on resume

    [ Upstream commit e0c2ce8217955537dd5434baeba061f209797119 ]

    Virtio devices might lose their state when the VMM is restarted
    after a suspend to disk (hibernation) cycle. This means that the
    guest page size register must be restored for the virtio_mmio legacy
    interface, since otherwise the virtio queues are not functional.

    This is particularly problematic for QEMU that currently still defaults
    to using the legacy interface for virtio_mmio. Write the guest page
    size register again in virtio_mmio_restore() to make legacy virtio_mmio
    devices work correctly after hibernation.

    Signed-off-by: Stephan Gerhold <stephan.gerhold@kernkonzept.com>
    Message-Id: <20220621110621.3638025-3-stephan.gerhold@kernkonzept.com>
```

```
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0723688d0c539d97a7a5bb1ebae907907cf0d19b
Author: Stephan Gerhold <stephan.gerhold@kernkonzept.com>
Date:   Tue Jun 21 13:06:20 2022 +0200

    virtio_mmio: Add missing PM calls to freeze/restore

    [ Upstream commit ed7ac37fde33ccd84e4bd2b9363c191f925364c7 ]

    Most virtio drivers provide freeze/restore callbacks to finish up
    device usage before suspend and to reinitialize the virtio device after
    resume. However, these callbacks are currently only called when using
    virtio_pci. virtio_mmio does not have any PM ops defined.

    This causes problems for example after suspend to disk (hibernation),
    since the virtio devices might lose their state after the VMM is
    restarted. Calling virtio_device_freeze()/restore() ensures that
    the virtio devices are re-initialized correctly.

    Fix this by implementing the dev_pm_ops for virtio_mmio,
    similar to virtio_pci_common.

    Signed-off-by: Stephan Gerhold <stephan.gerhold@kernkonzept.com>
    Message-Id: <20220621110621.3638025-2-stephan.gerhold@kernkonzept.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03303d4d9f99139ca69b3a49e1302caaf906cd2e
Author: Gayatri Kammela <gayatri.kammela@linux.intel.com>
Date:   Tue Jun 14 17:27:51 2022 -0700

    platform/x86: intel/pmc: Add Alder Lake N support to PMC core driver

    [ Upstream commit d63eae6747eb8b3192e89712f6553c6aa162f872 ]

    Add Alder Lake N (ADL-N) to the list of the platforms that Intel's
    PMC core driver supports. Alder Lake N reuses all the TigerLake PCH IPs.

    Cc: Srinivas Pandruvada <srinivas.pandruvada@intel.com>
    Cc: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Cc: David E. Box <david.e.box@linux.intel.com>
    Signed-off-by: Gayatri Kammela <gayatri.kammela@linux.intel.com>
    Reviewed-by: Rajneesh Bhardwaj <irenic.rajneesh@gmail.com>
    Link: https://lore.kernel.org/r/20220615002751.3371730-1-gayatri.kammela@linux.intel.com
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fbfa91978c0be7d92ee39e77a71dbe0b4814ef9f
Author: Parav Pandit <parav@nvidia.com>
Date:   Mon Jun 13 22:52:23 2022 +0300

    vduse: Tie vduse mgmtdev and its device

    [ Upstream commit 0e0348ac3f0a6e6606f1aa5acb1803ada913aa3d ]

    vduse devices are not backed by any real devices such as PCI. Hence it
    doesn't have any parent device linked to it.

    Kernel driver model in [1] suggests to avoid an empty device
    release callback.

    Hence tie the mgmtdevice object's life cycle to an allocate dummy struct
```

```
    device instead of static one.

    [1]
 https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/Documentation/core-
 api/kobject.rst?h=v5.18-rc7#n284

    Signed-off-by: Parav Pandit <parav@nvidia.com>
    Message-Id: <20220613195223.473966-1-parav@nvidia.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Reviewed-by: Xie Yongji <xieyongji@bytedance.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8a53aed793fb799d2c6f519c389ea84963d5c807
Author: Eli Cohen <elic@nvidia.com>
Date:   Mon Jun 13 10:59:58 2022 +0300

    vdpa/mlx5: Initialize CVQ vringh only once

    [ Upstream commit ace9252446ec615cd79a5f77d90edb25c0b9d024 ]

    Currently, CVQ vringh is initialized inside setup_virtqueues() which is
    called every time a memory update is done. This is undesirable since it
    resets all the context of the vring, including the available and used
    indices.

    Move the initialization to mlx5_vdpa_set_status() when
    VIRTIO_CONFIG_S_DRIVER_OK is set.

    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Message-Id: <20220613075958.511064-2-elic@nvidia.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Acked-by: Eugenio Pérez <eperezma@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 99d1c36bddd93919072b5a51a89297bbb5ad6a6f
Author: Nathan Lynch <nathanl@linux.ibm.com>
Date:   Thu Jun 23 13:25:09 2022 -0500

    powerpc/xive/spapr: correct bitmap allocation size

    [ Upstream commit 19fc5bb93c6bbdce8292b4d7eed04e2fa118d2fe ]

    kasan detects access beyond the end of the xibm->bitmap allocation:

    BUG: KASAN: slab-out-of-bounds in _find_first_zero_bit+0x40/0x140
    Read of size 8 at addr c00000001d1d0118 by task swapper/0/1

    CPU: 0 PID: 1 Comm: swapper/0 Not tainted 5.19.0-rc2-00001-g90df023b36dd #28
    Call Trace:
    [c00000001d98f770] [c0000000012baab8] dump_stack_lvl+0xac/0x108 (unreliable)
    [c00000001d98f7b0] [c00000000068faac] print_report+0x37c/0x710
    [c00000001d98f880] [c00000000006902c0] kasan_report+0x110/0x354
    [c00000001d98f950] [c000000000692324] __asan_load8+0xa4/0xe0
    [c00000001d98f970] [c0000000011c6ed0] _find_first_zero_bit+0x40/0x140
    [c00000001d98f9b0] [c0000000000dbfbc] xive_spapr_get_ipi+0xcc/0x260
    [c00000001d98fa70] [c0000000000d6d28] xive_setup_cpu_ipi+0x1e8/0x450
    [c00000001d98fb30] [c000000004032a20] pSeries_smp_probe+0x5c/0x118
    [c00000001d98fb60] [c000000004018b44] smp_prepare_cpus+0x944/0x9ac
    [c00000001d98fc90] [c000000004009f9c] kernel_init_freeable+0x2d4/0x640
    [c00000001d98fd90] [c0000000000131e8] kernel_init+0x28/0x1d0
    [c00000001d98fe10] [c00000000000cd54] ret_from_kernel_thread+0x5c/0x64

    Allocated by task 0:
```

```
      kasan_save_stack+0x34/0x70
      __kasan_kmalloc+0xb4/0xf0
      __kmalloc+0x268/0x540
      xive_spapr_init+0x4d0/0x77c
      pseries_init_irq+0x40/0x27c
      init_IRQ+0x44/0x84
      start_kernel+0x2a4/0x538
      start_here_common+0x1c/0x20

    The buggy address belongs to the object at c00000001d1d0118
     which belongs to the cache kmalloc-8 of size 8
    The buggy address is located 0 bytes inside of
     8-byte region [c00000001d1d0118, c00000001d1d0120)

    The buggy address belongs to the physical page:
    page:c00c000000074740 refcount:1 mapcount:0 mapping:0000000000000000
index:0xc00000001d1d0558 pfn:0x1d1d
    flags: 0x7ffff000000200(slab|node=0|zone=0|lastcpupid=0x7ffff)
    raw: 007ffff000000200 c00000001d0003c8 c00000001d0003c8 c00000001d010480
    raw: c00000001d1d0558 0000000001e1000a 00000001ffffffff 0000000000000000
    page dumped because: kasan: bad access detected

    Memory state around the buggy address:
     c00000001d1d0000: fc 00 fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     c00000001d1d0080: fc fc 00 fc fc fc fc fc fc fc fc fc fc fc fc fc
    >c00000001d1d0100: fc fc fc 02 fc fc fc fc fc fc fc fc fc fc fc fc
                           ^
     c00000001d1d0180: fc fc fc fc 04 fc fc fc fc fc fc fc fc fc fc fc
     c00000001d1d0200: fc fc fc fc fc 04 fc fc fc fc fc fc fc fc fc fc

    This happens because the allocation uses the wrong unit (bits) when it
    should pass (BITS_TO_LONGS(count) * sizeof(long)) or equivalent. With small
    numbers of bits, the allocated object can be smaller than sizeof(long),
    which results in invalid accesses.

    Use bitmap_zalloc() to allocate and initialize the irq bitmap, paired with
    bitmap_free() for consistency.

    Signed-off-by: Nathan Lynch <nathanl@linux.ibm.com>
    Reviewed-by: Cédric Le Goater <clg@kaod.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20220623182509.3985625-1-nathanl@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e2b1b8402efdf1a2fd8df71cbe6c6c07551b774
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Thu Jun 2 10:07:38 2022 +0900

    ksmbd: use SOCK_NONBLOCK type for kernel_accept()

    [ Upstream commit fe0fde09e1cb83effcf8fafa372533f438d93a1a ]

    I found that normally it is O_NONBLOCK but there are different value
    for some arch.

    /include/linux/net.h:
    #ifndef SOCK_NONBLOCK
    #define SOCK_NONBLOCK   O_NONBLOCK
    #endif

    /arch/alpha/include/asm/socket.h:
    #define SOCK_NONBLOCK   0x40000000

    Use SOCK_NONBLOCK instead of O_NONBLOCK for kernel_accept().
```

```
    Suggested-by: David Howells <dhowells@redhat.com>
    Signed-off-by: Namjae Jeon <linkinjeon@kerne.org>
    Reviewed-by: Hyunchul Lee <hyc.lee@gmail.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94c8cc503bec42e3d580d45716843ef9bcfe103a
Author: Israel Rukshin <israelr@nvidia.com>
Date:   Thu Jul 14 12:42:10 2022 +0000

    nvme: fix block device naming collision

    [ Upstream commit 6961b5e02876b3b47f030a1f1ee8fd3e631ac270 ]

    The issue exists when multipath is enabled and the namespace is
    shared, but all the other controller checks at nvme_is_unique_nsid()
    are false. The reason for this issue is that nvme_is_unique_nsid()
    returns false when is called from nvme_mpath_alloc_disk() due to an
    uninitialized value of head->shared. The patch fixes it by setting
    head->shared before nvme_mpath_alloc_disk() is called.

    Fixes: 5974ea7ce0f9 ("nvme: allow duplicate NSIDs for private namespaces")
    Signed-off-by: Israel Rukshin <israelr@nvidia.com>
    Reviewed-by: Keith Busch <kbusch@kernel.org>
    Reviewed-by: Max Gurtovoy <mgurtovoy@nvidia.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 799dcbf4d943401eb91c210e79d1f9d2e535ebfa
Author: Bjorn Andersson <bjorn.andersson@linaro.org>
Date:   Fri Jul 8 17:00:27 2022 -0700

    scsi: ufs: core: Drop loglevel of WriteBoost message

    [ Upstream commit 2ae57c995003a7840cb6b5ec5f0c06193695321b ]

    Commit '3b5f3c0d0548 ("scsi: ufs: core: Tidy up WB configuration code")'
    changed the log level of the write boost enable/disable notification from
    debug to info. This results in a lot of noise in the kernel log during
    normal operation.

    Drop it back to debug level to avoid this.

    Link: https://lore.kernel.org/r/20220709000027.3929970-1-bjorn.andersson@linaro.org
    Fixes: 3b5f3c0d0548 ("scsi: ufs: core: Tidy up WB configuration code")
    Reviewed-by: Alim Akhtar <alim.akhtar@samsung.com>
    Acked-by: Bean Huo <beanhuo@micron.com>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit afeb95a11a4cc4b45914b10d34dd71417b3f356b
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Jul 6 20:59:42 2022 +0800

    scsi: megaraid: Clear READ queue map's nr_queues

    [ Upstream commit 8312cd3a7b835ae3033a679e5f0014a40e7891c5 ]

    The megaraid SCSI driver sets set->nr_maps as 3 if poll_queues is > 0, and
    blk-mq actually initializes each map's nr_queues as nr_hw_queues.
    Consequently the driver has to clear READ queue map's nr_queues, otherwise
    the queue map becomes broken if poll_queues is set as non-zero.

    Link: https://lore.kernel.org/r/20220706125942.528533-1-ming.lei@redhat.com
```

```
    Fixes: 9e4bec5b2a23 ("scsi: megaraid_sas: mq_poll support")
    Cc: Kashyap Desai <kashyap.desai@broadcom.com>
    Cc: sumit.saxena@broadcom.com
    Cc: chandrakanth.patil@broadcom.com
    Cc: linux-block@vger.kernel.org
    Cc: Hannes Reinecke <hare@suse.de>
    Reported-by: Guangwu Zhang <guazhang@redhat.com>
    Tested-by: Guangwu Zhang <guazhang@redhat.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e40c724237218671af114d710175abbf9c91361f
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Mon Jun 27 14:50:53 2022 +0200

    s390/nospec: build expoline.o for modules_prepare target

    [ Upstream commit c4e789572557aa147b13bf7fe09cc99663ed0cf5 ]

    When CONFIG_EXPOLINE_EXTERN is used expoline thunks are generated
    from arch/s390/lib/expoline.S and postlinked into every module.
    This is also true for external modules. Add expoline.o build to
    the modules_prepare target.

    Fixes: 1d2ad084800e ("s390/nospec: add an option to use thunk-extern")
    Reported-by: Joe Lawrence <joe.lawrence@redhat.com>
    Tested-by: Sumanth Korikkar <sumanthk@linux.ibm.com>
    Acked-by: Sumanth Korikkar <sumanthk@linux.ibm.com>
    Tested-by: C. Erastus Toe <ctoe@redhat.com>
    Tested-by: Joe Lawrence <joe.lawrence@redhat.com>
    Link: https://lore.kernel.org/r/patch-1.thread-d13b6c.git-a2387a74dc49.your-ad-
here.call-01656331067-ext-4899@work.hours
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fa3302714c03e4e6c9b5aad5dacae33e75f76cf7
Author: Jiri Slaby <jirislaby@kernel.org>
Date:   Thu Jul 7 10:25:58 2022 +0200

    tty: use new tty_insert_flip_string_and_push_buffer() in pty_write()

    [ Upstream commit a501ab75e7624d133a5a3c7ec010687c8b961d23 ]

    There is a race in pty_write(). pty_write() can be called in parallel
    with e.g. ioctl(TIOCSTI) or ioctl(TCXONC) which also inserts chars to
    the buffer. Provided, tty_flip_buffer_push() in pty_write() is called
    outside the lock, it can commit inconsistent tail. This can lead to out
    of bounds writes and other issues. See the Link below.

    To fix this, we have to introduce a new helper called
    tty_insert_flip_string_and_push_buffer(). It does both
    tty_insert_flip_string() and tty_flip_buffer_commit() under the port
    lock. It also calls queue_work(), but outside the lock. See
    71a174b39f10 (pty: do tty_flip_buffer_push without port->lock in
    pty_write) for the reasons.

    Keep the helper internal-only (in drivers' tty.h). It is not intended to
    be used widely.

    Link: https://seclists.org/oss-sec/2022/q2/155
    Fixes: 71a174b39f10 (pty: do tty_flip_buffer_push without port->lock in pty_write)
    Cc: 一只狗 <chennbnbnb@gmail.com>
```

```
    Cc: Dan Carpenter <dan.carpenter@oracle.com>
    Suggested-by: Hillf Danton <hdanton@sina.com>
    Signed-off-by: Jiri Slaby <jslaby@suse.cz>
    Link: https://lore.kernel.org/r/20220707082558.9250-2-jslaby@suse.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1b922311ed063519f395f874a14ddd644a6868d
Author: Jiri Slaby <jirislaby@kernel.org>
Date:   Thu Jul 7 10:25:57 2022 +0200

    tty: extract tty_flip_buffer_commit() from tty_flip_buffer_push()

    [ Upstream commit 716b10580283fda66f2b88140e3964f8a7f9da89 ]

    We will need this new helper in the next patch.

    Cc: Hillf Danton <hdanton@sina.com>
    Cc: 一只狗 <chennbnbnb@gmail.com>
    Cc: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Jiri Slaby <jslaby@suse.cz>
    Link: https://lore.kernel.org/r/20220707082558.9250-1-jslaby@suse.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1bbe15676a37e4a67989ddd9cf357c0d1517387
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Mon Jun 6 13:43:53 2022 +0200

    tee: tee_get_drvdata(): fix description of return value

    [ Upstream commit e5ce073c8a1e01b215a5eb32ba48f8d17ded3bd5 ]

    This patch fixes the description of tee_get_drvdata()'s return value.
    It actually returns the driver_data pointer supplied to
    tee_device_alloc() since the TEE subsystem was added to the kernel.

    Fixes: 967c9cca2cc5 ("tee: generic TEE subsystem")
    Cc: Jens Wiklander <jens.wiklander@linaro.org>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Jens Wiklander <jens.wiklander@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 847f0c80c2ef9dcda7c47ee129d4253b5654c34c
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Wed Jun 29 12:40:12 2022 +0100

    ASoC: dt-bindings: Fix description for msm8916

    [ Upstream commit 94c65dffd4c4af052b3ea8934fbcb2fa8da276a8 ]

    For the existing msm8916 bindings the minimum reg/reg-names is 1 not 2.
    Similarly the minimum interrupt/interrupt-names is 1 not 2.

    Fixes: f3fc4fbfa2d2 ("ASoC: dt-bindings: Add SC7280 lpass cpu bindings")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20220629114012.3282945-1-bryan.odonoghue@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9987bb02667ffe5971881933119c6eaed7c3a9ed
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun Jun 19 20:40:16 2022 +0100
```

```
    ARM: 9212/1: domain: Modify Kconfig help text

    [ Upstream commit 2bf6204240fddb22cc4940b9e3f40c538390212e ]

    After the removal of set_fs() the reference to set_fs() is stale.
    Alter the helptext to reflect what the config option really does.

    Fixes: 8ac6f5d7f84b ("ARM: 9113/1: uaccess: remove set_fs() implementation")
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 82089fcb1051230b5dc7145d07034b407e4c1ce6
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun Jun 19 20:39:29 2022 +0100

    ARM: 9211/1: domain: drop modify_domain()

    [ Upstream commit cc45b836388f0ccc6831288a08f77a33845f10b0 ]

    This function/macro isn't used anywhere in the kernel.
    The only user was set_fs() and was deleted in the set_fs()
    removal patch set.

    Fixes: 8ac6f5d7f84b ("ARM: 9113/1: uaccess: remove set_fs() implementation")
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed388232ed26096648c451ba0d602af8d849389d
Author: Muchun Song <songmuchun@bytedance.com>
Date:   Thu Jun 9 18:40:32 2022 +0800

    mm: sysctl: fix missing numa_stat when !CONFIG_HUGETLB_PAGE

    [ Upstream commit 43b5240ca6b331089988105932481 86b1e3ae34a ]

    "numa_stat" should not be included in the scope of CONFIG_HUGETLB_PAGE, if
    CONFIG_HUGETLB_PAGE is not configured even if CONFIG_NUMA is configured,
    "numa_stat" is missed form /proc. Move it out of CONFIG_HUGETLB_PAGE to
    fix it.

    Fixes: 4518085e127d ("mm, sysctl: make NUMA stats configurable")
    Signed-off-by: Muchun Song <songmuchun@bytedance.com>
    Cc: <stable@vger.kernel.org>
    Acked-by: Michal Hocko <mhocko@suse.com>
    Acked-by: Mel Gorman <mgorman@techsingularity.net>
    Signed-off-by: Luis Chamberlain <mcgrof@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b00c5375a16413ada7537567d42cc0008f3378ae
Author: Tariq Toukan <tariqt@nvidia.com>
Date:   Thu Jul 14 10:07:54 2022 +0300

    net/tls: Check for errors in tls_device_init

    [ Upstream commit 3d8c51b25a235e283e37750943bbf356ef187230 ]

    Add missing error checks in tls_device_init.

    Fixes: e8f69799810c ("net/tls: Add generic NIC offload infrastructure")
    Reported-by: Jakub Kicinski <kuba@kernel.org>
    Reviewed-by: Maxim Mikityanskiy <maximmi@nvidia.com>
```

```
    Signed-off-by: Tariq Toukan <tariqt@nvidia.com>
    Link: https://lore.kernel.org/r/20220714070754.1428-1-tariqt@nvidia.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 80866cf468d062c8e2226e2f376103e971254707
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Fri Jul 8 14:51:47 2022 +0200

    KVM: x86: Fully initialize 'struct kvm_lapic_irq' in kvm_pv_kick_cpu_op()

    [ Upstream commit 8a414f943f8b5f94bbaafdec863d6f3dbef33f8a ]

    'vector' and 'trig_mode' fields of 'struct kvm_lapic_irq' are left
    uninitialized in kvm_pv_kick_cpu_op(). While these fields are normally
    not needed for APIC_DM_REMRD, they're still referenced by
    __apic_accept_irq() for trace_kvm_apic_accept_irq(). Fully initialize
    the structure to avoid consuming random stack memory.

    Fixes: a183b638b61c ("KVM: x86: make apic_accept_irq tracepoint more generic")
    Reported-by: syzbot+d6caa905917d353f0d07@syzkaller.appspotmail.com
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Reviewed-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20220708125147.593975-1-vkuznets@redhat.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 702419db7ded2e3686fbc05040b61d7428d611d5
Author: Chia-Lin Kao (AceLan) <acelan.kao@canonical.com>
Date:   Wed Jul 13 19:12:24 2022 +0800

    net: atlantic: remove aq_nic_deinit() when resume

    [ Upstream commit 2e15c51fefaffaf9f72255eaef4fada05055e4c5 ]

    aq_nic_deinit() has been called while suspending, so we don't have to call
    it again on resume.
    Actually, call it again leads to another hang issue when resuming from
    S3.

    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992345] Call Trace:
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992346] <TASK>
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992348] aq_nic_deinit+0xb4/0xd0
[atlantic]
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992356] aq_pm_thaw+0x7f/0x100
[atlantic]
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992362] pci_pm_resume+0x5c/0x90
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992366] ? pci_pm_thaw+0x80/0x80
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992368] dpm_run_callback+0x4e/0x120
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992371] device_resume+0xad/0x200
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992373] async_resume+0x1e/0x40
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992374]
async_run_entry_fn+0x33/0x120
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992377]
process_one_work+0x220/0x3c0
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992380] worker_thread+0x4d/0x3f0
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992382] ?
process_one_work+0x3c0/0x3c0
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992384] kthread+0x12a/0x150
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992386] ?
set_kthread_struct+0x40/0x40
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992387] ret_from_fork+0x22/0x30
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992391] </TASK>
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992392] ---[ end trace
1ec8c79604ed5e0d ]---
```

```
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992394] PM: dpm_run_callback():
 pci_pm_resume+0x0/0x90 returns -110
    Jul 8 03:09:44 u-Precision-7865-Tower kernel: [ 5910.992397] atlantic 0000:02:00.0: PM:
 failed to resume async: error -110

    Fixes: 1809c30b6e5a ("net: atlantic: always deep reset on pm op, fixing up my null deref
 regression")
    Signed-off-by: Chia-Lin Kao (AceLan) <acelan.kao@canonical.com>
    Link: https://lore.kernel.org/r/20220713111224.1535938-2-acelan.kao@canonical.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e151c507b7a7a77e78f666d9de00a38f3b2192c
Author: Chia-Lin Kao (AceLan) <acelan.kao@canonical.com>
Date:   Wed Jul 13 19:12:23 2022 +0800

    net: atlantic: remove deep parameter on suspend/resume functions

    [ Upstream commit 0f33250760384e05c36466b0a2f92f3c6007ba92 ]

    Below commit claims that atlantic NIC requires to reset the device on pm
    op, and had set the deep to true for all suspend/resume functions.
    commit 1809c30b6e5a ("net: atlantic: always deep reset on pm op, fixing up my null deref
 regression")
    So, we could remove deep parameter on suspend/resume functions without
    any functional change.

    Fixes: 1809c30b6e5a ("net: atlantic: always deep reset on pm op, fixing up my null deref
 regression")
    Signed-off-by: Chia-Lin Kao (AceLan) <acelan.kao@canonical.com>
    Link: https://lore.kernel.org/r/20220713111224.1535938-1-acelan.kao@canonical.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16662524ec5da801fb78a1afcaf6e782f1cf103a
Author: Íñigo Huguet <ihuguet@redhat.com>
Date:   Wed Jul 13 11:21:16 2022 +0200

    sfc: fix kernel panic when creating VF

    [ Upstream commit ada74c5539eba06cf8b47d068f92e0b3963a9a6e ]

    When creating VFs a kernel panic can happen when calling to
    efx_ef10_try_update_nic_stats_vf.

    When releasing a DMA coherent buffer, sometimes, I don't know in what
    specific circumstances, it has to unmap memory with vunmap. It is
    disallowed to do that in IRQ context or with BH disabled. Otherwise, we
    hit this line in vunmap, causing the crash:
      BUG_ON(in_interrupt());

    This patch reenables BH to release the buffer.

    Log messages when the bug is hit:
     kernel BUG at mm/vmalloc.c:2727!
     invalid opcode: 0000 [#1] PREEMPT SMP NOPTI
     CPU: 6 PID: 1462 Comm: NetworkManager Kdump: loaded Tainted: G          I      --------
 - ---  5.14.0-119.el9.x86_64 #1
     Hardware name: Dell Inc. PowerEdge R740/06WXJT, BIOS 2.8.2 08/27/2020
     RIP: 0010:vunmap+0x2e/0x30
     ...skip...
     Call Trace:
      __iommu_dma_free+0x96/0x100
      efx_nic_free_buffer+0x2b/0x40 [sfc]
      efx_ef10_try_update_nic_stats_vf+0x14a/0x1c0 [sfc]
```

```
       efx_ef10_update_stats_vf+0x18/0x40 [sfc]
       efx_start_all+0x15e/0x1d0 [sfc]
       efx_net_open+0x5a/0xe0 [sfc]
       __dev_open+0xe7/0x1a0
       __dev_change_flags+0x1d7/0x240
       dev_change_flags+0x21/0x60
       ...skip...
```

```
    Fixes: d778819609a2 ("sfc: DMA the VF stats only when requested")
    Reported-by: Ma Yuying <yuma@redhat.com>
    Signed-off-by: Íñigo Huguet <ihuguet@redhat.com>
    Acked-by: Edward Cree <ecree.xilinx@gmail.com>
    Link: https://lore.kernel.org/r/20220713092116.21238-1-ihuguet@redhat.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e980d6a20162396ea1858121892b396181c07b53
Author: Andrea Mayer <andrea.mayer@uniroma2.it>
Date:   Tue Jul 12 19:58:37 2022 +0200

    seg6: bpf: fix skb checksum in bpf_push_seg6_encap()

    [ Upstream commit 4889fbd98deaf243c3baadc54e296d71c6af1eb0 ]

    Both helper functions bpf_lwt_seg6_action() and bpf_lwt_push_encap() use
    the bpf_push_seg6_encap() to encapsulate the packet in an IPv6 with Segment
    Routing Header (SRH) or insert an SRH between the IPv6 header and the
    payload.
    To achieve this result, such helper functions rely on bpf_push_seg6_encap()
    which, in turn, leverages seg6_do_srh_{encap,inline}() to perform the
    required operation (i.e. encap/inline).

    This patch removes the initialization of the IPv6 header payload length
    from bpf_push_seg6_encap(), as it is now handled properly by
    seg6_do_srh_{encap,inline}() to prevent corruption of the skb checksum.

    Fixes: fe94cc290f53 ("bpf: Add IPv6 Segment Routing helpers")
    Signed-off-by: Andrea Mayer <andrea.mayer@uniroma2.it>
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a8e911a97f579d8e1997cb530f154c1e3b0da724
Author: Andrea Mayer <andrea.mayer@uniroma2.it>
Date:   Tue Jul 12 19:58:36 2022 +0200

    seg6: fix skb checksum in SRv6 End.B6 and End.B6.Encaps behaviors

    [ Upstream commit f048880fc77058d864aff5c674af7918b30f312a ]

    The SRv6 End.B6 and End.B6.Encaps behaviors rely on functions
    seg6_do_srh_{encap,inline}() to, respectively: i) encapsulate the
    packet within an outer IPv6 header with the specified Segment Routing
    Header (SRH); ii) insert the specified SRH directly after the IPv6
    header of the packet.

    This patch removes the initialization of the IPv6 header payload length
    from the input_action_end_b6{_encap}() functions, as it is now handled
    properly by seg6_do_srh_{encap,inline}() to avoid corruption of the skb
    checksum.

    Fixes: 140f04c33bbc ("ipv6: sr: implement several seg6local actions")
    Signed-off-by: Andrea Mayer <andrea.mayer@uniroma2.it>
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 281e3679e62403e74ba7862c98227b96c678fa88
Author: Andrea Mayer <andrea.mayer@uniroma2.it>
Date:   Tue Jul 12 19:58:35 2022 +0200

    seg6: fix skb checksum evaluation in SRH encapsulation/insertion

    [ Upstream commit df8386d13ea280d55beee1b95f61a59234a3798b ]

    Support for SRH encapsulation and insertion was introduced with
    commit 6c8702c60b88 ("ipv6: sr: add support for SRH encapsulation and
    injection with lwtunnels"), through the seg6_do_srh_encap() and
    seg6_do_srh_inline() functions, respectively.
    The former encapsulates the packet in an outer IPv6 header along with
    the SRH, while the latter inserts the SRH between the IPv6 header and
    the payload. Then, the headers are initialized/updated according to the
    operating mode (i.e., encap/inline).
    Finally, the skb checksum is calculated to reflect the changes applied
    to the headers.

    The IPv6 payload length ('payload_len') is not initialized
    within seg6_do_srh_{inline,encap}() but is deferred in seg6_do_srh(), i.e.
    the caller of seg6_do_srh_{inline,encap}().
    However, this operation invalidates the skb checksum, since the
    'payload_len' is updated only after the checksum is evaluated.

    To solve this issue, the initialization of the IPv6 payload length is
    moved from seg6_do_srh() directly into the seg6_do_srh_{inline,encap}()
    functions and before the skb checksum update takes place.

    Fixes: 6c8702c60b88 ("ipv6: sr: add support for SRH encapsulation and injection with
lwtunnels")
    Reported-by: Paolo Abeni <pabeni@redhat.com>
    Link: https://lore.kernel.org/all/20220705190727.69d532417be7438b15404ee1@uniroma2.it
    Signed-off-by: Andrea Mayer <andrea.mayer@uniroma2.it>
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 77be53ba0425ed1d4496f2b45708a863fe918730
Author: Xiubo Li <xiubli@redhat.com>
Date:   Mon Jul 11 12:11:21 2022 +0800

    netfs: do not unlock and put the folio twice

    [ Upstream commit fac47b43c760ea90e64b895dba60df0327be7775 ]

    check_write_begin() will unlock and put the folio when return
    non-zero.  So we should avoid unlocking and putting it twice in
    netfs layer.

    Change the way ->check_write_begin() works in the following two ways:

     (1) Pass it a pointer to the folio pointer, allowing it to unlock and put
         the folio prior to doing the stuff it wants to do, provided it clears
         the folio pointer.

     (2) Change the return values such that 0 with folio pointer set means
         continue, 0 with folio pointer cleared means re-get and all error
         codes indicating an error (no special treatment for -EAGAIN).

    [ bagasdotme: use Sphinx code text syntax for *foliop pointer ]

    Cc: stable@vger.kernel.org
    Link: https://tracker.ceph.com/issues/56423
    Link: https://lore.kernel.org/r/cf169f43-8ee7-8697-25da-0204d1b4343e@redhat.com
    Co-developed-by: David Howells <dhowells@redhat.com>
```

```
    Signed-off-by: Xiubo Li <xiubli@redhat.com>
    Signed-off-by: David Howells <dhowells@redhat.com>
    Signed-off-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Signed-off-by: Ilya Dryomov <idryomov@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e435c4aeeaa073091f7f3b7735af2ef5c97d63f2
Author: Íñigo Huguet <ihuguet@redhat.com>
Date:   Tue Jul 12 08:26:42 2022 +0200

    sfc: fix use after free when disabling sriov

    [ Upstream commit ebe41da5d47ac0fff877e57bd14c54dccf168827 ]

    Use after free is detected by kfence when disabling sriov. What was read
    after being freed was vf->pci_dev: it was freed from pci_disable_sriov
    and later read in efx_ef10_sriov_free_vf_vports, called from
    efx_ef10_sriov_free_vf_vswitching.

    Set the pointer to NULL at release time to not trying to read it later.

    Reproducer and dmesg log (note that kfence doesn't detect it every time):
    $ echo 1 > /sys/class/net/enp65s0f0np0/device/sriov_numvfs
    $ echo 0 > /sys/class/net/enp65s0f0np0/device/sriov_numvfs

     BUG: KFENCE: use-after-free read in efx_ef10_sriov_free_vf_vswitching+0x82/0x170 [sfc]

     Use-after-free read at 0x00000000ff3c1ba5 (in kfence-#224):
      efx_ef10_sriov_free_vf_vswitching+0x82/0x170 [sfc]
      efx_ef10_pci_sriov_disable+0x38/0x70 [sfc]
      efx_pci_sriov_configure+0x24/0x40 [sfc]
      sriov_numvfs_store+0xfe/0x140
      kernfs_fop_write_iter+0x11c/0x1b0
      new_sync_write+0x11f/0x1b0
      vfs_write+0x1eb/0x280
      ksys_write+0x5f/0xe0
      do_syscall_64+0x5c/0x80
      entry_SYSCALL_64_after_hwframe+0x44/0xae

     kfence-#224: 0x00000000edb8ef95-0x00000000671f5ce1, size=2792, cache=kmalloc-4k

     allocated by task 6771 on cpu 10 at 3137.860196s:
      pci_alloc_dev+0x21/0x60
      pci_iov_add_virtfn+0x2a2/0x320
      sriov_enable+0x212/0x3e0
      efx_ef10_sriov_configure+0x67/0x80 [sfc]
      efx_pci_sriov_configure+0x24/0x40 [sfc]
      sriov_numvfs_store+0xba/0x140
      kernfs_fop_write_iter+0x11c/0x1b0
      new_sync_write+0x11f/0x1b0
      vfs_write+0x1eb/0x280
      ksys_write+0x5f/0xe0
      do_syscall_64+0x5c/0x80
      entry_SYSCALL_64_after_hwframe+0x44/0xae

     freed by task 6771 on cpu 12 at 3170.991309s:
      device_release+0x34/0x90
      kobject_cleanup+0x3a/0x130
      pci_iov_remove_virtfn+0xd9/0x120
      sriov_disable+0x30/0xe0
      efx_ef10_pci_sriov_disable+0x57/0x70 [sfc]
      efx_pci_sriov_configure+0x24/0x40 [sfc]
      sriov_numvfs_store+0xfe/0x140
      kernfs_fop_write_iter+0x11c/0x1b0
      new_sync_write+0x11f/0x1b0
```

```
       vfs_write+0x1eb/0x280
       ksys_write+0x5f/0xe0
       do_syscall_64+0x5c/0x80
       entry_SYSCALL_64_after_hwframe+0x44/0xae

    Fixes: 3c5eb87605e85 ("sfc: create vports for VFs and assign random MAC addresses")
    Reported-by: Yanghang Liu <yanghliu@redhat.com>
    Signed-off-by: Íñigo Huguet <ihuguet@redhat.com>
    Acked-by: Martin Habets <habetsm.xilinx@gmail.com>
    Link: https://lore.kernel.org/r/20220712062642.6915-1-ihuguet@redhat.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68ea71d4eafcf5e6881286f551e9c35a750e0097
Author: Steve French <stfrench@microsoft.com>
Date:   Tue Jul 12 00:11:42 2022 -0500

    smb3: workaround negprot bug in some Samba servers

    [ Upstream commit 32f319183c439b239294cb2d70ada3564c4c7c39 ]

    Mount can now fail to older Samba servers due to a server
    bug handling padding at the end of the last negotiate
    context (negotiate contexts typically are rounded up to 8
    bytes by adding padding if needed). This server bug can
    be avoided by switching the order of negotiate contexts,
    placing a negotiate context at the end that does not
    require padding (prior to the recent netname context fix
    this was the case on the client).

    Fixes: 73130a7b1ac9 ("smb3: fix empty netname context on secondary channels")
    Reported-by: Julian Sikorski <belegdol@gmail.com>
    Tested-by: Julian Sikorski <belegdol+github@gmail.com>
    Reviewed-by: Shyam Prasad N <sprasad@microsoft.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f094271f6bcf183d29964c1f963d9e896e5a7462
Author: Michel Dänzer <mdaenzer@redhat.com>
Date:   Mon Jul 11 16:51:31 2022 +0200

    drm/amd/display: Ensure valid event timestamp for cursor-only commits

    [ Upstream commit 3283c83eb6fcfbda8ea03d7149d8e42e71c5d45e ]

    Requires enabling the vblank machinery for them.

    Bug: https://gitlab.freedesktop.org/drm/amd/-/issues/2030
    Signed-off-by: Michel Dänzer <mdaenzer@redhat.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17ece75cdbc2d816c54c3a43152a09db9fd918ae
Author: Yefim Barashkin <mr.b34r@kolabnow.com>
Date:   Mon Jul 11 14:35:11 2022 -0800

    drm/amd/pm: Prevent divide by zero

    [ Upstream commit 0638c98c17aa12fe914459c82cd178247e21fb2b ]

    divide error: 0000 [#1] SMP PTI
    CPU: 3 PID: 78925 Comm: tee Not tainted 5.15.50-1-lts #1
    Hardware name: MSI MS-7A59/Z270 SLI PLUS (MS-7A59), BIOS 1.90 01/30/2018
    RIP: 0010:smu_v11_0_set_fan_speed_rpm+0x11/0x110 [amdgpu]
```

```
    Speed is user-configurable through a file.
    I accidentally set it to zero, and the driver crashed.

    Reviewed-by: Evan Quan <evan.quan@amd.com>
    Reviewed-by: André Almeida <andrealmeid@igalia.com>
    Signed-off-by: Yefim Barashkin <mr.b34r@kolabnow.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19206b2bfccc3e834bcda885a5b84ac286109c2a
Author: Mario Kleiner <mario.kleiner.de@gmail.com>
Date:   Mon Jul 11 19:39:28 2022 +0200

    drm/amd/display: Only use depth 36 bpp linebuffers on DCN display engines.

    [ Upstream commit add61d3c31de6a4b5e11a2ab96aaf4c873481568 ]

    Various DCE versions had trouble with 36 bpp lb depth, requiring fixes,
    last time in commit 353ca0fa5630 ("drm/amd/display: Fix 10bit 4K display
    on CIK GPUs") for DCE-8. So far >= DCE-11.2 was considered ok, but now I
    found out that on DCE-11.2 it causes dithering when there shouldn't be
    any, so identity pixel passthrough with identity gamma LUTs doesn't work
    when it should. This breaks various important neuroscience applications,
    as reported to me by scientific users of Polaris cards under Ubuntu 22.04
    with Linux 5.15, and confirmed by testing it myself on DCE-11.2.

    Lets only use depth 36 for DCN engines, where my testing showed that it
    is both necessary for high color precision output, e.g., RGBA16 fb's,
    and not harmful, as far as more than one year in real-world use showed.

    DCE engines seem to work fine for high precision output at 30 bpp, so
    this ("famous last words") depth 30 should hopefully fix all known problems
    without introducing new ones.

    Successfully retested on DCE-11.2 Polaris and DCN-1.0 Raven Ridge on
    top of Linux 5.19.0-rc2 + drm-next.

    Fixes: 353ca0fa5630 ("drm/amd/display: Fix 10bit 4K display on CIK GPUs")
    Signed-off-by: Mario Kleiner <mario.kleiner.de@gmail.com>
    Tested-by: Mario Kleiner <mario.kleiner.de@gmail.com>
    Cc: stable@vger.kernel.org # 5.14.0
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Cc: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05c7d624b6a17fa89e9acd5ed75e01732dbf8f94
Author: Prike Liang <Prike.Liang@amd.com>
Date:   Mon Jul 11 16:03:08 2022 +0800

    drm/amdkfd: correct the MEC atomic support firmware checking for GC 10.3.7

    [ Upstream commit c0044865480a162146b9dfe7783e73a08e97b2b9 ]

    On the GC 10.3.7 platform the initial MEC release version #3 can support
    atomic operation,so need correct and set its MEC atomic support version to #3.

    Signed-off-by: Prike Liang <Prike.Liang@amd.com>
    Reviewed-by: Aaron Liu <aaron.liu@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org # 5.18.x
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5a0d38c6a1d6d6e7eb1c8f02e0077d117167749c
Author: Fangzhi Zuo <Jerry.Zuo@amd.com>
Date:   Wed Jul 6 15:52:46 2022 -0400

    drm/amd/display: Ignore First MST Sideband Message Return Error

    [ Upstream commit acea108fa067d140bd155161a79b1fcd967f4137 ]

    [why]
    First MST sideband message returns AUX_RET_ERROR_HPD_DISCON
    on certain intel platform. Aux transaction considered failure
    if HPD unexpected pulled low. The actual aux transaction success
    in such case, hence do not return error.

    [how]
    Not returning error when AUX_RET_ERROR_HPD_DISCON detected
    on the first sideband message.

    v2: squash in additional DMI entries
    v3: squash in static fix

    Signed-off-by: Fangzhi Zuo <Jerry.Zuo@amd.com>
    Acked-by: Solomon Chiu <solomon.chiu@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 830de9667b3ada0a75a3f098dfc7159709fe397b
Author: Jianglei Nie <niejianglei2021@163.com>
Date:   Tue Jul 12 09:10:37 2022 +0800

    ima: Fix potential memory leak in ima_init_crypto()

    [ Upstream commit 067d2521874135267e681c19d42761c601d503d6 ]

    On failure to allocate the SHA1 tfm, IMA fails to initialize and exits
    without freeing the ima_algo_array. Add the missing kfree() for
    ima_algo_array to avoid the potential memory leak.

    Signed-off-by: Jianglei Nie <niejianglei2021@163.com>
    Fixes: 6d94809af6b0 ("ima: Allocate and initialize tfm for each PCR bank")
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05f68241638e67457f7fab5066b08f6df181a3d4
Author: Coiby Xu <coxu@redhat.com>
Date:   Wed Jul 13 15:21:11 2022 +0800

    ima: force signature verification when CONFIG_KEXEC_SIG is configured

    [ Upstream commit af16df54b89dee72df253abc5e7b5e8a6d16c11c ]

    Currently, an unsigned kernel could be kexec'ed when IMA arch specific
    policy is configured unless lockdown is enabled. Enforce kernel
    signature verification check in the kexec_file_load syscall when IMA
    arch specific policy is configured.

    Fixes: 99d5cadfde2b ("kexec_file: split KEXEC_VERIFY_SIG into KEXEC_SIG and
KEXEC_SIG_FORCE")
    Reported-and-suggested-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Coiby Xu <coxu@redhat.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit dd91bc60f305610401b2196bedb573693d6c8e46
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Tue Jul 12 17:42:25 2022 +0300

    net: stmmac: fix leaks in probe

    [ Upstream commit 23aa6d5088e3bd65de77c5c307237b9937f8b48a ]

    These two error paths should clean up before returning.

    Fixes: 2bb4b98b60d7 ("net: stmmac: Add Ingenic SoCs MAC support.")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5cfdd61c807df5b2905a4de78514c2022db6954a
Author: Liang He <windhl@126.com>
Date:   Tue Jul 12 14:14:17 2022 +0800

    net: ftgmac100: Hold reference returned by of_get_child_by_name()

    [ Upstream commit 49b9f431ff0d845a36be0b3ede35ec324f2e5fee ]

    In ftgmac100_probe(), we should hold the refernece returned by
    of_get_child_by_name() and use it to call of_node_put() for
    reference balance.

    Fixes: 39bfab8844a0 ("net: ftgmac100: Add support for DT phy-handle property")
    Signed-off-by: Liang He <windhl@126.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae3054f6fbccc90f14ecd6cf9b2c09a2401c64fd
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:33 2022 -0700

    nexthop: Fix data-races around nexthop_compat_mode.

    [ Upstream commit bdf00bf24bef9be1ca641a6390fd5487873e0d2e ]

    While reading nexthop_compat_mode, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its readers.

    Fixes: 4f80116d3df3 ("net: ipv4: add sysctl for nexthop api compatibility mode")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f5042c5357ce931f1666faf6b26b27d12cbcc633
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:32 2022 -0700

    ipv4: Fix data-races around sysctl_ip_dynaddr.

    [ Upstream commit e49e4aff7ec19b2d0d0957ee30e93dade57dab9e ]

    While reading sysctl_ip_dynaddr, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its readers.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1ec3d6c2626ee6e1b36b7bd006873a271406ba61
```

```
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:31 2022 -0700

    tcp: Fix a data-race around sysctl_tcp_ecn_fallback.

    [ Upstream commit 12b8d9ca7e678abc48195294494f1815b555d658 ]

    While reading sysctl_tcp_ecn_fallback, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 492135557dc0 ("tcp: add rfc3168, section 6.1.1.1. fallback")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7af0cc1f99010a6e4fe54f31a4804146a77260be
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:30 2022 -0700

    tcp: Fix data-races around sysctl_tcp_ecn.

    [ Upstream commit 4785a66702f086cf2ea84bdbe6ec921f274bd9f2 ]

    While reading sysctl_tcp_ecn, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its readers.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab5adca2e17d6595f3fc0e25ccb6bcbe2e01ca4f
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:29 2022 -0700

    raw: Fix a data-race around sysctl_raw_l3mdev_accept.

    [ Upstream commit 1dace014928e6e385363032d359a04dee9158af0 ]

    While reading sysctl_raw_l3mdev_accept, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 6897445fb194 ("net: provide a sysctl raw_l3mdev_accept for raw socket lookup with
VRFs")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84492c1edafa7cb493767fd5479f5f984eb9ac10
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:28 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_ratemask.

    [ Upstream commit 1ebcb25ad6fc3d50fca87350acf451b9a66dd31e ]

    While reading sysctl_icmp_ratemask, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 83c68347b41af8e11a720fb826b9d824527ad902
```

```
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:27 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_ratelimit.

    [ Upstream commit 2a4eb714841f288cf51c7d942d98af6a8c6e4b01 ]

    While reading sysctl_icmp_ratelimit, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f9617844e4d5d6331dbce3fb19a24e5bda201e58
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:26 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_errors_use_inbound_ifaddr.

    [ Upstream commit d2efabce81db7eed1c98fa1a3f203f0edd738ac3 ]

    While reading sysctl_icmp_errors_use_inbound_ifaddr, it can be changed
    concurrently.  Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1c2fb7f93cb2 ("[IPV4]: Sysctl configurable icmp error source address.")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d9a9433f818e6a34fcd3b3cf1d657a3bde83bd24
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:25 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_ignore_bogus_error_responses.

    [ Upstream commit b04f9b7e85c7d7aecbada620e8759a662af068d3 ]

    While reading sysctl_icmp_ignore_bogus_error_responses, it can be changed
    concurrently.  Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48fda9af1df971076416bace9bfbb5d6d89ba6c1
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:24 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_echo_ignore_broadcasts.

    [ Upstream commit 66484bb98ed2dfa1dda37a32411483d8311ac269 ]

    While reading sysctl_icmp_echo_ignore_broadcasts, it can be changed
    concurrently.  Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05c615033174f1d19374f42285ccd8e9af13e427
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
```

```
Date:   Mon Jul 11 17:15:23 2022 -0700

    icmp: Fix data-races around sysctl_icmp_echo_enable_probe.

    [ Upstream commit 4a2f7083cc6cb72dade9a63699ca352fad26d1cd ]

    While reading sysctl_icmp_echo_enable_probe, it can be changed
    concurrently.  Thus, we need to add READ_ONCE() to its readers.

    Fixes: d329ea5bd884 ("icmp: add response to RFC 8335 PROBE messages")
    Fixes: 1fd07f33c3ea ("ipv6: ICMPV6: add response to ICMPV6 RFC 8335 PROBE messages")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dde0d75e461ea61b4c31b99cfbcb45514ebf0b6a
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:22 2022 -0700

    icmp: Fix a data-race around sysctl_icmp_echo_ignore_all.

    [ Upstream commit bb7bb35a63b4812da8e3aff587773678e31d23e3 ]

    While reading sysctl_icmp_echo_ignore_all, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 12d40638f02f75c565b7906045db811c56eb2325
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:21 2022 -0700

    tcp: Fix a data-race around sysctl_max_tw_buckets.

    [ Upstream commit 6f605b57f3782114e330e108ce1903ede22ec675 ]

    While reading sysctl_max_tw_buckets, it can be changed concurrently.
    Thus, we need to add READ_ONCE() to its reader.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 93793b9abeb1655c1e133c488759768aeaf21224
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:20 2022 -0700

    sysctl: Fix data-races in proc_dointvec_ms_jiffies().

    [ Upstream commit 7d1025e559782b58824b36cb8ad547a69f2e4b31 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_dointvec_ms_jiffies() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_dointvec_ms_jiffies() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
```

```
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f776daef0b5354615ec4b4234cd9539ca05f273
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Mon Jul 11 17:15:19 2022 -0700

    sysctl: Fix data-races in proc_dou8vec_minmax().

    [ Upstream commit 7dee5d7747a69aa2be41f04c6a7ecfe3ac8cdf18 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_dou8vec_minmax() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_dou8vec_minmax() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: cb9444130662 ("sysctl: add proc_dou8vec_minmax()")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1899873812fa7586a0bd0a488a01719c878dd11c
Author: Pavan Chebbi <pavan.chebbi@broadcom.com>
Date:   Mon Jul 11 22:26:18 2022 -0400

    bnxt_en: Fix bnxt_refclk_read()

    [ Upstream commit ddde5412fdaa5048bbca31529d46cb8da882870c ]

    The upper 32-bit PHC register is not latched when reading the lower
    32-bit PHC register.  Current code leaves a small window where we may
    not read correct higher order bits if the lower order bits are just about
    to wrap around.

    This patch fixes this by reading higher order bits twice and makes
    sure that final value is correctly paired with its lower 32 bits.

    Fixes: 30e96f487f64 ("bnxt_en: Do not read the PTP PHC during chip reset")
    Cc: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: Pavan Chebbi <pavan.chebbi@broadcom.com>
    Signed-off-by: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f0c89d59e3847932106cefec029d454960a6224b
Author: Vikas Gupta <vikas.gupta@broadcom.com>
Date:   Mon Jul 11 22:26:16 2022 -0400

    bnxt_en: fix livepatch query

    [ Upstream commit 619b9b1622c283cc5ca86f4c487db266a8f55dab ]

    In the livepatch query fw_target BNXT_FW_SRT_PATCH is
    applicable for P5 chips only.

    Fixes: 3c4153394e2c ("bnxt_en: implement firmware live patching")
    Reviewed-by: Saravanan Vajravel <saravanan.vajravel@broadcom.com>
    Reviewed-by: Somnath Kotur <somnath.kotur@broadcom.com>
    Signed-off-by: Vikas Gupta <vikas.gupta@broadcom.com>
    Signed-off-by: Michael Chan <michael.chan@broadcom.com>
```

```
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5769db185f180263eba24cc1dc16f666ab8100c1
Author: Michael Chan <michael.chan@broadcom.com>
Date:   Mon Jul 11 22:26:15 2022 -0400

    bnxt_en: Fix bnxt_reinit_after_abort() code path

    [ Upstream commit 4279414bff8af9898e8c53ae6c5bc17f68ad67b7 ]

    bnxt_reinit_after_abort() is called during ifup when a previous
    FW reset sequence has aborted or a previous ifup has failed after
    detecting FW reset.  In all cases, it is safe to assume that a
    previous FW reset has completed and the driver may not have fully
    reinitialized.

    Prior to this patch, it is assumed that the
    FUNC_DRV_IF_CHANGE_RESP_FLAGS_HOT_FW_RESET_DONE flag will always be
    set by the firmware in bnxt_hwrm_if_change().  This may not be true if
    the driver has already attempted to register with the firmware.  The
    firmware may not set the RESET_DONE flag again after the driver has
    registered, assuming that the driver has seen the flag already.

    Fix it to always go through the FW reset initialization path if
    the BNXT_STATE_FW_RESET_DET flag is set.  This flag is always set
    by the driver after successfully going through bnxt_reinit_after_abort().

    Fixes: 6882c36cf82e ("bnxt_en: attempt to reinitialize after aborted reset")
    Reviewed-by: Pavan Chebbi <pavan.chebbi@broadcom.com>
    Signed-off-by: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97928c06973922ef230c90c9ef330aab8e01122a
Author: Kashyap Desai <kashyap.desai@broadcom.com>
Date:   Mon Jul 11 22:26:14 2022 -0400

    bnxt_en: reclaim max resources if sriov enable fails

    [ Upstream commit c5b744d38c36a407a41e918602eec4d89730787b ]

    If bnxt_sriov_enable() fails after some resources have been reserved
    for the VFs, the current code is not unwinding properly and the
    reserved resources become unavailable afterwards.  Fix it by
    properly unwinding with a call to bnxt_hwrm_func_qcaps() to
    reset all maximum resources.

    Also, add the missing bnxt_ulp_sriov_cfg() call to let the RDMA
    driver know to abort.

    Fixes: c0c050c58d84 ("bnxt_en: New Broadcom ethernet driver.")
    Signed-off-by: Kashyap Desai <kashyap.desai@broadcom.com>
    Signed-off-by: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8997d2d6b8d764e12489f1af2a1ce1d7384ca2a
Author: Andrzej Hajda <andrzej.hajda@intel.com>
Date:   Fri Jun 24 13:35:28 2022 +0200

    drm/i915/selftests: fix subtraction overflow bug

    [ Upstream commit 333991c4e66b3d4b5613315f18016da80344f659 ]
```

```
    On some machines hole_end can be small enough to cause subtraction
    overflow. On the other side (addr + 2 * min_alignment) can overflow
    in case of mock tests. This patch should handle both cases.

    Fixes: e1c5f754067b59 ("drm/i915: Avoid overflow in computing pot_hole loop
termination")
    Closes: https://gitlab.freedesktop.org/drm/intel/-/issues/3674
    Signed-off-by: Andrzej Hajda <andrzej.hajda@intel.com>
    Reviewed-by: Andi Shyti <andi.shyti@linux.intel.com>
    Signed-off-by: Andi Shyti <andi.shyti@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220624113528.2159210-1-
andrzej.hajda@intel.com
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    (cherry picked from commit ab3edc679c552a466e4bf0b11af3666008bd65a2)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e204eff3888ae75fffc1f7cd6f6d976d4a0a4a8
Author: Chris Wilson <chris.p.wilson@intel.com>
Date:   Tue Jul 12 16:21:33 2022 +0100

    drm/i915/gt: Serialize TLB invalidates with GT resets

    [ Upstream commit a1c5a7bf79c1faa5633b918b5c0666545e84c4d1 ]

    Avoid trying to invalidate the TLB in the middle of performing an
    engine reset, as this may result in the reset timing out. Currently,
    the TLB invalidate is only serialised by its own mutex, forgoing the
    uncore lock, but we can take the uncore->lock as well to serialise
    the mmio access, thereby serialising with the GDRST.

    Tested on a NUC5i7RYB, BIOS RYBDWi35.86A.0380.2019.0517.1530 with
    i915 selftest/hangcheck.

    Cc: stable@vger.kernel.org  # v4.4 and upper
    Fixes: 7938d61591d3 ("drm/i915: Flush TLBs before releasing backing store")
    Reported-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Tested-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Reviewed-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Chris Wilson <chris.p.wilson@intel.com>
    Cc: Tvrtko Ursulin <tvrtko.ursulin@linux.intel.com>
    Reviewed-by: Andi Shyti <andi.shyti@linux.intel.com>
    Acked-by: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Link:
https://patchwork.freedesktop.org/patch/msgid/1e59a7c45dd919a530256b9ac721ac6ea86c0677.16576
39152.git.mchehab@kernel.org
    (cherry picked from commit 33da97894758737895e90c909f16786052680ef4)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8acf1575188989cf4b9553da0926c42b128b8c5f
Author: Chris Wilson <chris@chris-wilson.co.uk>
Date:   Tue Jul 12 16:21:32 2022 +0100

    drm/i915/gt: Serialize GRDOM access between multiple engine resets

    [ Upstream commit b24dcf1dc507f69ed3b5c66c2b6a0209ae80d4d4 ]

    Don't allow two engines to be reset in parallel, as they would both
    try to select a reset bit (and send requests to common registers)
    and wait on that register, at the same time. Serialize control of
    the reset requests/acks using the uncore->lock, which will also ensure
    that no other GT state changes at the same time as the actual reset.
```

```
    Cc: stable@vger.kernel.org # v4.4 and upper
    Reported-by: Mika Kuoppala <mika.kuoppala@linux.intel.com>
    Signed-off-by: Chris Wilson <chris@chris-wilson.co.uk>
    Acked-by: Mika Kuoppala <mika.kuoppala@linux.intel.com>
    Reviewed-by: Andi Shyti <andi.shyti@intel.com>
    Reviewed-by: Andrzej Hajda <andrzej.hajda@intel.com>
    Acked-by: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Link:
https://patchwork.freedesktop.org/patch/msgid/e0a2d894e77aed7c2e36b0d1abdc7dbac3011729.16576
39152.git.mchehab@kernel.org
    (cherry picked from commit 336561a914fc0c6f1218228718f633b31b7af1c3)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 46c5f0d16e1bcfb9cb9732b0bdb9f27a4d9e1e36
Author: Matthew Auld <matthew.auld@intel.com>
Date:   Mon Jul 11 09:58:59 2022 +0100

    drm/i915/ttm: fix sg_table construction

    [ Upstream commit aff1e0b09b54b64944b7fe32997229552737b9e9 ]

    If we encounter some monster sized local-memory page that exceeds the
    maximum sg length (UINT32_MAX), ensure that don't end up with some
    misaligned address in the entry that follows, leading to fireworks
    later. Also ensure we have some coverage of this in the selftests.

    v2(Chris):
      - Use round_down consistently to avoid udiv errors
    v3(Nirmoy):
      - Also update the max_segment in the selftest

    Fixes: f701b16d4cc5 ("drm/i915/ttm: add i915_sg_from_buddy_resource")
    Closes: https://gitlab.freedesktop.org/drm/intel/-/issues/6379
    Signed-off-by: Matthew Auld <matthew.auld@intel.com>
    Cc: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Cc: Nirmoy Das <nirmoy.das@linux.intel.com>
    Reviewed-by: Nirmoy Das <nirmoy.das@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220711085859.24198-1-
matthew.auld@intel.com
    (cherry picked from commit bc99f1209f19fefa3ee11e77464ccfae541f4291)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f7eec9e63bd86dde7b2c5515f5fe24b9f73a5234
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jul 8 12:41:04 2022 +0300

    drm/i915/selftests: fix a couple IS_ERR() vs NULL tests

    [ Upstream commit 896dcabd1f8f613c533d948df17408c41f8929f5 ]

    The shmem_pin_map() function doesn't return error pointers, it returns
    NULL.

    Fixes: be1cb55a07bf ("drm/i915/gt: Keep a no-frills swappable copy of the default
context state")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Matthew Auld <matthew.auld@intel.com>
    Signed-off-by: Matthew Auld <matthew.auld@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220708094104.GL2316@kadam
    (cherry picked from commit d50f5a109cf4ed50c5b575c1bb5fc3bd17b23308)
```

```
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 51eae602908cac76d2aad238f283722d5df8eaa0
Author: Douglas Anderson <dianders@chromium.org>
Date:   Fri Jul 8 17:09:52 2022 -0700

    tracing: Fix sleeping while atomic in kdb ftdump

    [ Upstream commit 495fcec8648cdfb483b5b9ab310f3839f07cb3b8 ]

    If you drop into kdb and type "ftdump" you'll get a sleeping while
    atomic warning from memory allocation in trace_find_next_entry().

    This appears to have been caused by commit ff895103a84a ("tracing:
    Save off entry when peeking at next entry"), which added the
    allocation in that path. The problematic commit was already fixed by
    commit 8e99cf91b99b ("tracing: Do not allocate buffer in
    trace_find_next_entry() in atomic") but that fix missed the kdb case.

    The fix here is easy: just move the assignment of the static buffer to
    the place where it should have been to begin with:
    trace_init_global_iter(). That function is called in two places, once
    is right before the assignment of the static buffer added by the
    previous fix and once is in kdb.

    Note that it appears that there's a second static buffer that we need
    to assign that was added in commit efbbdaa22bb7 ("tracing: Show real
    address for trace event arguments"), so we'll move that too.

    Link:
https://lkml.kernel.org/r/20220708170919.1.I75844e5038d9425add2ad853a608cb44bb39df40@changei
d

    Fixes: ff895103a84a ("tracing: Save off entry when peeking at next entry")
    Fixes: efbbdaa22bb7 ("tracing: Show real address for trace event arguments")
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6bd9f8470735eaee2ca83339e4ef8c3bd0d261ac
Author: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
Date:   Wed Jun 8 14:48:32 2022 -0700

    ice: change devlink code to read NVM in blocks

    [ Upstream commit 7b6f9462a3234c35cf808453d39a074a04e71de1 ]

    When creating a snapshot of the NVM the driver needs to read the entire
    contents from the NVM and store it. The NVM reads are protected by a lock
    that is shared between the driver and the firmware.

    If the driver takes too long to read the entire NVM (which can happen on
    some systems) then the firmware could reclaim the lock and cause subsequent
    reads from the driver to fail.

    We could fix this by increasing the timeout that we pass to the firmware,
    but we could end up in the same situation again if the system is slow.
    Instead have the driver break the reading of the NVM into blocks that are
    small enough that we have confidence that the read will complete within the
    timeout time, but large enough not to cause significant AQ overhead.

    Fixes: dce730f17825 ("ice: add a devlink region for dumping NVM contents")
    Signed-off-by: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
    Tested-by: Gurucharan <gurucharanx.g@intel.com> (A Contingent worker at Intel)
```

```
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b3a5ec835f091a40d3250cc73422c640e13016a1
Author: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
Date:   Wed Jun 8 14:09:52 2022 -0700

    ice: handle E822 generic device ID in PLDM header

    [ Upstream commit f52d166819a4d8e0d5cca07d8a8dd6397c96dcf1 ]

    The driver currently presumes that the record data in the PLDM header
    of the firmware image will match the device ID of the running device.
    This is true for E810 devices. It appears that for E822 devices that
    this is not guaranteed to be true.

    Fix this by adding a check for the generic E822 device.

    Fixes: d69ea414c9b4 ("ice: implement device flash update via devlink")
    Signed-off-by: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
    Tested-by: Gurucharan <gurucharanx.g@intel.com> (A Contingent worker at Intel)
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0820e797c63812e355c23754ee5319cff875feba
Author: Yevhen Orlov <yevhen.orlov@plvision.eu>
Date:   Sun Jul 10 15:20:21 2022 +0300

    net: marvell: prestera: fix missed deinit sequence

    [ Upstream commit f946964a9f79f8dcb5a6329265281eebfc23aee5 ]

    Add unregister_fib_notifier as rollback of register_fib_notifier.

    Fixes: 4394fbcb78cf ("net: marvell: prestera: handle fib notifications")
    Signed-off-by: Yevhen Orlov <yevhen.orlov@plvision.eu>
    Link: https://lore.kernel.org/r/20220710122021.7642-1-yevhen.orlov@plvision.eu
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4e556ecd16413050afb87495708a0166b2915cfc
Author: Jeff Layton <jlayton@kernel.org>
Date:   Mon Jul 11 14:30:14 2022 -0400

    lockd: fix nlm_close_files

    [ Upstream commit 1197eb5906a5464dbaea24cac296dfc38499cc00 ]

    This loop condition tries a bit too hard to be clever. Just test for
    the two indices we care about explicitly.

    Cc: J. Bruce Fields <bfields@fieldses.org>
    Fixes: 7f024fcd5c97 ("Keep read and write fds with each nlm_file")
    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58c3f48a540b821b3253e2dc31c8fa237293ffc9
Author: Jeff Layton <jlayton@kernel.org>
Date:   Mon Jul 11 14:30:13 2022 -0400

    lockd: set fl_owner when unlocking files

    [ Upstream commit aec158242b87a43d83322e99bc71ab4428e5ab79 ]
```

```
    Unlocking a POSIX lock on an inode with vfs_lock_file only works if
    the owner matches. Ensure we set it in the request.

    Cc: J. Bruce Fields <bfields@fieldses.org>
    Fixes: 7f024fcd5c97 ("Keep read and write fds with each nlm_file")
    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2bcf184c16f1a1f1c27ae75c9e997543428bd667
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Sun Jul 10 14:46:04 2022 -0400

    NFSD: Decode NFSv4 birth time attribute

    [ Upstream commit 5b2f3e0777da2a5dd62824bbe2fdab1d12caaf8f ]

    NFSD has advertised support for the NFSv4 time_create attribute
    since commit e377a3e698fb ("nfsd: Add support for the birth time
    attribute").

    Igor Mammedov reports that Mac OS clients attempt to set the NFSv4
    birth time attribute via OPEN(CREATE) and SETATTR if the server
    indicates that it supports it, but since the above commit was
    merged, those attempts now fail.

    Table 5 in RFC 8881 lists the time_create attribute as one that can
    be both set and retrieved, but the above commit did not add server
    support for clients to provide a time_create attribute. IMO that's
    a bug in our implementation of the NFSv4 protocol, which this commit
    addresses.

    Whether NFSD silently ignores the new birth time or actually sets it
    is another matter. I haven't found another filesystem service in the
    Linux kernel that enables users or clients to modify a file's birth
    time attribute.

    This commit reflects my (perhaps incorrect) understanding of whether
    Linux users can set a file's birth time. NFSD will now recognize a
    time_create attribute but it ignores its value. It clears the
    time_create bit in the returned attribute bitmask to indicate that
    the value was not used.

    Reported-by: Igor Mammedov <imammedo@redhat.com>
    Fixes: e377a3e698fb ("nfsd: Add support for the birth time attribute")
    Tested-by: Igor Mammedov <imammedo@redhat.com>
    Reviewed-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff151c477f6a529b0e4643d90f4f0b8eca85de04
Author: Paolo Abeni <pabeni@redhat.com>
Date:   Fri Jul 8 16:36:09 2022 -0700

    mptcp: fix subflow traversal at disconnect time

    [ Upstream commit 5c835bb142d4013c2ab24bff5ae9f6709a39cbcf ]

    At disconnect time the MPTCP protocol traverse the subflows
    list closing each of them. In some circumstances - MPJ subflow,
    passive MPTCP socket, the latter operation can remove the
    subflow from the list, invalidating the current iterator.

    Address the issue using the safe list traversing helper
    variant.
```

```
    Reported-by: van fantasy <g1042620637@gmail.com>
    Fixes: b29fcfb54cd7 ("mptcp: full disconnect implementation")
    Tested-by: Matthieu Baerts <matthieu.baerts@tessares.net>
    Reviewed-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7fb77b292abec73345852c179a444001730f6e7e
Author: Demi Marie Obenour <demi@invisiblethingslab.com>
Date:   Sun Jul 10 19:05:22 2022 -0400

    xen/gntdev: Ignore failure to unmap INVALID_GRANT_HANDLE

    [ Upstream commit 166d3863231667c4f64dee72b77d1102cdfad11f ]

    The error paths of gntdev_mmap() can call unmap_grant_pages() even
    though not all of the pages have been successfully mapped.  This will
    trigger the WARN_ON()s in __unmap_grant_pages_done().  The number of
    warnings can be very large; I have observed thousands of lines of
    warnings in the systemd journal.

    Avoid this problem by only warning on unmapping failure if the handle
    being unmapped is not INVALID_GRANT_HANDLE.  The handle field of any
    page that was not successfully mapped will be INVALID_GRANT_HANDLE, so
    this catches all cases where unmapping can legitimately fail.

    Fixes: dbe97cff7dd9 ("xen/gntdev: Avoid blocking in unmap_grant_pages()")
    Cc: stable@vger.kernel.org
    Suggested-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Demi Marie Obenour <demi@invisiblethingslab.com>
    Reviewed-by: Oleksandr Tyshchenko <oleksandr_tyshchenko@epam.com>
    Reviewed-by: Juergen Gross <jgross@suse.com>
    Link: https://lore.kernel.org/r/20220710230522.1563-1-demi@invisiblethingslab.com
    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 778e85514b617b3a6b14d7143d88244ed853a089
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jul 8 11:41:06 2022 +0300

    drm/i915/gvt: IS_ERR() vs NULL bug in intel_gvt_update_reg_whitelist()

    [ Upstream commit e87197fbd137c888fd6c871c72fe7e89445dd015 ]

    The shmem_pin_map() function returns NULL, it doesn't return error
    pointers.

    Fixes: 97ea656521c8 ("drm/i915/gvt: Parse default state to update reg whitelist")
    Reviewed-by: Andrzej Hajda <andrzej.hajda@intel.com>
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Zhenyu Wang <zhenyuw@linux.intel.com>
    Link: http://patchwork.freedesktop.org/patch/msgid/Ysftoia2BPUyqVcD@kili
    Acked-by: Zhenyu Wang <zhenyuw@linux.intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a541767afc81d739613557fd16a535236652cfad
Author: Anup Patel <apatel@ventanamicro.com>
Date:   Mon Jul 11 09:36:32 2022 +0530

    RISC-V: KVM: Fix SRCU deadlock caused by kvm_riscv_check_vcpu_requests()

    [ Upstream commit be82abe6a76ba8e76f25312566182b0f13c4fbf9 ]
```

```
    The kvm_riscv_check_vcpu_requests() is called with SRCU read lock held
    and for KVM_REQ_SLEEP request it will block the VCPU without releasing
    SRCU read lock. This causes KVM ioctls (such as KVM_IOEVENTFD) from
    other VCPUs of the same Guest/VM to hang/deadlock if there is any
    synchronize_srcu() or synchronize_srcu_expedited() in the path.

    To fix the above in kvm_riscv_check_vcpu_requests(), we should do SRCU
    read unlock before blocking the VCPU and do SRCU read lock after VCPU
    wakeup.

    Fixes: cce69aff689e ("RISC-V: KVM: Implement VCPU interrupts and requests handling")
    Reported-by: Bin Meng <bmeng.cn@gmail.com>
    Signed-off-by: Anup Patel <apatel@ventanamicro.com>
    Reviewed-by: Atish Patra <atishp@rivosinc.com>
    Tested-by: Heinrich Schuchardt <heinrich.schuchardt@canonical.com>
    Tested-by: Bin Meng <bmeng.cn@gmail.com>
    Signed-off-by: Anup Patel <anup@brainfault.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dbef6a21e49942736e5ce1c205802274410c1211
Author: Pablo Neira Ayuso <pablo@netfilter.org>
Date:   Tue Jul 5 11:41:59 2022 +0200

    netfilter: nf_tables: replace BUG_ON by element length check

    [ Upstream commit c39ba4de6b0a843bec5d46c2b6f2064428dada5e ]

    BUG_ON can be triggered from userspace with an element with a large
    userdata area. Replace it by length check and return EINVAL instead.
    Over time extensions have been growing in size.

    Pick a sufficiently old Fixes: tag to propagate this fix.

    Fixes: 7d7402642eaf ("netfilter: nf_tables: variable sized set element keys / data")
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4c43069bb1097dd6cc1cf0f7c43a36d1f7b3910b
Author: Eric Dumazet <edumazet@google.com>
Date:   Fri Jul 8 15:11:53 2022 +0000

    vlan: fix memory leak in vlan_newlink()

    [ Upstream commit 72a0b329114b1caa8e69dfa7cdad1dd3c69b8602 ]

    Blamed commit added back a bug I fixed in commit 9bbd917e0bec
    ("vlan: fix memory leak in vlan_dev_set_egress_priority")

    If a memory allocation fails in vlan_changelink() after other allocations
    succeeded, we need to call vlan_dev_free_egress_priority()
    to free all allocated memory because after a failed ->newlink()
    we do not call any methods like ndo_uninit() or dev->priv_destructor().

    In following example, if the allocation for last element 2000:2001 fails,
    we need to free eight prior allocations:

    ip link add link dummy0 dummy0.100 type vlan id 100 \
            egress-qos-map 1:2 2:3 3:4 4:5 5:6 6:7 7:8 8:9 2000:2001

    syzbot report was:

    BUG: memory leak
    unreferenced object 0xffff888117bd1060 (size 32):
    comm "syz-executor408", pid 3759, jiffies 4294956555 (age 34.090s)
```

```
    hex dump (first 32 bytes):
    09 00 00 00 00 a0 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    backtrace:
    [<ffffffff83fc60ad>] kmalloc include/linux/slab.h:600 [inline]
    [<ffffffff83fc60ad>] vlan_dev_set_egress_priority+0xed/0x170 net/8021q/vlan_dev.c:193
    [<ffffffff83fc6628>] vlan_changelink+0x178/0x1d0 net/8021q/vlan_netlink.c:128
    [<ffffffff83fc67c8>] vlan_newlink+0x148/0x260 net/8021q/vlan_netlink.c:185
    [<ffffffff838b1278>] rtnl_newlink_create net/core/rtnetlink.c:3363 [inline]
    [<ffffffff838b1278>] __rtnl_newlink+0xa58/0xdc0 net/core/rtnetlink.c:3580
    [<ffffffff838b1629>] rtnl_newlink+0x49/0x70 net/core/rtnetlink.c:3593
    [<ffffffff838ac66c>] rtnetlink_rcv_msg+0x21c/0x5c0 net/core/rtnetlink.c:6089
    [<ffffffff839f9c37>] netlink_rcv_skb+0x87/0x1d0 net/netlink/af_netlink.c:2501
    [<ffffffff839f8da7>] netlink_unicast_kernel net/netlink/af_netlink.c:1319 [inline]
    [<ffffffff839f8da7>] netlink_unicast+0x397/0x4c0 net/netlink/af_netlink.c:1345
    [<ffffffff839f9266>] netlink_sendmsg+0x396/0x710 net/netlink/af_netlink.c:1921
    [<ffffffff8384dbf6>] sock_sendmsg_nosec net/socket.c:714 [inline]
    [<ffffffff8384dbf6>] sock_sendmsg+0x56/0x80 net/socket.c:734
    [<ffffffff8384e15c>] ____sys_sendmsg+0x36c/0x390 net/socket.c:2488
    [<ffffffff838523cb>] ___sys_sendmsg+0x8b/0xd0 net/socket.c:2542
    [<ffffffff838525b8>] __sys_sendmsg net/socket.c:2571 [inline]
    [<ffffffff838525b8>] __do_sys_sendmsg net/socket.c:2580 [inline]
    [<ffffffff838525b8>] __se_sys_sendmsg net/socket.c:2578 [inline]
    [<ffffffff838525b8>] __x64_sys_sendmsg+0x78/0xf0 net/socket.c:2578
    [<ffffffff845ad8d5>] do_syscall_x64 arch/x86/entry/common.c:50 [inline]
    [<ffffffff845ad8d5>] do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80
    [<ffffffff8460006a>] entry_SYSCALL_64_after_hwframe+0x46/0xb0

    Fixes: 37aa50c539bc ("vlan: introduce vlan_dev_free_egress_priority")
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Xin Long <lucien.xin@gmail.com>
    Reviewed-by: Xin Long <lucien.xin@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e75554f74bce38b9e8e0cd8147c8a677e97d410a
Author: Baowen Zheng <baowen.zheng@corigine.com>
Date:   Fri Jul 8 11:07:18 2022 +0100

    nfp: fix issue of skb segments exceeds descriptor limitation

    [ Upstream commit 9c840d5f9aaef87e65db900bae21c70b059aba5f ]

    TCP packets will be dropped if the segments number in the tx skb
    exceeds limitation when sending iperf3 traffic with --zerocopy option.

    we make the following changes:

    Get nr_frags in nfp_nfdk_tx_maybe_close_block instead of passing from
    outside because it will be changed after skb_linearize operation.

    Fill maximum dma_len in first tx descriptor to make sure the whole
    head is included in the first descriptor.

    Fixes: c10d12e3dce8 ("nfp: add support for NFDK data path")
    Signed-off-by: Baowen Zheng <baowen.zheng@corigine.com>
    Reviewed-by: Louis Peens <louis.peens@corigine.com>
    Signed-off-by: Simon Horman <simon.horman@corigine.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ace115965b7d04bab14b90082d457a3ae63f783
Author: Pablo Neira Ayuso <pablo@netfilter.org>
Date:   Tue Jul 5 10:26:15 2022 +0200
```

```
    netfilter: nf_log: incorrect offset to network header

    [ Upstream commit 7a847c00eeba9744353ecdfad253143b9115678a ]

    NFPROTO_ARP is expecting to find the ARP header at the network offset.

    In the particular case of ARP, HTYPE= field shows the initial bytes of
    the ethernet header destination MAC address.

     netdev out: IN= OUT=bridge0 MACSRC=c2:76:e5:71:e1:de MACDST=36:b0:4a:e2:72:ea
MACPROTO=0806 ARP HTYPE=14000 PTYPE=0x4ae2 OPCODE=49782

    NFPROTO_NETDEV egress hook is also expecting to find the IP headers at
    the network offset.

    Fixes: 35b9395104d5 ("netfilter: add generic ARP packet logger")
    Reported-by: Tom Yan <tom.ty89@gmail.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7fc7c6d053cfca70bb81892f3f00937e5c459d5a
Author: William Zhang <william.zhang@broadcom.com>
Date:   Fri Jul 8 11:25:07 2022 -0700

    arm64: dts: broadcom: bcm4908: Fix cpu node for smp boot

    [ Upstream commit 8bd582ae9a71d7f14c4e0c735b2eacaf7516d626 ]

    Add spin-table enable-method and cpu-release-addr properties for
    cpu0 node. This is required by all ARMv8 SoC. Otherwise some
    bootloader like u-boot can not update cpu-release-addr and linux
    fails to start up secondary cpus.

    Fixes: 2961f69f151c ("arm64: dts: broadcom: add BCM4908 and Asus GT-AC5300 early DTS
files")
    Signed-off-by: William Zhang <william.zhang@broadcom.com>
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 658a67f68bca44d922bbd786080535a69e1ac000
Author: William Zhang <william.zhang@broadcom.com>
Date:   Fri Jul 8 11:25:06 2022 -0700

    arm64: dts: broadcom: bcm4908: Fix timer node for BCM4906 SoC

    [ Upstream commit b4a544e415e9be33b37d9bfa9d9f9f4d13f553d6 ]

    The cpu mask value in interrupt property inherits from bcm4908.dtsi
    which sets to four cpus. Correct the value to two cpus for dual core
    BCM4906 SoC.

    Fixes: c8b404fb05dc ("arm64: dts: broadcom: bcm4908: add BCM4906 Netgear R8000P DTS
files")
    Signed-off-by: William Zhang <william.zhang@broadcom.com>
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bee4dc387f64009011953191ea0e7bbab1642632
Author: Michal Suchanek <msuchanek@suse.de>
Date:   Fri Jul 8 19:45:29 2022 +0200

    ARM: dts: sunxi: Fix SPI NOR campatible on Orange Pi Zero

    [ Upstream commit 884b66976a7279ee889ba885fe364244d50b79e7 ]
```

```
    The device tree should include generic "jedec,spi-nor" compatible, and a
    manufacturer-specific one.
    The macronix part is what is shipped on the boards that come with a
    flash chip.

    Fixes: 45857ae95478 ("ARM: dts: orange-pi-zero: add node for SPI NOR")
    Signed-off-by: Michal Suchanek <msuchanek@suse.de>
    Acked-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Signed-off-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Link: https://lore.kernel.org/r/20220708174529.3360-1-msuchanek@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c0328380d0efeae8a100c39273e4860bf213bfea
Author: Ryan Wanner <Ryan.Wanner@microchip.com>
Date:   Thu Jul 7 14:58:12 2022 -0700

    ARM: dts: at91: sama5d2: Fix typo in i2s1 node

    [ Upstream commit 2fdf15b50a46e366740df4cccbe2343269b4ff55 ]

    Fix typo in i2s1 causing errors in dt binding validation.
    Change assigned-parrents to assigned-clock-parents
    to match i2s0 node formatting.

    Fixes: 1ca81883c557 ("ARM: dts: at91: sama5d2: add nodes for I2S controllers")
    Signed-off-by: Ryan Wanner <Ryan.Wanner@microchip.com>
    [claudiu.beznea: use imperative addressing in commit description, remove
     blank line after fixes tag, fix typo in commit message]
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20220707215812.193008-1-Ryan.Wanner@microchip.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 190cd4ff128373271e065afb20f1d2247b3f10c3
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:40:03 2022 -0700

    ipv4: Fix a data-race around sysctl_fib_sync_mem.

    [ Upstream commit 73318c4b7dbd0e781aaababff17376b2894745c0 ]

    While reading sysctl_fib_sync_mem, it can be changed concurrently.
    So, we need to add READ_ONCE() to avoid a data-race.

    Fixes: 9ab948a91b2c ("ipv4: Allow amount of dirty memory from fib resizing to be
controllable")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 798c2cf57c63ab39c8aac24d6a3d50f4fa5eeb06
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:40:02 2022 -0700

    icmp: Fix data-races around sysctl.

    [ Upstream commit 48d7ee321ea5182c6a70782aa186422a70e67e22 ]

    While reading icmp sysctl variables, they can be changed concurrently.
    So, we need to add READ_ONCE() to avoid data-races.

    Fixes: 4cdf507d5452 ("icmp: add a global rate limitation")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 59e26906b89cc35bb54476498772b45cbc32323f
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:40:01 2022 -0700

    cipso: Fix data-races around sysctl.

    [ Upstream commit dd44f04b9214adb68ef5684ae87a81ba03632250 ]

    While reading cipso sysctl variables, they can be changed concurrently.
    So, we need to add READ_ONCE() to avoid data-races.

    Fixes: 446fda4f2682 ("[NetLabel]: CIPSOv4 engine")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Acked-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 28a912218350ba3de991ed23dbdb51006ffece8b
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:40:00 2022 -0700

    net: Fix data-races around sysctl_mem.

    [ Upstream commit 310731e2f1611d1d13aae237abcf8e66d33345d5 ]

    While reading .sysctl_mem, it can be changed concurrently.
    So, we need to add READ_ONCE() to avoid data-races.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9b8ef81eb73a592e407e38c80ff0ddd042c514d
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:59 2022 -0700

    inetpeer: Fix data-races around sysctl.

    [ Upstream commit 3d32edf1f3c38d3301f6434e56316f293466d7fb ]

    While reading inetpeer sysctl variables, they can be changed
    concurrently.  So, we need to add READ_ONCE() to avoid data-races.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e293ebc8b3bfb4cc45614bcf789a766f9b13e158
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:58 2022 -0700

    tcp: Fix a data-race around sysctl_tcp_max_orphans.

    [ Upstream commit 47e6ab24e8c6e3ca10ceb5835413f401f90de4bf ]

    While reading sysctl_tcp_max_orphans, it can be changed concurrently.
    So, we need to add READ_ONCE() to avoid a data-race.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 1908c3d91b4927a03713287bcbe75bf8fa5c7ead
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:57 2022 -0700

    sysctl: Fix data races in proc_dointvec_jiffies().

    [ Upstream commit e877820877663fbae8cb9582ea597a7230b94df3 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_dointvec_jiffies() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_dointvec_jiffies() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8309d12bc3098ca8e56ab6c27a97f134c33a3dec
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:56 2022 -0700

    sysctl: Fix data races in proc_doulongvec_minmax().

    [ Upstream commit c31bcc8fb89fc2812663900589c6325ba35d9a65 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_doulongvec_minmax() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_doulongvec_minmax() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b60eddf98b9716651069dfda296c91311a7a6293
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:55 2022 -0700

    sysctl: Fix data races in proc_douintvec_minmax().

    [ Upstream commit 2d3b559df3ed39258737789aae2ae7973d205bc1 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_douintvec_minmax() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_douintvec_minmax() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: 61d9b56a8920 ("sysctl: add unsigned int range support")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
```

```
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ebc99cab28f52c3ed27f67c0c27e21a4964aebb
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:54 2022 -0700

    sysctl: Fix data races in proc_dointvec_minmax().

    [ Upstream commit f613d86d014b6375a4085901de39406598121e35 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_dointvec_minmax() to use READ_ONCE() and
    WRITE_ONCE() internally to fix data-races on the sysctl side.  For now,
    proc_dointvec_minmax() itself is tolerant to a data-race, but we still
    need to add annotations on the other subsystem's side.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 630c76850d554d7140232e71b5d1663e88cffb54
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:53 2022 -0700

    sysctl: Fix data races in proc_douintvec().

    [ Upstream commit 4762b532ec9539755aab61445d5da6e1926ccb99 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_douintvec() to use READ_ONCE() and WRITE_ONCE()
    internally to fix data-races on the sysctl side.  For now, proc_douintvec()
    itself is tolerant to a data-race, but we still need to add annotations on
    the other subsystem's side.

    Fixes: e7d316a02f68 ("sysctl: handle error writing UINT_MAX to u32 fields")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bfbb0be2c8a40a1a7421942fa3d567b2e956aa56
Author: Kuniyuki Iwashima <kuniyu@amazon.com>
Date:   Wed Jul 6 16:39:52 2022 -0700

    sysctl: Fix data races in proc_dointvec().

    [ Upstream commit 1f1be04b4d48a2475ea1aab46a99221bfc5c0968 ]

    A sysctl variable is accessed concurrently, and there is always a chance
    of data-race.  So, all readers and writers need some basic protection to
    avoid load/store-tearing.

    This patch changes proc_dointvec() to use READ_ONCE() and WRITE_ONCE()
    internally to fix data-races on the sysctl side.  For now, proc_dointvec()
    itself is tolerant to a data-race, but we still need to add annotations on
    the other subsystem's side.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
```

```
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 62acd1e2fe22e7ac4dccd17855fe65ca6c90a4b8
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Wed Jul 6 16:28:45 2022 +0300

    net: ocelot: fix wrong time_after usage

    [ Upstream commit f46fd3d7c3bd5d7bd5bb664135cf32ca9e97190b ]

    Accidentally noticed, that this driver is the only user of
    while (time_after(jiffies...)).

    It looks like typo, because likely this while loop will finish after 1st
    iteration, because time_after() returns true when 1st argument _is after_
    2nd one.

    There is one possible problem with this poll loop: the scheduler could put
    the thread to sleep, and it does not get woken up for
    OCELOT_FDMA_CH_SAFE_TIMEOUT_US. During that time, the hardware has done
    its thing, but you exit the while loop and return -ETIMEDOUT.

    Fix it by using sane poll API that avoids all problems described above

    Fixes: 753a026cfec1 ("net: ocelot: add FDMA support")
    Suggested-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Reviewed-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Link: https://lore.kernel.org/r/20220706132845.27968-1-paskripkin@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 79d60284bcd2d082a56f45ed716b93d78c24e12e
Author: Siddharth Vadapalli <s-vadapalli@ti.com>
Date:   Wed Jul 6 12:32:08 2022 +0530

    net: ethernet: ti: am65-cpsw: Fix devlink port register sequence

    [ Upstream commit 0680e20af5fbf41df8a11b11bd9a7c25b2ca0746 ]

    Renaming interfaces using udevd depends on the interface being registered
    before its netdev is registered. Otherwise, udevd reads an empty
    phys_port_name value, resulting in the interface not being renamed.

    Fix this by registering the interface before registering its netdev
    by invoking am65_cpsw_nuss_register_devlink() before invoking
    register_netdev() for the interface.

    Move the function call to devlink_port_type_eth_set(), invoking it after
    register_netdev() is invoked, to ensure that netlink notification for the
    port state change is generated after the netdev is completely initialized.

    Fixes: 58356eb31d60 ("net: ti: am65-cpsw-nuss: Add devlink support")
    Signed-off-by: Siddharth Vadapalli <s-vadapalli@ti.com>
    Link: https://lore.kernel.org/r/20220706070208.12207-1-s-vadapalli@ti.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfa4caf3e881ad6dd366c903c34f1c7f21b857ab
Author: Jon Hunter <jonathanh@nvidia.com>
Date:   Wed Jul 6 09:39:13 2022 +0100

    net: stmmac: dwc-qos: Disable split header for Tegra194
```

```
    [ Upstream commit 029c1c2059e9c4b38f97a06204cdecd10cfbeb8a ]

    There is a long-standing issue with the Synopsys DWC Ethernet driver
    for Tegra194 where random system crashes have been observed [0]. The
    problem occurs when the split header feature is enabled in the stmmac
    driver. In the bad case, a larger than expected buffer length is
    received and causes the calculation of the total buffer length to
    overflow. This results in a very large buffer length that causes the
    kernel to crash. Why this larger buffer length is received is not clear,
    however, the feedback from the NVIDIA design team is that the split
    header feature is not supported for Tegra194. Therefore, disable split
    header support for Tegra194 to prevent these random crashes from
    occurring.

    [0] https://lore.kernel.org/linux-tegra/b0b17697-f23e-8fa5-3757-604a86f3a095@nvidia.com/

    Fixes: 67afd6d1cfdf ("net: stmmac: Add Split Header support and enable it in XGMAC
cores")
    Signed-off-by: Jon Hunter <jonathanh@nvidia.com>
    Link: https://lore.kernel.org/r/20220706083913.13750-1-jonathanh@nvidia.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac5e76476d3c29704fbbc9783178a7751091cd86
Author: Florian Westphal <fw@strlen.de>
Date:   Wed Jul 6 16:50:04 2022 +0200

    netfilter: conntrack: fix crash due to confirmed bit load reordering

    [ Upstream commit 0ed8f619b412b52360ccdfaf997223ccd9319569 ]

    Kajetan Puchalski reports crash on ARM, with backtrace of:

    __nf_ct_delete_from_lists
    nf_ct_delete
    early_drop
    __nf_conntrack_alloc

    Unlike atomic_inc_not_zero, refcount_inc_not_zero is not a full barrier.
    conntrack uses SLAB_TYPESAFE_BY_RCU, i.e. it is possible that a 'newly'
    allocated object is still in use on another CPU:

    CPU1                                                    CPU2
                                                            encounter 'ct' during hlist walk
     delete_from_lists
     refcount drops to 0
     kmem_cache_free(ct);
     __nf_conntrack_alloc() // returns same object
                                                            refcount_inc_not_zero(ct); /* might fail
*/

                                                            /* If set, ct is public/in the hash
table */
                                                            test_bit(IPS_CONFIRMED_BIT, &ct-
>status);

    In case CPU1 already set refcount back to 1, refcount_inc_not_zero()
    will succeed.

    The expected possibilities for a CPU that obtained the object 'ct'
    (but no reference so far) are:

    1. refcount_inc_not_zero() fails.  CPU2 ignores the object and moves to
       the next entry in the list.  This happens for objects that are about
       to be free'd, that have been free'd, or that have been reallocated
```

```
        by __nf_conntrack_alloc(), but where the refcount has not been
        increased back to 1 yet.

     2. refcount_inc_not_zero() succeeds. CPU2 checks the CONFIRMED bit
        in ct->status.  If set, the object is public/in the table.

        If not, the object must be skipped; CPU2 calls nf_ct_put() to
        un-do the refcount increment and moves to the next object.

     Parallel deletion from the hlists is prevented by a
     'test_and_set_bit(IPS_DYING_BIT, &ct->status);' check, i.e. only one
     cpu will do the unlink, the other one will only drop its reference count.

     Because refcount_inc_not_zero is not a full barrier, CPU2 may try to
     delete an object that is not on any list:

     1. refcount_inc_not_zero() successful (refcount inited to 1 on other CPU)
     2. CONFIRMED test also successful (load was reordered or zeroing
        of ct->status not yet visible)
     3. delete_from_lists unlinks entry not on the hlist, because
        IPS_DYING_BIT is 0 (already cleared).

     2) is already wrong: CPU2 will handle a partially initited object
     that is supposed to be private to CPU1.

     Add needed barriers when refcount_inc_not_zero() is successful.

     It also inserts a smp_wmb() before the refcount is set to 1 during
     allocation.

     Because other CPU might still see the object, refcount_set(1)
     "resurrects" it, so we need to make sure that other CPUs will also observe
     the right content.  In particular, the CONFIRMED bit test must only pass
     once the object is fully initialised and either in the hash or about to be
     inserted (with locks held to delay possible unlink from early_drop or
     gc worker).

     I did not change flow_offload_alloc(), as far as I can see it should call
     refcount_inc(), not refcount_inc_not_zero(): the ct object is attached to
     the skb so its refcount should be >= 1 in all cases.

     v2: prefer smp_acquire__after_ctrl_dep to smp_rmb (Will Deacon).
     v3: keep smp_acquire__after_ctrl_dep close to refcount_inc_not_zero call
         add comment in nf_conntrack_netlink, no control dependency there
         due to locks.

     Cc: Peter Zijlstra <peterz@infradead.org>
     Link: https://lore.kernel.org/all/Yr7WTfd6AVTQkLjI@e126311.manchester.arm.com/
     Reported-by: Kajetan Puchalski <kajetan.puchalski@arm.com>
     Diagnosed-by: Will Deacon <will@kernel.org>
     Fixes: 719774377622 ("netfilter: conntrack: convert to refcount_t api")
     Signed-off-by: Florian Westphal <fw@strlen.de>
     Acked-by: Will Deacon <will@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 594cea2c09f7cd440d1ee1c4547d5bc6a646b0e4
Author: Florian Westphal <fw@strlen.de>
Date:   Mon Apr 11 13:01:18 2022 +0200

     netfilter: conntrack: remove the percpu dying list

     [ Upstream commit 1397af5bfd7d32b0cf2adb70a78c9a9e8f11d912 ]

     Its no longer needed. Entries that need event redelivery are placed
     on the new pernet dying list.
```

```
    The advantage is that there is no need to take additional spinlock on
    conntrack removal unless event redelivery failed or the conntrack entry
    was never added to the table in the first place (confirmed bit not set).

    The IPS_CONFIRMED bit now needs to be set as soon as the entry has been
    unlinked from the unconfirmed list, else the destroy function may
    attempt to unlink it a second time.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be4dd63974cf215fd10b3a411a7de4013f78293b
Author: Florian Westphal <fw@strlen.de>
Date:   Mon Apr 11 13:01:17 2022 +0200

    netfilter: conntrack: include ecache dying list in dumps

    [ Upstream commit 0d3cc504ba9cdcff76346306c37eb1ea01e60a86 ]

    The new pernet dying list includes conntrack entries that await
    delivery of the 'destroy' event via ctnetlink.

    The old percpu dying list will be removed soon.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 87be95810c083816e7fd928009a319ae9fb10c2a
Author: Florian Westphal <fw@strlen.de>
Date:   Mon Apr 11 13:01:16 2022 +0200

    netfilter: ecache: use dedicated list for event redelivery

    [ Upstream commit 2ed3bf188b33630cf9d93b996ebf001847a00b5a ]

    This disentangles event redelivery and the percpu dying list.

    Because entries are now stored on a dedicated list, all
    entries are in NFCT_ECACHE_DESTROY_FAIL state and all entries
    still have confirmed bit set -- the reference count is at least 1.

    The 'struct net' back-pointer can be removed as well.

    The pcpu dying list will be removed eventually, it has no functionality.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b900c5eaeef158deef762ea69f90675e844fa44f
Author: Florian Westphal <fw@strlen.de>
Date:   Wed Mar 23 14:22:03 2022 +0100

    netfilter: conntrack: split inner loop of list dumping to own function

    [ Upstream commit 49001a2e83a80f6d9c4287c46ffa41a03667bbd1 ]

    This allows code re-use in the followup patch.
    No functional changes intended.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
```

```
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 60d85bb77afd7b108e4b60b95886c34f6d3d39af
Author: Florian Westphal <fw@strlen.de>
Date:   Wed Mar 23 14:22:01 2022 +0100

    netfilter: ecache: move to separate structure

    [ Upstream commit 9027ce0b071a1bbd046682907fc2e23ca3592883 ]

    This makes it easier for a followup patch to only expose ecache
    related parts of nf_conntrack_net structure.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ba8d419b95a70d32823bf5755c06519efc6373f
Author: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
Date:   Thu Jun 30 09:56:38 2022 +0300

    ASoC: Intel: Skylake: Correct the handling of fmt_config flexible array

    [ Upstream commit fc976f5629afb4160ee77798b14a693eac903ffd ]

    The struct nhlt_format's fmt_config is a flexible array, it must not be
    used as normal array.
    When moving to the next nhlt_fmt_cfg we need to take into account the data
    behind the ->config.caps (indicated by ->config.size).

    The logic of the code also changed: it is no longer saves the _last_
    fmt_cfg for all found rates.

    Fixes: bc2bd45b1f7f3 ("ASoC: Intel: Skylake: Parse nhlt and register clock device")
    Signed-off-by: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
    Reviewed-by: Cezary Rojewski <cezary.rojewski@intel.com>
    Link: https://lore.kernel.org/r/20220630065638.11183-3-peter.ujfalusi@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6777a8e60e8c7ce7103be41224839d75cd0151c9
Author: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
Date:   Thu Jun 30 09:56:37 2022 +0300

    ASoC: Intel: Skylake: Correct the ssp rate discovery in skl_get_ssp_clks()

    [ Upstream commit 219af251bd1694bce1f627d238347d2eaf13de61 ]

    The present flag is only set once when one rate has been found to be saved.
    This will effectively going to ignore any rate discovered at later time and
    based on the code, this is not the intention.

    Fixes: bc2bd45b1f7f3 ("ASoC: Intel: Skylake: Parse nhlt and register clock device")
    Signed-off-by: Peter Ujfalusi <peter.ujfalusi@linux.intel.com>
    Reviewed-by: Cezary Rojewski <cezary.rojewski@intel.com>
    Link: https://lore.kernel.org/r/20220630065638.11183-2-peter.ujfalusi@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b174cca5647f14faca87bc52c85a68a93511682a
Author: Hector Martin <marcan@marcan.st>
Date:   Thu Jun 30 09:51:35 2022 +0200

    ASoC: tas2764: Fix amp gain register offset & default
```

```
    [ Upstream commit 1c4f29ec878bbf1cc0a1eb54ae7da5ff98e19641 ]

    The register default is 0x28 per the datasheet, and the amp gain field
    is supposed to be shifted left by one. With the wrong default, the ALSA
    controls lie about the power-up state. With the wrong shift, we get only
    half the gain we expect.

    Signed-off-by: Hector Martin <marcan@marcan.st>
    Fixes: 827ed8a0fa50 ("ASoC: tas2764: Add the driver for the TAS2764")
    Signed-off-by: Martin Povišer <povik+lin@cutebit.org>
    Link: https://lore.kernel.org/r/20220630075135.2221-4-povik+lin@cutebit.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc4a817fdcda4b70304357373a5acdd29ca9e92d
Author: Hector Martin <marcan@marcan.st>
Date:   Thu Jun 30 09:51:34 2022 +0200

    ASoC: tas2764: Correct playback volume range

    [ Upstream commit 3e99e5697e1f7120b5abc755e8a560b22612d6ed ]

    DVC value 0xc8 is -100dB and 0xc9 is mute; this needs to map to
    -100.5dB as far as the dB scale is concerned. Fix that and enable
    the mute flag, so alsamixer correctly shows the control as
    <0 dB .. -100 dB, mute>.

    Signed-off-by: Hector Martin <marcan@marcan.st>
    Fixes: 827ed8a0fa50 ("ASoC: tas2764: Add the driver for the TAS2764")
    Signed-off-by: Martin Povišer <povik+lin@cutebit.org>
    Link: https://lore.kernel.org/r/20220630075135.2221-3-povik+lin@cutebit.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a851ee8091ab383e5d0cd495410f6495959544d
Author: Martin Povišer <povik+lin@cutebit.org>
Date:   Thu Jun 30 09:51:33 2022 +0200

    ASoC: tas2764: Fix and extend FSYNC polarity handling

    [ Upstream commit d1a10f1b48202e2d183cce144c218a211e98d906 ]

    Fix setting of FSYNC polarity in case of LEFT_J and DSP_A/B formats.
    Do NOT set the SCFG field as was previously done, because that is not
    correct and is also in conflict with the "ASI1 Source" control which
    sets the same SCFG field!

    Also add support for explicit polarity inversion.

    Fixes: 827ed8a0fa50 ("ASoC: tas2764: Add the driver for the TAS2764")
    Signed-off-by: Martin Povišer <povik+lin@cutebit.org>
    Link: https://lore.kernel.org/r/20220630075135.2221-2-povik+lin@cutebit.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd486308d773d6d062a0140062458b48f8a0eb6b
Author: Martin Povišer <povik+lin@cutebit.org>
Date:   Thu Jun 30 09:51:32 2022 +0200

    ASoC: tas2764: Add post reset delays

    [ Upstream commit cd10bb89b0d57bca98eb75e0444854a1c129a14e ]

    Make sure there is at least 1 ms delay from reset to first command as
    is specified in the datasheet. This is a fix similar to commit
```

```
    307f31452078 ("ASoC: tas2770: Insert post reset delay").

    Fixes: 827ed8a0fa50 ("ASoC: tas2764: Add the driver for the TAS2764")
    Signed-off-by: Martin Povišer <povik+lin@cutebit.org>
    Link: https://lore.kernel.org/r/20220630075135.2221-1-povik+lin@cutebit.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f4f11e5efcf56ee8bdd9c92df6ec7021ea13082d
Author: Francesco Dolcini <francesco.dolcini@toradex.com>
Date:   Fri Jun 24 12:13:01 2022 +0200

    ASoC: sgtl5000: Fix noise on shutdown/remove

    [ Upstream commit 040e3360af3736348112d29425bf5d0be5b93115 ]

    Put the SGTL5000 in a silent/safe state on shutdown/remove, this is
    required since the SGTL5000 produces a constant noise on its output
    after it is configured and its clock is removed. Without this change
    this is happening every time the module is unbound/removed or from
    reboot till the clock is enabled again.

    The issue was experienced on both a Toradex Colibri/Apalis iMX6, but can
    be easily reproduced everywhere just playing something on the codec and
    after that removing/unbinding the driver.

    Fixes: 9b34e6cc3bc2 ("ASoC: Add Freescale SGTL5000 codec support")
    Signed-off-by: Francesco Dolcini <francesco.dolcini@toradex.com>
    Reviewed-by: Fabio Estevam <festevam@denx.de>
    Link: https://lore.kernel.org/r/20220624101301.441314-1-francesco.dolcini@toradex.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 640cea4c2839a821adfbb703b590a5928abe9286
Author: Huaxin Lu <luhuaxin1@huawei.com>
Date:   Tue Jul 5 13:14:17 2022 +0800

    ima: Fix a potential integer overflow in ima_appraise_measurement

    [ Upstream commit d2ee2cfc4aa85ff6a2a3b198a3a524ec54e3d999 ]

    When the ima-modsig is enabled, the rc passed to evm_verifyxattr() may be
    negative, which may cause the integer overflow problem.

    Fixes: 39b07096364a ("ima: Implement support for module-style appended signatures")
    Signed-off-by: Huaxin Lu <luhuaxin1@huawei.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 683feeb567732731238f620d181aede6b176da58
Author: Daniele Ceraolo Spurio <daniele.ceraolospurio@intel.com>
Date:   Tue Jun 21 16:30:05 2022 -0700

    drm/i915/guc: ADL-N should use the same GuC FW as ADL-S

    [ Upstream commit 25c95bf494067f7bd1dfa8064ef964abe88cafc2 ]

    The only difference between the ADL S and P GuC FWs is the HWConfig
    support. ADL-N does not support HWConfig, so we should use the same
    binary as ADL-S, otherwise the GuC might attempt to fetch a config
    table that does not exist. ADL-N is internally identified as an ADL-P,
    so we need to special-case it in the FW selection code.

    Fixes: 7e28d0b26759 ("drm/i915/adl-n: Enable ADL-N platform")
    Cc: John Harrison <John.C.Harrison@Intel.com>
```

```
    Cc: Tejas Upadhyay <tejas.upadhyay@intel.com>
    Cc: Anusha Srivatsa <anusha.srivatsa@intel.com>
    Cc: Jani Nikula <jani.nikula@intel.com>
    Signed-off-by: Daniele Ceraolo Spurio <daniele.ceraolospurio@intel.com>
    Reviewed-by: Matt Roper <matthew.d.roper@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220621233005.3952293-1-
daniele.ceraolospurio@intel.com
    (cherry picked from commit 971e4a9781742aaad1587e25fd5582b2dd595ef8)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a91522b4279bebb098106a19b91f82b9c3213be9
Author: Hangyu Hua <hbh25y@gmail.com>
Date:   Fri Jun 24 06:04:06 2022 -0700

    drm/i915: fix a possible refcount leak in intel_dp_add_mst_connector()

    [ Upstream commit 85144df9ff4652816448369de76897c57cbb1b93 ]

    If drm_connector_init fails, intel_connector_free will be called to take
    care of proper free. So it is necessary to drop the refcount of port
    before intel_connector_free.

    Fixes: 091a4f91942a ("drm/i915: Handle drm-layer errors in intel_dp_add_mst_connector")
    Signed-off-by: Hangyu Hua <hbh25y@gmail.com>
    Reviewed-by: José Roberto de Souza <jose.souza@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220624130406.17996-1-
jose.souza@intel.com
    Signed-off-by: José Roberto de Souza <jose.souza@intel.com>
    (cherry picked from commit cea9ed611e85d36a05db52b6457bf584b7d969e2)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ffcacab7145080187330accafae69e87a481eec
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Tue Jun 21 10:10:37 2022 -0400

    drm/amdgpu/display: disable prefer_shadow for generic fb helpers

    [ Upstream commit 3a4b1cc28fbdc2325b3e3ed7d8024995a75f9216 ]

    Seems to break hibernation.  Disable for now until we can root
    cause it.

    Fixes: 087451f372bf ("drm/amdgpu: use generic fb helpers instead of setting up AMD
own's.")
    Bug: https://bugzilla.kernel.org/show_bug.cgi?id=216119
    Acked-by: Evan Quan <evan.quan@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 27dccf616a0a82f4d8004b7ee04560e7de419e63
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Tue Jun 21 10:04:55 2022 -0400

    drm/amdgpu: keep fbdev buffers pinned during suspend

    [ Upstream commit f9a89117fbdc63c0d4ab63a8f3596a72c245bcfe ]

    Was dropped when we converted to the generic helpers.

    Fixes: 087451f372bf ("drm/amdgpu: use generic fb helpers instead of setting up AMD
own's.")
    Acked-by: Evan Quan <evan.quan@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
```

```
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16427298f3dc02ec90bdfa31c8ef9b384ea5534a
Author: Maxim Mikityanskiy <maximmi@nvidia.com>
Date:   Mon May 30 14:01:37 2022 +0300

    net/mlx5e: Ring the TX doorbell on DMA errors

    [ Upstream commit 5b759bf2f9d73db05369aef2344502095c4e5e73 ]

    TX doorbells may be postponed, because sometimes the driver knows that
    another packet follows (for example, when xmit_more is true, or when a
    MPWQE session is closed before transmitting a packet).

    However, the DMA mapping may fail for the next packet, in which case a
    new WQE is not posted, the doorbell isn't updated either, and the
    transmission of the previous packet will be delayed indefinitely.

    This commit fixes the described rare error flow by posting a NOP and
    ringing the doorbell on errors to flush all the previous packets. The
    MPWQE session is closed before that. DMA mapping in the MPWQE flow is
    moved to the beginning of mlx5e_sq_xmit_mpwqe, because empty sessions
    are not allowed. Stop room always has enough space for a NOP, because
    the actual TX WQE is not posted.

    Fixes: e586b3b0baee ("net/mlx5: Ethernet Datapath files")
    Signed-off-by: Maxim Mikityanskiy <maximmi@nvidia.com>
    Reviewed-by: Tariq Toukan <tariqt@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 987df1259671aa40f791141bfe0c7d388df27816
Author: Gal Pressman <gal@nvidia.com>
Date:   Mon Jun 27 15:05:53 2022 +0300

    net/mlx5e: Fix capability check for updating vnic env counters

    [ Upstream commit 452133dd580811f184e76b1402983182ee425298 ]

    The existing capability check for vnic env counters only checks for
    receive steering discards, although we need the counters update for the
    exposed internal queue oob counter as well. This could result in the
    latter counter not being updated correctly when the receive steering
    discards counter is not supported.
    Fix that by checking whether any counter is supported instead of only
    the steering counter capability.

    Fixes: 0cfafd4b4ddf ("net/mlx5e: Add device out of buffer counter")
    Signed-off-by: Gal Pressman <gal@nvidia.com>
    Reviewed-by: Tariq Toukan <tariqt@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8651049bdd77fa652bcf0f3157911a3a6fc4f2f
Author: Roi Dayan <roid@nvidia.com>
Date:   Tue Jun 21 10:43:55 2022 +0300

    net/mlx5e: CT: Use own workqueue instead of mlx5e priv

    [ Upstream commit 6c4e8fa03fde7e5b304594294e397a9ba92feaf6 ]

    Allocate a ct priv workqueue instead of using mlx5e priv one
    so flushing will only be of related CT entries.
    Also move flushing of the workqueue before rhashtable destroy
    otherwise entries won't be valid.
```

```
    Fixes: b069e14fff46 ("net/mlx5e: CT: Fix queued up restore put() executing after
relevant ft release")
    Signed-off-by: Roi Dayan <roid@nvidia.com>
    Reviewed-by: Oz Shlomo <ozsh@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33ef9ed40fcc3f7b37033823a9470ccf32206bcc
Author: Paul Blakey <paulb@nvidia.com>
Date:   Wed Jun 22 13:11:18 2022 +0300

    net/mlx5e: Fix enabling sriov while tc nic rules are offloaded

    [ Upstream commit 0c9d876545a56aebed30fa306d0460a4d28d271a ]

    There is a total of four 4M entries flow tables. In sriov disabled
    mode, ct, ct_nat and post_act take three of them. When adding the
    first tc nic rule in this mode, it will take another 4M table
    for the tc <chain,prio> table. If user then enables sriov, the legacy
    flow table tries to take another 4M and fails, and so enablement fails.

    To fix that, have legacy fdb take the next available maximum
    size from the fs ft pool.

    Fixes: 4a98544d1827 ("net/mlx5: Move chains ft pool to be used by all firmware
steering")
    Signed-off-by: Paul Blakey <paulb@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41ab41f695de241ec37bca9e36a396869f6d8862
Author: Tariq Toukan <tariqt@nvidia.com>
Date:   Mon Jun 6 21:21:10 2022 +0300

    net/mlx5e: kTLS, Fix build time constant test in RX

    [ Upstream commit 2ec6cf9b742a5c18982861322fa5de6510f8f57e ]

    Use the correct constant (TLS_DRIVER_STATE_SIZE_RX) in the comparison
    against the size of the private RX TLS driver context.

    Fixes: 1182f3659357 ("net/mlx5e: kTLS, Add kTLS RX HW offload support")
    Signed-off-by: Tariq Toukan <tariqt@nvidia.com>
    Reviewed-by: Maxim Mikityanskiy <maximmi@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3a19a4662126c3e8fb4cb2ae53838027ebd15c7
Author: Tariq Toukan <tariqt@nvidia.com>
Date:   Mon Jun 6 21:20:29 2022 +0300

    net/mlx5e: kTLS, Fix build time constant test in TX

    [ Upstream commit 6cc2714e85754a621219693ea8aa3077d6fca0cb ]

    Use the correct constant (TLS_DRIVER_STATE_SIZE_TX) in the comparison
    against the size of the private TX TLS driver context.

    Fixes: df8d866770f9 ("net/mlx5e: kTLS, Use kernel API to extract private offload
context")
    Signed-off-by: Tariq Toukan <tariqt@nvidia.com>
    Reviewed-by: Maxim Mikityanskiy <maximmi@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
```

```
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 079fe72a31c404ce66515dc90ebdf774aa59386e
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Mon Jun 13 15:05:41 2022 +0100

    ARM: 9210/1: Mark the FDT_FIXED sections as shareable

    [ Upstream commit 598f0a99fa8a35be44b27106b43ddc66417af3b1 ]

    commit 7a1be318f579 ("ARM: 9012/1: move device tree mapping out of linear
    region") use FDT_FIXED_BASE to map the whole FDT_FIXED_SIZE memory area
    which contains fdt. But it only reserves the exact physical memory that
    fdt occupied. Unfortunately, this mapping is non-shareable. An illegal or
    speculative read access can bring the RAM content from non-fdt zone into
    cache, PIPT makes it to be hit by subsequently read access through
    shareable mapping(such as linear mapping), and the cache consistency
    between cores is lost due to non-shareable property.

    |<---------FDT_FIXED_SIZE------>|
    |                               |
     -------------------------------
    | <non-fdt> | <fdt> | <non-fdt> |
     -------------------------------

    1. CoreA read <non-fdt> through MT_ROM mapping, the old data is loaded
       into the cache.
    2. CoreB write <non-fdt> to update data through linear mapping. CoreA
       received the notification to invalid the corresponding cachelines, but
       the property non-shareable makes it to be ignored.
    3. CoreA read <non-fdt> through linear mapping, cache hit, the old data
       is read.

    To eliminate this risk, add a new memory type MT_MEMORY_RO. Compared to
    MT_ROM, it is shareable and non-executable.

    Here's an example:
      list_del corruption. prev->next should be c0ecbf74, but was c08410dc
      kernel BUG at lib/list_debug.c:53!
      ... ...
      PC is at __list_del_entry_valid+0x58/0x98
      LR is at __list_del_entry_valid+0x58/0x98
      psr: 60000093
      sp : c0ecbf30  ip : 00000000  fp : 00000001
      r10: c08410d0  r9 : 00000001  r8 : c0825e0c
      r7 : 20000013  r6 : c08410d0  r5 : c0ecbf74  r4 : c0ecbf74
      r3 : c0825d08  r2 : 00000000  r1 : df7ce6f4  r0 : 00000044
      ... ...
      Stack: (0xc0ecbf30 to 0xc0ecc000)
      bf20:                                     c0ecbf74 c0164fd0 c0ecbf70 c0165170
      bf40: c0eca000 c0840c00 c0840c00 c0824500 c0825e0c c0189bbc c088f404 60000013
      bf60: 60000013 c0e85100 000004ec 00000000 c0ebcdc0 c0ecbf74 c0ecbf74 c0825d08
      ... ...                                           <  next     prev  >
      (__list_del_entry_valid) from (__list_del_entry+0xc/0x20)
      (__list_del_entry) from (finish_swait+0x60/0x7c)
      (finish_swait) from (rcu_gp_kthread+0x560/0xa20)
      (rcu_gp_kthread) from (kthread+0x14c/0x15c)
      (kthread) from (ret_from_fork+0x14/0x24)

    The faulty list node to be deleted is a local variable, its address is
    c0ecbf74. The dumped stack shows that 'prev' = c0ecbf74, but its value
    before lib/list_debug.c:53 is c08410dc. A large amount of printing results
    in swapping out the cacheline containing the old data(MT_ROM mapping is
    read only, so the cacheline cannot be dirty), and the subsequent dump
    operation obtains new data from the DDR.
```

```
    Fixes: 7a1be318f579 ("ARM: 9012/1: move device tree mapping out of linear region")
    Suggested-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Reviewed-by: Ard Biesheuvel <ardb@kernel.org>
    Reviewed-by: Kefeng Wang <wangkefeng.wang@huawei.com>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c10f57d82df1566af2ad3b5e6a81c35622889d08
Author: Ard Biesheuvel <ardb@kernel.org>
Date:   Tue May 31 09:53:42 2022 +0100

    ARM: 9209/1: Spectre-BHB: avoid pr_info() every time a CPU comes out of idle

    [ Upstream commit 0609e200246bfd3b7516091c491bec4308349055 ]

    Jon reports that the Spectre-BHB init code is filling up the kernel log
    with spurious notifications about which mitigation has been enabled,
    every time any CPU comes out of a low power state.

    Given that Spectre-BHB mitigations are system wide, only a single
    mitigation can be enabled, and we already print an error if two types of
    CPUs coexist in a single system that require different Spectre-BHB
    mitigations.

    This means that the pr_info() that describes the selected mitigation
    does not need to be emitted for each CPU anyway, and so we can simply
    emit it only once.

    In order to clarify the above in the log message, update it to describe
    that the selected mitigation will be enabled on all CPUs, including ones
    that are unaffected. If another CPU comes up later that is affected and
    requires a different mitigation, we report an error as before.

    Fixes: b9baf5c8c5c3 ("ARM: Spectre-BHB workaround")
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f65501ce53045740a86756471c48350e16f39add
Author: Cristian Ciocaltea <cristian.ciocaltea@collabora.com>
Date:   Wed Jul 6 13:06:22 2022 +0300

    spi: amd: Limit max transfer and message size

    [ Upstream commit 6ece49c56965544262523dae4a071ace3db63507 ]

    Enabling the SPI CS35L41 audio codec driver for Steam Deck [1]
    revealed a problem with the current AMD SPI controller driver
    implementation, consisting of an unrecoverable system hang.

    The issue can be prevented if we ensure the max transfer size
    and the max message size do not exceed the FIFO buffer size.

    According to the implementation of the downstream driver, the
    AMD SPI controller is not able to handle more than 70 bytes per
    transfer, which corresponds to the size of the FIFO buffer.

    Hence, let's fix this by setting the SPI limits mentioned above.

    [1] https://lore.kernel.org/r/20220621213819.262537-1-cristian.ciocaltea@collabora.com

    Reported-by: Anastasios Vacharakis <vacharakis@o2mail.de>
```

```
      Fixes: bbb336f39efc ("spi: spi-amd: Add AMD SPI controller driver support")
      Signed-off-by: Cristian Ciocaltea <cristian.ciocaltea@collabora.com>
      Link: https://lore.kernel.org/r/20220706100626.1234731-2-
  cristian.ciocaltea@collabora.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 0903e579544ad1e608e3b2c905d2d0cd1e501fb5
  Author: Conor Dooley <conor.dooley@microchip.com>
  Date:   Wed Jun 29 21:07:33 2022 +0100

      riscv: dts: microchip: hook up the mpfs' l2cache

      [ Upstream commit efa310ba00716d7a872bdc5fa1f5545edc9efd69 ]

      The initial PolarFire SoC devicetree must have been forked off from
      the fu540 one prior to the addition of l2cache controller support being
      added there. When the controller node was added to mpfs.dtsi, it was
      not hooked up to the CPUs & thus sysfs reports an incorrect cache
      configuration. Hook it up.

      Fixes: 0fa6107eca41 ("RISC-V: Initial DTS for Microchip ICICLE board")
      Reviewed-by: Sudeep Holla <sudeep.holla@arm.com>
      Reviewed-by: Daire McNamara <daire.mcnamara@microchip.com>
      Signed-off-by: Conor Dooley <conor.dooley@microchip.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 0125374851845eae4cb7932be2c1650ced2eda2b
  Author: Kris Bahnsen <kris@embeddedTS.com>
  Date:   Thu Jun 30 14:03:27 2022 -0700

      ARM: dts: imx6qdl-ts7970: Fix ngpio typo and count

      [ Upstream commit e95ea0f687e679fcb0a3a67d0755b81ee7d60db0 ]

      Device-tree incorrectly used "ngpio" which caused the driver to
      fallback to 32 ngpios.

      This platform has 62 GPIO registers.

      Fixes: 9ff8e9fccef9 ("ARM: dts: TS-7970: add basic device tree")
      Signed-off-by: Kris Bahnsen <kris@embeddedTS.com>
      Reviewed-by: Fabio Estevam <festevam@gmail.com>
      Signed-off-by: Shawn Guo <shawnguo@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 51904922902f04619635b1024ce3300955ccc599
  Author: Sean Anderson <sean.anderson@seco.com>
  Date:   Thu Apr 28 14:16:59 2022 -0400

      arm64: dts: ls1028a: Update SFP node to include clock

      [ Upstream commit 3c12e9da3098a30fc82dea01768d355c28e3692d ]

      The clocks property is now mandatory. Add it to avoid warning message.

      Signed-off-by: Sean Anderson <sean.anderson@seco.com>
      Reviewed-by: Michael Walle <michael@walle.cc>
      Fixes: eba5bea8f37f ("arm64: dts: ls1028a: add efuse node")
      Signed-off-by: Shawn Guo <shawnguo@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit ab5e910c989ad93b1eeb7f07aa8ed8dbe9b94bfd
  Author: Serge Semin <Sergey.Semin@baikalelectronics.ru>
  Date:   Fri Jun 24 17:18:45 2022 +0300
```

```
    reset: Fix devm bulk optional exclusive control getter

    [ Upstream commit a57f68ddc8865d59a19783080cc52fb4a11dc209 ]

    Most likely due to copy-paste mistake the device managed version of the
    denoted reset control getter has been implemented with invalid semantic,
    which can be immediately spotted by having "WARN_ON(shared && acquired)"
    warning in the system log as soon as the method is called. Anyway let's
    fix it by altering the boolean arguments passed to the
    __devm_reset_control_bulk_get() method from
    - shared = true, optional = false, acquired = true
    to
    + shared = false, optional = true, acquired = true
    That's what they were supposed to be in the first place (see the non-devm
    version of the same method: reset_control_bulk_get_optional_exclusive()).

    Fixes: 48d71395896d ("reset: Add reset_control_bulk API")
    Signed-off-by: Serge Semin <Sergey.Semin@baikalelectronics.ru>
    Reviewed-by: Dmitry Osipenko <dmitry.osipenko@collabora.com>
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Link: https://lore.kernel.org/r/20220624141853.7417-2-Sergey.Semin@baikalelectronics.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a92d44b412e75dd66543843165e46637457f22cc
Author: Xiu Jianfeng <xiujianfeng@huawei.com>
Date:   Fri May 27 19:17:26 2022 +0800

    Revert "evm: Fix memleak in init_desc"

    commit 51dd64bb99e4478fc5280171acd8e1b529eadaf7 upstream.

    This reverts commit ccf11dbaa07b328fa469415c362d33459c140a37.

    Commit ccf11dbaa07b ("evm: Fix memleak in init_desc") said there is
    memleak in init_desc. That may be incorrect, as we can see, tmp_tfm is
    saved in one of the two global variables hmac_tfm or evm_tfm[hash_algo],
    then if init_desc is called next time, there is no need to alloc tfm
    again, so in the error path of kmalloc desc or crypto_shash_init(desc),
    It is not a problem without freeing tmp_tfm.

    And also that commit did not reset the global variable to NULL after
    freeing tmp_tfm and this makes *tfm a dangling pointer which may cause a
    UAF issue.

    Reported-by: Guozihua (Scott) <guozihua@huawei.com>
    Signed-off-by: Xiu Jianfeng <xiujianfeng@huawei.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c41ab0ecda85f984a889626184ef4487e36cde5
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Jun 20 09:01:43 2022 +0200

    sh: convert nommu io{re,un}map() to static inline functions

    commit d684e0a52d36f8939eda30a0f31ee235ee4ee741 upstream.

    Recently, nommu iounmap() was converted from a static inline function to a
    macro again, basically reverting commit 4580ba4ad2e6b8dd ("sh: Convert
    iounmap() macros to inline functions").  With -Werror, this leads to build
    failures like:

        drivers/iio/adc/xilinx-ams.c: In function `ams_iounmap_ps':
        drivers/iio/adc/xilinx-ams.c:1195:14: error: unused variable `ams' [-Werror=unused-
```

```
 variable]
         1195 |  struct ams *ams = data;
              |             ^~~
```

    Fix this by replacing the macros for ioremap() and iounmap() by static
    inline functions, based on <asm-generic/io.h>.

    Link:
 https://lkml.kernel.org/r/8d1b1766260961799b04035e7bc39a7f59729f72.1655708312.git.geert+rene
 sas@glider.be
    Fixes: 13f1fc870dd74713 ("sh: move the ioremap implementation out of line")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Acked-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4eb6ad83932e841ec6a93daff363bf00ba78c75b
Author: Ryusuke Konishi <konishi.ryusuke@gmail.com>
Date: Thu Jun 23 17:54:01 2022 +0900

    nilfs2: fix incorrect masking of permission flags for symlinks

    commit 5924e6ec1585445f251ea92713eb15beb732622a upstream.

    The permission flags of newly created symlinks are wrongly dropped on
    nilfs2 with the current umask value even though symlinks should have 777
    (rwxrwxrwx) permissions:

     $ umask
     0022
     $ touch file && ln -s file symlink; ls -l file symlink
     -rw-r--r--. 1 root root 0 Jun 23 16:29 file
     lrwxr-xr-x. 1 root root 4 Jun 23 16:29 symlink -> file

    This fixes the bug by inserting a missing check that excludes
    symlinks.

    Link: https://lkml.kernel.org/r/1655974441-5612-1-git-send-email-
 konishi.ryusuke@gmail.com
    Signed-off-by: Ryusuke Konishi <konishi.ryusuke@gmail.com>
    Reported-by: Tommy Pettersson <ptp@lysator.liu.se>
    Reported-by: Ciprian Craciun <ciprian.craciun@gmail.com>
    Tested-by: Ryusuke Konishi <konishi.ryusuke@gmail.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e711a68a0e66ad64d2234d80ecbf10ec49c0fa00
Author: Dave Chinner <dchinner@redhat.com>
Date: Wed Jul 13 17:49:15 2022 +1000

    fs/remap: constrain dedupe of EOF blocks

    commit 5750676b64a561f7ec920d7c6ba130fc9c7378f3 upstream.

    If dedupe of an EOF block is not constrainted to match against only
    other EOF blocks with the same EOF offset into the block, it can
    match against any other block that has the same matching initial
    bytes in it, even if the bytes beyond EOF in the source file do
    not match.

    Fix this by constraining the EOF block matching to only match
    against other EOF blocks that have identical EOF offsets and data.

```
    This allows "whole file dedupe" to continue to work without allowing
    eof blocks to randomly match against partial full blocks with the
    same data.

    Reported-by: Ansgar Lößer <ansgar.loesser@tu-darmstadt.de>
    Fixes: 1383a7ed6749 ("vfs: check file ranges before cloning files")
    Link: https://lore.kernel.org/linux-fsdevel/a7c93559-4ba1-df2f-7a85-55a143696405@tu-
darmstadt.de/
    Signed-off-by: Dave Chinner <dchinner@redhat.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f036392edd9c49090781d8cca26ad6557a63bae4
Author: Dmitry Osipenko <dmitry.osipenko@collabora.com>
Date:   Thu Jun 30 23:06:01 2022 +0300

    drm/panfrost: Fix shrinker list corruption by madvise IOCTL

    commit 9fc33eaaa979d112d10fea729edcd2a2e21aa912 upstream.

    Calling madvise IOCTL twice on BO causes memory shrinker list corruption
    and crashes kernel because BO is already on the list and it's added to
    the list again, while BO should be removed from the list before it's
    re-added. Fix it.

    Cc: stable@vger.kernel.org
    Fixes: 013b65101315 ("drm/panfrost: Add madvise and shrinker support")
    Acked-by: Alyssa Rosenzweig <alyssa.rosenzweig@collabora.com>
    Reviewed-by: Steven Price <steven.price@arm.com>
    Signed-off-by: Dmitry Osipenko <dmitry.osipenko@collabora.com>
    Signed-off-by: Steven Price <steven.price@arm.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220630200601.1884120-3-
dmitry.osipenko@collabora.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3d1ca4114fe644c3a9ba229d02b593ae7b6cc77f
Author: Dmitry Osipenko <dmitry.osipenko@collabora.com>
Date:   Thu Jun 30 23:06:00 2022 +0300

    drm/panfrost: Put mapping instead of shmem obj on panfrost_mmu_map_fault_addr() error

    commit fb6e0637ab7ebd8e61fe24f4d663c4bae99cfa62 upstream.

    When panfrost_mmu_map_fault_addr() fails, the BO's mapping should be
    unreferenced and not the shmem object which backs the mapping.

    Cc: stable@vger.kernel.org
    Fixes: bdefca2d8dc0 ("drm/panfrost: Add the panfrost_gem_mapping concept")
    Reviewed-by: Steven Price <steven.price@arm.com>
    Signed-off-by: Dmitry Osipenko <dmitry.osipenko@collabora.com>
    Signed-off-by: Steven Price <steven.price@arm.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220630200601.1884120-2-
dmitry.osipenko@collabora.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 02fc3a6ed7873e85c037ea9b67614780713966d7
Author: Christoph Hellwig <hch@lst.de>
Date:   Thu Jun 30 18:03:19 2022 +0200

    btrfs: zoned: fix a leaked bioc in read_zone_info

    commit 2963457829decf0c824a443238d251151ed18ff5 upstream.

    The bioc would leak on the normal completion path and also on the RAID56
    check (but that one won't happen in practice due to the invalid
```

```
    combination with zoned mode).

    Fixes: 7db1c5d14dcd ("btrfs: zoned: support dev-replace in zoned filesystems")
    CC: stable@vger.kernel.org # 5.16+
    Reviewed-by: Anand Jain <anand.jain@oracle.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    [ update changelog ]
    Reviewed-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 722ab9a4632bbf6dcc93aee03d7a121ac85d7f4c
Author: Filipe Manana <fdmanana@suse.com>
Date:   Mon Jul 4 12:42:03 2022 +0100

    btrfs: return -EAGAIN for NOWAIT dio reads/writes on compressed and inline extents

    commit a4527e1853f8ff6e0b7c2dadad6268bd38427a31 upstream.

    When doing a direct IO read or write, we always return -ENOTBLK when we
    find a compressed extent (or an inline extent) so that we fallback to
    buffered IO. This however is not ideal in case we are in a NOWAIT context
    (io_uring for example), because buffered IO can block and we currently
    have no support for NOWAIT semantics for buffered IO, so if we need to
    fallback to buffered IO we should first signal the caller that we may
    need to block by returning -EAGAIN instead.

    This behaviour can also result in short reads being returned to user
    space, which although it's not incorrect and user space should be able
    to deal with partial reads, it's somewhat surprising and even some popular
    applications like QEMU (Link tag #1) and MariaDB (Link tag #2) don't
    deal with short reads properly (or at all).

    The short read case happens when we try to read from a range that has a
    non-compressed and non-inline extent followed by a compressed extent.
    After having read the first extent, when we find the compressed extent we
    return -ENOTBLK from btrfs_dio_iomap_begin(), which results in iomap to
    treat the request as a short read, returning 0 (success) and waiting for
    previously submitted bios to complete (this happens at
    fs/iomap/direct-io.c:__iomap_dio_rw()). After that, and while at
    btrfs_file_read_iter(), we call filemap_read() to use buffered IO to
    read the remaining data, and pass it the number of bytes we were able to
    read with direct IO. Than at filemap_read() if we get a page fault error
    when accessing the read buffer, we return a partial read instead of an
    -EFAULT error, because the number of bytes previously read is greater
    than zero.

    So fix this by returning -EAGAIN for NOWAIT direct IO when we find a
    compressed or an inline extent.

    Reported-by: Dominique MARTINET <dominique.martinet@atmark-techno.com>
    Link: https://lore.kernel.org/linux-btrfs/YrrFGO4A1jS0GI0G@atmark-techno.com/
    Link: https://jira.mariadb.org/browse/MDEV-27900?
focusedCommentId=216582&page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-
tabpanel#comment-216582
    Tested-by: Dominique MARTINET <dominique.martinet@atmark-techno.com>
    CC: stable@vger.kernel.org # 5.10+
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0e41774b564befa6d271e8d5086bf870d617a4e6
Author: Tejun Heo <tj@kernel.org>
Date:   Mon Jun 13 12:19:50 2022 -1000
```

cgroup: Use separate src/dst nodes when preloading css_sets for migration

commit 07fd5b6cdf3cc30bfde8fe0f644771688be04447 upstream.

Each cset (css_set) is pinned by its tasks. When we're moving tasks around
across csets for a migration, we need to hold the source and destination
csets to ensure that they don't go away while we're moving tasks about. This
is done by linking cset->mg_preload_node on either the
mgctx->preloaded_src_csets or mgctx->preloaded_dst_csets list. Using the
same cset->mg_preload_node for both the src and dst lists was deemed okay as
a cset can't be both the source and destination at the same time.

Unfortunately, this overloading becomes problematic when multiple tasks are
involved in a migration and some of them are identity noop migrations while
others are actually moving across cgroups. For example, this can happen with
the following sequence on cgroup1:

```
 #1> mkdir -p /sys/fs/cgroup/misc/a/b
 #2> echo $$ > /sys/fs/cgroup/misc/a/cgroup.procs
 #3> RUN_A_COMMAND_WHICH_CREATES_MULTIPLE_THREADS &
 #4> PID=$!
 #5> echo $PID > /sys/fs/cgroup/misc/a/b/tasks
 #6> echo $PID > /sys/fs/cgroup/misc/a/cgroup.procs
```

the process including the group leader back into a. In this final migration,
non-leader threads would be doing identity migration while the group leader
is doing an actual one.

After #3, let's say the whole process was in cset A, and that after #4, the
leader moves to cset B. Then, during #6, the following happens:

1. cgroup_migrate_add_src() is called on B for the leader.

2. cgroup_migrate_add_src() is called on A for the other threads.

3. cgroup_migrate_prepare_dst() is called. It scans the src list.

4. It notices that B wants to migrate to A, so it tries to A to the dst
   list but realizes that its ->mg_preload_node is already busy.

5. and then it notices A wants to migrate to A as it's an identity
   migration, it culls it by list_del_init()'ing its ->mg_preload_node and
   putting references accordingly.

6. The rest of migration takes place with B on the src list but nothing on
   the dst list.

This means that A isn't held while migration is in progress. If all tasks
leave A before the migration finishes and the incoming task pins it, the
cset will be destroyed leading to use-after-free.

This is caused by overloading cset->mg_preload_node for both src and dst
preload lists. We wanted to exclude the cset from the src list but ended up
inadvertently excluding it from the dst list too.

This patch fixes the issue by separating out cset->mg_preload_node into
->mg_src_preload_node and ->mg_dst_preload_node, so that the src and dst
preloadings don't interfere with each other.

Signed-off-by: Tejun Heo <tj@kernel.org>
Reported-by: Mukesh Ojha <quic_mojha@quicinc.com>
Reported-by: shisiyuan <shisiyuan19870131@gmail.com>
Link: http://lkml.kernel.org/r/1654187688-27411-1-git-send-email-shisiyuan@xiaomi.com
Link: https://www.spinics.net/lists/cgroups/msg33313.html

```
    Fixes: f817de98513d ("cgroup: prepare migration path for unified hierarchy")
    Cc: stable@vger.kernel.org # v3.16+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 41ecab279a70dced9005f4d55fa0fcde8938603f
Author: Felix Fietkau <nbd@nbd.name>
Date:   Sat Jul 2 16:52:27 2022 +0200

    wifi: mac80211: fix queue selection for mesh/OCB interfaces

    commit 50e2ab39291947b6c6c7025cf01707c270fcde59 upstream.

    When using iTXQ, the code assumes that there is only one vif queue for
    broadcast packets, using the BE queue. Allowing non-BE queue marking
    violates that assumption and txq->ac == skb_queue_mapping is no longer
    guaranteed. This can cause issues with queue handling in the driver and
    also causes issues with the recent ATF change, resulting in an AQL
    underflow warning.

    Cc: stable@vger.kernel.org
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Link: https://lore.kernel.org/r/20220702145227.39356-1-nbd@nbd.name
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 05c214703f286006db285078e70cec4416e5bf7b
Author: Ard Biesheuvel <ardb@kernel.org>
Date:   Thu Jun 30 16:46:54 2022 +0100

    ARM: 9214/1: alignment: advance IT state after emulating Thumb instruction

    commit e5c46fde75e43c15a29b40e5fc5641727f97ae47 upstream.

    After emulating a misaligned load or store issued in Thumb mode, we have
    to advance the IT state by hand, or it will get out of sync with the
    actual instruction stream, which means we'll end up applying the wrong
    condition code to subsequent instructions. This might corrupt the
    program state rather catastrophically.

    So borrow the it_advance() helper from the probing code, and use it on
    CPSR if the emulated instruction is Thumb.

    Cc: <stable@vger.kernel.org>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9bac88b10f0f8e442fc1b8b4a72b29d708208805
Author: Dmitry Osipenko <dmitry.osipenko@collabora.com>
Date:   Tue Jun 28 08:55:45 2022 +0100

    ARM: 9213/1: Print message about disabled Spectre workarounds only once

    commit e4ced82deb5fb17222fb82e092c3f8311955b585 upstream.

    Print the message about disabled Spectre workarounds only once. The
    message is printed each time CPU goes out from idling state on NVIDIA
    Tegra boards, causing storm in KMSG that makes system unusable.

    Cc: stable@vger.kernel.org
    Signed-off-by: Dmitry Osipenko <dmitry.osipenko@collabora.com>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 681da5bff05fbbbba7bf15c0004a371cb57f2de9
Author: Nicolas Dichtel <nicolas.dichtel@6wind.com>
Date:   Wed Jul 13 13:48:52 2022 +0200

    ip: fix dflt addr selection for connected nexthop

    commit 747c14307214b55dbd8250e1ab44cad8305756f1 upstream.

    When a nexthop is added, without a gw address, the default scope was set
    to 'host'. Thus, when a source address is selected, 127.0.0.1 may be chosen
    but rejected when the route is used.

    When using a route without a nexthop id, the scope can be configured in the
    route, thus the problem doesn't exist.

    To explain more deeply: when a user creates a nexthop, it cannot specify
    the scope. To create it, the function nh_create_ipv4() calls fib_check_nh()
    with scope set to 0. fib_check_nh() calls fib_check_nh_nongw() wich was
    setting scope to 'host'. Then, nh_create_ipv4() calls
    fib_info_update_nhc_saddr() with scope set to 'host'. The src addr is
    chosen before the route is inserted.

    When a 'standard' route (ie without a reference to a nexthop) is added,
    fib_create_info() calls fib_info_update_nhc_saddr() with the scope set by
    the user. iproute2 set the scope to 'link' by default.

    Here is a way to reproduce the problem:
    ip netns add foo
    ip -n foo link set lo up
    ip netns add bar
    ip -n bar link set lo up
    sleep 1

    ip -n foo link add name eth0 type dummy
    ip -n foo link set eth0 up
    ip -n foo address add 192.168.0.1/24 dev eth0

    ip -n foo link add name veth0 type veth peer name veth1 netns bar
    ip -n foo link set veth0 up
    ip -n bar link set veth1 up

    ip -n bar address add 192.168.1.1/32 dev veth1
    ip -n bar route add default dev veth1

    ip -n foo nexthop add id 1 dev veth0
    ip -n foo route add 192.168.1.1 nhid 1

    Try to get/use the route:
    > $ ip -n foo route get 192.168.1.1
    > RTNETLINK answers: Invalid argument
    > $ ip netns exec foo ping -c1 192.168.1.1
    > ping: connect: Invalid argument

    Try without nexthop group (iproute2 sets scope to 'link' by dflt):
    ip -n foo route del 192.168.1.1
    ip -n foo route add 192.168.1.1 dev veth0

    Try to get/use the route:
    > $ ip -n foo route get 192.168.1.1
    > 192.168.1.1 dev veth0 src 192.168.0.1 uid 0
    >     cache
    > $ ip netns exec foo ping -c1 192.168.1.1
    > PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
    > 64 bytes from 192.168.1.1: icmp_seq=1 ttl=64 time=0.039 ms
    >
```

```
      > --- 192.168.1.1 ping statistics ---
      > 1 packets transmitted, 1 received, 0% packet loss, time 0ms
      > rtt min/avg/max/mdev = 0.039/0.039/0.039/0.000 ms

      CC: stable@vger.kernel.org
      Fixes: 597cfe4fc339 ("nexthop: Add support for IPv4 nexthops")
      Reported-by: Edwin Brossette <edwin.brossette@6wind.com>
      Signed-off-by: Nicolas Dichtel <nicolas.dichtel@6wind.com>
      Link: https://lore.kernel.org/r/20220713114853.29406-1-nicolas.dichtel@6wind.com
      Signed-off-by: Paolo Abeni <pabeni@redhat.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 157269e332865dec1ed64f243d122b9f066f6639
  Author: Steven Rostedt (Google) <rostedt@goodmis.org>
  Date:   Wed Jul 6 10:50:40 2022 -0400

      net: sock: tracing: Fix sock_exceed_buf_limit not to dereference stale pointer

      commit 820b8963adaea34a87abbecb906d1f54c0aabfb7 upstream.

      The trace event sock_exceed_buf_limit saves the prot->sysctl_mem pointer
      and then dereferences it in the TP_printk() portion. This is unsafe as the
      TP_printk() portion is executed at the time the buffer is read. That is,
      it can be seconds, minutes, days, months, even years later. If the proto
      is freed, then this dereference will can also lead to a kernel crash.

      Instead, save the sysctl_mem array into the ring buffer and have the
      TP_printk() reference that instead. This is the proper and safe way to
      read pointers in trace events.

      Link: https://lore.kernel.org/all/20220706052130.16368-12-kuniyu@amazon.com/

      Cc: stable@vger.kernel.org
      Fixes: 3847ce32aea9f ("core: add tracepoints for queueing skb to rcvbuf")
      Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
      Acked-by: Kuniyuki Iwashima <kuniyu@amazon.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 4d453eb5e1eec89971aa5b3262857ee26cfdffd3
  Author: Zheng Yejian <zhengyejian1@huawei.com>
  Date:   Mon Jul 11 09:47:31 2022 +0800

      tracing/histograms: Fix memory leak problem

      commit 7edc3945bdce9c39198a10d6129377a5c53559c2 upstream.

      This reverts commit 46bbe5c671e06f070428b9be142cc4ee5cedebac.

      As commit 46bbe5c671e0 ("tracing: fix double free") said, the
      "double free" problem reported by clang static analyzer is:
        > In parse_var_defs() if there is a problem allocating
        > var_defs.expr, the earlier var_defs.name is freed.
        > This free is duplicated by free_var_defs() which frees
        > the rest of the list.

      However, if there is a problem allocating N-th var_defs.expr:
        + in parse_var_defs(), the freed 'earlier var_defs.name' is
          actually the N-th var_defs.name;
        + then in free_var_defs(), the names from 0th to (N-1)-th are freed;

                                IF ALLOCATING PROBLEM HAPPENED HERE!!! -+
                                                                       \
                                                                        |
                  0th           1th                 (N-1)-th      N-th    V
```

```
              +-------------+-------------+-----+-------------+-----------
    var_defs: | name | expr | name | expr | ... | name | expr | name | ///
              +-------------+-------------+-----+-------------+-----------
```

    These two frees don't act on same name, so there was no "double free"
    problem before. Conversely, after that commit, we get a "memory leak"
    problem because the above "N-th var_defs.name" is not freed.

    If enable CONFIG_DEBUG_KMEMLEAK and inject a fault at where the N-th
    var_defs.expr allocated, then execute on shell like:
      $ echo 'hist:key=call_site:val=$v1,$v2:v1=bytes_req,v2=bytes_alloc' > \
    /sys/kernel/debug/tracing/events/kmem/kmalloc/trigger

    Then kmemleak reports:
      unreferenced object 0xffff8fb100ef3518 (size 8):
        comm "bash", pid 196, jiffies 4295681690 (age 28.538s)
        hex dump (first 8 bytes):
          76 31 00 00 b1 8f ff ff   v1......
        backtrace:
          [<0000000038fe4895>] kstrdup+0x2d/0x60
          [<00000000c99c049a>] event_hist_trigger_parse+0x206f/0x20e0
          [<00000000ae70d2cc>] trigger_process_regex+0xc0/0x110
          [<0000000066737a4c>] event_trigger_write+0x75/0xd0
          [<000000007341e40c>] vfs_write+0xbb/0x2a0
          [<0000000087fde4c2>] ksys_write+0x59/0xd0
          [<00000000581e9cdf>] do_syscall_64+0x3a/0x80
          [<00000000cf3b065c>] entry_SYSCALL_64_after_hwframe+0x46/0xb0

    Link: https://lkml.kernel.org/r/20220711014731.69520-1-zhengyejian1@huawei.com

    Cc: stable@vger.kernel.org
    Fixes: 46bbe5c671e0 ("tracing: fix double free")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Suggested-by: Steven Rostedt <rostedt@goodmis.org>
    Reviewed-by: Tom Zanussi <tom.zanussi@linux.intel.com>
    Signed-off-by: Zheng Yejian <zhengyejian1@huawei.com>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3920c5843e4ef47e5b0f6c8d78829fe4e07421ec
Author: Baolin Wang <baolin.wang@linux.alibaba.com>
Date: Mon Jun 20 10:34:42 2022 +0800

    mm/damon: use set_huge_pte_at() to make huge pte old

    commit ed1523a895ffdabcab6e067af18685ed00f5ce15 upstream.

    The huge_ptep_set_access_flags() can not make the huge pte old according
    to the discussion [1], that means we will always mornitor the young state
    of the hugetlb though we stopped accessing the hugetlb, as a result DAMON
    will get inaccurate accessing statistics.

    So changing to use set_huge_pte_at() to make the huge pte old to fix this
    issue.

    [1] https://lore.kernel.org/all/Yqy97gXI4Nqb7dYo@arm.com/

    Link: https://lkml.kernel.org/r/1655692482-28797-1-git-send-email-
baolin.wang@linux.alibaba.com
    Fixes: 49f4203aae06 ("mm/damon: add access checking for hugetlb pages")
    Signed-off-by: Baolin Wang <baolin.wang@linux.alibaba.com>
    Reviewed-by: SeongJae Park <sj@kernel.org>
    Acked-by: Mike Kravetz <mike.kravetz@oracle.com>
    Reviewed-by: Muchun Song <songmuchun@bytedance.com>
    Cc: <stable@vger.kernel.org>

```
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3d637c7df81ef8cbed8e80cf57c9c619ea6f9fb8
Author: Gowans, James <jgowans@amazon.com>
Date:   Thu Jun 23 05:24:03 2022 +0000

    mm: split huge PUD on wp_huge_pud fallback

    commit 14c99d65941538aa33edd8dc7b1bbbb593c324a2 upstream.

    Currently the implementation will split the PUD when a fallback is taken
    inside the create_huge_pud function.  This isn't where it should be done:
    the splitting should be done in wp_huge_pud, just like it's done for PMDs.
    Reason being that if a callback is taken during create, there is no PUD
    yet so nothing to split, whereas if a fallback is taken when encountering
    a write protection fault there is something to split.

    It looks like this was the original intention with the commit where the
    splitting was introduced, but somehow it got moved to the wrong place
    between v1 and v2 of the patch series.  Rebase mistake perhaps.

    Link:
https://lkml.kernel.org/r/6f48d622eb8bce1ae5dd75327b0b73894a2ec407.camel@amazon.com
    Fixes: 327e9fd48972 ("mm: Split huge pages on write-notify or COW")
    Signed-off-by: James Gowans <jgowans@amazon.com>
    Reviewed-by: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: Jan H. Schönherr <jschoenh@amazon.de>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6ebbd46d2ef5f6ab815907e1137bdea6e1a3ec3f
Author: Muchun Song <songmuchun@bytedance.com>
Date:   Mon Jun 20 10:30:19 2022 +0800

    mm: sparsemem: fix missing higher order allocation splitting

    commit 39d35edee4537487e5178f258e23518272a66413 upstream.

    Higher order allocations for vmemmap pages from buddy allocator must be
    able to be treated as indepdenent small pages as they can be freed
    individually by the caller.  There is no problem for higher order vmemmap
    pages allocated at boot time since each individual small page will be
    initialized at boot time.  However, it will be an issue for memory hotplug
    case since those higher order vmemmap pages are allocated from buddy
    allocator without initializing each individual small page's refcount.  The
    system will panic in put_page_testzero() when CONFIG_DEBUG_VM is enabled
    if the vmemmap page is freed.

    Link: https://lkml.kernel.org/r/20220620023019.94257-1-songmuchun@bytedance.com
    Fixes: d8d55f5616cf ("mm: sparsemem: use page table lock to protect kernel pmd
operations")
    Signed-off-by: Muchun Song <songmuchun@bytedance.com>
    Reviewed-by: Mike Kravetz <mike.kravetz@oracle.com>
    Cc: Xiongchun Duan <duanxiongchun@bytedance.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 43c5ac008f03ba2b8c957a0b7b7936acbe481e7a
Author: Axel Rasmussen <axelrasmussen@google.com>
Date:   Fri Jun 10 10:38:12 2022 -0700

    mm: userfaultfd: fix UFFDIO_CONTINUE on fallocated shmem pages
```

```
    commit 73f37dbcfe1763ee2294c7717a1f571e27d17fd8 upstream.

    When fallocate() is used on a shmem file, the pages we allocate can end up
    with !PageUptodate.

    Since UFFDIO_CONTINUE tries to find the existing page the user wants to
    map with SGP_READ, we would fail to find such a page, since
    shmem_getpage_gfp returns with a "NULL" pagep for SGP_READ if it discovers
    !PageUptodate.  As a result, UFFDIO_CONTINUE returns -EFAULT, as it would
    do if the page wasn't found in the page cache at all.

    This isn't the intended behavior.  UFFDIO_CONTINUE is just trying to find
    if a page exists, and doesn't care whether it still needs to be cleared or
    not.  So, instead of SGP_READ, pass in SGP_NOALLOC.  This is the same,
    except for one critical difference: in the !PageUptodate case, SGP_NOALLOC
    will clear the page and then return it.  With this change, UFFDIO_CONTINUE
    works properly (succeeds) on a shmem file which has been fallocated, but
    otherwise not modified.

    Link: https://lkml.kernel.org/r/20220610173812.1768919-1-axelrasmussen@google.com
    Fixes: 153132571f02 ("userfaultfd/shmem: support UFFDIO_CONTINUE for shmem")
    Signed-off-by: Axel Rasmussen <axelrasmussen@google.com>
    Acked-by: Peter Xu <peterx@redhat.com>
    Cc: Hugh Dickins <hughd@google.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f45f4f82cb48c6ce1d838cf3bace9993c90aaee6
Author: Oleg Nesterov <oleg@redhat.com>
Date:   Mon Jul 11 18:16:25 2022 +0200

    fix race between exit_itimers() and /proc/pid/timers

    commit d5b36a4dbd06c5e8e36ca8ccc552f679069e2946 upstream.

    As Chris explains, the comment above exit_itimers() is not correct,
    we can race with proc_timers_seq_ops. Change exit_itimers() to clear
    signal->posix_timers with ->siglock held.

    Cc: <stable@vger.kernel.org>
    Reported-by: chris@accessvector.net
    Signed-off-by: Oleg Nesterov <oleg@redhat.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f0b5c819b062df8bf5f2acf4697e3871cb3722da
Author: Juergen Gross <jgross@suse.com>
Date:   Wed Jul 13 15:53:22 2022 +0200

    xen/netback: avoid entering xenvif_rx_next_skb() with an empty rx queue

    commit 94e8100678889ab428e68acadf042de723f094b9 upstream.

    xenvif_rx_next_skb() is expecting the rx queue not being empty, but
    in case the loop in xenvif_rx_action() is doing multiple iterations,
    the availability of another skb in the rx queue is not being checked.

    This can lead to crashes:

    [40072.537261] BUG: unable to handle kernel NULL pointer dereference at 0000000000000080
    [40072.537407] IP: xenvif_rx_skb+0x23/0x590 [xen_netback]
    [40072.537534] PGD 0 P4D 0
    [40072.537644] Oops: 0000 [#1] SMP NOPTI
```

```
    [40072.537749] CPU: 0 PID: 12505 Comm: v1-c40247-q2-gu Not tainted 4.12.14-122.121-
default #1 SLE12-SP5
    [40072.537867] Hardware name: HP ProLiant DL580 Gen9/ProLiant DL580 Gen9, BIOS U17
11/23/2021
    [40072.537999] task: ffff880433b38100 task.stack: ffffc90043d40000
    [40072.538112] RIP: e030:xenvif_rx_skb+0x23/0x590 [xen_netback]
    [40072.538217] RSP: e02b:ffffc90043d43de0 EFLAGS: 00010246
    [40072.538319] RAX: 0000000000000000 RBX: ffffc90043cd7cd0 RCX: 00000000000000f7
    [40072.538430] RDX: 0000000000000000 RSI: 0000000000000006 RDI: ffffc90043d43df8
    [40072.538531] RBP: 000000000000003f R08: 000077ff80000000 R09: 0000000000000008
    [40072.538644] R10: 0000000000007ff0 R11: 00000000000008f6 R12: ffffc90043ce2708
    [40072.538745] R13: 0000000000000000 R14: ffffc90043d43ed0 R15: ffff88043ea748c0
    [40072.538861] FS: 0000000000000000(0000) GS:ffff880484600000(0000)
knlGS:0000000000000000
    [40072.538988] CS: e033 DS: 0000 ES: 0000 CR0: 0000000080050033
    [40072.539088] CR2: 0000000000000080 CR3: 0000000407ac8000 CR4: 0000000000040660
    [40072.539211] Call Trace:
    [40072.539319] xenvif_rx_action+0x71/0x90 [xen_netback]
    [40072.539429] xenvif_kthread_guest_rx+0x14a/0x29c [xen_netback]

    Fix that by stopping the loop in case the rx queue becomes empty.

    Cc: stable@vger.kernel.org
    Fixes: 98f6d57ced73 ("xen-netback: process guest rx packets in batches")
    Signed-off-by: Juergen Gross <jgross@suse.com>
    Reviewed-by: Jan Beulich <jbeulich@suse.com>
    Reviewed-by: Paul Durrant <paul@xen.org>
    Link: https://lore.kernel.org/r/20220713135322.19616-1-jgross@suse.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c70092121f3c748a1f8a365401aa687e990ea471
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Wed Jul 13 17:41:33 2022 +0800

    ALSA: hda/realtek - Enable the headset-mic on a Xiaomi's laptop

    commit 9b043a8f386485c74c0f8eea2c287d5bdbdf3279 upstream.

    The headset on this machine is not defined, after applying the quirk
    ALC256_FIXUP_ASUS_HEADSET_MIC, the headset-mic works well

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220713094133.9894-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c639c0221efa7548dac3f8573c805456b3faacb
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Wed Jul 13 14:33:32 2022 +0800

    ALSA: hda/realtek - Fix headset mic problem for a HP machine with alc221

    commit 4ba5c853d7945b3855c3dcb293f7f9f019db641e upstream.

    On a HP 288 Pro G2 MT (X9W02AV), the front mic could not be detected.
    In order to get it working, the pin configuration needs to be set
    correctly, and the ALC221_FIXUP_HP_288PRO_MIC_NO_PRESENCE fixup needs
    to be applied.

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220713063332.30095-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
```

```
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 44978dd619c27b1191f3b5f38700a8103aceeed0
Author: Jeremy Szu <jeremy.szu@canonical.com>
Date:   Wed Jul 13 10:27:04 2022 +0800

    ALSA: hda/realtek: fix mute/micmute LEDs for HP machines

    commit 61d307855eb1a2ae849da445edd5389db8a58a5c upstream.

    The HP ProBook 440/450 G9 and EliteBook 640/650 G9 have multiple
    motherboard design and they are using different subsystem ID of audio
    codec. Add the same quirk for other MBs.

    Signed-off-by: Jeremy Szu <jeremy.szu@canonical.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220713022706.22892-1-jeremy.szu@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b65a9d237bd61a5cdfdd334ccceed2f86cdb61f
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Tue Jul 12 17:22:22 2022 +0800

    ALSA: hda/realtek - Fix headset mic problem for a HP machine with alc671

    commit dbe75d314748e08fc6e4576d153d8a69621ee5ca upstream.

    On a HP 288 Pro G6, the front mic could not be detected.In order to
    get it working, the pin configuration needs to be set correctly, and
    the ALC671_FIXUP_HP_HEADSET_MIC2 fixup needs to be applied.

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220712092222.21738-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c6cce94bdcc269f14f0800fa530cafaa7b96fe87
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Mon Jul 11 16:15:27 2022 +0800

    ALSA: hda/realtek: Fix headset mic for Acer SF313-51

    commit 5f3fe25e70559fa3b096ab17e13316c93ddb7020 upstream.

    The issue on Acer SWIFT SF313-51 is that headset microphone
    doesn't work. The following quirk fixed headset microphone issue.
    Note that the fixup of SF314-54/55 (ALC256_FIXUP_ACER_HEADSET_MIC)
    was not successful on my SF313-51.

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220711081527.6254-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a9ff4345d5ad35276a51f6b6dec7af8207e20599
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Mon Jul 11 18:17:44 2022 +0800

    ALSA: hda/conexant: Apply quirk for another HP ProDesk 600 G3 model

    commit d16d69bf5a25d91c6d8f3e29711be12551bf56cd upstream.
```

```
    There is another HP ProDesk 600 G3 model with the PCI SSID 103c:82b4
    that requires the quirk HP_MIC_NO_PRESENCE. Add the corresponding
    entry to the quirk table.

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220711101744.25189-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cd52154b924f2ea05069d4296045d9fd56a8da23
Author: Meng Tang <tangmeng@uniontech.com>
Date:   Tue Jul 12 14:00:05 2022 +0800

    ALSA: hda - Add fixup for Dell Latitidue E5430

    commit 841bdf85c226803a78a9319af9b2caa9bf3e2eda upstream.

    Another Dell model, another fixup entry: Latitude E5430 needs the same
    fixup as other Latitude E series as workaround for noise problems.

    Signed-off-by: Meng Tang <tangmeng@uniontech.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20220712060005.20176-1-tangmeng@uniontech.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f0f7faf210c9c993604b839816124a8aaadd4b2
Author: Juergen Gross <jgross@suse.com>
Date:   Thu Jun 30 09:14:39 2022 +0200

    x86/xen: Use clear_bss() for Xen PV guests

    commit 96e8fc5818686d4a1591bb6907e7fdb64ef29884 upstream.

    Instead of clearing the bss area in assembly code, use the clear_bss()
    function.

    This requires to pass the start_info address as parameter to
    xen_start_kernel() in order to avoid the xen_start_info being zeroed
    again.

    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Jan Beulich <jbeulich@suse.com>
    Reviewed-by: Boris Ostrovsky <boris.ostrovsky@oracle.com>
    Link: https://lore.kernel.org/r/20220630071441.28576-2-jgross@suse.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b24c8369355afb1e7a3caa6eed1209748db00d20
Author: Chanho Park <chanho61.park@samsung.com>
Date:   Mon Jun 27 15:51:13 2022 +0900

    tty: serial: samsung_tty: set dma burst_size to 1

    commit f7e35e4bf1e8dc2c8cbd5e0955dc1bd58558dae0 upstream.

    The src_maxburst and dst_maxburst have been changed to 1 but the settings
    of the UCON register aren't changed yet. They should be changed as well
    according to the dmaengine slave config.

    Fixes: aa2f80e752c7 ("serial: samsung: fix maxburst parameter for DMA transactions")
    Cc: stable <stable@kernel.org>
    Cc: Marek Szyprowski <m.szyprowski@samsung.com>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
```

```
    Signed-off-by: Chanho Park <chanho61.park@samsung.com>
    Link: https://lore.kernel.org/r/20220627065113.139520-1-chanho61.park@samsung.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a25d361f141196ad6fd2956bc1861e72bc275879
Author: Bartosz Golaszewski <brgl@bgdev.pl>
Date:   Tue Jul 12 09:40:55 2022 +0200

    gpio: sim: fix the chip_name configfs item

    commit 7329b071729645e243b6207e76bca2f4951c991b upstream.

    The chip_name configs attribute always displays the device name of the
    first GPIO bank because the logic of the relevant function is simply
    wrong.

    Fix it by correctly comparing the bank's swnode against the GPIO
    device's children.

    Fixes: cb8c474e79be ("gpio: sim: new testing module")
    Cc: stable@vger.kernel.org
    Reported-by: Kent Gibson <warthog618@gmail.com>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Reviewed-by: Kent Gibson <warthog618@gmail.com>
    Tested-by: Kent Gibson <warthog618@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1bbc49745687b0b1edf47856f7e5bc8cb37ddd4e
Author: Thinh Nguyen <Thinh.Nguyen@synopsys.com>
Date:   Mon Jun 27 18:41:19 2022 -0700

    usb: dwc3: gadget: Fix event pending check

    commit 7441b273388b9a59d8387a03ffbbca9d5af6348c upstream.

    The DWC3_EVENT_PENDING flag is used to protect against invalid call to
    top-half interrupt handler, which can occur when there's a delay in
    software detection of the interrupt line deassertion.

    However, the clearing of this flag was done prior to unmasking the
    interrupt line, creating opportunity where the top-half handler can
    come. This breaks the serialization and creates a race between the
    top-half and bottom-half handler, resulting in losing synchronization
    between the controller and the driver when processing events.

    To fix this, make sure the clearing of the DWC3_EVENT_PENDING is done at
    the end of the bottom-half handler.

    Fixes: d325a1de49d6 ("usb: dwc3: gadget: Prevent losing events in event cache")
    Cc: stable@vger.kernel.org
    Signed-off-by: Thinh Nguyen <Thinh.Nguyen@synopsys.com>
    Link:
https://lore.kernel.org/r/8670aaf1cf52e7d1e6df2a827af2d77263b93b75.1656380429.git.Thinh.Nguy
en@synopsys.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fd54a40491ca5b874aa6a720d9257063929fb1ab
Author: Linyu Yuan <quic_linyyuan@quicinc.com>
Date:   Fri Jul 1 16:08:54 2022 +0800

    usb: typec: add missing uevent when partner support PD

    commit 6fb9e1d94789e8ee5a258a23bc588693f743fd6c upstream.
```

```
    System like Android allow user control power role from UI, it is possible
    to implement application base on typec uevent to refresh UI, but found
    there is chance that UI show different state from typec attribute file.

    In typec_set_pwr_opmode(), when partner support PD, there is no uevent
    send to user space which cause the problem.

    Fix it by sending uevent notification when change power mode to PD.

    Fixes: bdecb33af34f ("usb: typec: API for controlling USB Type-C Multiplexers")
    Cc: stable@vger.kernel.org
    Signed-off-by: Linyu Yuan <quic_linyyuan@quicinc.com>
    Link: https://lore.kernel.org/r/1656662934-10226-1-git-send-email-
quic_linyyuan@quicinc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8185962cbe6f5b89b992409cc104f9075dffe84a
Author: Lucien Buchmann <lucien.buchmann@gmx.net>
Date:   Sat Jun 25 02:17:44 2022 +0200

    USB: serial: ftdi_sio: add Belimo device ids

    commit 7c239a071d1f04b7137789810807b4108d475c72 upstream.

    Those two product ids are known.

    Signed-off-by: Lucien Buchmann <lucien.buchmann@gmx.net>
    Cc: stable@vger.kernel.org
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```