```
commit cdc56cf3e0ca6bd3119fc841de1c37b0e6e5b802
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Fri Mar 10 09:29:57 2023 +0100

    Linux 6.2.3

    Link: https://lore.kernel.org/r/20230307170022.094103862@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Luna Jernberg <droidbittin@gmail.com>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Conor Dooley <conor.dooley@microchip.com>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Link: https://lore.kernel.org/r/20230308091912.362228731@linuxfoundation.org
    Tested-by: Luna Jernberg <droidbittin@gmail.com>
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Jiri Slaby <jirislaby@kernel.org>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cb6aedc1fd9d808d7319db2f953f4886dd46c627
Author: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
Date:   Tue Mar 7 14:37:07 2023 -0300

    net: avoid double iput when sock_alloc_file fails

    commit 649c15c7691e9b13cbe9bf6c65c365350e056067 upstream.

    When sock_alloc_file fails to allocate a file, it will call sock_release.
    __sys_socket_file should then not call sock_release again, otherwise there
    will be a double free.

    [   89.319884] ------------[ cut here ]------------
    [   89.320286] kernel BUG at fs/inode.c:1764!
    [   89.320656] invalid opcode: 0000 [#1] PREEMPT SMP NOPTI
    [   89.321051] CPU: 7 PID: 125 Comm: iou-sqp-124 Not tainted 6.2.0+ #361
    [   89.321535] RIP: 0010:iput+0x1ff/0x240
    [   89.321808] Code: d1 83 e1 03 48 83 f9 02 75 09 48 81 fa 00 10 00 00 77 05 83 e2 01
75 1f 4c 89 ef e8 fb d2 ba 00 e9 80 fe ff ff c3 cc cc cc cc <0f> 0b 0f 0b e9 d0 fe ff ff 0f
0b eb 8d 49 8d b4 24 08 01 00 00 48
    [   89.322760] RSP: 0018:ffffbdd60068bd50 EFLAGS: 00010202
    [   89.323036] RAX: 0000000000000000 RBX: ffff9d7ad3cacac0 RCX: 0000000000001107
    [   89.323412] RDX: 000000000003af00 RSI: 0000000000000000 RDI: ffff9d7ad3cacb40
    [   89.323785] RBP: ffffbdd60068bd68 R08: ffffffffffffffff R09: ffffffffab606438
    [   89.324157] R10: ffffffffacb3dfa0 R11: 6465686361657256 R12: ffff9d7ad3cacb40
    [   89.324529] R13: 0000000080000001 R14: 0000000080000001 R15: 0000000000000002
    [   89.324904] FS:  00007f7b28516740(0000) GS:ffff9d7aeb1c0000(0000)
knlGS:0000000000000000
    [   89.325328] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [   89.325629] CR2: 00007f0af52e96c0 CR3: 0000000002a02006 CR4: 0000000000770ee0
    [   89.326004] PKRU: 55555554
    [   89.326161] Call Trace:
    [   89.326298]  <TASK>
    [   89.326419]  __sock_release+0xb5/0xc0
    [   89.326632]  __sys_socket_file+0xb2/0xd0
    [   89.326844]  io_socket+0x88/0x100
    [   89.327039]  ? io_issue_sqe+0x6a/0x430
    [   89.327258]  io_issue_sqe+0x67/0x430
    [   89.327450]  io_submit_sqes+0x1fe/0x670
    [   89.327661]  io_sq_thread+0x2e6/0x530
```

```
[   89.327859]  ? __pfx_autoremove_wake_function+0x10/0x10
[   89.328145]  ? __pfx_io_sq_thread+0x10/0x10
[   89.328367]  ret_from_fork+0x29/0x50
[   89.328576] RIP: 0033:0x0
[   89.328732] Code: Unable to access opcode bytes at 0xfffffffffffffd6.
[   89.329073] RSP: 002b:0000000000000000 EFLAGS: 00000202 ORIG_RAX: 00000000000001a9
[   89.329477] RAX: 0000000000000000 RBX: 0000000000000000 RCX: 00007f7b28637a3d
[   89.329845] RDX: 00007fff4e4318a8 RSI: 00007fff4e4318b0 RDI: 0000000000000400
[   89.330216] RBP: 00007fff4e431830 R08: 00007fff4e431711 R09: 00007fff4e4318b0
[   89.330584] R10: 0000000000000000 R11: 0000000000000202 R12: 00007fff4e441b38
[   89.330950] R13: 0000563835e3e725 R14: 0000563835e40d10 R15: 00007f7b28784040
[   89.331318]  </TASK>
[   89.331441] Modules linked in:
[   89.331617] ---[ end trace 0000000000000000 ]---
```

Fixes: da214a475f8b ("net: add __sys_socket_file()")
Signed-off-by: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
Reviewed-by: Jens Axboe <axboe@kernel.dk>
Reviewed-by: Eric Dumazet <edumazet@google.com>
Reviewed-by: Kuniyuki Iwashima <kuniyu@amazon.com>
Link: https://lore.kernel.org/r/20230307173707.468744-1-cascardo@canonical.com
Signed-off-by: Jakub Kicinski <kuba@kernel.org>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2cbe0f05664de37bbc084b79001f1436e19b5da6
Author: Pankaj Raghav <p.raghav@samsung.com>
Date: Fri Feb 17 17:44:44 2023 +0530

brd: use radix_tree_maybe_preload instead of radix_tree_preload

commit 0aa2988e4fd23c0c8b33999d7b47dfbc5e6bf24b upstream.

Unconditionally calling radix_tree_preload_end() results in a OOPS
message as the preload is only conditionally called for
gfpflags_allow_blocking().

```
[   20.267323] BUG: using smp_processor_id() in preemptible [00000000] code: fio/416
[   20.267837] caller is brd_insert_page.part.0+0xbe/0x190 [brd]
[   20.269436] Call Trace:
[   20.269598]  <TASK>
[   20.269742]  dump_stack_lvl+0x32/0x50
[   20.269982]  check_preemption_disabled+0xd1/0xe0
[   20.270289]  brd_insert_page.part.0+0xbe/0x190 [brd]
[   20.270664]  brd_submit_bio+0x33f/0xf40 [brd]
```

Use radix_tree_maybe_preload() which does preload only if
gfpflags_allow_blocking() is true but also takes the lock. Therefore,
unconditionally calling radix_tree_preload_end() should not create any
issues and the message disappears.

Fixes: 6ded703c56c2 ("brd: check for REQ_NOWAIT and set correct page allocation mask")
Signed-off-by: Pankaj Raghav <p.raghav@samsung.com>
Link: https://lore.kernel.org/r/20230217121442.33914-1-p.raghav@samsung.com
Signed-off-by: Jens Axboe <axboe@kernel.dk>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2d3b582b80c7dcb8ca12b36a4ac0e836828ecdeb
Author: Michal Schmidt <mschmidt@redhat.com>
Date: Fri Feb 24 01:41:45 2023 +0100

qede: avoid uninitialized entries in coal_entry array

commit aaa3c08ee0653beaa649d4adfb27ad562641cfd8 upstream.

Even after commit 908d4bb7c54c ("qede: fix interrupt coalescing

```
      configuration"), some entries of the coal_entry array may theoretically
      be used uninitialized:

       1. qede_alloc_fp_array() allocates QEDE_MAX_RSS_CNT entries for
          coal_entry. The initial allocation uses kcalloc, so everything is
          initialized.
       2. The user sets a small number of queues (ethtool -L).
          coal_entry is reallocated for the actual small number of queues.
       3. The user sets a bigger number of queues.
          coal_entry is reallocated bigger. The added entries are not
          necessarily initialized.

      In practice, the reallocations will actually keep using the originally
      allocated region of memory, but we should not rely on it.

      The reallocation is unnecessary. coal_entry can always have
      QEDE_MAX_RSS_CNT entries.

      Fixes: 908d4bb7c54c ("qede: fix interrupt coalescing configuration")
      Signed-off-by: Michal Schmidt <mschmidt@redhat.com>
      Nacked-by: Manish Chopra <manishc@marvell.com>
      Acked-by: Manish Chopra <manishc@marvell.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 58be98f457f96132c994b3e00bf477588f3d49ba
  Author: Jani Nikula <jani.nikula@intel.com>
  Date:   Wed Jan 4 12:05:17 2023 +0200

      drm/edid: fix parsing of 3D modes from HDMI VSDB

      commit 72794d16bd535a984e6653a18f5862405b49b5f9 upstream.

      Commit 537d9ed2f6c1 ("drm/edid: convert add_cea_modes() to use cea db
      iter") inadvertently moved the do_hdmi_vsdb_modes() call within the db
      iteration loop, always passing NULL as the CTA VDB to
      do_hdmi_vsdb_modes(), skipping a lot of stereo modes.

      Move the call back outside of the loop.

      This does mean only one CTA VDB and HDMI VSDB combination will be
      handled, but it's an unlikely scenario to have more than one of either
      block, and it was not accounted for before the regression either.

      Fixes: 537d9ed2f6c1 ("drm/edid: convert add_cea_modes() to use cea db iter")
      Cc: <stable@vger.kernel.org> # v6.0+
      Cc: Ville Syrjälä <ville.syrjala@linux.intel.com>
      Signed-off-by: Jani Nikula <jani.nikula@intel.com>
      Reviewed-by: Ville Syrjälä <ville.syrjala@linux.intel.com>
      Link:
  https://patchwork.freedesktop.org/patch/msgid/cf159b8816191ed595a3cb954acaf189c4528cc7.16728
  26282.git.jani.nikula@intel.com
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 86e58683796f0dfda6a0f84d251a69cc0c17181a
  Author: Jani Nikula <jani.nikula@intel.com>
  Date:   Wed Jan 4 12:05:16 2023 +0200

      drm/edid: fix AVI infoframe aspect ratio handling

      commit 1cbc1f0d324ba6c4d1b10ac6362b5e0b029f63d5 upstream.

      We try to avoid sending VICs defined in the later specs in AVI
      infoframes to sinks that conform to the earlier specs, to not upset
      them, and use 0 for the VIC instead. However, we do this detection and
```

```
    conversion to 0 too early, as we'll need the actual VIC to figure out
    the aspect ratio.

    In particular, for a mode with 64:27 aspect ratio, 0 for VIC fails the
    AVI infoframe generation altogether with -EINVAL.

    Separate the VIC lookup from the "filtering", and postpone the
    filtering, to use the proper VIC for aspect ratio handling, and the 0
    VIC for the infoframe video code as needed.

    Reported-by: William Tseng <william.tseng@intel.com>
    Closes: https://gitlab.freedesktop.org/drm/intel/-/issues/6153
    References: https://lore.kernel.org/r/20220920062316.43162-1-william.tseng@intel.com
    Cc: <stable@vger.kernel.org>
    Cc: Ville Syrjälä <ville.syrjala@linux.intel.com>
    Signed-off-by: Jani Nikula <jani.nikula@intel.com>
    Reviewed-by: Ville Syrjälä <ville.syrjala@linux.intel.com>
    Link:
https://patchwork.freedesktop.org/patch/msgid/c3e78cc6d01ed237f71ad0038826b08d83d75eef.16728
26282.git.jani.nikula@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e1078b270d218f8d58efb4d78ea25a4d16ba3490
Author: Noralf Trønnes <noralf@tronnes.org>
Date:   Wed Nov 30 20:26:49 2022 +0100

    drm/gud: Fix UBSAN warning

    commit 951df98024f7272f85df5044eca7374f5b5b24ef upstream.

    UBSAN complains about invalid value for bool:

    [  101.165172] [drm] Initialized gud 1.0.0 20200422 for 2-3.2:1.0 on minor 1
    [  101.213360] gud 2-3.2:1.0: [drm] fb1: guddrmfb frame buffer device
    [  101.213426] usbcore: registered new interface driver gud
    [  101.989431]
================================================================================
    [  101.989441] UBSAN: invalid-load in linux/include/linux/iosys-map.h:253:9
    [  101.989447] load of value 121 is not a valid value for type '_Bool'
    [  101.989451] CPU: 1 PID: 455 Comm: kworker/1:6 Not tainted 5.18.0-rc5-gud-5.18-rc5 #3
    [  101.989456] Hardware name: Hewlett-Packard HP EliteBook 820 G1/1991, BIOS L71 Ver.
01.44 04/12/2018
    [  101.989459] Workqueue: events_long gud_flush_work [gud]
    [  101.989471] Call Trace:
    [  101.989474]  <TASK>
    [  101.989479]  dump_stack_lvl+0x49/0x5f
    [  101.989488]  dump_stack+0x10/0x12
    [  101.989493]  ubsan_epilogue+0x9/0x3b
    [  101.989498]  __ubsan_handle_load_invalid_value.cold+0x44/0x49
    [  101.989504]  dma_buf_vmap.cold+0x38/0x3d
    [  101.989511]  ? find_busiest_group+0x48/0x300
    [  101.989520]  drm_gem_shmem_vmap+0x76/0x1b0 [drm_shmem_helper]
    [  101.989528]  drm_gem_shmem_object_vmap+0x9/0xb [drm_shmem_helper]
    [  101.989535]  drm_gem_vmap+0x26/0x60 [drm]
    [  101.989594]  drm_gem_fb_vmap+0x47/0x150 [drm_kms_helper]
    [  101.989630]  gud_prep_flush+0xc1/0x710 [gud]
    [  101.989639]  ? _raw_spin_lock+0x17/0x40
    [  101.989648]  gud_flush_work+0x1e0/0x430 [gud]
    [  101.989653]  ? __switch_to+0x11d/0x470
    [  101.989664]  process_one_work+0x21f/0x3f0
    [  101.989673]  worker_thread+0x200/0x3e0
    [  101.989679]  ? rescuer_thread+0x390/0x390
    [  101.989684]  kthread+0xfd/0x130
    [  101.989690]  ? kthread_complete_and_exit+0x20/0x20
    [  101.989696]  ret_from_fork+0x22/0x30
```

```
     [  101.989706]  </TASK>
     [  101.989708]
 ================================================================================

     The source of this warning is in iosys_map_clear() called from
     dma_buf_vmap(). It conditionally sets values based on map->is_iomem. The
     iosys_map variables are allocated uninitialized on the stack leading to
     ->is_iomem having all kinds of values and not only 0/1.

     Fix this by zeroing the iosys_map variables.

     Fixes: 40e1a70b4aed ("drm: Add GUD USB Display driver")
     Cc: <stable@vger.kernel.org> # v5.18+
     Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
     Reviewed-by: Thomas Zimmermann <tzimmermann@suse.de>
     Signed-off-by: Noralf Trønnes <noralf@tronnes.org>
     Link: https://patchwork.freedesktop.org/patch/msgid/20221122-gud-shadow-plane-v2-1-
 435037990a83@tronnes.org
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit f8c6d7f805855095b79beb510b5746c5890ae6e1
 Author: John Harrison <John.C.Harrison@Intel.com>
 Date:   Wed Feb 15 17:11:01 2023 -0800

     drm/i915: Don't use BAR mappings for ring buffers with LLC

     commit 85636167e3206c3fbd52254fc432991cc4e90194 upstream.

     Direction from hardware is that ring buffers should never be mapped
     via the BAR on systems with LLC. There are too many caching pitfalls
     due to the way BAR accesses are routed. So it is safest to just not
     use it.

     Signed-off-by: John Harrison <John.C.Harrison@Intel.com>
     Fixes: 9d80841ea4c9 ("drm/i915: Allow ringbuffers to be bound anywhere")
     Cc: Chris Wilson <chris@chris-wilson.co.uk>
     Cc: Joonas Lahtinen <joonas.lahtinen@linux.intel.com>
     Cc: Jani Nikula <jani.nikula@linux.intel.com>
     Cc: Rodrigo Vivi <rodrigo.vivi@intel.com>
     Cc: Tvrtko Ursulin <tvrtko.ursulin@linux.intel.com>
     Cc: intel-gfx@lists.freedesktop.org
     Cc: <stable@vger.kernel.org> # v4.9+
     Tested-by: Jouni Högander <jouni.hogander@intel.com>
     Reviewed-by: Daniele Ceraolo Spurio <daniele.ceraolospurio@intel.com>
     Link: https://patchwork.freedesktop.org/patch/msgid/20230216011101.1909009-3-
 John.C.Harrison@Intel.com
     (cherry picked from commit 65c08339db1ada87afd6cfe7db8e60bb4851d919)
     Signed-off-by: Jani Nikula <jani.nikula@intel.com>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit f8354e04e5faac765fef967c24772e358c426cef
 Author: John Harrison <John.C.Harrison@Intel.com>
 Date:   Wed Feb 15 17:11:00 2023 -0800

     drm/i915: Don't use stolen memory for ring buffers with LLC

     commit 690e0ec8e63da9a29b39fedc6ed5da09c7c82651 upstream.

     Direction from hardware is that stolen memory should never be used for
     ring buffer allocations on platforms with LLC. There are too many
     caching pitfalls due to the way stolen memory accesses are routed. So
     it is safest to just not use it.

     Signed-off-by: John Harrison <John.C.Harrison@Intel.com>
     Fixes: c58b735fc762 ("drm/i915: Allocate rings from stolen")
```

```
    Cc: Chris Wilson <chris@chris-wilson.co.uk>
    Cc: Joonas Lahtinen <joonas.lahtinen@linux.intel.com>
    Cc: Jani Nikula <jani.nikula@linux.intel.com>
    Cc: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Cc: Tvrtko Ursulin <tvrtko.ursulin@linux.intel.com>
    Cc: intel-gfx@lists.freedesktop.org
    Cc: <stable@vger.kernel.org> # v4.9+
    Tested-by: Jouni Högander <jouni.hogander@intel.com>
    Reviewed-by: Daniele Ceraolo Spurio <daniele.ceraolospurio@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230216011101.1909009-2-
John.C.Harrison@Intel.com
    (cherry picked from commit f54c1f6c697c4297f7ed94283c184acc338a5cf8)
    Signed-off-by: Jani Nikula <jani.nikula@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b6f66ec601b6add736265d49dfd96be18b724e36
Author: Mark Hawrylak <mark.hawrylak@gmail.com>
Date:   Sun Feb 19 16:02:00 2023 +1100

    drm/radeon: Fix eDP for single-display iMac11,2

    commit 05eacc198c68cbb35a7281ce4011f8899ee1cfb8 upstream.

    Apple iMac11,2 (mid 2010) also with Radeon HD-4670 that has the same
    issue as iMac10,1 (late 2009) where the internal eDP panel stays dark on
    driver load.  This patch treats iMac11,2 the same as iMac10,1,
    so the eDP panel stays active.

    Additional steps:
    Kernel boot parameter radeon.nomodeset=0 required to keep the eDP
    panel active.

    This patch is an extension of
    commit 564d8a2cf3ab ("drm/radeon: Fix eDP for single-display iMac10,1 (v2)")
    Link: https://lore.kernel.org/all/lsq.1507553064.833262317@decadent.org.uk/
    Signed-off-by: Mark Hawrylak <mark.hawrylak@gmail.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 21c7c846279fad6e245d32f343336ae021853b6c
Author: Mavroudis Chatzilaridis <mavchatz@protonmail.com>
Date:   Wed Feb 1 18:51:25 2023 +0000

    drm/i915/quirks: Add inverted backlight quirk for HP 14-r206nv

    commit 5e438bf7f9a1705ebcae5fa89cdbfbc6932a7871 upstream.

    This laptop uses inverted backlight PWM. Thus, without this quirk,
    backlight brightness decreases as the brightness value increases and
    vice versa.

    Closes: https://gitlab.freedesktop.org/drm/intel/-/issues/8013
    Cc: stable@vger.kernel.org
    Signed-off-by: Mavroudis Chatzilaridis <mavchatz@protonmail.com>
    Reviewed-by: Jani Nikula <jani.nikula@intel.com>
    Signed-off-by: Jani Nikula <jani.nikula@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230201184947.8835-1-
mavchatz@protonmail.com
    (cherry picked from commit 83e7d6fd330d413cb2064e680ffea91b0512a520)
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 55449229e88fab284c9e6554005c4c3296ea98b9
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Sun Feb 19 23:04:04 2023 -0600
```

```
    drm/amd: Fix initialization for nbio 7.5.1

    commit 65a24000808f70ac69bd2a96381fa0c7341f20c0 upstream.

    A mistake has been made in the BIOS for some ASICs with NBIO 7.5.1
    where some NBIO registers aren't properly setup.

    Ensure that they're set during initialization.

    Tested-by: Richard Gong <richard.gong@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org # 6.1.x
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 98bd3240f16ec1fc8890f3102acd1e49edb69b5b
Author: Steve Sistare <steven.sistare@oracle.com>
Date:   Tue Jan 31 08:58:06 2023 -0800

    vfio/type1: restore locked_vm

    commit 90fdd158a695d70403163f9a0e4efc5b20f3fd3e upstream.

    When a vfio container is preserved across exec or fork-exec, the new
    task's mm has a locked_vm count of 0.  After a dma vaddr is updated using
    VFIO_DMA_MAP_FLAG_VADDR, locked_vm remains 0, and the pinned memory does
    not count against the task's RLIMIT_MEMLOCK.

    To restore the correct locked_vm count, when VFIO_DMA_MAP_FLAG_VADDR is
    used and the dma's mm has changed, add the dma's locked_vm count to
    the new mm->locked_vm, subject to the rlimit, and subtract it from the
    old mm->locked_vm.

    Fixes: c3cbab24db38 ("vfio/type1: implement interfaces to update vaddr")
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve Sistare <steven.sistare@oracle.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Link: https://lore.kernel.org/r/1675184289-267876-5-git-send-email-
steven.sistare@oracle.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 273fa49906a02bf9fdadbb15acfd86882d915a55
Author: Steve Sistare <steven.sistare@oracle.com>
Date:   Tue Jan 31 08:58:05 2023 -0800

    vfio/type1: track locked_vm per dma

    commit 18e292705ba21cc9b3227b9ad5b1c28973605ee5 upstream.

    Track locked_vm per dma struct, and create a new subroutine, both for use
    in a subsequent patch.  No functional change.

    Fixes: c3cbab24db38 ("vfio/type1: implement interfaces to update vaddr")
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve Sistare <steven.sistare@oracle.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Link: https://lore.kernel.org/r/1675184289-267876-4-git-send-email-
steven.sistare@oracle.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit b0790dff0760b7734cf0961f497ad64628ca550b
Author: Steve Sistare <steven.sistare@oracle.com>
Date:   Tue Jan 31 08:58:04 2023 -0800

    vfio/type1: prevent underflow of locked_vm via exec()

    commit 046eca5018f8a5dd1dc2cedf87fb5843b9ea3026 upstream.

    When a vfio container is preserved across exec, the task does not change,
    but it gets a new mm with locked_vm=0, and loses the count from existing
    dma mappings.  If the user later unmaps a dma mapping, locked_vm underflows
    to a large unsigned value, and a subsequent dma map request fails with
    ENOMEM in __account_locked_vm.

    To avoid underflow, grab and save the mm at the time a dma is mapped.
    Use that mm when adjusting locked_vm, rather than re-acquiring the saved
    task's mm, which may have changed.  If the saved mm is dead, do nothing.

    locked_vm is incremented for existing mappings in a subsequent patch.

    Fixes: 73fa0d10d077 ("vfio: Type1 IOMMU implementation")
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve Sistare <steven.sistare@oracle.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Link: https://lore.kernel.org/r/1675184289-267876-3-git-send-email-
steven.sistare@oracle.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 346fb41f157aa7e507bd7828d4ce431255aa59e4
Author: Steve Sistare <steven.sistare@oracle.com>
Date:   Tue Jan 31 08:58:03 2023 -0800

    vfio/type1: exclude mdevs from VFIO_UPDATE_VADDR

    commit ef3a3f6a294ba65fd906a291553935881796f8a5 upstream.

    Disable the VFIO_UPDATE_VADDR capability if mediated devices are present.
    Their kernel threads could be blocked indefinitely by a misbehaving
    userland while trying to pin/unpin pages while vaddrs are being updated.

    Do not allow groups to be added to the container while vaddr's are invalid,
    so we never need to block user threads from pinning, and can delete the
    vaddr-waiting code in a subsequent patch.

    Fixes: c3cbab24db38 ("vfio/type1: implement interfaces to update vaddr")
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve Sistare <steven.sistare@oracle.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Link: https://lore.kernel.org/r/1675184289-267876-2-git-send-email-
steven.sistare@oracle.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bc348a6cc8f92e9979fbfbb3c6b48fd51b4b1dbf
Author: Jacob Pan <jacob.jun.pan@linux.intel.com>
Date:   Thu Feb 16 21:08:15 2023 +0800

    iommu/vt-d: Fix PASID directory pointer coherency

    commit 194b3348bdbb7db65375c72f3f774aee4cc6614e upstream.
```

```
    On platforms that do not support IOMMU Extended capability bit 0
    Page-walk Coherency, CPU caches are not snooped when IOMMU is accessing
    any translation structures. IOMMU access goes only directly to
    memory. Intel IOMMU code was missing a flush for the PASID table
    directory that resulted in the unrecoverable fault as shown below.

    This patch adds clflush calls whenever allocating and updating
    a PASID table directory to ensure cache coherency.

    On the reverse direction, there's no need to clflush the PASID directory
    pointer when we deactivate a context entry in that IOMMU hardware will
    not see the old PASID directory pointer after we clear the context entry.
    PASID directory entries are also never freed once allocated.

     DMAR: DRHD: handling fault status reg 3
     DMAR: [DMA Read NO_PASID] Request device [00:0d.2] fault addr 0x1026a4000
           [fault reason 0x51] SM: Present bit in Directory Entry is clear
     DMAR: Dump dmar1 table entries for IOVA 0x1026a4000
     DMAR: scalable mode root entry: hi 0x0000000102448001, low 0x0000000101b3e001
     DMAR: context entry: hi 0x0000000000000000, low 0x0000000101b4d401
     DMAR: pasid dir entry: 0x0000000101b4e001
     DMAR: pasid table entry[0]: 0x0000000000000109
     DMAR: pasid table entry[1]: 0x0000000000000001
     DMAR: pasid table entry[2]: 0x0000000000000000
     DMAR: pasid table entry[3]: 0x0000000000000000
     DMAR: pasid table entry[4]: 0x0000000000000000
     DMAR: pasid table entry[5]: 0x0000000000000000
     DMAR: pasid table entry[6]: 0x0000000000000000
     DMAR: pasid table entry[7]: 0x0000000000000000
     DMAR: PTE not present at level 4

    Cc: <stable@vger.kernel.org>
    Fixes: 0bbeb01a4faf ("iommu/vt-d: Manage scalalble mode PASID tables")
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reported-by: Sukumar Ghorai <sukumar.ghorai@intel.com>
    Signed-off-by: Ashok Raj <ashok.raj@intel.com>
    Signed-off-by: Jacob Pan <jacob.jun.pan@linux.intel.com>
    Link: https://lore.kernel.org/r/20230209212843.1788125-1-jacob.jun.pan@linux.intel.com
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 863b125439c2491d88d9ca981841e737d701344e
Author: Jacob Pan <jacob.jun.pan@linux.intel.com>
Date:   Thu Feb 16 21:08:14 2023 +0800

    iommu/vt-d: Avoid superfluous IOTLB tracking in lazy mode

    commit 16a75bbe480c3598b3af57a2504ea89b1e32c3ac upstream.

    Intel IOMMU driver implements IOTLB flush queue with domain selective
    or PASID selective invalidations. In this case there's no need to track
    IOVA page range and sync IOTLBs, which may cause significant performance
    hit.

    This patch adds a check to avoid IOVA gather page and IOTLB sync for
    the lazy path.

    The performance difference on Sapphire Rapids 100Gb NIC is improved by
    the following (as measured by iperf send):

    w/o this fix~48 Gbits/s. with this fix ~54 Gbits/s

    Cc: <stable@vger.kernel.org>
    Fixes: 2a2b8eaa5b25 ("iommu: Handle freelists when using deferred flushing in iommu
```

```
drivers")
    Reviewed-by: Robin Murphy <robin.murphy@arm.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Tested-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Signed-off-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Signed-off-by: Jacob Pan <jacob.jun.pan@linux.intel.com>
    Link: https://lore.kernel.org/r/20230209175330.1783556-1-jacob.jun.pan@linux.intel.com
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c44adefdcf472f946f0632f4e0ddcbf3e00b8516
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date:   Mon Feb 13 14:02:42 2023 -0400

    iommufd: Do not add the same hwpt to the ioas->hwpt_list twice

    commit b4ff830eca097df51af10a9be29e8cc817327919 upstream.

    The hwpt is added to the hwpt_list only during its creation, it is never
    added again. This hunk is some missed leftover from rework. Adding it
    twice will corrupt the linked list in some cases.

    It effects HWPT specific attachment, which is something the test suite
    cannot cover until we can create a legitimate struct device with a
    non-system iommu "driver" (ie we need the bus removed from the iommu code)

    Cc: stable@vger.kernel.org
    Fixes: e8d57210035b ("iommufd: Add kAPI toward external drivers for physical devices")
    Link: https://lore.kernel.org/r/1-v1-4336b5cb2fe4+1d7-iommufd_hwpt_jgg@nvidia.com
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reported-by: Kevin Tian <kevin.tian@intel.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7adcec686e4d699c169d34c722132b2bce5232cb
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date:   Mon Feb 13 10:32:21 2023 -0400

    iommufd: Make sure to zero vfio_iommu_type1_info before copying to user

    commit b3551ead616318ea155558cdbe7e91495b8d9b33 upstream.

    Missed a zero initialization here. Most of the struct is filled with
    a copy_from_user(), however minsz for that copy is smaller than the
    actual struct by 8 bytes, thus we don't fill the padding.

    Cc: stable@vger.kernel.org # 6.1+
    Fixes: d624d6652a65 ("iommufd: vfio container FD ioctl compatibility")
    Link: https://lore.kernel.org/r/0-v1-a74499ece799+1a-
iommufd_get_info_leak_jgg@nvidia.com
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reported-by: syzbot+cb1e0978f6bf46b83a58@syzkaller.appspotmail.com
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6144b1cdeb65bdd50a1957e74037f1a2ec0a8456
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed Dec 28 21:47:04 2022 +0530

    bus: mhi: ep: Save channel state locally during suspend and resume

    commit 8a1c24bb908f9ecbc4be0fea014df67d43161551 upstream.

    During suspend and resume, the channel state needs to be saved locally.
```

```
    Otherwise, the endpoint may access the channels while they were being
    suspended and causing access violations.

    Fix it by saving the channel state locally during suspend and resume.

    Cc: <stable@vger.kernel.org> # 5.19
    Fixes: e4b7b5f0f30a ("bus: mhi: ep: Add support for suspending and resuming channels")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com)
    Link: https://lore.kernel.org/r/20221228161704.255268-7-manivannan.sadhasivam@linaro.org
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b619b60153656c40c008f63ad950671bb54b8df2
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed Dec 28 21:47:03 2022 +0530

    bus: mhi: ep: Move chan->lock to the start of processing queued ch ring

    commit 8d6a1fea53864cd9545741f48f4ae4df804db557 upstream.

    There is a good chance that while the channel ring gets processed, the STOP
    or RESET command for the channel might be received from the MHI host. In
    those cases, the entire channel ring processing needs to be protected by
    chan->lock to prevent the race where the corresponding channel ring might
    be reset.

    While at it, let's also add a sanity check to make sure that the ring is
    started before processing it. Because, if the STOP/RESET command gets
    processed while mhi_ep_ch_ring_worker() waited for chan->lock, the ring
    would've been reset.

    Cc: <stable@vger.kernel.org> # 5.19
    Fixes: 03c0bb8ec983 ("bus: mhi: ep: Add support for processing channel rings")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com>
    Link: https://lore.kernel.org/r/20221228161704.255268-6-manivannan.sadhasivam@linaro.org
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 860ad591056d7e4dc30bc130b6ec6e6d70930c85
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed Dec 28 21:47:01 2022 +0530

    bus: mhi: ep: Only send -ENOTCONN status if client driver is available

    commit e6cebcc27519dcf1652e604c73b9fd4f416987c0 upstream.

    For the STOP and RESET commands, only send the channel disconnect status
    -ENOTCONN if client driver is available. Otherwise, it will result in
    null pointer dereference.

    Cc: <stable@vger.kernel.org> # 5.19
    Fixes: e827569062a8 ("bus: mhi: ep: Add support for processing command rings")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com>
    Link: https://lore.kernel.org/r/20221228161704.255268-4-manivannan.sadhasivam@linaro.org
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0081032082b5b45ca902b3c3d6986cb5cca69ff2
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 15 09:20:33 2023 +0100

    PCI/DPC: Await readiness of secondary bus after reset
```

```
    commit 53b54ad074de1896f8b021615f65b27f557ce874 upstream.

    pci_bridge_wait_for_secondary_bus() is called after a Secondary Bus
    Reset, but not after a DPC-induced Hot Reset.

    As a result, the delays prescribed by PCIe r6.0 sec 6.6.1 are not
    observed and devices on the secondary bus may be accessed before
    they're ready.

    One affected device is Intel's Ponte Vecchio HPC GPU.  It comprises a
    PCIe switch whose upstream port is not immediately ready after reset.
    Because its config space is restored too early, it remains in
    D0uninitialized, its subordinate devices remain inaccessible and DPC
    recovery fails with messages such as:

      i915 0000:8c:00.0: can't change power state from D3cold to D0 (config space
inaccessible)
      intel_vsec 0000:8e:00.1: can't change power state from D3cold to D0 (config space
inaccessible)
      pcieport 0000:89:02.0: AER: device recovery failed

    Fix it.

    Link:
https://lore.kernel.org/r/9f5ff00e1593d8d9a4b452398b98aa14d23fca11.1673769517.git.lukas@wunn
er.de
    Tested-by: Ravi Kishore Koppuravuri <ravi.kishore.koppuravuri@intel.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 07a966891b8ac07809a4b73b5a4fe867108def5d
Author: Damien Le Moal <damien.lemoal@opensource.wdc.com>
Date:   Sat Jan 28 10:39:51 2023 +0900

    PCI: Avoid FLR for AMD FCH AHCI adapters

    commit 63ba51db24ed1b8f8088a897290eb6c036c5435d upstream.

    PCI passthrough to VMs does not work with AMD FCH AHCI adapters: the guest
    OS fails to correctly probe devices attached to the controller due to FIS
    communication failures:

      ata4: softreset failed (1st FIS failed)
      ...
      ata4.00: qc timeout after 5000 msecs (cmd 0xec)
      ata4.00: failed to IDENTIFY (I/O error, err_mask=0x4)

    Forcing the "bus" reset method before unbinding & binding the adapter to
    the vfio-pci driver solves this issue, e.g.:

      echo "bus" > /sys/bus/pci/devices/<ID>/reset_method

    gives a working guest OS, indicating that the default FLR reset method
    doesn't work correctly.

    Apply quirk_no_flr() to AMD FCH AHCI devices to work around this issue.

    Link: https://lore.kernel.org/r/20230128013951.523247-1-damien.lemoal@opensource.wdc.com
    Reported-by: Niklas Cassel <niklas.cassel@wdc.com>
    Signed-off-by: Damien Le Moal <damien.lemoal@opensource.wdc.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
```

```
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3f1719324e15ab237a76d9f725b59d658275bff1
Author: Lukas Wunner <lukas@wunner.de>
Date:   Fri Jan 20 10:19:02 2023 +0100

    PCI: hotplug: Allow marking devices as disconnected during bind/unbind

    commit 74ff8864cc842be994853095dba6db48e716400a upstream.

    On surprise removal, pciehp_unconfigure_device() and acpiphp's
    trim_stale_devices() call pci_dev_set_disconnected() to mark removed
    devices as permanently offline.  Thereby, the PCI core and drivers know
    to skip device accesses.

    However pci_dev_set_disconnected() takes the device_lock and thus waits for
    a concurrent driver bind or unbind to complete.  As a result, the driver's
    ->probe and ->remove hooks have no chance to learn that the device is gone.

    That doesn't make any sense, so drop the device_lock and instead use atomic
    xchg() and cmpxchg() operations to update the device state.

    As a byproduct, an AB-BA deadlock reported by Anatoli is fixed which occurs
    on surprise removal with AER concurrently performing a bus reset.

    AER bus reset:

      INFO: task irq/26-aerdrv:95 blocked for more than 120 seconds.
      Tainted: G        W          6.2.0-rc3-custom-norework-jan11+
      schedule
      rwsem_down_write_slowpath
      down_write_nested
      pciehp_reset_slot                     # acquires reset_lock
      pci_reset_hotplug_slot
      pci_slot_reset                        # acquires device_lock
      pci_bus_error_reset
      aer_root_reset
      pcie_do_recovery
      aer_process_err_devices
      aer_isr

    pciehp surprise removal:

      INFO: task irq/26-pciehp:96 blocked for more than 120 seconds.
      Tainted: G        W          6.2.0-rc3-custom-norework-jan11+
      schedule_preempt_disabled
      __mutex_lock
      mutex_lock_nested
      pci_dev_set_disconnected              # acquires device_lock
      pci_walk_bus
      pciehp_unconfigure_device
      pciehp_disable_slot
      pciehp_handle_presence_or_link_change
      pciehp_ist                            # acquires reset_lock

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=215590
    Fixes: a6bd101b8f84 ("PCI: Unify device inaccessible")
    Link:
https://lore.kernel.org/r/3dc88ea82bdc0e37d9000e413d5ebce481cbd629.1674205689.git.lukas@wunn
er.de
    Reported-by: Anatoli Antonovitch <anatoli.antonovitch@amd.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Cc: stable@vger.kernel.org # v4.20+
```

```
    Cc: Keith Busch <kbusch@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 145cf271d5a0be12316f48655bdc31a664ea031d
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 15 09:20:32 2023 +0100

    PCI: Unify delay handling for reset and resume

    commit ac91e6980563ed53afadd925fa6585ffd2bc4a2c upstream.

    Sheng Bi reports that pci_bridge_secondary_bus_reset() may fail to wait
    for devices on the secondary bus to become accessible after reset:

    Although it does call pci_dev_wait(), it erroneously passes the bridge's
    pci_dev rather than that of a child.  The bridge of course is always
    accessible while its secondary bus is reset, so pci_dev_wait() returns
    immediately.

    Sheng Bi proposes introducing a new pci_bridge_secondary_bus_wait()
    function which is called from pci_bridge_secondary_bus_reset():

    https://lore.kernel.org/linux-pci/20220523171517.32407-1-windy.bi.enflame@gmail.com/

    However we already have pci_bridge_wait_for_secondary_bus() which does
    almost exactly what we need.  So far it's only called on resume from
    D3cold (which implies a Fundamental Reset per PCIe r6.0 sec 5.8).
    Re-using it for Secondary Bus Resets is a leaner and more rational
    approach than introducing a new function.

    That only requires a few minor tweaks:

    - Amend pci_bridge_wait_for_secondary_bus() to await accessibility of
      the first device on the secondary bus by calling pci_dev_wait() after
      performing the prescribed delays.  pci_dev_wait() needs two parameters,
      a reset reason and a timeout, which callers must now pass to
      pci_bridge_wait_for_secondary_bus().  The timeout is 1 sec for resume
      (PCIe r6.0 sec 6.6.1) and 60 sec for reset (commit 821cdad5c46c ("PCI:
      Wait up to 60 seconds for device to become ready after FLR")).
      Introduce a PCI_RESET_WAIT macro for the 1 sec timeout.

    - Amend pci_bridge_wait_for_secondary_bus() to return 0 on success or
      -ENOTTY on error for consumption by pci_bridge_secondary_bus_reset().

    - Drop an unnecessary 1 sec delay from pci_reset_secondary_bus() which
      is now performed by pci_bridge_wait_for_secondary_bus().  A static
      delay this long is only necessary for Conventional PCI, so modern
      PCIe systems benefit from shorter reset times as a side effect.

    Fixes: 6b2f1351af56 ("PCI: Wait for device to become ready after secondary bus reset")
    Link:
https://lore.kernel.org/r/da77c92796b99ec568bd070cbe4725074a117038.1673769517.git.lukas@wunn
er.de
    Reported-by: Sheng Bi <windy.bi.enflame@gmail.com>
    Tested-by: Ravi Kishore Koppuravuri <ravi.kishore.koppuravuri@intel.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Reviewed-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Cc: stable@vger.kernel.org # v4.17+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1ec8738fe38a094ce0e7a10f7bb318d9694ede7b
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 15 09:20:31 2023 +0100
```

```
    PCI/PM: Observe reset delay irrespective of bridge_d3

    commit 8ef0217227b42e2c34a18de316cee3da16c9bf1e upstream.

    If a PCI bridge is suspended to D3cold upon entering system sleep,
    resuming it entails a Fundamental Reset per PCIe r6.0 sec 5.8.

    The delay prescribed after a Fundamental Reset in PCIe r6.0 sec 6.6.1
    is sought to be observed by:

      pci_pm_resume_noirq()
        pci_pm_bridge_power_up_actions()
          pci_bridge_wait_for_secondary_bus()

    However, pci_bridge_wait_for_secondary_bus() bails out if the bridge_d3
    flag is not set.  That flag indicates whether a bridge is allowed to
    suspend to D3cold at *runtime*.

    Hence *no* delay is observed on resume from system sleep if runtime
    D3cold is forbidden.  That doesn't make any sense, so drop the bridge_d3
    check from pci_bridge_wait_for_secondary_bus().

    The purpose of the bridge_d3 check was probably to avoid delays if a
    bridge remained in D0 during suspend.  However the sole caller of
    pci_bridge_wait_for_secondary_bus(), pci_pm_bridge_power_up_actions(),
    is only invoked if the previous power state was D3cold.  Hence the
    additional bridge_d3 check seems superfluous.

    Fixes: ad9001f2f411 ("PCI/PM: Add missing link delays required by the PCIe spec")
    Link:
https://lore.kernel.org/r/eb37fa345285ec8bacabbf06b020b803f77bdd3d.1673769517.git.lukas@wunn
er.de
    Tested-by: Ravi Kishore Koppuravuri <ravi.kishore.koppuravuri@intel.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Reviewed-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Cc: stable@vger.kernel.org # v5.5+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5475b05b5fe492d8f727f50ea6abed6fe9ca8885
Author: H. Nikolaus Schaller <hns@goldelico.com>
Date:   Sun Jan 29 19:57:04 2023 +0100

    MIPS: DTS: CI20: fix otg power gpio

    commit 0cb4228f6cc9ed0ca2be0d9ddf29168a8e3a3905 upstream.

    According to schematics it is PF15 and not PF14 (MIC_SW_EN).
    Seems as if it was hidden and not noticed during testing since
    there is no sound DT node.

    Fixes: 158c774d3c64 ("MIPS: Ingenic: Add missing nodes for Ingenic SoCs and boards.")
    Cc: stable@vger.kernel.org
    Signed-off-by: H. Nikolaus Schaller <hns@goldelico.com>
    Acked-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8a0fc187e9a2a540482643b38f686ea54ce775f2
Author: Guo Ren <guoren@kernel.org>
Date:   Thu Jan 12 04:05:59 2023 -0500

    riscv: ftrace: Reduce the detour code size to half
```

```
    commit 6724a76cff85ee271bbbff42ac527e4643b2ec52 upstream.

    Use a temporary register to reduce the size of detour code from 16 bytes to
    8 bytes. The previous implementation is from 'commit afc76b8b8011 ("riscv:
    Using PATCHABLE_FUNCTION_ENTRY instead of MCOUNT")'.

    Before the patch:
    <func_prolog>:
     0: REG_S  ra, -SZREG(sp)
     4: auipc  ra, ?
     8: jalr   ?(ra)
    12: REG_L  ra, -SZREG(sp)
     (func_boddy)

    After the patch:
    <func_prolog>:
     0: auipc  t0, ?
     4: jalr   t0, ?(t0)
     (func_boddy)

    This patch not just reduces the size of detour code, but also fixes an
    important issue:

    An Ftrace callback registered with FTRACE_OPS_FL_IPMODIFY flag can
    actually change the instruction pointer, e.g. to "replace" the given
    kernel function with a new one, which is needed for livepatching, etc.

    In this case, the trampoline (ftrace_regs_caller) would not return to
    <func_prolog+12> but would rather jump to the new function. So, "REG_L
    ra, -SZREG(sp)" would not run and the original return address would not
    be restored. The kernel is likely to hang or crash as a result.

    This can be easily demonstrated if one tries to "replace", say,
    cmdline_proc_show() with a new function with the same signature using
    instruction_pointer_set(&fregs->regs, new_func_addr) in the Ftrace
    callback.

    Link: https://lore.kernel.org/linux-riscv/20221122075440.1165172-1-
suagrfillet@gmail.com/
    Link: https://lore.kernel.org/linux-riscv/d7d5730b-ebef-68e5-5046-
e763e1ee6164@yadro.com/
    Co-developed-by: Song Shuai <suagrfillet@gmail.com>
    Signed-off-by: Song Shuai <suagrfillet@gmail.com>
    Signed-off-by: Guo Ren <guoren@linux.alibaba.com>
    Signed-off-by: Guo Ren <guoren@kernel.org>
    Cc: Evgenii Shatokhin <e.shatokhin@yadro.com>
    Reviewed-by: Evgenii Shatokhin <e.shatokhin@yadro.com>
    Link: https://lore.kernel.org/r/20230112090603.1295340-4-guoren@kernel.org
    Cc: stable@vger.kernel.org
    Fixes: 10626c32e382 ("riscv/ftrace: Add basic support")
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 10509799f471bd4655ff30acf4cf544d3df72f7f
Author: Guo Ren <guoren@kernel.org>
Date:   Thu Jan 12 04:05:58 2023 -0500

    riscv: ftrace: Remove wasted nops for !RISCV_ISA_C

    commit 409c8fb20c66df7150e592747412438c04aeb11f upstream.

    When CONFIG_RISCV_ISA_C=n, -fpatchable-function-entry=8 would generate
    more nops than we expect. Because it treat nop opcode as 0x00000013
    instead of 0x0001.
```

```
    Dump of assembler code for function dw_pcie_free_msi:
       0xffffffff806fce94 <+0>:     sd      ra,-8(sp)
       0xffffffff806fce98 <+4>:     auipc   ra,0xff90f
       0xffffffff806fce9c <+8>:     jalr    -684(ra) # 0xffffffff8000bbec
    <ftrace_caller>
       0xffffffff806fcea0 <+12>:    ld      ra,-8(sp)
       0xffffffff806fcea4 <+16>:    nop /* wasted */
       0xffffffff806fcea8 <+20>:    nop /* wasted */
       0xffffffff806fceac <+24>:    nop /* wasted */
       0xffffffff806fceb0 <+28>:    nop /* wasted */
       0xffffffff806fceb4 <+0>:     addi    sp,sp,-48
       0xffffffff806fceb8 <+4>:     sd      s0,32(sp)
       0xffffffff806fcebc <+8>:     sd      s1,24(sp)
       0xffffffff806fcec0 <+12>:    sd      s2,16(sp)
       0xffffffff806fcec4 <+16>:    sd      s3,8(sp)
       0xffffffff806fcec8 <+20>:    sd      ra,40(sp)
       0xffffffff806fcecc <+24>:    addi    s0,sp,48

    Signed-off-by: Guo Ren <guoren@linux.alibaba.com>
    Signed-off-by: Guo Ren <guoren@kernel.org>
    Link: https://lore.kernel.org/r/20230112090603.1295340-3-guoren@kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ed157545d06e97651eedb151a6f5f264a8156013
Author: Björn Töpel <bjorn@rivosinc.com>
Date:   Tue Feb 14 17:25:15 2023 +0100

    riscv, mm: Perform BPF exhandler fixup on page fault

    commit 416721ff05fddc58ca531b6f069de250301de6e5 upstream.

    Commit 21855cac82d3 ("riscv/mm: Prevent kernel module to access user
    memory without uaccess routines") added early exits/deaths for page
    faults stemming from accesses to user-space without using proper
    uaccess routines (where sstatus.SUM is set).

    Unfortunatly, this is too strict for some BPF programs, which relies
    on BPF exhandler fixups. These BPF programs loads "BTF pointers". A
    BTF pointers could either be a valid kernel pointer or NULL, but not a
    userspace address.

    Resolve the problem by calling the fixup handler in the early exit
    path.

    Fixes: 21855cac82d3 ("riscv/mm: Prevent kernel module to access user memory without
uaccess routines")
    Signed-off-by: Björn Töpel <bjorn@rivosinc.com>
    Link: https://lore.kernel.org/r/20230214162515.184827-1-bjorn@kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 20a7510e781084364691b4962de31de758194cc9
Author: Andy Chiu <andy.chiu@sifive.com>
Date:   Thu Jan 12 04:05:57 2023 -0500

    riscv: ftrace: Fixup panic by disabling preemption

    commit 8547649981e6631328cd64f583667501ae385531 upstream.

    In RISCV, we must use an AUIPC + JALR pair to encode an immediate,
    forming a jump that jumps to an address over 4K. This may cause errors
```

```
    if we want to enable kernel preemption and remove dependency from
    patching code with stop_machine(). For example, if a task was switched
    out on auipc. And, if we changed the ftrace function before it was
    switched back, then it would jump to an address that has updated 11:0
    bits mixing with previous XLEN:12 part.

    p: patched area performed by dynamic ftrace
    ftrace_prologue:
    p|      REG_S   ra, -SZREG(sp)
    p|      auipc   ra, 0x? ------------> preempted
                                          ...
                                  change ftrace function
                                          ...
    p|      jalr    -?(ra) <------------- switched back
    p|      REG_L   ra, -SZREG(sp)
    func:
            xxx
            ret

    Fixes: afc76b8b8011 ("riscv: Using PATCHABLE_FUNCTION_ENTRY instead of MCOUNT")
    Signed-off-by: Andy Chiu <andy.chiu@sifive.com>
    Signed-off-by: Guo Ren <guoren@kernel.org>
    Link: https://lore.kernel.org/r/20230112090603.1295340-2-guoren@kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 60312d09c622a5b8a3e9b6cc3811fb9af3111fc1
Author: Andy Chiu <andy.chiu@sifive.com>
Date:   Mon Feb 6 04:04:40 2023 -0500

    riscv: jump_label: Fixup unaligned arch_static_branch function

    commit 9ddfc3cd806081ce1f6c9c2f988cbb031f35d28f upstream.

    Runtime code patching must be done at a naturally aligned address, or we
    may execute on a partial instruction.

    We have encountered problems traced back to static jump functions during
    the test. We switched the tracer randomly for every 1~5 seconds on a
    dual-core QEMU setup and found the kernel sucking at a static branch
    where it jumps to itself.

    The reason is that the static branch was 2-byte but not 4-byte aligned.
    Then, the kernel would patch the instruction, either J or NOP, with two
    half-word stores if the machine does not have efficient unaligned
    accesses. Thus, moments exist where half of the NOP mixes with the other
    half of the J when transitioning the branch. In our particular case, on
    a little-endian machine, the upper half of the NOP was mixed with the
    lower part of the J when enabling the branch, resulting in a jump that
    jumped to itself. Conversely, it would result in a HINT instruction when
    disabling the branch, but it might not be observable.

    ARM64 does not have this problem since all instructions must be 4-byte
    aligned.

    Fixes: ebc00dde8a97 ("riscv: Add jump-label implementation")
    Link: https://lore.kernel.org/linux-riscv/20220913094252.3555240-6-andy.chiu@sifive.com/
    Reviewed-by: Greentime Hu <greentime.hu@sifive.com>
    Signed-off-by: Andy Chiu <andy.chiu@sifive.com>
    Signed-off-by: Guo Ren <guoren@kernel.org>
    Link: https://lore.kernel.org/r/20230206090440.1255001-1-guoren@kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 8ef3d23d4d821632167ef76a2e6b99c2bb9ada89
Author: Sergey Matyukevich <sergey.matyukevich@syntacore.com>
Date:   Mon Jan 30 00:18:18 2023 +0300

    riscv: mm: fix regression due to update_mmu_cache change

    commit b49f700668fff7565b945dce823def79bff59bb0 upstream.

    This is a partial revert of the commit 4bd1d80efb5a ("riscv: mm: notify
    remote harts about mmu cache updates"). Original commit included two
    loosely related changes serving the same purpose of fixing stale TLB
    entries causing user-space application crash:
    - introduce deferred per-ASID TLB flush for CPUs not running the task
    - switch to per-ASID TLB flush on all CPUs running the task in update_mmu_cache

    According to report and discussion in [1], the second part caused a
    regression on Renesas RZ/Five SoC. For now restore the old behavior
    of the update_mmu_cache.

    [1] https://lore.kernel.org/linux-riscv/20220829205219.283543-1-geomatsi@gmail.com/

    Fixes: 4bd1d80efb5a ("riscv: mm: notify remote harts about mmu cache updates")
    Reported-by: "Lad, Prabhakar" <prabhakar.csengg@gmail.com>
    Signed-off-by: Sergey Matyukevich <sergey.matyukevich@syntacore.com>
    Link: trailer, so that it can be parsed with git's trailer functionality?
    Reviewed-by: Conor Dooley <conor.dooley@microchip.com>
    Link: https://lore.kernel.org/r/20230129211818.686557-1-geomatsi@gmail.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 424818a16b7385f5ae8d1575596b25be6b372505
Author: Mattias Nissler <mnissler@rivosinc.com>
Date:   Wed Feb 15 14:48:28 2023 +0000

    riscv: Avoid enabling interrupts in die()

    commit 130aee3fd9981297ff9354e5d5609cd59aafbbea upstream.

    While working on something else, I noticed that the kernel would start
    accepting interrupts again after crashing in an interrupt handler. Since
    the kernel is already in inconsistent state, enabling interrupts is
    dangerous and opens up risk of kernel state deteriorating further.
    Interrupts do get enabled via what looks like an unintended side effect of
    spin_unlock_irq, so switch to the more cautious
    spin_lock_irqsave/spin_unlock_irqrestore instead.

    Fixes: 76d2a0493a17 ("RISC-V: Init and Halt Code")
    Signed-off-by: Mattias Nissler <mnissler@rivosinc.com>
    Reviewed-by: Björn Töpel <bjorn@kernel.org>
    Link: https://lore.kernel.org/r/20230215144828.3370316-1-mnissler@rivosinc.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fec6331a42d38c5a98dffa5ed123bbac0aa9b69f
Author: Conor Dooley <conor.dooley@microchip.com>
Date:   Fri Feb 10 18:59:45 2023 +0000

    RISC-V: add a spin_shadow_stack declaration

    commit eb9be8310c58c166f9fae3b71c0ad9d6741b4897 upstream.

    The patchwork automation reported a sparse complaint that
```

```
    spin_shadow_stack was not declared and should be static:
    ../arch/riscv/kernel/traps.c:335:15: warning: symbol 'spin_shadow_stack' was not
declared. Should it be static?

    However, this is used in entry.S and therefore shouldn't be static.
    The same applies to the shadow_stack that this pseudo spinlock is
    trying to protect, so do like its charge and add a declaration to
    thread_info.h

    Signed-off-by: Conor Dooley <conor.dooley@microchip.com>
    Fixes: 7e1864332fbc ("riscv: fix race when vmap stack overflow")
    Reviewed-by: Guo Ren <guoren@kernel.org>
    Link: https://lore.kernel.org/r/20230210185945.915806-1-conor@kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2fb1fa8425cce2dc4dce298275d22d7077694b73
Author: Tomas Henzl <thenzl@redhat.com>
Date:   Thu Feb 2 17:24:51 2023 +0100

    scsi: ses: Fix slab-out-of-bounds in ses_intf_remove()

    commit 578797f0c8cbc2e3ec5fc0dab87087b4c7073686 upstream.

    A fix for:

    BUG: KASAN: slab-out-of-bounds in ses_intf_remove+0x23f/0x270 [ses]
    Read of size 8 at addr ffff88a10d32e5d8 by task rmmod/12013

    When edev->components is zero, accessing edev->component[0] members is
    wrong.

    Link: https://lore.kernel.org/r/20230202162451.15346-5-thenzl@redhat.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Tomas Henzl <thenzl@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4b8cae410472653a59e15af62c57c49b8e0a1201
Author: Tomas Henzl <thenzl@redhat.com>
Date:   Thu Feb 2 17:24:50 2023 +0100

    scsi: ses: Fix possible desc_ptr out-of-bounds accesses

    commit 801ab13d50cf3d26170ee073ea8bb4eececb76ab upstream.

    Sanitize possible desc_ptr out-of-bounds accesses in
    ses_enclosure_data_process().

    Link: https://lore.kernel.org/r/20230202162451.15346-4-thenzl@redhat.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Tomas Henzl <thenzl@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b91ef85a32fdba45fcbad87dd526d73d3b6d857d
Author: Tomas Henzl <thenzl@redhat.com>
Date:   Thu Feb 2 17:24:49 2023 +0100

    scsi: ses: Fix possible addl_desc_ptr out-of-bounds accesses

    commit db95d4df71cb55506425b6e4a5f8d68e3a765b63 upstream.

    Sanitize possible addl_desc_ptr out-of-bounds accesses in
```

```
    ses_enclosure_data_process().

    Link: https://lore.kernel.org/r/20230202162451.15346-3-thenzl@redhat.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Tomas Henzl <thenzl@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 799e8dd2022d2e13f0c5c1906b40ceca07a23349
Author: Tomas Henzl <thenzl@redhat.com>
Date:   Thu Feb 2 17:24:48 2023 +0100

    scsi: ses: Fix slab-out-of-bounds in ses_enclosure_data_process()

    commit 9b4f5028e493cb353a5c8f5c45073eeea0303abd upstream.

    A fix for:

    BUG: KASAN: slab-out-of-bounds in ses_enclosure_data_process+0x949/0xe30 [ses]
    Read of size 1 at addr ffff88a1b043a451 by task systemd-udevd/3271

    Checking after (and before in next loop) addl_desc_ptr[1] is sufficient, we
    expect the size to be sanitized before first access to addl_desc_ptr[1].
    Make sure we don't walk beyond end of page.

    Link: https://lore.kernel.org/r/20230202162451.15346-2-thenzl@redhat.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Tomas Henzl <thenzl@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f182ad02024d3f45374a9e0c9d76f28b776d762d
Author: James Bottomley <jejb@linux.ibm.com>
Date:   Sat Nov 28 15:27:21 2020 -0800

    scsi: ses: Don't attach if enclosure has no components

    commit 3fe97ff3d94934649abb0652028dd7296170c8d0 upstream.

    An enclosure with no components can't usefully be operated by the driver
    (since effectively it has nothing to manage), so report the problem and
    don't attach. Not attaching also fixes an oops which could occur if the
    driver tries to manage a zero component enclosure.

    [mkp: Switched to KERN_WARNING since this scenario is common]

    Link:
https://lore.kernel.org/r/c5deac044ac409e32d9ad9968ce0dcbc996bfc7a.camel@linux.ibm.com
    Cc: stable@vger.kernel.org
    Reported-by: Ding Hui <dinghui@sangfor.com.cn>
    Signed-off-by: James Bottomley <James.Bottomley@HansenPartnership.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4f5b18fd8505e427138541c2c4237f21f0de8f93
Author: Saurav Kashyap <skashyap@marvell.com>
Date:   Mon Dec 19 03:07:46 2022 -0800

    scsi: qla2xxx: Remove increment of interface err cnt

    commit d676a9e3d9efb7e93df460bcf4c445496c16314f upstream.

    Residual underrun is not an interface error, hence no need to increment
    that count.
```

```
    Fixes: dbf1f53cfd23 ("scsi: qla2xxx: Implementation to get and manage host, target stats
and initiator port")
    Cc: stable@vger.kernel.org
    Signed-off-by: Saurav Kashyap <skashyap@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2cb341ea2e7a35c0047fb64b2b320acbfc8c750e
Author: Quinn Tran <qutran@marvell.com>
Date:   Mon Dec 19 03:07:45 2022 -0800

    scsi: qla2xxx: Fix erroneous link down

    commit 3fbc74feb642deb688cc97f76d40b7287ddd4cb1 upstream.

    If after an adapter reset the appearance of link is not recovered, the
    devices are not rediscovered.  This is result of a race condition between
    adapter reset (abort_isp) and the topology scan.  During adapter reset, the
    ABORT_ISP_ACTIVE flag is set.  Topology scan usually occurred after adapter
    reset.  In this case, the topology scan came earlier than usual where it
    ran into problem due to ABORT_ISP_ACTIVE flag was still set.

    kernel: qla2xxx [0000:13:00.0]-1005:1: Cmd 0x6a aborted with timeout since ISP Abort is
pending
    kernel: qla2xxx [0000:13:00.0]-28a0:1: MBX_GET_PORT_NAME failed, No FL Port.
    kernel: qla2xxx [0000:13:00.0]-286b:1: qla2x00_configure_loop: exiting normally. local
port wwpn 51402ec0123d9a80 id 012300)
    kernel: qla2xxx [0000:13:00.0]-8017:1: ADAPTER RESET SUCCEEDED nexus=1:0:15.

    Allow adapter reset to complete before any scan can start.

    Cc: stable@vger.kernel.org
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e5000a69f4362e9130c9b151c8b49e860f121379
Author: Quinn Tran <qutran@marvell.com>
Date:   Mon Dec 19 03:07:44 2022 -0800

    scsi: qla2xxx: Remove unintended flag clearing

    commit 7e8a936a2d0f98dd6e5d05d4838affabe606cabc upstream.

    FCF_ASYNC_SENT flag is used in session management. This flag is cleared in
    task management path by accident.  Remove unintended flag clearing.

    Fixes: 388a49959ee4 ("scsi: qla2xxx: Fix panic from use after free in
qla2x00_async_tm_cmd")
    Cc: stable@vger.kernel.org
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ad6af23593594402c826eefdf43ae174e5f0f202
Author: Arun Easi <aeasi@marvell.com>
Date:   Mon Dec 19 03:07:40 2022 -0800

    scsi: qla2xxx: Fix DMA-API call trace on NVMe LS requests
```

```
    commit c75e6aef5039830cce5d4cf764dd204522f89e6b upstream.

    The following message and call trace was seen with debug kernels:

    DMA-API: qla2xxx 0000:41:00.0: device driver failed to check map
    error [device address=0x00000002a3ff38d8] [size=1024 bytes] [mapped as
    single]
    WARNING: CPU: 0 PID: 2930 at kernel/dma/debug.c:1017
             check_unmap+0xf42/0x1990

    Call Trace:
            debug_dma_unmap_page+0xc9/0x100
            qla_nvme_ls_unmap+0x141/0x210 [qla2xxx]

    Remove DMA mapping from the driver altogether, as it is already done by FC
    layer. This prevents the warning.

    Fixes: c85ab7d9e27a ("scsi: qla2xxx: Fix missed DMA unmap for NVMe ls requests")
    Cc: stable@vger.kernel.org
    Signed-off-by: Arun Easi <aeasi@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d75af67e3ecba4b79469583fd4e6f659d9398e04
Author: Shreyas Deodhar <sdeodhar@marvell.com>
Date:   Mon Dec 19 03:07:38 2022 -0800

    scsi: qla2xxx: Check if port is online before sending ELS

    commit 0c227dc22ca18856055983f27594feb2e0149965 upstream.

    CT Ping and ELS cmds fail for NVMe targets.  Check if port is online before
    sending ELS instead of sending login.

    Cc: stable@vger.kernel.org
    Signed-off-by: Shreyas Deodhar <sdeodhar@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec66d4f04245eec347e81e4230e330b73452de86
Author: Quinn Tran <qutran@marvell.com>
Date:   Mon Dec 19 03:07:39 2022 -0800

    scsi: qla2xxx: Fix link failure in NPIV environment

    commit b1ae65c082f74536ec292b15766f2846f0238373 upstream.

    User experienced symptoms of adapter failure in NPIV environment. NPIV
    hosts were allowed to trigger chip reset back to back due to NPIV link
    state being slow to come online.

    Fix link failure in NPIV environment by removing NPIV host from directly
    being able to perform chip reset.

     kernel: qla2xxx [0000:04:00.1]-6009:261: Loop down - aborting ISP.
     kernel: qla2xxx [0000:04:00.1]-6009:262: Loop down - aborting ISP.
     kernel: qla2xxx [0000:04:00.1]-6009:281: Loop down - aborting ISP.
     kernel: qla2xxx [0000:04:00.1]-6009:285: Loop down - aborting ISP

    Fixes: 0d6e61bc6a4f ("[SCSI] qla2xxx: Correct various NPIV issues.")
```

```
    Cc: stable@vger.kernel.org
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 99decfb04b854716498db8af806d4778f90091f1
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Jan 27 16:04:13 2023 -0800

    scsi: aacraid: Allocate cmd_priv with scsicmd

    commit 7ab734fc759828707dae22fe48b1eb4dcf70beea upstream.

    The aac_priv() helper assumes that the private cmd area immediately follows
    struct scsi_cmnd. Allocate this space as part of scsicmd, else there is a
    risk of heap overflow. Seen with GCC 13:

    ../drivers/scsi/aacraid/aachba.c: In function 'aac_probe_container':
    ../drivers/scsi/aacraid/aachba.c:841:26: warning: array subscript 16 is outside array
bounds of 'void[392]' [-Warray-bounds=]
      841 |         status = cmd_priv->status;
          |                          ^~
    In file included from ../include/linux/resource_ext.h:11,
                     from ../include/linux/pci.h:40,
                     from ../drivers/scsi/aacraid/aachba.c:22:
    In function 'kmalloc',
        inlined from 'kzalloc' at ../include/linux/slab.h:720:9,
        inlined from 'aac_probe_container' at ../drivers/scsi/aacraid/aachba.c:821:30:
    ../include/linux/slab.h:580:24: note: at offset 392 into object of size 392 allocated by
'kmalloc_trace'
      580 |                 return kmalloc_trace(
          |                        ^~~~~~~~~~~~~~
      581 |                                 kmalloc_caches[kmalloc_type(flags)][index],
          |                                 ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
      582 |                                 flags, size);
          |                                 ~~~~~~~~~~~~

    Fixes: 76a3451b64c6 ("scsi: aacraid: Move the SCSI pointer to private command data")
    Link: https://lore.kernel.org/r/20230128000409.never.976-kees@kernel.org
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: Hannes Reinecke <hare@suse.de>
    Cc: Himanshu Madhani <himanshu.madhani@oracle.com>
    Cc: Adaptec OEM Raid Solutions <aacraid@microsemi.com>
    Cc: "James E.J. Bottomley" <jejb@linux.ibm.com>
    Cc: "Martin K. Petersen" <martin.petersen@oracle.com>
    Cc: linux-scsi@vger.kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Vegard Nossum <vegard.nossum@oracle.com>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 63cd11165e5e0ea2012254c764003eda1f9adb7d
Author: Gavrilov Ilia <Ilia.Gavrilov@infotecs.ru>
Date:   Thu Feb 2 08:26:56 2023 +0000

    iommu/amd: Add a length limitation for the ivrs_acpihid command-line parameter

    commit b6b26d86c61c441144c72f842f7469bb686e1211 upstream.

    The 'acpiid' buffer in the parse_ivrs_acpihid function may overflow,
```

```
    because the string specifier in the format string sscanf()
    has no width limitation.

    Found by InfoTeCS on behalf of Linux Verification Center
    (linuxtesting.org) with SVACE.

    Fixes: ca3bf5d47cec ("iommu/amd: Introduces ivrs_acpihid kernel parameter")
    Cc: stable@vger.kernel.org
    Signed-off-by: Ilia.Gavrilov <Ilia.Gavrilov@infotecs.ru>
    Reviewed-by: Kim Phillips <kim.phillips@amd.com>
    Link: https://lore.kernel.org/r/20230202082719.1513849-1-Ilia.Gavrilov@infotecs.ru
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 43520b214e37bd5a0c67b941160c79851b4e5972
Author: Masami Hiramatsu (Google) <mhiramat@kernel.org>
Date:   Tue Feb 21 08:49:16 2023 +0900

    tracing/eprobe: Fix to add filter on eprobe description in README file

    commit 133921530c42960c07d25d12677f9e131a2b0cdf upstream.

    Fix to add a description of the filter on eprobe in README file. This
    is required to identify the kernel supports the filter on eprobe or not.

    Link:
https://lore.kernel.org/all/167309833728.640500.12232259238201433587.stgit@devnote3/

    Fixes: 752be5c5c910 ("tracing/eprobe: Add eprobe filter support")
    Cc: stable@vger.kernel.org
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Reviewed-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb79a6a801aa68b80f48b7e3e5837d921212e240
Author: Antonio Alvarez Feijoo <antonio.feijoo@suse.com>
Date:   Wed Feb 22 08:27:35 2023 +0900

    tools/bootconfig: fix single & used for logical condition

    commit cf8c59a3756b2735c409a9b3ac1e4ec556546a7a upstream.

    A single & will create a background process and return true, so the grep
    command will run even if the file checked in the first condition does not
    exist.

    Link: https://lore.kernel.org/all/20230112114215.17103-1-antonio.feijoo@suse.com/

    Fixes: 1eaad3ac3f39 ("tools/bootconfig: Use per-group/all enable option in ftrace2bconf
script")
    Signed-off-by: Antonio Alvarez Feijoo <antonio.feijoo@suse.com>
    Cc: stable@vger.kernel.org
    Acked-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 09b1bf25f7f7a8f2bf8cd4278bba9c3172db8013
Author: Mukesh Ojha <quic_mojha@quicinc.com>
Date:   Tue Feb 14 17:36:43 2023 +0530

    ring-buffer: Handle race between rb_move_tail and rb_check_pages

    commit 8843e06f67b14f71c044bf6267b2387784c7e198 upstream.

    It seems a data race between ring_buffer writing and integrity check.
```

That is, RB_FLAG of head_page is been updating, while at same time
RB_FLAG was cleared when doing integrity check rb_check_pages():

```
  rb_check_pages()            rb_handle_head_page():
  --------                    --------
  rb_head_page_deactivate()
                              rb_head_page_set_normal()
  rb_head_page_activate()
```

We do intergrity test of the list to check if the list is corrupted and
it is still worth doing it. So, let's refactor rb_check_pages() such that
we no longer clear and set flag during the list sanity checking.

[1] and [2] are the test to reproduce and the crash report respectively.

1:
``` read_trace.sh
  while true;
  do
    # the "trace" file is closed after read
    head -1 /sys/kernel/tracing/trace > /dev/null
  done
```
``` repro.sh
  sysctl -w kernel.panic_on_warn=1
  # function tracer will writing enough data into ring_buffer
  echo function > /sys/kernel/tracing/current_tracer
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
  ./read_trace.sh &
```

2:
```
------------[ cut here ]------------
WARNING: CPU: 9 PID: 62 at kernel/trace/ring_buffer.c:2653
rb_move_tail+0x450/0x470
Modules linked in:
CPU: 9 PID: 62 Comm: ksoftirqd/9 Tainted: G        W          6.2.0-rc6+
Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS
rel-1.15.0-0-g2dd4b9b3f840-prebuilt.qemu.org 04/01/2014
RIP: 0010:rb_move_tail+0x450/0x470
Code: ff ff 4c 89 c8 f0 4d 0f b1 02 48 89 c2 48 83 e2 fc 49 39 d0 75 24
83 e0 03 83 f8 02 0f 84 e1 fb ff ff 48 8b 57 10 f0 ff 42 08 <0f> 0b 83
f8 02 0f 84 ce fb ff ff e9 db
RSP: 0018:ffffb5564089bd00 EFLAGS: 00000203
RAX: 0000000000000000 RBX: ffff9db385a2bf81 RCX: ffffb5564089bd18
RDX: ffff9db281110100 RSI: 0000000000000fe4 RDI: ffff9db380145400
RBP: ffff9db385a2bf80 R08: ffff9db385a2bfc0 R09: ffff9db385a2bfc2
R10: ffff9db385a6c000 R11: ffff9db385a2bf80 R12: 0000000000000000
R13: 00000000000003e8 R14: ffff9db281110100 R15: ffffffffbb006108
FS:  0000000000000000(0000) GS:ffff9db3bdcc0000(0000)
knlGS:0000000000000000
CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
CR2: 00005602323024c8 CR3: 0000000022e0c000 CR4: 00000000000006e0
Call Trace:
 <TASK>
 ring_buffer_lock_reserve+0x136/0x360
 ? __do_softirq+0x287/0x2df
 ? __pfx_rcu_softirq_qs+0x10/0x10
 trace_function+0x21/0x110
```

```
      ? __pfx_rcu_softirq_qs+0x10/0x10
      ? __do_softirq+0x287/0x2df
      function_trace_call+0xf6/0x120
      0xffffffffc038f097
      ? rcu_softirq_qs+0x5/0x140
      rcu_softirq_qs+0x5/0x140
      __do_softirq+0x287/0x2df
      run_ksoftirqd+0x2a/0x30
      smpboot_thread_fn+0x188/0x220
      ? __pfx_smpboot_thread_fn+0x10/0x10
      kthread+0xe7/0x110
      ? __pfx_kthread+0x10/0x10
      ret_from_fork+0x2c/0x50
      </TASK>
     ---[ end trace 0000000000000000 ]---
```

[crash report and test reproducer credit goes to Zheng Yejian]

Link: https://lore.kernel.org/linux-trace-kernel/1676376403-16462-1-git-send-email-quic_mojha@quicinc.com

Cc: <mhiramat@kernel.org>
Cc: stable@vger.kernel.org
Fixes: 1039221cc278 ("ring-buffer: Do not disable recording when there is an iterator")
Reported-by: Zheng Yejian <zhengyejian1@huawei.com>
Signed-off-by: Mukesh Ojha <quic_mojha@quicinc.com>
Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec9a6d473f979174f1ea923dd496046e1b73376a
Author: Tong Tiangen <tongtiangen@huawei.com>
Date: Sun Jan 29 04:06:51 2023 +0000

memory tier: release the new_memtier in find_create_memory_tier()

commit 93419139fa14124c1c507d804f2b28866ebee28d upstream.

In find_create_memory_tier(), if failed to register device, then we should release new_memtier from the tier list and put device instead of memtier.

Link: https://lkml.kernel.org/r/20230129040651.1329208-1-tongtiangen@huawei.com
Fixes: 9832fb87834e ("mm/demotion: expose memory tier details via sysfs")
Signed-off-by: Tong Tiangen <tongtiangen@huawei.com>
Cc: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
Cc: Hanjun Guo <guohanjun@huawei.com>
Cc: Kefeng Wang <wangkefeng.wang@huawei.com>
Cc: Guohanjun <guohanjun@huawei.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 72667d71c17d981aeb1298ba43d5a8dda9ca35bd
Author: Steven Rostedt <rostedt@goodmis.org>
Date: Wed Jan 18 16:37:25 2023 -0500

ktest.pl: Add RUN_TIMEOUT option with default unlimited

commit 4e7d2a8f0b52abf23b1dc13b3d88bc0923383cd5 upstream.

There is a disconnect between the run_command function and the wait_for_input. The wait_for_input has a default timeout of 2 minutes. But if that happens, the run_command loop will exit out to the waitpid() of the executing command. This fails in that it no longer monitors the command, and also, the ssh to the test box can hang when its finished, as it's waiting for the pipe it's writing to to flush, but the loop that

```
    reads that pipe has already exited, leaving the command stuck, and the
    test hangs.

    Instead, make the default "wait_for_input" of the run_command infinite,
    and allow the user to override it if they want with a default timeout
    option "RUN_TIMEOUT".

    But this fixes the hang that happens when the pipe is full and the ssh
    session never exits.

    Cc: stable@vger.kernel.org
    Fixes: 6e98d1b4415fe ("ktest: Add timeout to ssh command")
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f7838a04f65f28a58c8763df494eda1278344b09
Author: Steven Rostedt <rostedt@goodmis.org>
Date:   Wed Jan 18 11:31:25 2023 -0500

    ktest.pl: Fix missing "end_monitor" when machine check fails

    commit e8bf9b98d40dbdf4e39362e3b85a70c61da68cb7 upstream.

    In the "reboot" command, it does a check of the machine to see if it is
    still alive with a simple "ssh echo" command. If it fails, it will assume
    that a normal "ssh reboot" is not possible and force a power cycle.

    In this case, the "start_monitor" is executed, but the "end_monitor" is
    not, and this causes the screen will not be given back to the console. That
    is, after the test, a "reset" command needs to be performed, as "echo" is
    turned off.

    Cc: stable@vger.kernel.org
    Fixes: 6474ace999edd ("ktest.pl: Powercycle the box on reboot if no connection can be
made")
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6b51f0f3c1aa1fd13325e2a01b0d746d2e17b68d
Author: Masami Hiramatsu (Google) <mhiramat@kernel.org>
Date:   Tue Feb 21 08:49:16 2023 +0900

    kprobes: Fix to handle forcibly unoptimized kprobes on freeing_list

    commit 4fbd2f83fda0ca44a2ec6421ca3508b355b31858 upstream.

    Since forcibly unoptimized kprobes will be put on the freeing_list directly
    in the unoptimize_kprobe(), do_unoptimize_kprobes() must continue to check
    the freeing_list even if unoptimizing_list is empty.

    This bug can happen if a kprobe is put in an instruction which is in the
    middle of the jump-replaced instruction sequence of an optprobe, *and* the
    optprobe is recently unregistered and queued on unoptimizing_list.
    In this case, the optprobe will be unoptimized forcibly (means immediately)
    and put it into the freeing_list, expecting the optprobe will be handled in
    do_unoptimize_kprobe().
    But if there is no other optprobes on the unoptimizing_list, current code
    returns from the do_unoptimize_kprobe() soon and does not handle the
    optprobe which is on the freeing_list. Then the optprobe will hit the
    WARN_ON_ONCE() in the do_free_cleaned_kprobes(), because it is not handled
    in the latter loop of the do_unoptimize_kprobe().

    To solve this issue, do not return from do_unoptimize_kprobes() immediately
    even if unoptimizing_list is empty.
```

```
    Moreover, this change affects another case. kill_optimized_kprobes() expects
    kprobe_optimizer() will just free the optprobe on freeing_list.
    So I changed it to just do list_move() to freeing_list if optprobes are on
    unoptimizing list. And the do_unoptimize_kprobe() will skip
    arch_disarm_kprobe() if the probe on freeing_list has gone flag.

    Link: https://lore.kernel.org/all/Y8URdIfVr3pq2X8w@xpf.sh.intel.com/
    Link:
https://lore.kernel.org/all/167448024501.3253718.13037333683110512967.stgit@devnote3/

    Fixes: e4add247789e ("kprobes: Fix optimize_kprobe()/unoptimize_kprobe() cancellation
logic")
    Reported-by: Pengfei Xu <pengfei.xu@intel.com>
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Cc: stable@vger.kernel.org
    Acked-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0b3657bb31979a37ad0e0a0b23b22d750834f899
Author: Steven Rostedt <rostedt@goodmis.org>
Date:   Wed Jan 18 16:32:13 2023 -0500

    ktest.pl: Give back console on Ctrt^C on monitor

    commit 83d29d439cd3ef23041570d55841f814af2ecac0 upstream.

    When monitoring the console output, the stdout is being redirected to do
    so. If Ctrl^C is hit during this mode, the stdout is not back to the
    console, the user does not see anything they type (no echo).

    Add "end_monitor" to the SIGINT interrupt handler to give back the console
    on Ctrl^C.

    Cc: stable@vger.kernel.org
    Fixes: 9f2cdcbbb90e7 ("ktest: Give console process a dedicated tty")
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5a2a2bec925f8e303657eae5c34f65b48ede4218
Author: Yin Fengwei <fengwei.yin@intel.com>
Date:   Fri Dec 23 21:52:07 2022 +0800

    mm/thp: check and bail out if page in deferred queue already

    commit 81e506bec9be1eceaf5a2c654e28ba5176ef48d8 upstream.

    Kernel build regression with LLVM was reported here:
    https://lore.kernel.org/all/Y1GCYXGtEVZbcv%2F5@dev-arch.thelio-3990X/ with
    commit f35b5d7d676e ("mm: align larger anonymous mappings on THP
    boundaries").  And the commit f35b5d7d676e was reverted.

    It turned out the regression is related with madvise(MADV_DONTNEED)
    was used by ld.lld. But with none PMD_SIZE aligned parameter len.
    trace-bpfcc captured:
    531607  531732  ld.lld          do_madvise.part.0 start: 0x7feca9000000, len: 0x7fb000,
behavior: 0x4
    531607  531793  ld.lld          do_madvise.part.0 start: 0x7fec86a00000, len: 0x7fb000,
behavior: 0x4

    If the underneath physical page is THP, the madvise(MADV_DONTNEED) can
    trigger split_queue_lock contention raised significantly. perf showed
    following data:
        14.85%     0.00%  ld.lld           [kernel.kallsyms]           [k]
           entry_SYSCALL_64_after_hwframe
               11.52%
```

```
                    entry_SYSCALL_64_after_hwframe
                    do_syscall_64
                    __x64_sys_madvise
                    do_madvise.part.0
                    zap_page_range
                    unmap_single_vma
                    unmap_page_range
                    page_remove_rmap
                    deferred_split_huge_page
                    __lock_text_start
                    native_queued_spin_lock_slowpath

    If THP can't be removed from rmap as whole THP, partial THP will be
    removed from rmap by removing sub-pages from rmap.  Even the THP head page
    is added to deferred queue already, the split_queue_lock will be acquired
    and check whether the THP head page is in the queue already.  Thus, the
    contention of split_queue_lock is raised.

    Before acquire split_queue_lock, check and bail out early if the THP
    head page is in the queue already. The checking without holding
    split_queue_lock could race with deferred_split_scan, but it doesn't
    impact the correctness here.

    Test result of building kernel with ld.lld:
    commit 7b5a0b664ebe (parent commit of f35b5d7d676e):
    time -f "\t%E real,\t%U user,\t%S sys" make LD=ld.lld -skj96 allmodconfig all
            6:07.99 real,   26367.77 user,  5063.35 sys

    commit f35b5d7d676e:
    time -f "\t%E real,\t%U user,\t%S sys" make LD=ld.lld -skj96 allmodconfig all
            7:22.15 real,   26235.03 user,  12504.55 sys

    commit f35b5d7d676e with the fixing patch:
    time -f "\t%E real,\t%U user,\t%S sys" make LD=ld.lld -skj96 allmodconfig all
            6:08.49 real,   26520.15 user,  5047.91 sys

    Link: https://lkml.kernel.org/r/20221223135207.2275317-1-fengwei.yin@intel.com
    Signed-off-by: Yin Fengwei <fengwei.yin@intel.com>
    Tested-by: Nathan Chancellor <nathan@kernel.org>
    Acked-by: David Rientjes <rientjes@google.com>
    Reviewed-by: "Huang, Ying" <ying.huang@intel.com>
    Cc: Feng Tang <feng.tang@intel.com>
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: Rik van Riel <riel@surriel.com>
    Cc: Xing Zhengjun <zhengjun.xing@linux.intel.com>
    Cc: Yang Shi <shy828301@gmail.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 706377987d0a0b349e9bf433ffbc746f969a0a08
Author: Johannes Weiner <hannes@cmpxchg.org>
Date:   Wed Dec 7 14:00:39 2022 +0100

    mm: memcontrol: deprecate charge moving

    commit da34a8484d162585e22ed8c1e4114aa2f60e3567 upstream.

    Charge moving mode in cgroup1 allows memory to follow tasks as they
    migrate between cgroups.  This is, and always has been, a questionable
    thing to do - for several reasons.

    First, it's expensive.  Pages need to be identified, locked and isolated
    from various MM operations, and reassigned, one by one.
```

```
    Second, it's unreliable.  Once pages are charged to a cgroup, there isn't
    always a clear owner task anymore.  Cache isn't moved at all, for example.
    Mapped memory is moved - but if trylocking or isolating a page fails,
    it's arbitrarily left behind.  Frequent moving between domains may leave a
    task's memory scattered all over the place.

    Third, it isn't really needed.  Launcher tasks can kick off workload tasks
    directly in their target cgroup.  Using dedicated per-workload groups
    allows fine-grained policy adjustments - no need to move tasks and their
    physical pages between control domains.  The feature was never
    forward-ported to cgroup2, and it hasn't been missed.

    Despite it being a niche usecase, the maintenance overhead of supporting
    it is enormous.  Because pages are moved while they are live and subject
    to various MM operations, the synchronization rules are complicated.
    There are lock_page_memcg() in MM and FS code, which non-cgroup people
    don't understand.  In some cases we've been able to shift code and cgroup
    API calls around such that we can rely on native locking as much as
    possible.  But that's fragile, and sometimes we need to hold MM locks for
    longer than we otherwise would (pte lock e.g.).

    Mark the feature deprecated. Hopefully we can remove it soon.

    And backport into -stable kernels so that people who develop against
    earlier kernels are warned about this deprecation as early as possible.

    [akpm@linux-foundation.org: fix memory.rst underlining]
    Link: https://lkml.kernel.org/r/Y5COd+qXwk/S+n8N@cmpxchg.org
    Signed-off-by: Johannes Weiner <hannes@cmpxchg.org>
    Acked-by: Shakeel Butt <shakeelb@google.com>
    Acked-by: Hugh Dickins <hughd@google.com>
    Acked-by: Michal Hocko <mhocko@suse.com>
    Cc: Muchun Song <songmuchun@bytedance.com>
    Cc: Roman Gushchin <roman.gushchin@linux.dev>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dd4f59bb3b26746d583ee62137d9acf361c9f79b
Author: John Ogness <john.ogness@linutronix.de>
Date:   Thu Dec 29 14:49:39 2022 +0106

    docs: gdbmacros: print newest record

    commit f2e4cca2f670c8e52fbb551a295f2afc9aa2bd72 upstream.

    @head_id points to the newest record, but the printing loop
    exits when it increments to this value (before printing).

    Exit the printing loop after the newest record has been printed.

    The python-based function in scripts/gdb/linux/dmesg.py already
    does this correctly.

    Fixes: e60768311af8 ("scripts/gdb: update for lockless printk ringbuffer")
    Cc: stable@vger.kernel.org
    Signed-off-by: John Ogness <john.ogness@linutronix.de>
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Signed-off-by: Petr Mladek <pmladek@suse.com>
    Link: https://lore.kernel.org/r/20221229134339.197627-1-john.ogness@linutronix.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8f24eef598ce7cce0bbefe0ec642bcc031d0f528
Author: Yan Zhao <yan.y.zhao@intel.com>
Date:   Wed Feb 22 15:49:38 2023 +0800
```

```
    vfio: Fix NULL pointer dereference caused by uninitialized group->iommufd

    commit d649c34cb916b015fdcb487e51409fcc5caeca8d upstream.

    group->iommufd is not initialized for the iommufd_ctx_put()

    [20018.331541] BUG: kernel NULL pointer dereference, address: 0000000000000000
    [20018.377508] RIP: 0010:iommufd_ctx_put+0x5/0x10 [iommufd]
    ...
    [20018.476483] Call Trace:
    [20018.479214]  <TASK>
    [20018.481555]  vfio_group_fops_unl_ioctl+0x506/0x690 [vfio]
    [20018.487586]  __x64_sys_ioctl+0x6a/0xb0
    [20018.491773]  ? trace_hardirqs_on+0xc5/0xe0
    [20018.496347]  do_syscall_64+0x67/0x90
    [20018.500340]  entry_SYSCALL_64_after_hwframe+0x4b/0xb5

    Fixes: 9eefba8002c2 ("vfio: Move vfio group specific code into group.c")
    Cc: stable@vger.kernel.org
    Signed-off-by: Yan Zhao <yan.y.zhao@intel.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Reviewed-by: Yi Liu <yi.l.liu@intel.com>
    Link: https://lore.kernel.org/r/20230222074938.13681-1-yan.y.zhao@intel.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3e64bb920ae53665a8d2e44bedbb0680e6d8c692
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Wed Jan 4 16:31:10 2023 +0800

    remoteproc/mtk_scp: Move clk ops outside send_lock

    commit e46ceea3148163166ef9b7bcac578e72dd30c064 upstream.

    Clocks are properly reference counted and do not need to be inside the
    lock range.

    Right now this triggers a false-positive lockdep warning on MT8192 based
    Chromebooks, through a combination of mtk-scp that has a cros-ec-rpmsg
    sub-device, the (actual) cros-ec I2C adapter registration, I2C client
    (not on cros-ec) probe doing i2c transfers and enabling clocks.

    This is a false positive because the cros-ec-rpmsg under mtk-scp does
    not have an I2C adapter, and also each I2C adapter and cros-ec instance
    have their own mutex.

    Move the clk operations outside of the send_lock range.

    Fixes: 63c13d61eafe ("remoteproc/mediatek: add SCP support for mt8183")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230104083110.736377-1-wenst@chromium.org
    [Fixed "Fixes:" tag line]
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2a65f2293e658004a0126587b13861eb4079e648
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Wed Dec 21 09:30:11 2022 +0100

    media: ipu3-cio2: Fix PM runtime usage_count in driver unbind

    commit 909d3096ac99fa2289f9b8945a3eab2269947a0a upstream.
```

```
    Get the PM runtime usage_count and forbid PM runtime at driver unbind. The
    opposite is being done in probe() already.

    Fixes: commit c2a6a07afe4a ("media: intel-ipu3: cio2: add new MIPI-CSI2 driver")
    Cc: stable@vger.kernel.org # for >= 4.16
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Reviewed-by: Bingbu Cao <bingbu.cao@intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8263e14a99902c23af3559352efa08fb021167d2
Author: Elvira Khabirova <lineprinter0@gmail.com>
Date:   Sat Feb 18 23:43:59 2023 +0100

    mips: fix syscall_get_nr

    commit 85cc91e2ba4262a602ec65e2b76c4391a9e60d3d upstream.

    The implementation of syscall_get_nr on mips used to ignore the task
    argument and return the syscall number of the calling thread instead of
    the target thread.

    The bug was exposed to user space by commit 201766a20e30f ("ptrace: add
    PTRACE_GET_SYSCALL_INFO request") and detected by strace test suite.

    Link: https://github.com/strace/strace/issues/235
    Fixes: c2d9f1775731 ("MIPS: Fix syscall_get_nr for the syscall exit tracing.")
    Cc: <stable@vger.kernel.org> # v3.19+
    Co-developed-by: Dmitry V. Levin <ldv@strace.io>
    Signed-off-by: Dmitry V. Levin <ldv@strace.io>
    Signed-off-by: Elvira Khabirova <lineprinter0@gmail.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e2f26c695e4e9fbc9c6c6dce97d3f92d3492b3aa
Author: Dan Williams <dan.j.williams@intel.com>
Date:   Thu Feb 16 00:36:02 2023 -0800

    dax/kmem: Fix leak of memory-hotplug resources

    commit e686c32590f40bffc45f105c04c836ffad3e531a upstream.

    While experimenting with CXL region removal the following corruption of
    /proc/iomem appeared.

    Before:
    f010000000-f04fffffff : CXL Window 0
      f010000000-f02fffffff : region4
        f010000000-f02fffffff : dax4.0
          f010000000-f02fffffff : System RAM (kmem)

    After (modprobe -r cxl_test):
    f010000000-f02fffffff : **redacted binary garbage**
      f010000000-f02fffffff : System RAM (kmem)

    ...and testing further the same is visible with persistent memory
    assigned to kmem:

    Before:
    480000000-243fffffff : Persistent Memory
      480000000-57e1fffff : namespace3.0
      580000000-243fffffff : dax3.0
        580000000-243fffffff : System RAM (kmem)
```

```
    After (ndctl disable-region all):
    480000000-243fffffff : Persistent Memory
      580000000-243fffffff : ***redacted binary garbage***
        580000000-243fffffff : System RAM (kmem)

    The corrupted data is from a use-after-free of the "dax4.0" and "dax3.0"
    resources, and it also shows that the "System RAM (kmem)" resource is
    not being removed. The bug does not appear after "modprobe -r kmem", it
    requires the parent of "dax4.0" and "dax3.0" to be removed which
    re-parents the leaked "System RAM (kmem)" instances. Those in turn
    reference the freed resource as a parent.

    First up for the fix is release_mem_region_adjustable() needs to
    reliably delete the resource inserted by add_memory_driver_managed().
    That is thwarted by a check for IORESOURCE_SYSRAM that predates the
    dax/kmem driver, from commit:

    65c78784135f ("kernel, resource: check for IORESOURCE_SYSRAM in
release_mem_region_adjustable")

    That appears to be working around the behavior of HMM's
    "MEMORY_DEVICE_PUBLIC" facility that has since been deleted. With that
    check removed the "System RAM (kmem)" resource gets removed, but
    corruption still occurs occasionally because the "dax" resource is not
    reliably removed.

    The dax range information is freed before the device is unregistered, so
    the driver can not reliably recall (another use after free) what it is
    meant to release. Lastly if that use after free got lucky, the driver
    was covering up the leak of "System RAM (kmem)" due to its use of
    release_resource() which detaches, but does not free, child resources.
    The switch to remove_resource() forces remove_memory() to be responsible
    for the deletion of the resource added by add_memory_driver_managed().

    Fixes: c2f3011ee697 ("device-dax: add an allocation interface for device-dax instances")
    Cc: <stable@vger.kernel.org>
    Cc: Oscar Salvador <osalvador@suse.de>
    Cc: David Hildenbrand <david@redhat.com>
    Cc: Pavel Tatashin <pasha.tatashin@soleen.com>
    Reviewed-by: Vishal Verma <vishal.l.verma@intel.com>
    Reviewed-by: Pasha Tatashin <pasha.tatashin@soleen.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/167653656244.3147810.5705900882794040229.stgit@dwillia2-
xfh.jf.intel.com
    Signed-off-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 16850c9cec70b07ff9670cb5da7b0734c2e7fd69
Author: Al Viro <viro@zeniv.linux.org.uk>
Date:   Fri Jan 6 19:25:59 2023 -0500

    alpha: fix FEN fault handling

    commit 977a3009547dad4a5bc95d91be4a58c9f7eedac0 upstream.

    Type 3 instruction fault (FPU insn with FPU disabled) is handled
    by quietly enabling FPU and returning.  Which is fine, except that
    we need to do that both for fault in userland and in the kernel;
    the latter *can* legitimately happen - all it takes is this:

    .global _start
    _start:
            call_pal 0xae
            lda $0, 0
            ldq $0, 0($0)
```

```
    - call_pal CLRFEN to clear "FPU enabled" flag and arrange for
    a signal delivery (SIGSEGV in this case).

    Fixed by moving the handling of type 3 into the common part of
    do_entIF(), before we check for kernel vs. user mode.

    Incidentally, the check for kernel mode is unidiomatic; the normal
    way to do that is !user_mode(regs).  The difference is that
    the open-coded variant treats any of bits 63..3 of regs->ps being
    set as "it's user mode" while the normal approach is to check just
    the bit 3.  PS is a 4-bit register and regs->ps always will have
    bits 63..4 clear, so the open-coded variant here is actually equivalent
    to !user_mode(regs).  Harder to follow, though...

    Cc: stable@vger.kernel.org
    Reviewed-by: Richard Henderson <rth@twiddle.net>
    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 034245d771e392728c7e6d23efa85587bcacc0dd
Author: Dhruva Gole <d-gole@ti.com>
Date:   Thu Feb 23 15:22:02 2023 +0530

    spi: spi-sn-f-ospi: fix duplicate flag while assigning to mode_bits

    commit 078a5517d22342eb0474046d3e891427a2552e3c upstream.

    Replace the SPI_TX_OCTAL flag that appeared two time with SPI_RX_OCTAL
    in the chain of '|' operators while assigning to mode_bits

    Fixes: 1b74dd64c861 ("spi: Add Socionext F_OSPI SPI flash controller driver")

    Reported-by: David Binderman <dcb314@hotmail.com>
    Link:
https://lore.kernel.org/all/DB6P189MB0568F3BE9384315F5C8C1A3E9CA49@DB6P189MB0568.EURP189.PRO
D.OUTLOOK.COM/

    Cc: stable@vger.kernel.org
    Signed-off-by: Dhruva Gole <d-gole@ti.com>
    Link: https://lore.kernel.org/r/20230223095202.924626-1-d-gole@ti.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a354667993183b54430a34937bd3718380c63780
Author: Marc Zyngier <maz@kernel.org>
Date:   Mon Feb 20 19:01:01 2023 +0000

    genirq/msi: Take the per-device MSI lock before validating the control structure

    commit 0af2795f936f1ea1f9f1497447145dfcc7ed2823 upstream.

    Calling msi_ctrl_valid() ultimately results in calling
    msi_get_device_domain(), which requires holding the device MSI lock.

    However, in msi_domain_populate_irqs() the lock is taken right after having
    called msi_ctrl_valid(), which is just a tad too late.

    Take the lock before invoking msi_ctrl_valid().

    Fixes: 40742716f294 ("genirq/msi: Make msi_add_simple_msi_descs() device domain aware")
    Reported-by: "Russell King (Oracle)" <linux@armlinux.org.uk>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
```

```
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/Y/Opu6ETe3ZzZ/8E@shell.armlinux.org.uk
    Link: https://lore.kernel.org/r/20230220190101.314446-1-maz@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e0b0354738f4cced09c0fb6a30021c5400390701
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Mar 1 22:07:48 2023 +0100

    genirq/msi, platform-msi: Ensure that MSI descriptors are unreferenced

    commit 0fb7fb713461e44b12e72c292bf90ee300f40710 upstream.

    Miquel reported a warning in the MSI core which is triggered when
    interrupts are freed via platform_msi_device_domain_free().

    This code got reworked to use core functions for freeing the MSI
    descriptors, but nothing took care to clear the msi_desc->irq entry, which
    then triggers the warning in msi_free_msi_desc() which uses desc->irq to
    validate that the descriptor has been torn down. The same issue exists in
    msi_domain_populate_irqs().

    Up to the point that msi_free_msi_descs() grew a warning for this case,
    this went un-noticed.

    Provide the counterpart of msi_domain_populate_irqs() and invoke it in
    platform_msi_device_domain_free() before freeing the interrupts and MSI
    descriptors and also in the error path of msi_domain_populate_irqs().

    Fixes: 2f2940d16823 ("genirq/msi: Remove filter from msi_free_descs_free_range()")
    Reported-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/87mt4wkwnv.ffs@tglx
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6dcf132fe236045bd7f50c008660ea086d09af1f
Author: Naoya Horiguchi <naoya.horiguchi@nec.com>
Date:   Tue Feb 21 17:59:05 2023 +0900

    mm/hwpoison: convert TTU_IGNORE_HWPOISON to TTU_HWPOISON

    commit 6da6b1d4a7df8c35770186b53ef65d388398e139 upstream.

    After a memory error happens on a clean folio, a process unexpectedly
    receives SIGBUS when it accesses the error page.  This SIGBUS killing is
    pointless and simply degrades the level of RAS of the system, because the
    clean folio can be dropped without any data lost on memory error handling
    as we do for a clean pagecache.

    When memory_failure() is called on a clean folio, try_to_unmap() is called
    twice (one from split_huge_page() and one from hwpoison_user_mappings()).
    The root cause of the issue is that pte conversion to hwpoisoned entry is
    now done in the first call of try_to_unmap() because PageHWPoison is
    already set at this point, while it's actually expected to be done in the
    second call.  This behavior disturbs the error handling operation like
    removing pagecache, which results in the malfunction described above.

    So convert TTU_IGNORE_HWPOISON into TTU_HWPOISON and set TTU_HWPOISON only
    when we really intend to convert pte to hwpoison entry.  This can prevent
    other callers of try_to_unmap() from accidentally converting to hwpoison
    entries.

    Link: https://lkml.kernel.org/r/20230221085905.1465385-1-naoya.horiguchi@linux.dev
```

```
    Fixes: a42634a6c07d ("readahead: Use a folio in read_pages()")
    Signed-off-by: Naoya Horiguchi <naoya.horiguchi@nec.com>
    Cc: David Hildenbrand <david@redhat.com>
    Cc: Hugh Dickins <hughd@google.com>
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: Miaohe Lin <linmiaohe@huawei.com>
    Cc: Minchan Kim <minchan@kernel.org>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ece25d35038d72624a1bc08c36e85645fd3f37ce
Author: Guilherme G. Piccoli <gpiccoli@igalia.com>
Date:   Sun Feb 26 13:08:38 2023 -0300

    panic: fix the panic_print NMI backtrace setting

    commit b905039e428d639adeebb719b76f98865ea38d4d upstream.

    Commit 8d470a45d1a6 ("panic: add option to dump all CPUs backtraces in
    panic_print") introduced a setting for the "panic_print" kernel parameter
    to allow users to request a NMI backtrace on panic.  Problem is that the
    panic_print handling happens after the secondary CPUs are already
    disabled, hence this option ended-up being kind of a no-op - kernel skips
    the NMI trace in idling CPUs, which is the case of offline CPUs.

    Fix it by checking the NMI backtrace bit in the panic_print prior to the
    CPU disabling function.

    Link: https://lkml.kernel.org/r/20230226160838.414257-1-gpiccoli@igalia.com
    Fixes: 8d470a45d1a6 ("panic: add option to dump all CPUs backtraces in panic_print")
    Signed-off-by: Guilherme G. Piccoli <gpiccoli@igalia.com>
    Cc: <stable@vger.kernel.org>
    Cc: Baoquan He <bhe@redhat.com>
    Cc: Dave Young <dyoung@redhat.com>
    Cc: Feng Tang <feng.tang@intel.com>
    Cc: HATAYAMA Daisuke <d.hatayama@jp.fujitsu.com>
    Cc: Hidehiro Kawai <hidehiro.kawai.ez@hitachi.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Michael Kelley <mikelley@microsoft.com>
    Cc: Petr Mladek <pmladek@suse.com>
    Cc: Vivek Goyal <vgoyal@redhat.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 15a6cdd02c8193e99b8b022e40abdb8822b3b2b5
Author: Matthias Kaehlcke <mka@chromium.org>
Date:   Thu Feb 23 00:33:30 2023 +0000

    regulator: core: Use ktime_get_boottime() to determine how long a regulator was off

    commit 80d2c29e09e663761c2778167a625b25ffe01b6f upstream.

    For regulators with 'off-on-delay-us' the regulator framework currently
    uses ktime_get() to determine how long the regulator has been off
    before re-enabling it (after a delay if needed). A problem with using
    ktime_get() is that it doesn't account for the time the system is
    suspended. As a result a regulator with a longer 'off-on-delay' (e.g.
    500ms) that was switched off during suspend might still incurr in a
    delay on resume before it is re-enabled, even though the regulator
    might have been off for hours. ktime_get_boottime() accounts for
    suspend time, use it instead of ktime_get().

    Fixes: a8ce7bd89689 ("regulator: core: Fix off_on_delay handling")
```

```
    Cc: stable@vger.kernel.org    # 5.13+
    Signed-off-by: Matthias Kaehlcke <mka@chromium.org>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Link:
https://lore.kernel.org/r/20230223003301.v2.1.I9719661b8eb0a73b8c416f9c26cf5bd8c0563f99@chan
geid
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c79e1d1207f53d5b10d95c2b347afd2a543e8e4b
Author: Xiubo Li <xiubli@redhat.com>
Date:   Mon Feb 13 13:56:20 2023 +0800

    ceph: update the time stamps and try to drop the suid/sgid

    commit e027253c4b77d395798600a90b6a96fe4adf4d5e upstream.

    The fallocate will try to clear the suid/sgid if a unprevileged user
    changed the file.

    There is no POSIX item requires that we should clear the suid/sgid
    in fallocate code path but this is the default behaviour for most of
    the filesystems and the VFS layer. And also the same for the write
    code path, which have already support it.

    And also we need to update the time stamps since the fallocate will
    change the file contents.

    Cc: stable@vger.kernel.org
    Link: https://tracker.ceph.com/issues/58054
    Signed-off-by: Xiubo Li <xiubli@redhat.com>
    Reviewed-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Ilya Dryomov <idryomov@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cc8c0dd2984503ed09efa37bcafcef3d3da104e8
Author: Ilya Dryomov <idryomov@gmail.com>
Date:   Fri Feb 24 18:48:54 2023 +0100

    rbd: avoid use-after-free in do_rbd_add() when rbd_dev_create() fails

    commit f7c4d9b133c7a04ca619355574e96b6abf209fba upstream.

    If getting an ID or setting up a work queue in rbd_dev_create() fails,
    use-after-free on rbd_dev->rbd_client, rbd_dev->spec and rbd_dev->opts
    is triggered in do_rbd_add().  The root cause is that the ownership of
    these structures is transfered to rbd_dev prematurely and they all end
    up getting freed when rbd_dev_create() calls rbd_dev_free() prior to
    returning to do_rbd_add().

    Found by Linux Verification Center (linuxtesting.org) with SVACE, an
    incomplete patch submitted by Natalia Petrova <n.petrova@fintech.ru>.

    Cc: stable@vger.kernel.org
    Fixes: 1643dfa4c2c8 ("rbd: introduce a per-device ordered workqueue")
    Signed-off-by: Ilya Dryomov <idryomov@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 833f8dbbdd7b23dbdb17ef4a94513b20ce390c33
Author: Alexander Mikhalitsyn <alexander@mihalicyn.com>
Date:   Thu Jan 26 11:23:18 2023 +0100

    fuse: add inode/permission checks to fileattr_get/fileattr_set

    commit 1cc4606d19e3710bfab3f6704b87ff9580493c69 upstream.
```

```
    It looks like these checks were accidentally lost during the conversion to
    fileattr API.

    Fixes: 72227eac177d ("fuse: convert to fileattr")
    Cc: <stable@vger.kernel.org> # v5.13
    Signed-off-by: Alexander Mikhalitsyn <aleksandr.mikhalitsyn@canonical.com>
    Signed-off-by: Miklos Szeredi <mszeredi@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d1c9727ba8190a575f14ba3cb136f3bc2b27f7e0
Author: Peter Collingbourne <pcc@google.com>
Date:   Tue Feb 14 21:09:11 2023 -0800

    arm64: Reset KASAN tag in copy_highpage with HW tags only

    commit e74a68468062d7ebd8ce17069e12ccc64cc6a58c upstream.

    During page migration, the copy_highpage function is used to copy the
    page data to the target page. If the source page is a userspace page
    with MTE tags, the KASAN tag of the target page must have the match-all
    tag in order to avoid tag check faults during subsequent accesses to the
    page by the kernel. However, the target page may have been allocated in
    a number of ways, some of which will use the KASAN allocator and will
    therefore end up setting the KASAN tag to a non-match-all tag. Therefore,
    update the target page's KASAN tag to match the source page.

    We ended up unintentionally fixing this issue as a result of a bad
    merge conflict resolution between commit e059853d14ca ("arm64: mte:
    Fix/clarify the PG_mte_tagged semantics") and commit 20794545c146 ("arm64:
    kasan: Revert "arm64: mte: reset the page tag in page->flags""), which
    preserved a tag reset for PG_mte_tagged pages which was considered to be
    unnecessary at the time. Because SW tags KASAN uses separate tag storage,
    update the code to only reset the tags when HW tags KASAN is enabled.

    Signed-off-by: Peter Collingbourne <pcc@google.com>
    Link: https://linux-review.googlesource.com/id/If303d8a709438d3ff5af5fd85706505830f52e0c
    Reported-by: "Kuan-Ying Lee (李冠穎)" <Kuan-Ying.Lee@mediatek.com>
    Cc: <stable@vger.kernel.org> # 6.1
    Fixes: 20794545c146 ("arm64: kasan: Revert "arm64: mte: reset the page tag in page-
>flags"")
    Reviewed-by: Andrey Konovalov <andreyknvl@gmail.com>
    Link: https://lore.kernel.org/r/20230215050911.1433132-1-pcc@google.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e614cc200e87ac616e0e3a9b6a557db1a3c5b59
Author: Catalin Marinas <catalin.marinas@arm.com>
Date:   Wed Feb 22 17:52:32 2023 +0000

    arm64: mm: hugetlb: Disable HUGETLB_PAGE_OPTIMIZE_VMEMMAP

    commit 060a2c92d1b627c86c5c42ca69baf00457c00c5a upstream.

    Revert the HUGETLB_PAGE_FREE_VMEMMAP selection from commit 1e63ac088f20
    ("arm64: mm: hugetlb: enable HUGETLB_PAGE_FREE_VMEMMAP for arm64") but
    keep the flush_dcache_page() compound_head() change as it aligns with
    the corresponding check in the __sync_icache_dcache() function.

    The original config option was renamed in commit 47010c040dec ("mm:
    hugetlb_vmemmap: cleanup CONFIG_HUGETLB_PAGE_FREE_VMEMMAP*") to
    HUGETLB_PAGE_OPTIMIZE_VMEMMAP and the flush_dcache_page() check was
    further simplified by commit 2da1c30929a2 ("mm: hugetlb_vmemmap: delete
    hugetlb_optimize_vmemmap_enabled()").
```

```
    The reason for the revert is that the generic vmemmap_remap_pte()
    function changes both the permissions (writeable to read-only) and the
    output address (pfn) of the vmemmap ptes. This is deemed UNPREDICTABLE
    by the Arm architecture without a break-before-make sequence (make the
    PTE invalid, TLBI, write the new valid PTE). However, such sequence is
    not possible since the vmemmap may be concurrently accessed by the
    kernel. Disable the optimisation until a better solution is found.

    Fixes: 1e63ac088f20 ("arm64: mm: hugetlb: enable HUGETLB_PAGE_FREE_VMEMMAP for arm64")
    Cc: <stable@vger.kernel.org> # 5.19.x
    Cc: Muchun Song <muchun.song@linux.dev>
    Cc: Will Deacon <will@kernel.org>
    Cc: Anshuman Khandual <anshuman.khandual@arm.com>
    Link: https://lore.kernel.org/r/Y9pZALdn3pKiJUeQ@arm.com
    Reviewed-by: Anshuman Khandual <anshuman.khandual@arm.com>
    Link: https://lore.kernel.org/r/20230222175232.540851-1-catalin.marinas@arm.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7521da2eb42d65f89f511b7912d3757cf3d9168a
Author: Sudeep Holla <sudeep.holla@arm.com>
Date:   Thu Feb 23 13:57:42 2023 +0000

    arm64: acpi: Fix possible memory leak of ffh_ctxt

    commit 1b561d3949f8478c5403c9752b5533211a757226 upstream.

    Allocated 'ffh_ctxt' memory leak is possible if the SMCCC version
    and conduit checks fail and -EOPNOTSUPP is returned without freeing the
    allocated memory.

    Fix the same by moving the allocation after the SMCCC version and
    conduit checks.

    Fixes: 1d280ce099db ("arm64: Add architecture specific ACPI FFH Opregion callbacks")
    Cc: <stable@vger.kernel.org> # 6.2.x
    Cc: Will Deacon <will@kernel.org>
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <error27@gmail.com>
    Suggested-by: Dan Carpenter <error27@gmail.com>
    Link: https://lore.kernel.org/r/202302191417.dAl9NuE8-lkp@intel.com/
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Link: https://lore.kernel.org/r/20230223135742.2952091-1-sudeep.holla@arm.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 897e7a13a58aca72161b158e15baea6d4e48a7cb
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:40 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Odroid HC1

    commit 2e3d0e20d8456f876607a8af61fdb83dfbf98cb6 upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.  This was not critical before,
    but since rework of thermal Devicetree initialization in the
    commit 3fd6d6e2b4e8 ("thermal/of: Rework the thermal device tree
    initialization"), this leads to errors registering thermal zones other
    than first one:

      thermal_sys: cpu0-thermal: Failed to read thermal-sensors cells: -2
      thermal_sys: Failed to find thermal zone for tmu id=0
      exynos-tmu 10064000.tmu: Failed to register sensor: -2
      exynos-tmu: probe of 10064000.tmu failed with error -2
```

```
    Fixes: 1ac49427b566 ("ARM: dts: exynos: Add support for Hardkernel's Odroid HC1 board")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230209105841.779596-5-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e2fb42a97b6773fc59cc3e84595a191befd40fb
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:39 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Odroid XU

    commit 9372eca505e7a19934d750b4b4c89a3652738e66 upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.  Since thermal-sensors property
    is already defined in included exynosi5410.dtsi, drop it from
    exynos5410-odroidxu.dts to fix the error and remoev redundancy.

    Fixes: 88644b4c750b ("ARM: dts: exynos: Configure PWM, usb3503, PMIC and thermal on
Odroid XU board")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230209105841.779596-4-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 676c5672ab4f0b1cec2af1142b3dacefec4e350a
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:38 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Exynos5250

    commit 33e2c595e2e4016991ead44933a29d1ef93d5f26 upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.

    Cc: <stable@vger.kernel.org>
    Fixes: 9843a2236003 ("ARM: dts: Provide dt bindings identical for Exynos TMU")
    Link: https://lore.kernel.org/r/20230209105841.779596-3-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 044a0326fcfe2ffb1e055895091a5c718c8fc38f
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:41 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Odroid XU3 family

    commit a3583e92d188ec6c58c7f603ac5e72dd8a11c21a upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.  This was not critical before,
    but since rework of thermal Devicetree initialization in the
    commit 3fd6d6e2b4e8 ("thermal/of: Rework the thermal device tree
    initialization"), this leads to errors registering thermal zones other
    than first one:

      thermal_sys: cpu0-thermal: Failed to read thermal-sensors cells: -2
      thermal_sys: Failed to find thermal zone for tmu id=0
      exynos-tmu 10064000.tmu: Failed to register sensor: -2
      exynos-tmu: probe of 10064000.tmu failed with error -2

    Fixes: f1722d7dd8b8 ("ARM: dts: Define default thermal-zones for exynos5422")
```

```
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230209105841.779596-6-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 360c3fdec7a2c9ed5053347a532d6183112eec67
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:36 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Exynos4

    commit 8e4505e617a80f601e2f53a917611777f128f925 upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.

    Fixes: 328829a6ad70 ("ARM: dts: define default thermal-zones for exynos4")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230209105841.779596-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d77be13b9e32a7598ab638af5b786aca19080dbb
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 9 11:58:37 2023 +0100

    ARM: dts: exynos: correct TMU phandle in Exynos4210

    commit 408ab6786dbf6dd696488054c9559681112ef994 upstream.

    TMU node uses 0 as thermal-sensor-cells, thus thermal zone referencing
    it must not have an argument to phandle.  Since thermal-sensors property is
    already defined in included exynos4-cpu-thermal.dtsi, drop it from
    exynos4210.dtsi to fix the error and remoev redundancy.

    Fixes: 9843a2236003 ("ARM: dts: Provide dt bindings identical for Exynos TMU")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230209105841.779596-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ad08f81f8047192eefcfcd2be6c1946c0cf47655
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Mon Jan 23 18:49:30 2023 +0530

    ARM: dts: qcom: sdx55: Add Qcom SMMU-500 as the fallback for IOMMU node

    commit af4ab377543853b690cc85b4c46cf976ab560dc2 upstream.

    SDX55 uses the Qcom version of the SMMU-500 IP. So use "qcom,smmu-500"
    compatible as the fallback to the SoC specific compatible.

    Cc: <stable@vger.kernel.org> # 5.12
    Fixes: a2bdfdfba2af ("ARM: dts: qcom: sdx55: Enable ARM SMMU")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230123131931.263024-3-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8acd379b922538e42ca500ad8bf699f4142bab33
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Mon Jan 23 18:49:31 2023 +0530

    ARM: dts: qcom: sdx65: Add Qcom SMMU-500 as the fallback for IOMMU node
```

```
    commit 157178a7912e00a0aa0371dc9041952c1a21d112 upstream.

    SDX65 uses the Qcom version of the SMMU-500 IP. So use "qcom,smmu-500"
    compatible as the fallback to the SoC specific compatible.

    Cc: <stable@vger.kernel.org> # 5.19
    Fixes: 98187f7b74bf ("ARM: dts: qcom: sdx65: Enable ARM SMMU")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230123131931.263024-4-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 837fffc028f9e6a1a5746a5e000a22acb4463404
Author: Mika Westerberg <mika.westerberg@linux.intel.com>
Date:   Wed Feb 15 13:00:40 2023 +0200

    spi: intel: Check number of chip selects after reading the descriptor

    commit 574fbb95cd9d88bdc9c9c4c64223a38a61d7de9a upstream.

    The flash decriptor contains the number of flash components that we use
    to figure out how many flash chips there are connected. Therefore we
    need to read it first before deciding how many chip selects the
    controller has.

    Reported-by: Marcin Witkowski <marcin.witkowski@intel.com>
    Fixes: 3f03c618bebb ("spi: intel: Add support for second flash chip")
    Cc: stable@vger.kernel.org
    Signed-off-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Link: https://lore.kernel.org/r/20230215110040.42186-1-mika.westerberg@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8014342094442e93eabd7a445d5402b5255273f8
Author: Zev Weiss <zev@bewilderbeest.net>
Date:   Mon Jan 2 13:28:57 2023 -0800

    hwmon: (nct6775) Fix incorrect parenthesization in nct6775_write_fan_div()

    commit 2fbb848b65cde5b876cce52ebcb34de4aaa5a94a upstream.

    Commit 4ef2774511dc ("hwmon: (nct6775) Convert register access to
    regmap API") fumbled the shifting & masking of the fan_div values such
    that odd-numbered fan divisors would always be set to zero.  Fix it so
    that we actually OR in the bits we meant to.

    Signed-off-by: Zev Weiss <zev@bewilderbeest.net>
    Fixes: 4ef2774511dc ("hwmon: (nct6775) Convert register access to regmap API")
    Cc: stable@kernel.org # v5.19+
    Link: https://lore.kernel.org/r/20230102212857.5670-1-zev@bewilderbeest.net
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 984d16e83c53bca3072281b5541e1db64cc1d30a
Author: Zev Weiss <zev@bewilderbeest.net>
Date:   Wed Feb 1 18:18:25 2023 -0800

    hwmon: (peci/cputemp) Fix off-by-one in coretemp_label allocation

    commit f00093608fa790580da309bb9feb5108fbe7c331 upstream.

    The find_last_bit() call produces the index of the highest-numbered
    core in core_mask; because cores are numbered from zero, the number of
```

```
    elements we need to allocate is one more than that.

    Signed-off-by: Zev Weiss <zev@bewilderbeest.net>
    Cc: stable@kernel.org # v5.18
    Fixes: bf3608f338e9 ("hwmon: peci: Add cputemp driver")
    Reviewed-by: Iwona Winiarska <iwona.winiarska@intel.com>
    Link: https://lore.kernel.org/r/20230202021825.21486-1-zev@bewilderbeest.net
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 844ea6c307e8e06fb5839bc6559b4384f438a9b1
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Sun Jan 22 14:03:31 2023 -0500

    dm flakey: fix a bug with 32-bit highmem systems

    commit 8eb29c4fbf9661e6bd4dd86197a37ffe0ecc9d50 upstream.

    The function page_address does not work with 32-bit systems with high
    memory. Use bvec_kmap_local/kunmap_local instead.

    Cc: stable@vger.kernel.org
    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Reviewed-by: Sweet Tea Dorminy <sweettea-kernel@dorminy.me>
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 63d31617883d64b43b0e2d529f0751f40713ecae
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Sun Jan 22 14:02:57 2023 -0500

    dm flakey: don't corrupt the zero page

    commit f50714b57aecb6b3dc81d578e295f86d9c73f078 upstream.

    When we need to zero some range on a block device, the function
    __blkdev_issue_zero_pages submits a write bio with the bio vector pointing
    to the zero page. If we use dm-flakey with corrupt bio writes option, it
    will corrupt the content of the zero page which results in crashes of
    various userspace programs. Glibc assumes that memory returned by mmap is
    zeroed and it uses it for calloc implementation; if the newly mapped
    memory is not zeroed, calloc will return non-zeroed memory.

    Fix this bug by testing if the page is equal to ZERO_PAGE(0) and
    avoiding the corruption in this case.

    Cc: stable@vger.kernel.org
    Fixes: a00f5276e266 ("dm flakey: Properly corrupt multi-page bios.")
    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Reviewed-by: Sweet Tea Dorminy <sweettea-kernel@dorminy.me>
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ed56ad5cacb7a3aeb611494d5d66e2399d2bfecc
Author: Joe Thornber <ejt@redhat.com>
Date:   Thu Jan 26 09:59:10 2023 +0000

    dm cache: free background tracker's queued work in btracker_destroy

    commit 95ab80a8a0fef2ce0cc494a306dd283948066ce7 upstream.

    Otherwise the kernel can BUG with:

    [ 2245.426978]
================================================================================
```

```
    [ 2245.435155] BUG bt_work (Tainted: G    B   W        ): Objects remaining in bt_work
on __kmem_cache_shutdown()
    [ 2245.445233] -------------------------------------------------------------------------
----
    [ 2245.445233]
    [ 2245.454879] Slab 0x00000000b0ce2b30 objects=64 used=2 fp=0x000000000a3c6a4e
flags=0x17ffffc0000200(slab|node=0|zone=2|lastcpupid=0x1fffff)
    [ 2245.467300] CPU: 7 PID: 10805 Comm: lvm Kdump: loaded Tainted: G    B   W
6.0.0-rc2 #19
    [ 2245.476078] Hardware name: Dell Inc. PowerEdge R7525/0590KW, BIOS 2.5.6 10/06/2021
    [ 2245.483646] Call Trace:
    [ 2245.486100]  <TASK>
    [ 2245.488206]  dump_stack_lvl+0x34/0x48
    [ 2245.491878]  slab_err+0x95/0xcd
    [ 2245.495028]  __kmem_cache_shutdown.cold+0x31/0x136
    [ 2245.499821]  kmem_cache_destroy+0x49/0x130
    [ 2245.503928]  btracker_destroy+0x12/0x20 [dm_cache]
    [ 2245.508728]  smq_destroy+0x15/0x60 [dm_cache_smq]
    [ 2245.513435]  dm_cache_policy_destroy+0x12/0x20 [dm_cache]
    [ 2245.518834]  destroy+0xc0/0x110 [dm_cache]
    [ 2245.522933]  dm_table_destroy+0x5c/0x120 [dm_mod]
    [ 2245.527649]  __dm_destroy+0x10e/0x1c0 [dm_mod]
    [ 2245.532102]  dev_remove+0x117/0x190 [dm_mod]
    [ 2245.536384]  ctl_ioctl+0x1a2/0x290 [dm_mod]
    [ 2245.540579]  dm_ctl_ioctl+0xa/0x20 [dm_mod]
    [ 2245.544773]  __x64_sys_ioctl+0x8a/0xc0
    [ 2245.548524]  do_syscall_64+0x5c/0x90
    [ 2245.552104]  ? syscall_exit_to_user_mode+0x12/0x30
    [ 2245.556897]  ? do_syscall_64+0x69/0x90
    [ 2245.560648]  ? do_syscall_64+0x69/0x90
    [ 2245.564394]  entry_SYSCALL_64_after_hwframe+0x63/0xcd
    [ 2245.569447] RIP: 0033:0x7fe52583ec6b
    ...
    [ 2245.646771] ------------[ cut here ]------------
    [ 2245.651395] kmem_cache_destroy bt_work: Slab cache still has objects when called from
btracker_destroy+0x12/0x20 [dm_cache]
    [ 2245.651408] WARNING: CPU: 7 PID: 10805 at mm/slab_common.c:478
kmem_cache_destroy+0x128/0x130

    Found using: lvm2-testsuite --only "cache-single-split.sh"

    Ben bisected and found that commit 0495e337b703 ("mm/slab_common:
    Deleting kobject in kmem_cache_destroy() without holding
    slab_mutex/cpu_hotplug_lock") first exposed dm-cache's incomplete
    cleanup of its background tracker work objects.

    Reported-by: Benjamin Marzinski <bmarzins@redhat.com>
    Tested-by: Benjamin Marzinski <bmarzins@redhat.com>
    Cc: stable@vger.kernel.org # 6.0+
    Signed-off-by: Joe Thornber <ejt@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7bcf6d686071117b80605e4512a90f821a6a4541
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Sun Jan 22 14:03:56 2023 -0500

    dm flakey: fix logic when corrupting a bio

    commit aa56b9b75996ff4c76a0a4181c2fa0206c3d91cc upstream.

    If "corrupt_bio_byte" is set to corrupt reads and corrupt_bio_flags is
    used, dm-flakey would erroneously return all writes as errors. Likewise,
    if "corrupt_bio_byte" is set to corrupt writes, dm-flakey would return
    errors for all reads.
```

```
    Fix the logic so that if fc->corrupt_bio_byte is non-zero, dm-flakey
    will not abort reads on writes with an error.

    Cc: stable@vger.kernel.org
    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Reviewed-by: Sweet Tea Dorminy <sweettea-kernel@dorminy.me>
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6921f8116f0e11a024197c423274754372c71926
Author: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
Date:   Wed Feb 1 12:39:41 2023 -0800

    thermal: intel: powerclamp: Fix cur_state for multi package system

    commit 8e47363588377e1bdb65e2b020b409cfb44dd260 upstream.

    The powerclamp cooling device cur_state shows actual idle observed by
    package C-state idle counters. But the implementation is not sufficient
    for multi package or multi die system. The cur_state value is incorrect.
    On these systems, these counters must be read from each package/die and
    somehow aggregate them. But there is no good method for aggregation.

    It was not a problem when explicit CPU model addition was required to
    enable intel powerclamp. In this way certain CPU models could have
    been avoided. But with the removal of CPU model check with the
    availability of Package C-state counters, the driver is loaded on most
    of the recent systems.

    For multi package/die systems, just show the actual target idle state,
    the system is trying to achieve. In powerclamp this is the user set
    state minus one.

    Also there is no use of starting a worker thread for polling package
    C-state counters and applying any compensation for multiple package
    or multiple die systems.

    Fixes: b721ca0d1927 ("thermal/powerclamp: remove cpu whitelist")
    Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Cc: 4.14+ <stable@vger.kernel.org> # 4.14+
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 205b02a83c76f779f0f0f32000657ca6557bac17
Author: Manish Chopra <manishc@marvell.com>
Date:   Thu Feb 16 03:54:47 2023 -0800

    qede: fix interrupt coalescing configuration

    commit 908d4bb7c54caa58253a363d63e797a468eaf321 upstream.

    On default driver load device gets configured with unexpected
    higher interrupt coalescing values instead of default expected
    values as memory allocated from krealloc() is not supposed to
    be zeroed out and may contain garbage values.

    Fix this by allocating the memory of required size first with
    kcalloc() and then use krealloc() to resize and preserve the
    contents across down/up of the interface.

    Signed-off-by: Manish Chopra <manishc@marvell.com>
    Fixes: b0ec5489c480 ("qede: preserve per queue stats across up/down of interface")
    Cc: stable@vger.kernel.org
    Cc: Bhaskar Upadhaya <bupadhaya@marvell.com>
```

```
    Cc: David S. Miller <davem@davemloft.net>
    Link: https://bugzilla.redhat.com/show_bug.cgi?id=2160054
    Signed-off-by: Alok Prasad <palok@marvell.com>
    Signed-off-by: Ariel Elior <aelior@marvell.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f16394792f81fcd6b989174168e429dde686877b
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Feb 6 20:33:06 2023 +0100

    cpuidle: add ARCH_SUSPEND_POSSIBLE dependencies

    commit 7787943a3a8ade6594a68db28c166adbb1d3708c upstream.

    Some ARMv4 processors don't support suspend, which leads
    to a build failure with the tegra and qualcomm cpuidle driver:

    WARNING: unmet direct dependencies detected for ARM_CPU_SUSPEND
      Depends on [n]: ARCH_SUSPEND_POSSIBLE [=n]
      Selected by [y]:
      - ARM_TEGRA_CPUIDLE [=y] && CPU_IDLE [=y] && (ARM [=y] || ARM64) && (ARCH_TEGRA [=n]
|| COMPILE_TEST [=y]) && !ARM64 && MMU [=y]

    arch/arm/kernel/sleep.o: in function `__cpu_suspend':
    (.text+0x68): undefined reference to `cpu_sa110_suspend_size'
    (.text+0x68): undefined reference to `cpu_fa526_suspend_size'

    Add an explicit dependency to make randconfig builds avoid
    this combination.

    Fixes: faae6c9f2e68 ("cpuidle: tegra: Enable compile testing")
    Fixes: a871be6b8eee ("cpuidle: Convert Qualcomm SPM driver to a generic CPUidle driver")
    Link: https://lore.kernel.org/all/20211013160125.772873-1-arnd@kernel.org/
    Cc: All applicable <stable@vger.kernel.org>
    Reviewed-by: Dmitry Osipenko <digetx@gmail.com>
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Acked-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bf3c348c5fdcf00a7eeed04a1b83e454d2dca2e5
Author: Marc Bornand <dev.mbornand@systemb.ch>
Date:   Wed Feb 15 08:47:53 2023 +0000

    wifi: cfg80211: Set SSID if it is not already set

    commit c38c701851011c94ce3be1ccb3593678d2933fd8 upstream.

    When a connection was established without going through
    NL80211_CMD_CONNECT, the ssid was never set in the wireless_dev struct.
    Now we set it in __cfg80211_connect_result() when it is not already set.

    When using a userspace configuration that does not call
    cfg80211_connect() (can be checked with breakpoints in the kernel),
    this patch should allow `networkctl status device_name` to output the
    SSID instead of null.

    Cc: stable@vger.kernel.org
    Reported-by: Yohan Prod'homme <kernel@zoddo.fr>
    Fixes: 7b0a0e3c3a88 (wifi: cfg80211: do some rework towards MLO link APIs)
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216711
    Signed-off-by: Marc Bornand <dev.mbornand@systemb.ch>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit f4b6a138efb8a32507b8946104e32cb926308da7
Author: Alexander Wetzel <alexander@wetzel-home.de>
Date:   Tue Jan 24 15:18:56 2023 +0100

    wifi: cfg80211: Fix use after free for wext

    commit 015b8cc5e7c4d7bb671f1984d7b7338c310b185b upstream.

    Key information in wext.connect is not reset on (re)connect and can hold
    data from a previous connection.

    Reset key data to avoid that drivers or mac80211 incorrectly detect a
    WEP connection request and access the freed or already reused memory.

    Additionally optimize cfg80211_sme_connect() and avoid an useless
    schedule of conn_work.

    Fixes: fffd0934b939 ("cfg80211: rework key operation")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230124141856.356646-1-alexander@wetzel-home.de
    Signed-off-by: Alexander Wetzel <alexander@wetzel-home.de>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2388cfc52822d23f9704a56f57b7ed45aa9c9051
Author: Len Brown <len.brown@intel.com>
Date:   Wed Feb 1 12:32:01 2023 -0600

    wifi: ath11k: allow system suspend to survive ath11k

    commit 7c15430822e71e90203d87e6d0cfe83fa058b0dc upstream.

    When ath11k runs into internal errors upon suspend,
    it returns an error code to pci_pm_suspend, which
    aborts the entire system suspend.

    The driver should not abort system suspend, but should
    keep its internal errors to itself, and allow the system
    to suspend.  Otherwise, a user can suspend a laptop
    by closing the lid and sealing it into a case, assuming
    that is will suspend, rather than heating up and draining
    the battery when in transit.

    In practice, the ath11k device seems to have plenty of transient
    errors, and subsequent suspend cycles after this failure
    often succeed.

    https://bugzilla.kernel.org/show_bug.cgi?id=216968

    Fixes: d1b0c33850d29 ("ath11k: implement suspend for QCA6390 PCI devices")

    Signed-off-by: Len Brown <len.brown@intel.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230201183201.14431-1-len.brown@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5368db4d7e0b65caf08929e5332d19d70186508a
Author: Bitterblue Smith <rtl8821cerfe2@gmail.com>
Date:   Sun Jan 8 17:08:16 2023 +0200

    wifi: rtl8xxxu: Use a longer retry limit of 48

    commit 2a86aa9a1892d60ef2e3f310f5b42b8b05546d65 upstream.
```

```
    The Realtek rate control algorithm goes back and forth a lot between
    the highest and the lowest rate it's allowed to use. This is due to
    a lot of frames being dropped because the retry limits set by
    IEEE80211_CONF_CHANGE_RETRY_LIMITS are too low. (Experimentally, they
    are 4 for long frames and 7 for short frames.)

    The vendor drivers hardcode the value 48 for both retry limits (for
    station mode), which makes dropped frames very rare and thus the rate
    control is more stable.

    Because most Realtek chips handle the rate control in the firmware,
    which can't be modified, ignore the limits set by
    IEEE80211_CONF_CHANGE_RETRY_LIMITS and use the value 48 (set during
    chip initialisation), same as the vendor drivers.

    Cc: stable@vger.kernel.org
    Signed-off-by: Bitterblue Smith <rtl8821cerfe2@gmail.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/477d745b-6bac-111d-403c-487fc19aa30d@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 976fe246a57db503d2a0841c8cb813f2f66b78af
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Thu Feb 16 13:36:33 2023 +0800

    wifi: rtw88: use RTW_FLAG_POWERON flag to prevent to power on/off twice

    commit 4a267bc5ea8f159b614d0549030216d0434eccca upstream.

    Use power state to decide whether we can enter or leave IPS accurately,
    and then prevent to power on/off twice.

    The commit 6bf3a083407b ("wifi: rtw88: add flag check before enter or leave IPS")
    would like to prevent this as well, but it still can't entirely handle all
    cases. The exception is that WiFi gets connected and does suspend/resume,
    it will power on twice and cause it failed to power on after resuming,
    like:

      rtw_8723de 0000:03:00.0: failed to poll offset=0x6 mask=0x2 value=0x2
      rtw_8723de 0000:03:00.0: mac power on failed
      rtw_8723de 0000:03:00.0: failed to power on mac
      rtw_8723de 0000:03:00.0: leave idle state failed
      rtw_8723de 0000:03:00.0: failed to leave ips state
      rtw_8723de 0000:03:00.0: failed to leave idle state
      rtw_8723de 0000:03:00.0: failed to send h2c command

    To fix this, introduce new flag RTW_FLAG_POWERON to reflect power state,
    and call rtw_mac_pre_system_cfg() to configure registers properly between
    power-off/-on.

    Reported-by: Paul Gover <pmw.gover@yahoo.co.uk>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=217016
    Fixes: 6bf3a083407b ("wifi: rtw88: add flag check before enter or leave IPS")
    Cc: <Stable@vger.kernel.org>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230216053633.20366-1-pkshih@realtek.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 06d2e5a15125a86ac782587873d96400caabf8c9
Author: Mike Snitzer <snitzer@kernel.org>
Date:   Thu Feb 16 12:10:05 2023 -0500
```

```
    dm: add cond_resched() to dm_wq_requeue_work()

    commit f77692d65d54665d81815349cc727baa85e8b71d upstream.

    Otherwise the while() loop in dm_wq_requeue_work() can result in a
    "dead loop" on systems that have preemption disabled. This is
    particularly problematic on single cpu systems.

    Fixes: 8b211aaccb915 ("dm: add two stage requeue mechanism")
    Cc: stable@vger.kernel.org
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eb1d07e21f8852df2c2bd6fad726ad9bae4a98e3
Author: Pingfan Liu <piliu@redhat.com>
Date:   Wed Feb 15 19:23:40 2023 +0800

    dm: add cond_resched() to dm_wq_work()

    commit 0ca44fcef241768fd25ee763b3d203b9852f269b upstream.

    Otherwise the while() loop in dm_wq_work() can result in a "dead
    loop" on systems that have preemption disabled. This is particularly
    problematic on single cpu systems.

    Cc: stable@vger.kernel.org
    Signed-off-by: Pingfan Liu <piliu@redhat.com>
    Acked-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4eb29b0236a71cb147294f5487f7a4217b9b6382
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Tue Feb 7 08:33:06 2023 -0500

    dm: send just one event on resize, not two

    commit 7533afa1d27ba1234146d31d2402c195cf195962 upstream.

    Device mapper sends an uevent when the device is suspended, using the
    function set_capacity_and_notify. However, this causes a race condition
    with udev.

    Udev skips scanning dm devices that are suspended. If we send an uevent
    while we are suspended, udev will be racing with device mapper resume
    code. If the device mapper resume code wins the race, udev will process
    the uevent after the device is resumed and it will properly scan the
    device.

    However, if udev wins the race, it will receive the uevent, find out that
    the dm device is suspended and skip scanning the device. This causes bugs
    such as systemd unmounting the device - see
    https://bugzilla.redhat.com/show_bug.cgi?id=2158628

    This commit fixes this race.

    We replace the function set_capacity_and_notify with set_capacity, so that
    the uevent is not sent at this point. In do_resume, we detect if the
    capacity has changed and we pass a boolean variable need_resize_uevent to
    dm_kobject_uevent. dm_kobject_uevent adds "RESIZE=1" to the uevent if
    need_resize_uevent is set.

    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Tested-by: Peter Rajnoha <prajnoha@redhat.com>
    Cc: stable@vger.kernel.org
```

```
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 99341b8aee7b5b4255b339345bbcaa35867dfd0c
Author: Louis Rannou <lrannou@baylibre.com>
Date:   Fri Feb 3 09:07:54 2023 +0200

    mtd: spi-nor: Fix shift-out-of-bounds in spi_nor_set_erase_type

    commit f0f0cfdc3a024e21161714f2e05f0df3b84d42ad upstream.

    spi_nor_set_erase_type() was used either to set or to mask out an erase
    type. When we used it to mask out an erase type a shift-out-of-bounds
    was hit:
    UBSAN: shift-out-of-bounds in drivers/mtd/spi-nor/core.c:2237:24
    shift exponent 4294967295 is too large for 32-bit type 'int'

    The setting of the size_{shift, mask} and of the opcode are unnecessary
    when the erase size is zero, as throughout the code just the erase size
    is considered to determine whether an erase type is supported or not.
    Setting the opcode to 0xFF was wrong too as nobody guarantees that 0xFF
    is an unused opcode. Thus when masking out an erase type, just set the
    erase size to zero. This will fix the shift-out-of-bounds.

    Fixes: 5390a8df769e ("mtd: spi-nor: add support to non-uniform SFDP SPI NOR flash
memories")
    Cc: stable@vger.kernel.org
    Reported-by: Alexander Stein <Alexander.Stein@tq-group.com>
    Signed-off-by: Louis Rannou <lrannou@baylibre.com>
    Tested-by: Alexander Stein <Alexander.Stein@tq-group.com>
    Link: https://lore.kernel.org/r/20230203070754.50677-1-tudor.ambarus@linaro.org
    [ta: refine changes, new commit message, fix compilation error]
    Signed-off-by: Tudor Ambarus <tudor.ambarus@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9b2152bbf1634d6573e095a169d5201ac148b33b
Author: Tudor Ambarus <tudor.ambarus@linaro.org>
Date:   Tue Jan 10 18:47:02 2023 +0200

    mtd: spi-nor: spansion: Consider reserved bits in CFR5 register

    commit 3f592a869f87723314f0cb1ac232bd3bf8245be8 upstream.

    CFR5[6] is reserved bit and must be always 1. Set it to comply with flash
    requirements. While fixing SPINOR_REG_CYPRESS_CFR5V_OCT_DTR_{EN, DS}
    definition, stop using magic numbers and describe the missing bit fields
    in CFR5 register. This is useful for both readability and future possible
    addition of Octal STR mode support.

    Fixes: c3266af101f2 ("mtd: spi-nor: spansion: add support for Cypress Semper flash")
    Cc: stable@vger.kernel.org
    Reported-by: Takahiro Kuwano <Takahiro.Kuwano@infineon.com>
    Signed-off-by: Tudor Ambarus <tudor.ambarus@linaro.org>
    Reviewed-by: Dhruva Gole <d-gole@ti.com>
    Reviewed-by: Pratyush Yadav <ptyadav@amazon.de>
    Tested-by: Dhruva Gole <d-gole@ti.com>
    Link: https://lore.kernel.org/linux-mtd/20230110164703.83413-1-tudor.ambarus@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7205650530206c70d854328ae8b7c6fffb6b210c
Author: Takahiro Kuwano <Takahiro.Kuwano@infineon.com>
Date:   Mon Dec 26 13:01:58 2022 +0900

    mtd: spi-nor: sfdp: Fix index value for SCCR dwords
```

```
    commit ad9679f3811899fd1c21dc7bdd715e8e1cfb46b9 upstream.

    Array index for SCCR 22th DOWRD should be 21.

    Fixes: 981a8d60e01f ("mtd: spi-nor: Parse SFDP SCCR Map")
    Signed-off-by: Takahiro Kuwano <Takahiro.Kuwano@infineon.com>
    Signed-off-by: Tudor Ambarus <tudor.ambarus@linaro.org>
    Reviewed-by: Michael Walle <michael@walle.cc>
    Cc: stable@vger.kernel.org
    Link:
https://lore.kernel.org/r/d8a2a77c2c95cf776e7dcae6392d29fdcf5d6307.1672026365.git.Takahiro.K
uwano@infineon.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bee57c20fc0ca5ef9b9a53a0335eab2ac9e9cae1
Author: Dmitry Torokhov <dmitry.torokhov@gmail.com>
Date:   Fri Feb 3 16:43:29 2023 -0800

    Input: exc3000 - properly stop timer on shutdown

    commit 79c81d137d36f9635bbcbc3916c0cccb418a61dd upstream.

    We need to stop the timer on driver unbind or probe failures, otherwise
    we get UAF/Oops.

    Fixes: 7e577a17f2ee ("Input: add I2C attached EETI EXC3000 multi touch driver")
    Reported-by: "Stahl, Michael" <mstahl@moba.de>
    Link: https://lore.kernel.org/r/Y9dK57BFqtlf8NmN@google.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 18c65667fa9104780eeaa0dc1bc240f0c2094772
Author: Dan Williams <dan.j.williams@intel.com>
Date:   Mon Feb 13 17:01:05 2023 -0800

    cxl/pmem: Fix nvdimm registration races

    commit f57aec443c24d2e8e1f3b5b4856aea12ddda4254 upstream.

    A loop of the form:

        while true; do modprobe cxl_pci; modprobe -r cxl_pci; done

    ...fails with the following crash signature:

        BUG: kernel NULL pointer dereference, address: 0000000000000040
        [..]
        RIP: 0010:cxl_internal_send_cmd+0x5/0xb0 [cxl_core]
        [..]
        Call Trace:
         <TASK>
         cxl_pmem_ctl+0x121/0x240 [cxl_pmem]
         nvdimm_get_config_data+0xd6/0x1a0 [libnvdimm]
         nd_label_data_init+0x135/0x7e0 [libnvdimm]
         nvdimm_probe+0xd6/0x1c0 [libnvdimm]
         nvdimm_bus_probe+0x7a/0x1e0 [libnvdimm]
         really_probe+0xde/0x380
         __driver_probe_device+0x78/0x170
         driver_probe_device+0x1f/0x90
         __device_attach_driver+0x85/0x110
         bus_for_each_drv+0x7d/0xc0
         __device_attach+0xb4/0x1e0
         bus_probe_device+0x9f/0xc0
         device_add+0x445/0x9c0
```

```
          nd_async_device_register+0xe/0x40 [libnvdimm]
          async_run_entry_fn+0x30/0x130

    ...namely that the bottom half of async nvdimm device registration runs
    after the CXL has already torn down the context that cxl_pmem_ctl()
    needs. Unlike the ACPI NFIT case that benefits from launching multiple
    nvdimm device registrations in parallel from those listed in the table,
    CXL is already marked PROBE_PREFER_ASYNCHRONOUS. So provide for a
    synchronous registration path to preclude this scenario.

    Fixes: 21083f51521f ("cxl/pmem: Register 'pmem' / cxl_nvdimm devices")
    Cc: <stable@vger.kernel.org>
    Reported-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a458a8c1d1fc4e10a1813786132b09a3863ad3f2
Author: Jun Nie <jun.nie@linaro.org>
Date:   Tue Jan 3 09:45:17 2023 +0800

    ext4: refuse to create ea block when umounted

    commit f31173c19901a96bb2ebf6bcfec8a08df7095c91 upstream.

    The ea block expansion need to access s_root while it is
    already set as NULL when umount is triggered. Refuse this
    request to avoid panic.

    Reported-by: syzbot+2dacb8f015bf1420155f@syzkaller.appspotmail.com
    Link: https://syzkaller.appspot.com/bug?id=3613786cb88c93aa1c6a279b1df6a7b201347d08
    Link: https://lore.kernel.org/r/20230103014517.495275-3-jun.nie@linaro.org
    Cc: stable@kernel.org
    Signed-off-by: Jun Nie <jun.nie@linaro.org>
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d2efaf8c870c7067b8d1779773134f3481cd8f68
Author: Jun Nie <jun.nie@linaro.org>
Date:   Tue Jan 3 09:45:16 2023 +0800

    ext4: optimize ea_inode block expansion

    commit 1e9d62d252812575ded7c620d8fc67c32ff06c16 upstream.

    Copy ea data from inode entry when expanding ea block if possible.
    Then remove the ea entry if expansion success. Thus memcpy to a
    temporary buffer may be avoided.

    If the expansion fails, we do not need to recovery the removed ea
    entry neither in this way.

    Reported-by: syzbot+2dacb8f015bf1420155f@syzkaller.appspotmail.com
    Link: https://syzkaller.appspot.com/bug?id=3613786cb88c93aa1c6a279b1df6a7b201347d08
    Link: https://lore.kernel.org/r/20230103014517.495275-2-jun.nie@linaro.org
    Cc: stable@kernel.org
    Signed-off-by: Jun Nie <jun.nie@linaro.org>
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1ae96656b8f13c28e07bfc4fce48e1cb04578890
Author: Zhihao Cheng <chengzhihao1@huawei.com>
Date:   Tue Jan 10 09:53:27 2023 +0800

    jbd2: fix data missing when reusing bh which is ready to be checkpointed
```

```
    commit e6b9bd7290d334451ce054e98e752abc055e0034 upstream.

    Following process will make data lost and could lead to a filesystem
    corrupted problem:

    1. jh(bh) is inserted into T1->t_checkpoint_list, bh is dirty, and
       jh->b_transaction = NULL
    2. T1 is added into journal->j_checkpoint_transactions.
    3. Get bh prepare to write while doing checkpoing:
               PA                                  PB
       do_get_write_access             jbd2_log_do_checkpoint
        spin_lock(&jh->b_state_lock)
         if (buffer_dirty(bh))
          clear_buffer_dirty(bh)   // clear buffer dirty
           set_buffer_jbddirty(bh)
                                       transaction =
                                       journal->j_checkpoint_transactions
                                       jh = transaction->t_checkpoint_list
                                       if (!buffer_dirty(bh))
                                         __jbd2_journal_remove_checkpoint(jh)
                                         // bh won't be flushed
                                       jbd2_cleanup_journal_tail
        __jbd2_journal_file_buffer(jh, transaction, BJ_Reserved)
    4. Aborting journal/Power-cut before writing latest bh on journal area.

    In this way we get a corrupted filesystem with bh's data lost.

    Fix it by moving the clearing of buffer_dirty bit just before the call
    to __jbd2_journal_file_buffer(), both bit clearing and jh->b_transaction
    assignment are under journal->j_list_lock locked, so that
    jbd2_log_do_checkpoint() will wait until jh's new transaction fininshed
    even bh is currently not dirty. And journal_shrink_one_cp_list() won't
    remove jh from checkpoint list if the buffer head is reused in
    do_get_write_access().

    Fetch a reproducer in [Link].

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216898
    Cc: <stable@kernel.org>
    Signed-off-by: Zhihao Cheng <chengzhihao1@huawei.com>
    Signed-off-by: zhanchengbin <zhanchengbin1@huawei.com>
    Suggested-by: Jan Kara <jack@suse.cz>
    Reviewed-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20230110015327.1181863-1-chengzhihao1@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f500c21ae286d754e88137cd822f06e698a9a9f1
Author: Łukasz Stelmach <l.stelmach@samsung.com>
Date:   Thu Feb 23 08:47:48 2023 +0100

    ALSA: hda/realtek: Add quirk for HP EliteDesk 800 G6 Tower PC

    commit ea24b9953bcd3889f77a66e7f1d7e86e995dd9c3 upstream.

    HP EliteDesk 800 G6 Tower PC (103c:870c) requires a quirk for enabling
    headset-mic.

    Signed-off-by: Łukasz Stelmach <l.stelmach@samsung.com>
    Cc: <stable@vger.kernel.org>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=217008
    Link: https://lore.kernel.org/r/20230223074749.1026060-1-l.stelmach@samsung.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 536ee36a9610c2f54d75a748de289ab5e26a17c1
Author: Dmitry Fomin <fomindmitriyfoma@mail.ru>
Date:   Sat Feb 25 21:43:21 2023 +0300

    ALSA: ice1712: Do not left ice->gpio_mutex locked in aureon_add_controls()

    commit 951606a14a8901e3551fe4d8d3cedd73fe954ce1 upstream.

    If snd_ctl_add() fails in aureon_add_controls(), it immediately returns
    and leaves ice->gpio_mutex locked. ice->gpio_mutex locks in
    snd_ice1712_save_gpio_status and unlocks in
    snd_ice1712_restore_gpio_status(ice).

    It seems that the mutex is required only for aureon_cs8415_get(),
    so snd_ice1712_restore_gpio_status(ice) can be placed
    just after that. Compile tested only.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Dmitry Fomin <fomindmitriyfoma@mail.ru>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230225184322.6286-1-fomindmitriyfoma@mail.ru
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dca31ef86d3c2ef454b89a334159d60cebf5caa0
Author: andrew.yang <andrew.yang@mediatek.com>
Date:   Wed Feb 22 14:42:20 2023 +0800

    mm/damon/paddr: fix missing folio_put()

    commit 3f98c9a62c338bbe06a215c9491e6166ea39bf82 upstream.

    damon_get_folio() would always increase folio _refcount and
    folio_isolate_lru() would increase folio _refcount if the folio's lru flag
    is set.

    If an unevictable folio isolated successfully, there will be two more
    _refcount.  The one from folio_isolate_lru() will be decreased in
    folio_puback_lru(), but the other one from damon_get_folio() will be left
    behind.  This causes a pin page.

    Whatever the case, the _refcount from damon_get_folio() should be
    decreased.

    Link: https://lkml.kernel.org/r/20230222064223.6735-1-andrew.yang@mediatek.com
    Fixes: 57223ac29584 ("mm/damon/paddr: support the pageout scheme")
    Signed-off-by: andrew.yang <andrew.yang@mediatek.com>
    Reviewed-by: SeongJae Park <sj@kernel.org>
    Cc: <stable@vger.kernel.org>    [5.16.x]
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: SeongJae Park <sj@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2b1501f058245573a3aa6bf234d205dde1196184
Author: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
Date:   Wed Feb 1 15:59:44 2023 +0000

    crypto: qat - fix out-of-bounds read

    commit f6044cc3030e139f60c281386f28bda6e3049d66 upstream.

    When preparing an AER-CTR request, the driver copies the key provided by
    the user into a data structure that is accessible by the firmware.
```

```
    If the target device is QAT GEN4, the key size is rounded up by 16 since
    a rounded up size is expected by the device.
    If the key size is rounded up before the copy, the size used for copying
    the key might be bigger than the size of the region containing the key,
    causing an out-of-bounds read.

    Fix by doing the copy first and then update the keylen.

    This is to fix the following warning reported by KASAN:

            [  138.150574] BUG: KASAN: global-out-of-bounds in
qat_alg_skcipher_init_com.isra.0+0x197/0x250 [intel_qat]
            [  138.150641] Read of size 32 at addr ffffffff88c402c0 by task
cryptomgr_test/2340

            [  138.150651] CPU: 15 PID: 2340 Comm: cryptomgr_test Not tainted 6.2.0-rc1+ #45
            [  138.150659] Hardware name: Intel Corporation ArcherCity/ArcherCity, BIOS
EGSDCRB1.86B.0087.D13.2208261706 08/26/2022
            [  138.150663] Call Trace:
            [  138.150668]  <TASK>
            [  138.150922]  kasan_check_range+0x13a/0x1c0
            [  138.150931]  memcpy+0x1f/0x60
            [  138.150940]  qat_alg_skcipher_init_com.isra.0+0x197/0x250 [intel_qat]
            [  138.151006]  qat_alg_skcipher_init_sessions+0xc1/0x240 [intel_qat]
            [  138.151073]  crypto_skcipher_setkey+0x82/0x160
            [  138.151085]  ? prepare_keybuf+0xa2/0xd0
            [  138.151095]  test_skcipher_vec_cfg+0x2b8/0x800

    Fixes: 67916c951689 ("crypto: qat - add AES-CTR support for QAT GEN4 devices")
    Cc: <stable@vger.kernel.org>
    Reported-by: Vladis Dronov <vdronov@redhat.com>
    Signed-off-by: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
    Reviewed-by: Fiona Trahe <fiona.trahe@intel.com>
    Reviewed-by: Vladis Dronov <vdronov@redhat.com>
    Tested-by: Vladis Dronov <vdronov@redhat.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 79f4dbde3043fc27f5023b8801af9ab444368bbc
Author: Marc Zyngier <maz@kernel.org>
Date:   Mon Feb 13 11:42:49 2023 +0100

    irqdomain: Fix domain registration race

    commit 8932c32c3053accd50702b36e944ac2016cd103c upstream.

    Hierarchical domains created using irq_domain_create_hierarchy() are
    currently added to the domain list before having been fully initialised.

    This specifically means that a racing allocation request might fail to
    allocate irq data for the inner domains of a hierarchy in case the
    parent domain pointer has not yet been set up.

    Note that this is not really any issue for irqchip drivers that are
    registered early (e.g. via IRQCHIP_DECLARE() or IRQCHIP_ACPI_DECLARE())
    but could potentially cause trouble with drivers that are registered
    later (e.g. modular drivers using IRQCHIP_PLATFORM_DRIVER_BEGIN(),
    gpiochip drivers, etc.).

    Fixes: afb7da83b9f4 ("irqdomain: Introduce helper function irq_domain_add_hierarchy()")
    Cc: stable@vger.kernel.org      # 3.19
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    [ johan: add commit message ]
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
```

```
    Link: https://lore.kernel.org/r/20230213104302.17307-8-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a6e3d281c27944a88b44b29d3606ef0b594c557f
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:48 2023 +0100

    irqdomain: Fix mapping-creation race

    commit 601363cc08da25747feb87c55573dd54de91d66a upstream.

    Parallel probing of devices that share interrupts (e.g. when a driver
    uses asynchronous probing) can currently result in two mappings for the
    same hardware interrupt to be created due to missing serialisation.

    Make sure to hold the irq_domain_mutex when creating mappings so that
    looking for an existing mapping before creating a new one is done
    atomically.

    Fixes: 765230b5f084 ("driver-core: add asynchronous probing support for drivers")
    Fixes: b62b2cf5759b ("irqdomain: Fix handling of type settings for existing mappings")
    Link: https://lore.kernel.org/r/YuJXMHoT4ijUxnRb@hovoldconsulting.com
    Cc: stable@vger.kernel.org      # 4.8
    Cc: Dmitry Torokhov <dtor@chromium.org>
    Cc: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-7-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit adeaa9a9e0374208db69c5cfe98e1dd1f44b9f88
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:47 2023 +0100

    irqdomain: Refactor __irq_domain_alloc_irqs()

    commit d55f7f4c58c07beb5050a834bf57ae2ede599c7e upstream.

    Refactor __irq_domain_alloc_irqs() so that it can be called internally
    while holding the irq_domain_mutex.

    This will be used to fix a shared-interrupt mapping race, hence the
    Fixes tag.

    Fixes: b62b2cf5759b ("irqdomain: Fix handling of type settings for existing mappings")
    Cc: stable@vger.kernel.org      # 4.8
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-6-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b1f10867a10b774d3a524056a48f318808fe469d
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:45 2023 +0100

    irqdomain: Drop bogus fwspec-mapping error handling

    commit e3b7ab025e931accdc2c12acf9b75c6197f1c062 upstream.

    In case a newly allocated IRQ ever ends up not having any associated
    struct irq_data it would not even be possible to dispose the mapping.
```

```
    Replace the bogus disposal with a WARN_ON().

    This will also be used to fix a shared-interrupt mapping race, hence the
    CC-stable tag.

    Fixes: 1e2a7d78499e ("irqdomain: Don't set type when mapping an IRQ")
    Cc: stable@vger.kernel.org      # 4.8
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-4-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a9419361372175d7ad0102f61defec39690e2311
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:46 2023 +0100

    irqdomain: Look for existing mapping only once

    commit 6e6f75c9c98d2d246d90411ff2b6f0cd271f4cba upstream.

    Avoid looking for an existing mapping twice when creating a new mapping
    using irq_create_fwspec_mapping() by factoring out the actual allocation
    which is shared with irq_create_mapping_affinity().

    The new helper function will also be used to fix a shared-interrupt
    mapping race, hence the Fixes tag.

    Fixes: b62b2cf5759b ("irqdomain: Fix handling of type settings for existing mappings")
    Cc: stable@vger.kernel.org      # 4.8
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-5-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4c002a3c32c7fbf5c8ca4a488d53f25a536bb025
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:44 2023 +0100

    irqdomain: Fix disassociation race

    commit 3f883c38f5628f46b30bccf090faec054088e262 upstream.

    The global irq_domain_mutex is held when mapping interrupts from
    non-hierarchical domains but currently not when disposing them.

    This specifically means that updates of the domain mapcount is racy
    (currently only used for statistics in debugfs).

    Make sure to hold the global irq_domain_mutex also when disposing
    mappings from non-hierarchical domains.

    Fixes: 9dc6be3d4193 ("genirq/irqdomain: Add map counter")
    Cc: stable@vger.kernel.org      # 4.13
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-3-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit b364e1f44f38ed56c745ad0675899fb2b48b7762
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Feb 13 11:42:43 2023 +0100

    irqdomain: Fix association race

    commit b06730a571a9ff1ba5bd6b20bf9e50e5a12f1ec6 upstream.

    The sanity check for an already mapped virq is done outside of the
    irq_domain_mutex-protected section which means that an (unlikely) racing
    association may not be detected.

    Fix this by factoring out the association implementation, which will
    also be used in a follow-on change to fix a shared-interrupt mapping
    race.

    Fixes: ddaf144c61da ("irqdomain: Refactor irq_domain_associate_many()")
    Cc: stable@vger.kernel.org      # 3.11
    Tested-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230213104302.17307-2-johan+linaro@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 70ece74e138297944f4de52f064b585df170a334
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:46 2023 -0500

    selftests: seccomp: Fix incorrect kernel headers search path

    commit 07d42dd854446ba3177ad7a217870a5b4edee165 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9930b568dccf197a2cd32183905ebe4df195dff9
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:49 2023 -0500

    selftests: vm: Fix incorrect kernel headers search path

    commit 8eb3751c73bec746f61fb6bada60d1074d92b8c3 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
```

```
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3664b27db883e706649d181ff1d9c3e5039d5967
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:27 2023 -0500

    selftests: dmabuf-heaps: Fix incorrect kernel headers search path

    commit f80f09b59fdd45753dd80ac623981ad00ece4c2d upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c5f32e5cc8d759d92b1884fee076377dcb89e5c2
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:28 2023 -0500

    selftests: drivers: Fix incorrect kernel headers search path

    commit 07f0148aafe8c95a3a76cd59e9e75b4d78d1d31d upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bf5f0963a5736b27ec34fb41f9c6e6d0ea1128a1
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:30 2023 -0500

    selftests: futex: Fix incorrect kernel headers search path

    commit 24c55275ba0d538def2b1220002d0e808a85d50f upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit f1e57244c66f8b96a89af8a1f968fd8272ea2104
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:32 2023 -0500

    selftests: ipc: Fix incorrect kernel headers search path

    commit ecf9fdb5c2a9d63c732acccb6318feb73dd1589f upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2f905e7926267c31d17ba2df0af95d1834d029e6
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:39 2023 -0500

    selftests: perf_events: Fix incorrect kernel headers search path

    commit 465cbb1b9a9fd5f6907adb2d761facaf1a46bfbe upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e96c8d292a1e2c5490392d1fa76cfedad7bed03c
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:36 2023 -0500

    selftests: mount_setattr: Fix incorrect kernel headers search path

    commit 5d11f2d0eb39d2b5c5e8f05e1f650c4a4de69918 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 079509269354c056bcc2bb87fc88fdeca55dab2d
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
```

```
Date:   Fri Jan 27 08:57:37 2023 -0500

    selftests: move_mount_set_group: Fix incorrect kernel headers search path

    commit 65c68af0131bfef8e395c325735b6c40638cb931 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f14636c8baaeeeddbdcca138fbf79b340cc552ea
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:44 2023 -0500

    selftests: rseq: Fix incorrect kernel headers search path

    commit 2279bfc03211045c8f43a76b01889a5ca86acd5a upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1326589b1d3a30ad68a2cbe309bb90e5d635c968
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:47 2023 -0500

    selftests: sync: Fix incorrect kernel headers search path

    commit 5ad0c8e42c13623bd996e19ce76f2596e16eb0db upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit acaac31d8df314f7342b8e43e494deeb544a6d46
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:43 2023 -0500
```

```
    selftests: ptp: Fix incorrect kernel headers search path

    commit 01ede99e9de16e7a1ed689c99f41022aa878f2f4 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9671a5ffa753f628f2f12df10e00aa0413268565
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:48 2023 -0500

    selftests: user_events: Fix incorrect kernel headers search path

    commit f3886fd28987c119a98493f625cb9940b5f1c9a0 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d38e9b0a7b682bad1c56f678cf5b8aef8c029d28
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:29 2023 -0500

    selftests: filesystems: Fix incorrect kernel headers search path

    commit c2d3cf3653a8ff6e4b402d55e7f84790ac08a8ad upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 25bbf791f03128e2e78cd5c078492a84a2f57447
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:31 2023 -0500

    selftests: gpio: Fix incorrect kernel headers search path
```

```
    commit 8bb9c1808628babcc7b99ec2439bf102379bd4ac upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9d8a289cc0e7a5ac9dff354abb201f441ba1cc0d
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:34 2023 -0500

    selftests: media_tests: Fix incorrect kernel headers search path

    commit f2f9592b736087f695230410fb8dc1afd3cafbbb upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cbe62ad09bc58046635b831197a4c845dd81efb1
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:33 2023 -0500

    selftests: kcmp: Fix incorrect kernel headers search path

    commit 5d74231a2caad259f6669d8d6112814cef6bcd60 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1a7651cd32d1c9fa0e4bd437fedba9fa66c54df6
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:35 2023 -0500

    selftests: membarrier: Fix incorrect kernel headers search path

    commit 498bb027726371ba4a94686d251f9be1d437573e upstream.
```

```
    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 79ceed1e0341b10d9b56bc457eb273cb803fd34e
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:41 2023 -0500

    selftests: pidfd: Fix incorrect kernel headers search path

    commit 3f7d71768795c386019f2295c1986d00035c9f0f upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 81755c53491a270c7c770e670b860a07f5bdc833
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:24 2023 -0500

    selftests: clone3: Fix incorrect kernel headers search path

    commit 612cf4d283414a5ee2733db6608d917deb45fa46 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Acked-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 424cb47446d447373049bc2dd763ca42cb9653d5
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:22 2023 -0500

    selftests: arm64: Fix incorrect kernel headers search path

    commit 7482c19173b7eb044d476b3444d7ee55bc669d03 upstream.
```

```
    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Acked-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c742054509ef5885f7608055ab43b71ebc86bf3
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:40 2023 -0500

    selftests: pid_namespace: Fix incorrect kernel headers search path

    commit e81ff69f66969a16a98a2e0977c1860f1c182c74 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7b8a9d564cc2614961a86cc322d725f8a271e3f7
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:25 2023 -0500

    selftests: core: Fix incorrect kernel headers search path

    commit 145df2fdc38f24b3e52e4c2a59b02d874a074fbd upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org> # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e13dade91152b0cb95547a699e554fa2fe8cd6a4
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:45 2023 -0500

    selftests: sched: Fix incorrect kernel headers search path

    commit 0d2cace5af50806a6b32ab73d367b80e46acda0f upstream.
```

```
    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ea2808276d9b781763447a93e5ac6d0f44bffbbc
Author: Masami Hiramatsu (Google) <mhiramat@kernel.org>
Date:   Tue Feb 21 08:49:16 2023 +0900

    selftests/ftrace: Fix eprobe syntax test case to check filter support

    commit a457e944df92789ab31aaf35fae9db064e3c51c4 upstream.

    Fix eprobe syntax test case to check whether the kernel supports the filter
    on eprobe for filter syntax test command. Without this fix, this test case
    will fail if the kernel supports eprobe but doesn't support the filter on
    eprobe.

    Link: https://lore.kernel.org/all/167309834742.640500.379128668288448035.stgit@devnote3/

    Fixes: 9e14bae7d049 ("selftests/ftrace: Add eprobe syntax error testcase")
    Cc: stable@vger.kernel.org
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Reviewed-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 03379280558e6d57322d308ec71efb40e0975ab1
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:42 2023 -0500

    selftests/powerpc: Fix incorrect kernel headers search path

    commit 4f11410bf6da87defe8fd59b0413f0d9f71744da upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Cc: stable@vger.kernel.org # v5.18+
    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20230127135755.79929-22-mathieu.desnoyers@efficios.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eb57afe405fd3010a2bad8abd0d37b3a4ab1c59c
Author: Roberto Sassu <roberto.sassu@huawei.com>
Date:   Tue Jan 31 18:42:43 2023 +0100

    ima: Align ima_file_mmap() parameters with mmap_file LSM hook

    commit 4971c268b85e1c7a734a61622fc0813c86e2362e upstream.

    Commit 98de59bfe4b2f ("take calculation of final prot in
```

```
    security_mmap_file() into a helper") moved the code to update prot, to be
    the actual protections applied to the kernel, to a new helper called
    mmap_prot().

    However, while without the helper ima_file_mmap() was getting the updated
    prot, with the helper ima_file_mmap() gets the original prot, which
    contains the protections requested by the application.

    A possible consequence of this change is that, if an application calls
    mmap() with only PROT_READ, and the kernel applies PROT_EXEC in addition,
    that application would have access to executable memory without having this
    event recorded in the IMA measurement list. This situation would occur for
    example if the application, before mmap(), calls the personality() system
    call with READ_IMPLIES_EXEC as the first argument.

    Align ima_file_mmap() parameters with those of the mmap_file LSM hook, so
    that IMA can receive both the requested prot and the final prot. Since the
    requested protections are stored in a new variable, and the final
    protections are stored in the existing variable, this effectively restores
    the original behavior of the MMAP_CHECK hook.

    Cc: stable@vger.kernel.org
    Fixes: 98de59bfe4b2 ("take calculation of final prot in security_mmap_file() into a
helper")
    Signed-off-by: Roberto Sassu <roberto.sassu@huawei.com>
    Reviewed-by: Stefan Berger <stefanb@linux.ibm.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 51bf199731929ccd2b78f503bf2957556f0d0a1d
Author: Matt Bobrowski <mattbobrowski@google.com>
Date:   Wed Jan 4 03:41:44 2023 +0000

    ima: fix error handling logic when file measurement failed

    commit 6dc387d52eb67f45d68caa263704fa4e39ef8e76 upstream.

    Restore the error handling logic so that when file measurement fails,
    the respective iint entry is not left with the digest data being
    populated with zeroes.

    Fixes: 54f03916fb89 ("ima: permit fsverity's file digests in the IMA measurement list")
    Cc: stable@vger.kernel.org      # 5.19
    Signed-off-by: Matt Bobrowski <mattbobrowski@google.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 57a0078b1a07dff4397744d911fc89f1d23b117e
Author: Jens Axboe <axboe@kernel.dk>
Date:   Thu Feb 16 08:01:08 2023 -0700

    brd: check for REQ_NOWAIT and set correct page allocation mask

    commit 6ded703c56c21bfb259725d4f1831a5feb563e9b upstream.

    If REQ_NOWAIT is set, then do a non-blocking allocation if the operation
    is a write and we need to insert a new page. Currently REQ_NOWAIT cannot
    be set as the queue isn't marked as supporting nowait, this change is in
    preparation for allowing that.

    radix_tree_preload() warns on attempting to call it with an allocation
    mask that doesn't allow blocking. While that warning could arguably
    be removed, we need to handle radix insertion failures anyway as they
    are more likely if we cannot block to get memory.
```

```
    Remove legacy BUG_ON()'s and turn them into proper errors instead, one
    for the allocation failure and one for finding a page that doesn't
    match the correct index.

    Cc: stable@vger.kernel.org # 5.10+
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 87f1bf26501a51aeabc7aeb55b92afc550260d61
Author: Jens Axboe <axboe@kernel.dk>
Date:   Thu Feb 16 07:57:32 2023 -0700

    brd: return 0/-error from brd_insert_page()

    commit db0ccc44a20b4bb3039c0f6885a1f9c3323c7673 upstream.

    It currently returns a page, but callers just check for NULL/page to
    gauge success. Clean this up and return the appropriate error directly
    instead.

    Cc: stable@vger.kernel.org # 5.10+
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c3435c7e7e06d8dea390c63149d5f0447f640127
Author: Jens Axboe <axboe@kernel.dk>
Date:   Wed Feb 15 16:43:47 2023 -0700

    brd: mark as nowait compatible

    commit 67205f80be9910207481406c47f7d85e703fb2e9 upstream.

    By default, non-mq drivers do not support nowait. This causes io_uring
    to use a slower path as the driver cannot be trust not to block. brd
    can safely set the nowait flag, as worst case all it does is a NOIO
    allocation.

    For io_uring, this makes a substantial difference. Before:

    submitter=0, tid=453, file=/dev/ram0, node=-1
    polled=0, fixedbufs=1/0, register_files=1, buffered=0, QD=128
    Engine=io_uring, sq_ring=128, cq_ring=128
    IOPS=440.03K, BW=1718MiB/s, IOS/call=32/31
    IOPS=428.96K, BW=1675MiB/s, IOS/call=32/32
    IOPS=442.59K, BW=1728MiB/s, IOS/call=32/31
    IOPS=419.65K, BW=1639MiB/s, IOS/call=32/32
    IOPS=426.82K, BW=1667MiB/s, IOS/call=32/31

    and after:

    submitter=0, tid=354, file=/dev/ram0, node=-1
    polled=0, fixedbufs=1/0, register_files=1, buffered=0, QD=128
    Engine=io_uring, sq_ring=128, cq_ring=128
    IOPS=3.37M, BW=13.15GiB/s, IOS/call=32/31
    IOPS=3.45M, BW=13.46GiB/s, IOS/call=32/31
    IOPS=3.43M, BW=13.42GiB/s, IOS/call=32/32
    IOPS=3.43M, BW=13.39GiB/s, IOS/call=32/31
    IOPS=3.43M, BW=13.38GiB/s, IOS/call=32/31

    or about an 8x in difference. Now that brd is prepared to deal with
    REQ_NOWAIT reads/writes, mark it as supporting that.

    Cc: stable@vger.kernel.org # 5.10+
```

```
    Link: https://lore.kernel.org/linux-block/20230203103005.31290-1-p.raghav@samsung.com/
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 36261d4eacaa6de2cd705e69468571ec78082491
Author: Tom Lendacky <thomas.lendacky@amd.com>
Date:   Wed Feb 22 10:39:39 2023 -0600

    virt/sev-guest: Return -EIO if certificate buffer is not large enough

    commit dd093fb08e8f8a958fec4eef36f9f09eac047f60 upstream.

    Commit

      47894e0fa6a5 ("virt/sev-guest: Prevent IV reuse in the SNP guest driver")

    changed the behavior associated with the return value when the caller
    does not supply a large enough certificate buffer. Prior to the commit a
    value of -EIO was returned. Now, 0 is returned.  This breaks the
    established ABI with the user.

    Change the code to detect the buffer size error and return -EIO.

    Fixes: 47894e0fa6a5 ("virt/sev-guest: Prevent IV reuse in the SNP guest driver")
    Reported-by: Larry Dewey <larry.dewey@amd.com>
    Signed-off-by: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Tested-by: Larry Dewey <larry.dewey@amd.com>
    Cc: <stable@kernel.org>
    Link:
https://lore.kernel.org/r/2afbcae6daf13f7ad5a4296692e0a0fe1bc1e4ee.1677083979.git.thomas.len
dacky@amd.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b3add485357e03bb8ea339d29aa0af111166b5aa
Author: KP Singh <kpsingh@kernel.org>
Date:   Mon Feb 27 07:05:41 2023 +0100

    Documentation/hw-vuln: Document the interaction between IBRS and STIBP

    commit e02b50ca442e88122e1302d4dbc1b71a4808c13f upstream.

    Explain why STIBP is needed with legacy IBRS as currently implemented
    (KERNEL_IBRS) and why STIBP is not needed when enhanced IBRS is enabled.

    Fixes: 7c693f54c873 ("x86/speculation: Add spectre_v2=ibrs option to support Kernel
IBRS")
    Signed-off-by: KP Singh <kpsingh@kernel.org>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230227060541.1939092-2-kpsingh@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ead3c8e54d28fa1d5454b1f8a21b96b4a969b1cb
Author: KP Singh <kpsingh@kernel.org>
Date:   Mon Feb 27 07:05:40 2023 +0100

    x86/speculation: Allow enabling STIBP with legacy IBRS

    commit 6921ed9049bc7457f66c1596c5b78aec0dae4a9d upstream.

    When plain IBRS is enabled (not enhanced IBRS), the logic in
    spectre_v2_user_select_mitigation() determines that STIBP is not needed.

    The IBRS bit implicitly protects against cross-thread branch target
    injection. However, with legacy IBRS, the IBRS bit is cleared on
```

```
    returning to userspace for performance reasons which leaves userspace
    threads vulnerable to cross-thread branch target injection against which
    STIBP protects.

    Exclude IBRS from the spectre_v2_in_ibrs_mode() check to allow for
    enabling STIBP (through seccomp/prctl() by default or always-on, if
    selected by spectre_v2_user kernel cmdline parameter).

      [ bp: Massage. ]

    Fixes: 7c693f54c873 ("x86/speculation: Add spectre_v2=ibrs option to support Kernel
IBRS")
    Reported-by: José Oliveira <joseloliveira11@gmail.com>
    Reported-by: Rodrigo Branco <rodrigo@kernelhacking.com>
    Signed-off-by: KP Singh <kpsingh@kernel.org>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230220120127.1975241-1-kpsingh@kernel.org
    Link: https://lore.kernel.org/r/20230221184908.2349578-1-kpsingh@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eaf5dea1eb8c2928554b3ca717575cbe232b843c
Author: Borislav Petkov (AMD) <bp@alien8.de>
Date:   Thu Jan 26 16:26:17 2023 +0100

    x86/microcode/AMD: Fix mixed steppings support

    commit 7ff6edf4fef38ab404ee7861f257e28eaaeed35f upstream.

    The AMD side of the loader has always claimed to support mixed
    steppings. But somewhere along the way, it broke that by assuming that
    the cached patch blob is a single one instead of it being one per
    *node*.

    So turn it into a per-node one so that each node can stash the blob
    relevant for it.

      [ NB: Fixes tag is not really the exactly correct one but it is good
        enough. ]

    Fixes: fe055896c040 ("x86/microcode: Merge the early microcode loader")
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Cc: <stable@kernel.org> # 2355370cd941 ("x86/microcode/amd: Remove
load_microcode_amd()'s bsp parameter")
    Cc: <stable@kernel.org> # a5ad92134bd1 ("x86/microcode/AMD: Add a @cpu parameter to the
reloading functions")
    Link: https://lore.kernel.org/r/20230130161709.11615-4-bp@alien8.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4515c059e2690bd3f50960373c88edd83832194b
Author: Borislav Petkov (AMD) <bp@alien8.de>
Date:   Thu Jan 26 00:08:03 2023 +0100

    x86/microcode/AMD: Add a @cpu parameter to the reloading functions

    commit a5ad92134bd153a9ccdcddf09a95b088f36c3cce upstream.

    Will be used in a subsequent change.

    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230130161709.11615-3-bp@alien8.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 034cf5ae1863ea1f3cf5d6c3a86ecdc61b976f4f
Author: Borislav Petkov (AMD) <bp@alien8.de>
```

```
Date:   Tue Jan 17 23:59:24 2023 +0100

    x86/microcode/amd: Remove load_microcode_amd()'s bsp parameter

    commit 2355370cd941cbb20882cc3f34460f9f2b8f9a18 upstream.

    It is always the BSP.

    No functional changes.

    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230130161709.11615-2-bp@alien8.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit edcbb2800a36500ec18656f41685ed23c535c214
Author: Yang Jihong <yangjihong1@huawei.com>
Date:   Tue Feb 21 08:49:16 2023 +0900

    x86/kprobes: Fix arch_check_optimized_kprobe check within optimized_kprobe range

    commit f1c97a1b4ef709e3f066f82e3ba3108c3b133ae6 upstream.

    When arch_prepare_optimized_kprobe calculating jump destination address,
    it copies original instructions from jmp-optimized kprobe (see
    __recover_optprobed_insn), and calculated based on length of original
    instruction.

    arch_check_optimized_kprobe does not check KPROBE_FLAG_OPTIMATED when
    checking whether jmp-optimized kprobe exists.
    As a result, setup_detour_execution may jump to a range that has been
    overwritten by jump destination address, resulting in an inval opcode error.

    For example, assume that register two kprobes whose addresses are
    <func+9> and <func+11> in "func" function.
    The original code of "func" function is as follows:

       0xffffffff816cb5e9 <+9>:     push   %r12
       0xffffffff816cb5eb <+11>:    xor    %r12d,%r12d
       0xffffffff816cb5ee <+14>:    test   %rdi,%rdi
       0xffffffff816cb5f1 <+17>:    setne  %r12b
       0xffffffff816cb5f5 <+21>:    push   %rbp

    1.Register the kprobe for <func+11>, assume that is kp1, corresponding optimized_kprobe
is op1.
      After the optimization, "func" code changes to:

       0xffffffff816cc079 <+9>:     push   %r12
       0xffffffff816cc07b <+11>:    jmp    0xffffffffa0210000
       0xffffffff816cc080 <+16>:    incl   0xf(%rcx)
       0xffffffff816cc083 <+19>:    xchg   %eax,%ebp
       0xffffffff816cc084 <+20>:    (bad)
       0xffffffff816cc085 <+21>:    push   %rbp

    Now op1->flags == KPROBE_FLAG_OPTIMATED;

    2. Register the kprobe for <func+9>, assume that is kp2, corresponding optimized_kprobe
is op2.

    register_kprobe(kp2)
      register_aggr_kprobe
        alloc_aggr_kprobe
          __prepare_optimized_kprobe
            arch_prepare_optimized_kprobe
              __recover_optprobed_insn    // copy original bytes from kp1-
>optinsn.copied_insn,
```

```
                                        // jump address = <func+14>

    3. disable kp1:

    disable_kprobe(kp1)
      __disable_kprobe
        ...
        if (p == orig_p || aggr_kprobe_disabled(orig_p)) {
          ret = disarm_kprobe(orig_p, true)       // add op1 in unoptimizing_list, not
unoptimized
          orig_p->flags |= KPROBE_FLAG_DISABLED;  // op1->flags ==  KPROBE_FLAG_OPTIMATED |
KPROBE_FLAG_DISABLED
        ...

    4. unregister kp2
    __unregister_kprobe_top
      ...
      if (!kprobe_disabled(ap) && !kprobes_all_disarmed) {
        optimize_kprobe(op)
          ...
          if (arch_check_optimized_kprobe(op) < 0) // because op1 has KPROBE_FLAG_DISABLED,
here not return
            return;
          p->kp.flags |= KPROBE_FLAG_OPTIMIZED;   //  now op2 has KPROBE_FLAG_OPTIMIZED
      }

    "func" code now is:

       0xffffffff816cc079 <+9>:     int3
       0xffffffff816cc07a <+10>:    push   %rsp
       0xffffffff816cc07b <+11>:    jmp    0xffffffffa0210000
       0xffffffff816cc080 <+16>:    incl   0xf(%rcx)
       0xffffffff816cc083 <+19>:    xchg   %eax,%ebp
       0xffffffff816cc084 <+20>:    (bad)
       0xffffffff816cc085 <+21>:    push   %rbp

    5. if call "func", int3 handler call setup_detour_execution:

      if (p->flags & KPROBE_FLAG_OPTIMIZED) {
        ...
        regs->ip = (unsigned long)op->optinsn.insn + TMPL_END_IDX;
        ...
      }

    The code for the destination address is

       0xffffffffa021072c:  push   %r12
       0xffffffffa021072e:  xor    %r12d,%r12d
       0xffffffffa0210731:  jmp    0xffffffff816cb5ee <func+14>

    However, <func+14> is not a valid start instruction address. As a result, an error
occurs.

    Link: https://lore.kernel.org/all/20230216034247.32348-3-yangjihong1@huawei.com/

    Fixes: f66c0447cca1 ("kprobes: Set unoptimized flag after unoptimizing code")
    Signed-off-by: Yang Jihong <yangjihong1@huawei.com>
    Cc: stable@vger.kernel.org
    Acked-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ce84d9983273de0e09267318431c8287f6b4a3a6
Author: Yang Jihong <yangjihong1@huawei.com>
Date:   Tue Feb 21 08:49:16 2023 +0900
```

```
    x86/kprobes: Fix __recover_optprobed_insn check optimizing logic

    commit 868a6fc0ca2407622d2833adefe1c4d284766c4c upstream.

    Since the following commit:

      commit f66c0447cca1 ("kprobes: Set unoptimized flag after unoptimizing code")

    modified the update timing of the KPROBE_FLAG_OPTIMIZED, a optimized_kprobe
    may be in the optimizing or unoptimizing state when op.kp->flags
    has KPROBE_FLAG_OPTIMIZED and op->list is not empty.

    The __recover_optprobed_insn check logic is incorrect, a kprobe in the
    unoptimizing state may be incorrectly determined as unoptimizing.
    As a result, incorrect instructions are copied.

    The optprobe_queued_unopt function needs to be exported for invoking in
    arch directory.

    Link: https://lore.kernel.org/all/20230216034247.32348-2-yangjihong1@huawei.com/

    Fixes: f66c0447cca1 ("kprobes: Set unoptimized flag after unoptimizing code")
    Cc: stable@vger.kernel.org
    Signed-off-by: Yang Jihong <yangjihong1@huawei.com>
    Acked-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1698fda30976dd6eaae6fe1a71662e70cb31aa8a
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Nov 30 23:36:50 2022 +0000

    x86/reboot: Disable SVM, not just VMX, when stopping CPUs

    commit a2b07fa7b93321c059af0c6d492cc9a4f1e390aa upstream.

    Disable SVM and more importantly force GIF=1 when halting a CPU or
    rebooting the machine.  Similar to VMX, SVM allows software to block
    INITs via CLGI, and thus can be problematic for a crash/reboot.  The
    window for failure is smaller with SVM as INIT is only blocked while
    GIF=0, i.e. between CLGI and STGI, but the window does exist.

    Fixes: fba4f472b33a ("x86/reboot: Turn off KVM when halting a CPU")
    Cc: stable@vger.kernel.org
    Reviewed-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lore.kernel.org/r/20221130233650.1404148-5-seanjc@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3a4afb752482dcbd4a54ead94611147d4d7e505c
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Nov 30 23:36:49 2022 +0000

    x86/reboot: Disable virtualization in an emergency if SVM is supported

    commit d81f952aa657b76cea381384bef1fea35c5fd266 upstream.

    Disable SVM on all CPUs via NMI shootdown during an emergency reboot.
    Like VMX, SVM can block INIT, e.g. if the emergency reboot is triggered
    between CLGI and STGI, and thus can prevent bringing up other CPUs via
    INIT-SIPI-SIPI.

    Cc: stable@vger.kernel.org
    Reviewed-by: Thomas Gleixner <tglx@linutronix.de>
```

```
    Link: https://lore.kernel.org/r/20221130233650.1404148-4-seanjc@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b01c0d04d3bceb14717d43ea548426c060cfb1ad
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Nov 30 23:36:47 2022 +0000

    x86/crash: Disable virt in core NMI crash handler to avoid double shootdown

    commit 26044aff37a5455b19a91785086914fd33053ef4 upstream.

    Disable virtualization in crash_nmi_callback() and rework the
    emergency_vmx_disable_all() path to do an NMI shootdown if and only if a
    shootdown has not already occurred.   NMI crash shootdown fundamentally
    can't support multiple invocations as responding CPUs are deliberately
    put into halt state without unblocking NMIs.  But, the emergency reboot
    path doesn't have any work of its own, it simply cares about disabling
    virtualization, i.e. so long as a shootdown occurred, emergency reboot
    doesn't care who initiated the shootdown, or when.

    If "crash_kexec_post_notifiers" is specified on the kernel command line,
    panic() will invoke crash_smp_send_stop() and result in a second call to
    nmi_shootdown_cpus() during native_machine_emergency_restart().

    Invoke the callback _before_ disabling virtualization, as the current
    VMCS needs to be cleared before doing VMXOFF.  Note, this results in a
    subtle change in ordering between disabling virtualization and stopping
    Intel PT on the responding CPUs.  While VMX and Intel PT do interact,
    VMXOFF and writes to MSR_IA32_RTIT_CTL do not induce faults between one
    another, which is all that matters when panicking.

    Harden nmi_shootdown_cpus() against multiple invocations to try and
    capture any such kernel bugs via a WARN instead of hanging the system
    during a crash/dump, e.g. prior to the recent hardening of
    register_nmi_handler(), re-registering the NMI handler would trigger a
    double list_add() and hang the system if CONFIG_BUG_ON_DATA_CORRUPTION=y.

     list_add double add: new=ffffffff82220800, prev=ffffffff8221cfe8,
next=ffffffff82220800.
     WARNING: CPU: 2 PID: 1319 at lib/list_debug.c:29 __list_add_valid+0x67/0x70
     Call Trace:
      __register_nmi_handler+0xcf/0x130
      nmi_shootdown_cpus+0x39/0x90
      native_machine_emergency_restart+0x1c9/0x1d0
      panic+0x237/0x29b

    Extract the disabling logic to a common helper to deduplicate code, and
    to prepare for doing the shootdown in the emergency reboot path if SVM
    is supported.

    Note, prior to commit ed72736183c4 ("x86/reboot: Force all cpus to exit
    VMX root if VMX is supported"), nmi_shootdown_cpus() was subtly protected
    against a second invocation by a cpu_vmx_enabled() check as the kdump
    handler would disable VMX if it ran first.

    Fixes: ed72736183c4 ("x86/reboot: Force all cpus to exit VMX root if VMX is supported")
    Cc: stable@vger.kernel.org
    Reported-by: Guilherme G. Piccoli <gpiccoli@igalia.com>
    Cc: Vitaly Kuznetsov <vkuznets@redhat.com>
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Link: https://lore.kernel.org/all/20220427224924.592546-2-gpiccoli@igalia.com
    Tested-by: Guilherme G. Piccoli <gpiccoli@igalia.com>
    Reviewed-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lore.kernel.org/r/20221130233650.1404148-2-seanjc@google.com
```

```
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 28b136dbf10c39867ca1fa5b7c8c75fb2cdaf087
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Nov 30 23:36:48 2022 +0000

    x86/virt: Force GIF=1 prior to disabling SVM (for reboot flows)

    commit 6a3236580b0b1accc3976345e723104f74f6f8e6 upstream.

    Set GIF=1 prior to disabling SVM to ensure that INIT is recognized if the
    kernel is disabling SVM in an emergency, e.g. if the kernel is about to
    jump into a crash kernel or may reboot without doing a full CPU RESET.
    If GIF is left cleared, the new kernel (or firmware) will be unabled to
    awaken APs.  Eat faults on STGI (due to EFER.SVME=0) as it's possible
    that SVM could be disabled via NMI shootdown between reading EFER.SVME
    and executing STGI.

    Link: https://lore.kernel.org/all/cbcb6f35-e5d7-c1c9-4db9-fe5cc4de579a@amd.com
    Cc: stable@vger.kernel.org
    Cc: Andrew Cooper <Andrew.Cooper3@citrix.com>
    Cc: Tom Lendacky <thomas.lendacky@amd.com>
    Reviewed-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lore.kernel.org/r/20221130233650.1404148-3-seanjc@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 52b6f404674749e69d7564bbeaaae5541ee9d3ac
Author: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
Date:   Fri Jan 27 08:57:50 2023 -0500

    selftests: x86: Fix incorrect kernel headers search path

    commit ac5ec90e94fe8eddb4499e51398640fa6a89d657 upstream.

    Use $(KHDR_INCLUDES) as lookup path for kernel headers. This prevents
    building against kernel headers from the build environment in scenarios
    where kernel headers are installed into a specific output directory
    (O=...).

    Signed-off-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: linux-kselftest@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: <stable@vger.kernel.org>  # 5.18+
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 59dc47c1516aa72261d0c8ad1df587ab1d251ec7
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Tue Feb 21 23:33:15 2023 -0800

    KVM: SVM: hyper-v: placate modpost section mismatch error

    commit 45dd9bc75d9adc9483f0c7d662ba6e73ed698a0b upstream.

    modpost reports section mismatch errors/warnings:
    WARNING: modpost: vmlinux.o: section mismatch in reference: svm_hv_hardware_setup
(section: .text) -> (unknown) (section: .init.data)
    WARNING: modpost: vmlinux.o: section mismatch in reference: svm_hv_hardware_setup
(section: .text) -> (unknown) (section: .init.data)
    WARNING: modpost: vmlinux.o: section mismatch in reference: svm_hv_hardware_setup
(section: .text) -> (unknown) (section: .init.data)
```

```
    This "(unknown) (section: .init.data)" all refer to svm_x86_ops.

    Tag svm_hv_hardware_setup() with __init to fix a modpost warning as the
    non-stub implementation accesses __initdata (svm_x86_ops), i.e. would
    generate a use-after-free if svm_hv_hardware_setup() were actually invoked
    post-init.  The helper is only called from svm_hardware_setup(), which is
    also __init, i.e. lack of __init is benign other than the modpost warning.

    Fixes: 1e0c7d40758b ("KVM: SVM: hyper-v: Remote TLB flush for SVM")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Vineeth Pillai <viremana@linux.microsoft.com>
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Cc: kvm@vger.kernel.org
    Cc: stable@vger.kernel.org
    Reviewed-by: Sean Christopherson <seanjc@google.com>
    Reviewed-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20230222073315.9081-1-rdunlap@infradead.org>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 57e604e2e92e336df597cb7cf9445d426ac47921
Author: Peter Gonda <pgonda@google.com>
Date:   Tue Feb 7 09:13:54 2023 -0800

    KVM: SVM: Fix potential overflow in SEV's send|receive_update_data()

    commit f94f053aa3a5d6ff17951870483d9eb9e13de2e2 upstream.

    KVM_SEV_SEND_UPDATE_DATA and KVM_SEV_RECEIVE_UPDATE_DATA have an integer
    overflow issue. Params.guest_len and offset are both 32 bits wide, with a
    large params.guest_len the check to confirm a page boundary is not
    crossed can falsely pass:

        /* Check if we are crossing the page boundary *
        offset = params.guest_uaddr & (PAGE_SIZE - 1);
        if ((params.guest_len + offset > PAGE_SIZE))

    Add an additional check to confirm that params.guest_len itself is not
    greater than PAGE_SIZE.

    Note, this isn't a security concern as overflow can happen if and only if
    params.guest_len is greater than 0xfffff000, and the FW spec says these
    commands fail with lengths greater than 16KB, i.e. the PSP will detect
    KVM's goof.

    Fixes: 15fb7de1a7f5 ("KVM: SVM: Add KVM_SEV_RECEIVE_UPDATE_DATA command")
    Fixes: d3d1af85e2c7 ("KVM: SVM: Add KVM_SEND_UPDATE_DATA command")
    Reported-by: Andy Nguyen <theflow@google.com>
    Suggested-by: Thomas Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Peter Gonda <pgonda@google.com>
    Cc: David Rientjes <rientjes@google.com>
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Cc: Sean Christopherson <seanjc@google.com>
    Cc: kvm@vger.kernel.org
    Cc: stable@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Reviewed-by: Tom Lendacky <thomas.lendacky@amd.com>
    Link: https://lore.kernel.org/r/20230207171354.4012821-1-pgonda@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bbbd67a327df41e1284461fefa0958fbc6f6d42b
Author: Sean Christopherson <seanjc@google.com>
Date:   Sat Jan 7 01:10:21 2023 +0000
```

```
    KVM: x86: Inject #GP on x2APIC WRMSR that sets reserved bits 63:32

    commit ab52be1b310bcb39e6745d34a8f0e8475d67381a upstream.

    Reject attempts to set bits 63:32 for 32-bit x2APIC registers, i.e. all
    x2APIC registers except ICR.  Per Intel's SDM:

      Non-zero writes (by WRMSR instruction) to reserved bits to these
      registers will raise a general protection fault exception

    Opportunistically fix a typo in a nearby comment.

    Reported-by: Marc Orr <marcorr@google.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20230107011025.565472-3-seanjc@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4e01634d74511cc07502c3819ef3adaeee162227
Author: Sean Christopherson <seanjc@google.com>
Date:   Sat Jan 7 01:10:20 2023 +0000

    KVM: x86: Inject #GP if WRMSR sets reserved bits in APIC Self-IPI

    commit ba5838abb05334e4abfdff1490585c7f365e0424 upstream.

    Inject a #GP if the guest attempts to set reserved bits in the x2APIC-only
    Self-IPI register.  Bits 7:0 hold the vector, all other bits are reserved.

    Reported-by: Marc Orr <marcorr@google.com>
    Cc: Ben Gardon <bgardon@google.com>
    Cc: Venkatesh Srinivas <venkateshs@chromium.org>
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20230107011025.565472-2-seanjc@google.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e943f98967ee6d4cc8b1a0351a8f07aef12bd7ba
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:40 2023 +0000

    KVM: SVM: Don't put/load AVIC when setting virtual APIC mode

    commit e0bead97e7590da888148feb9e9133bc278c534b upstream.

    Move the VMCB updates from avic_refresh_apicv_exec_ctrl() into
    avic_set_virtual_apic_mode() and invert the dependency being said
    functions to avoid calling avic_vcpu_{load,put}() and
    avic_set_pi_irte_mode() when "only" setting the virtual APIC mode.

    avic_set_virtual_apic_mode() is invoked from common x86 with preemption
    enabled, which makes avic_vcpu_{load,put}() unhappy.  Luckily, calling
    those and updating IRTE stuff is unnecessary as the only reason
    avic_set_virtual_apic_mode() is called is to handle transitions between
    xAPIC and x2APIC that don't also toggle APICv activation.  And if
    activation doesn't change, there's no need to fiddle with the physical
    APIC ID table or update IRTE.

    The "full" refresh is guaranteed to be called if activation changes in
    this case as the only call to the "set" path is:

            kvm_vcpu_update_apicv(vcpu);
            static_call_cond(kvm_x86_set_virtual_apic_mode)(vcpu);
```

```
    and kvm_vcpu_update_apicv() invokes the refresh if activation changes:

            if (apic->apicv_active == activate)
                    goto out;

            apic->apicv_active = activate;
            kvm_apic_update_apicv(vcpu);
            static_call(kvm_x86_refresh_apicv_exec_ctrl)(vcpu);

    Rename the helper to reflect that it is also called during "refresh".

      WARNING: CPU: 183 PID: 49186 at arch/x86/kvm/svm/avic.c:1081 avic_vcpu_put+0xde/0xf0
[kvm_amd]
      CPU: 183 PID: 49186 Comm: stable Tainted: G           O       6.0.0-smp--fcddbca45f0a-
sink #34
      Hardware name: Google, Inc. Arcadia_IT_80/Arcadia_IT_80, BIOS 10.48.0 01/27/2022
      RIP: 0010:avic_vcpu_put+0xde/0xf0 [kvm_amd]
       avic_refresh_apicv_exec_ctrl+0x142/0x1c0 [kvm_amd]
       avic_set_virtual_apic_mode+0x5a/0x70 [kvm_amd]
       kvm_lapic_set_base+0x149/0x1a0 [kvm]
       kvm_set_apic_base+0x8f/0xd0 [kvm]
       kvm_set_msr_common+0xa3a/0xdc0 [kvm]
       svm_set_msr+0x364/0x6b0 [kvm_amd]
       __kvm_set_msr+0xb8/0x1c0 [kvm]
       kvm_emulate_wrmsr+0x58/0x1d0 [kvm]
       msr_interception+0x1c/0x30 [kvm_amd]
       svm_invoke_exit_handler+0x31/0x100 [kvm_amd]
       svm_handle_exit+0xfc/0x160 [kvm_amd]
       vcpu_enter_guest+0x21bb/0x23e0 [kvm]
       vcpu_run+0x92/0x450 [kvm]
       kvm_arch_vcpu_ioctl_run+0x43e/0x6e0 [kvm]
       kvm_vcpu_ioctl+0x559/0x620 [kvm]

    Fixes: 05c4fe8c1bd9 ("KVM: SVM: Refresh AVIC configuration when changing APIC mode")
    Cc: stable@vger.kernel.org
    Cc: Suravee Suthikulpanit <suravee.suthikulpanit@amd.com>
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-8-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ac46aca8d1a6690924310f71e4270fc56273abb7
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:37 2023 +0000

    KVM: SVM: Process ICR on AVIC IPI delivery failure due to invalid target

    commit 5aede752a839904059c2b5d68be0dc4501c6c15f upstream.

    Emulate ICR writes on AVIC IPI failures due to invalid targets using the
    same logic as failures due to invalid types.  AVIC acceleration fails if
    _any_ of the targets are invalid, and crucially VM-Exits before sending
    IPIs to targets that _are_ valid.  In logical mode, the destination is a
    bitmap, i.e. a single IPI can target multiple logical IDs.  Doing nothing
    causes KVM to drop IPIs if at least one target is valid and at least one
    target is invalid.

    Fixes: 18f40c53e10f ("svm: Add VMEXIT handlers for AVIC")
    Cc: stable@vger.kernel.org
    Reviewed-by: Paolo Bonzini <pbonzini@redhat.com>
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-5-seanjc@google.com>
```

```
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 164e2971ad31d8ebb452408816c01d8716a2389c
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:36 2023 +0000

    KVM: SVM: Flush the "current" TLB when activating AVIC

    commit 0ccf3e7cb95a2db8ddb2a44812037ffba8166dc9 upstream.

    Flush the TLB when activating AVIC as the CPU can insert into the TLB
    while AVIC is "locally" disabled.  KVM doesn't treat "APIC hardware
    disabled" as VM-wide AVIC inhibition, and so when a vCPU has its APIC
    hardware disabled, AVIC is not guaranteed to be inhibited.  As a result,
    KVM may create a valid NPT mapping for the APIC base, which the CPU can
    cache as a non-AVIC translation.

    Note, Intel handles this in vmx_set_virtual_apic_mode().

    Reviewed-by: Paolo Bonzini <pbonzini@redhat.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Message-Id: <20230106011306.85230-4-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c60d9137974bcc5be9831f9a378d7912db8d1216
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:39 2023 +0000

    KVM: x86: Don't inhibit APICv/AVIC if xAPIC ID mismatch is due to 32-bit ID

    commit f651a008954803d7bb2d85b7042d0fd46133d782 upstream.

    Truncate the vcpu_id, a.k.a. x2APIC ID, to an 8-bit value when comparing
    it against the xAPIC ID to avoid false positives (sort of) on systems
    with >255 CPUs, i.e. with IDs that don't fit into a u8.  The intent of
    APIC_ID_MODIFIED is to inhibit APICv/AVIC when the xAPIC is changed from
    it's original value,

    The mismatch isn't technically a false positive, as architecturally the
    xAPIC IDs do end up being aliased in this scenario, and neither APICv
    nor AVIC correctly handles IPI virtualization when there is aliasing.
    However, KVM already deliberately does not honor the aliasing behavior
    that results when an x2APIC ID gets truncated to an xAPIC ID.  I.e. the
    resulting APICv/AVIC behavior is aligned with KVM's existing behavior
    when KVM's x2APIC hotplug hack is effectively enabled.

    If/when KVM provides a way to disable the hotplug hack, APICv/AVIC can
    piggyback whatever logic disables the optimized APIC map (which is what
    provides the hotplug hack), i.e. so that KVM's optimized map and APIC
    virtualization yield the same behavior.

    For now, fix the immediate problem of APIC virtualization being disabled
    for large VMs, which is a much more pressing issue than ensuring KVM
    honors architectural behavior for APIC ID aliasing.

    Fixes: 3743c2f02517 ("KVM: x86: inhibit APICv/AVIC on changes to APIC ID or APIC base")
    Reported-by: Suravee Suthikulpanit <suravee.suthikulpanit@amd.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-7-seanjc@google.com>
```

```
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f911843eb5af2fb931e5dbd9ad72fc0272cd5a49
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:38 2023 +0000

    KVM: x86: Don't inhibit APICv/AVIC on xAPIC ID "change" if APIC is disabled

    commit a58a66afc464d6d2ec294cd3102f36f3652e7ce4 upstream.

    Don't inhibit APICv/AVIC due to an xAPIC ID mismatch if the APIC is
    hardware disabled.  The ID cannot be consumed while the APIC is disabled,
    and the ID is guaranteed to be set back to the vcpu_id when the APIC is
    hardware enabled (architectural behavior correctly emulated by KVM).

    Fixes: 3743c2f02517 ("KVM: x86: inhibit APICv/AVIC on changes to APIC ID or APIC base")
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-6-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7e65b89d558afcd4c8334e493045fccc32ea0fb6
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:34 2023 +0000

    KVM: x86: Blindly get current x2APIC reg value on "nodecode write" traps

    commit 0a19807b464fb10aa79b9dd7f494bc317438fada upstream.

    When emulating a x2APIC write in response to an APICv/AVIC trap, get the
    the written value from the vAPIC page without checking that reads are
    allowed for the target register.  AVIC can generate trap-like VM-Exits on
    writes to EOI, and so KVM needs to get the written value from the backing
    page without running afoul of EOI's write-only behavior.

    Alternatively, EOI could be special cased to always write '0', e.g. so
    that the sanity check could be preserved, but x2APIC on AMD is actually
    supposed to disallow non-zero writes (not emulated by KVM), and the
    sanity check was a byproduct of how the KVM code was written, i.e. wasn't
    added to guard against anything in particular.

    Fixes: 70c8327c11c6 ("KVM: x86: Bug the VM if an accelerated x2APIC trap occurs on a
"bad" reg")
    Fixes: 1bd9dfec9fd4 ("KVM: x86: Do not block APIC write for non ICR registers")
    Reported-by: Alejandro Jimenez <alejandro.j.jimenez@oracle.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-2-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6e255bc8f63e1d5d4c81f649ddd690873d942f99
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 6 01:12:35 2023 +0000

    KVM: x86: Purge "highest ISR" cache when updating APICv state

    commit 97a71c444a147ae41c7d0ab5b3d855d7f762f3ed upstream.

    Purge the "highest ISR" cache when updating APICv state on a vCPU.  The
    cache must not be used when APICv is active as hardware may emulate EOIs
```

```
    (and other operations) without exiting to KVM.

    This fixes a bug where KVM will effectively block IRQs in perpetuity due
    to the "highest ISR" never getting reset if APICv is activated on a vCPU
    while an IRQ is in-service.  Hardware emulates the EOI and KVM never gets
    a chance to update its cache.

    Fixes: b26a695a1d78 ("kvm: lapic: Introduce APICv update helper function")
    Cc: stable@vger.kernel.org
    Cc: Suravee Suthikulpanit <suravee.suthikulpanit@amd.com>
    Cc: Maxim Levitsky <mlevitsk@redhat.com>
    Reviewed-by: Paolo Bonzini <pbonzini@redhat.com>
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20230106011306.85230-3-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 906ea0086b70c146e1871a773430d7574391bf84
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Nov 30 23:08:45 2022 +0000

    KVM: Register /dev/kvm as the _very_ last thing during initialization

    commit 2b01281273738bf2d6551da48d65db2df3f28998 upstream.

    Register /dev/kvm, i.e. expose KVM to userspace, only after all other
    setup has completed.  Once /dev/kvm is exposed, userspace can start
    invoking KVM ioctls, creating VMs, etc...  If userspace creates a VM
    before KVM is done with its configuration, bad things may happen, e.g.
    KVM will fail to properly migrate vCPU state if a VM is created before
    KVM has registered preemption notifiers.

    Cc: stable@vger.kernel.org
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20221130230934.1014142-2-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3ba95cc671c025d0d2a1c7d5e2930f0ff0980cf4
Author: Alexandru Matei <alexandru.matei@uipath.com>
Date:   Tue Jan 24 00:12:08 2023 +0200

    KVM: VMX: Fix crash due to uninitialized current_vmcs

    commit 93827a0a36396f2fd6368a54a020f420c8916e9b upstream.

    KVM enables 'Enlightened VMCS' and 'Enlightened MSR Bitmap' when running as
    a nested hypervisor on top of Hyper-V. When MSR bitmap is updated,
    evmcs_touch_msr_bitmap function uses current_vmcs per-cpu variable to mark
    that the msr bitmap was changed.

    vmx_vcpu_create() modifies the msr bitmap via vmx_disable_intercept_for_msr
    -> vmx_msr_bitmap_l01_changed which in the end calls this function. The
    function checks for current_vmcs if it is null but the check is
    insufficient because current_vmcs is not initialized. Because of this, the
    code might incorrectly write to the structure pointed by current_vmcs value
    left by another task. Preemption is not disabled, the current task can be
    preempted and moved to another CPU while current_vmcs is accessed multiple
    times from evmcs_touch_msr_bitmap() which leads to crash.

    The manipulation of MSR bitmaps by callers happens only for vmcs01 so the
    solution is to use vmx->vmcs01.vmcs instead of current_vmcs.

      BUG: kernel NULL pointer dereference, address: 0000000000000338
```

```
       PGD 4e1775067 P4D 0
       Oops: 0002 [#1] PREEMPT SMP NOPTI
       ...
       RIP: 0010:vmx_msr_bitmap_l01_changed+0x39/0x50 [kvm_intel]
       ...
       Call Trace:
        vmx_disable_intercept_for_msr+0x36/0x260 [kvm_intel]
        vmx_vcpu_create+0xe6/0x540 [kvm_intel]
        kvm_arch_vcpu_create+0x1d1/0x2e0 [kvm]
        kvm_vm_ioctl_create_vcpu+0x178/0x430 [kvm]
        kvm_vm_ioctl+0x53f/0x790 [kvm]
        __x64_sys_ioctl+0x8a/0xc0
        do_syscall_64+0x5c/0x90
        entry_SYSCALL_64_after_hwframe+0x63/0xcd

    Fixes: ceef7d10dfb6 ("KVM: x86: VMX: hyper-v: Enlightened MSR-Bitmap support")
    Cc: stable@vger.kernel.org
    Suggested-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Alexandru Matei <alexandru.matei@uipath.com>
    Link: https://lore.kernel.org/r/20230123221208.4964-1-alexandru.matei@uipath.com
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fb436dd6914325075f07d19851ab277b7a693ae7
Author: Sean Christopherson <seanjc@google.com>
Date:   Mon Dec 19 17:19:24 2022 +0000

    KVM: Destroy target device if coalesced MMIO unregistration fails

    commit b1cb1fac22abf102ffeb29dd3eeca208a3869d54 upstream.

    Destroy and free the target coalesced MMIO device if unregistering said
    device fails.  As clearly noted in the code, kvm_io_bus_unregister_dev()
    does not destroy the target device.

      BUG: memory leak
      unreferenced object 0xffff888112a54880 (size 64):
        comm "syz-executor.2", pid 5258, jiffies 4297861402 (age 14.129s)
        hex dump (first 32 bytes):
          38 c7 67 15 00 c9 ff ff 38 c7 67 15 00 c9 ff ff  8.g.....8.g.....
          e0 c7 e1 83 ff ff ff ff 00 30 67 15 00 c9 ff ff  .........0g.....
        backtrace:
          [<0000000006995a8a>] kmalloc include/linux/slab.h:556 [inline]
          [<0000000006995a8a>] kzalloc include/linux/slab.h:690 [inline]
          [<0000000006995a8a>] kvm_vm_ioctl_register_coalesced_mmio+0x8e/0x3d0
arch/x86/kvm/../../../virt/kvm/coalesced_mmio.c:150
          [<00000000022550c2>] kvm_vm_ioctl+0x47d/0x1600
arch/x86/kvm/../../../virt/kvm/kvm_main.c:3323
          [<000000008a75102f>] vfs_ioctl fs/ioctl.c:46 [inline]
          [<000000008a75102f>] file_ioctl fs/ioctl.c:509 [inline]
          [<000000008a75102f>] do_vfs_ioctl+0xbab/0x1160 fs/ioctl.c:696
          [<0000000080e3f669>] ksys_ioctl+0x76/0xa0 fs/ioctl.c:713
          [<0000000059ef4888>] __do_sys_ioctl fs/ioctl.c:720 [inline]
          [<0000000059ef4888>] __se_sys_ioctl fs/ioctl.c:718 [inline]
          [<0000000059ef4888>] __x64_sys_ioctl+0x6f/0xb0 fs/ioctl.c:718
          [<000000006444fa05>] do_syscall_64+0x9f/0x4e0 arch/x86/entry/common.c:290
          [<000000009a4ed50b>] entry_SYSCALL_64_after_hwframe+0x49/0xbe

      BUG: leak checking failed

    Fixes: 5d3c4c79384a ("KVM: Stop looking for coalesced MMIO zones if the bus is
destroyed")
    Cc: stable@vger.kernel.org
    Reported-by: 柳菁峰 <liujingfeng@qianxin.com>
    Reported-by: Michal Luczaj <mhal@rbox.co>
```

```
    Link: https://lore.kernel.org/r/20221219171924.67989-1-seanjc@google.com
    Link: https://lore.kernel.org/all/20230118220003.1239032-1-mhal@rbox.co
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6930ca3f5ef99578a8518e06166450db2fac406c
Author: Hou Tao <houtao1@huawei.com>
Date:   Tue Jan 31 15:07:19 2023 +0800

    md: don't update recovery_cp when curr_resync is ACTIVE

    commit 1d1f25bfda432a6b61bd0205d426226bbbd73504 upstream.

    Don't update recovery_cp when curr_resync is MD_RESYNC_ACTIVE, otherwise
    md may skip the resync of the first 3 sectors if the resync procedure is
    interrupted before the first calling of ->sync_request() as shown below:

    md_do_sync thread          control thread
      // setup resync
      mddev->recovery_cp = 0
      j = 0
      mddev->curr_resync = MD_RESYNC_ACTIVE

                                 // e.g., set array as idle
                                 set_bit(MD_RECOVERY_INTR, &&mddev_recovery)
      // resync loop
      // check INTR before calling sync_request
      !test_bit(MD_RECOVERY_INTR, &mddev->recovery

      // resync interrupted
      // update recovery_cp from 0 to 3
      // the resync of three 3 sectors will be skipped
      mddev->recovery_cp = 3

    Fixes: eac58d08d493 ("md: Use enum for overloaded magic numbers used by mddev-
>curr_resync")
    Cc: stable@vger.kernel.org # 6.0+
    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Reviewed-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5d6cc0541e279a45bf6ad69041b0b3b0633cf7bd
Author: Jan Kara <jack@suse.cz>
Date:   Mon Jan 23 14:18:47 2023 +0100

    udf: Fix file corruption when appending just after end of preallocated extent

    commit 36ec52ea038b18a53e198116ef7d7e70c87db046 upstream.

    When we append new block just after the end of preallocated extent, the
    code in inode_getblk() wrongly determined we're going to use the
    preallocated extent which resulted in adding block into a wrong logical
    offset in the file. Sequence like this manifests it:

    xfs_io -f -c "pwrite 0x2cacf 0xd122" -c "truncate 0x2dd6f" \
      -c "pwrite 0x27fd9 0x69a9" -c "pwrite 0x32981 0x7244" <file>

    The code that determined the use of preallocated extent is actually
    stale because udf_do_extend_file() does not create preallocation anymore
    so after calling that function we are sure there's no usable
    preallocation. Just remove the faulty condition.

    CC: stable@vger.kernel.org
    Fixes: 16d055656814 ("udf: Discard preallocation before extending file with a hole")
```

```
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9e3b5ef7d02eaa6553e79b4af9bd99227280f245
Author: Jan Kara <jack@suse.cz>
Date:   Tue Jan 3 10:03:35 2023 +0100

    udf: Detect system inodes linked into directory hierarchy

    commit 85a37983ec69cc9fcd188bc37c4de15ee326355a upstream.

    When UDF filesystem is corrupted, hidden system inodes can be linked
    into directory hierarchy which is an avenue for further serious
    corruption of the filesystem and kernel confusion as noticed by syzbot
    fuzzed images. Refuse to access system inodes linked into directory
    hierarchy and vice versa.

    CC: stable@vger.kernel.org
    Reported-by: syzbot+38695a20b8addcbc1084@syzkaller.appspotmail.com
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bc170fa03fb88d29fc139da1c1ed3c290097d4ae
Author: Jan Kara <jack@suse.cz>
Date:   Tue Jan 3 09:56:56 2023 +0100

    udf: Preserve link count of system files

    commit fc8033a34a3ca7d23353e645e6dde5d364ac5f12 upstream.

    System files in UDF filesystem have link count 0. To not confuse VFS we
    fudge the link count to be 1 when reading such inodes however we forget
    to restore the link count of 0 when writing such inodes. Fix that.

    CC: stable@vger.kernel.org
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6d18cedc1ef0caeb1567cab660079e48844ff6d6
Author: Jan Kara <jack@suse.cz>
Date:   Mon Jan 2 20:14:47 2023 +0100

    udf: Do not update file length for failed writes to inline files

    commit 256fe4162f8b5a1625b8603ca5f7ff79725bfb47 upstream.

    When write to inline file fails (or happens only partly), we still
    updated length of inline data as if the whole write succeeded. Fix the
    update of length of inline data to happen only if the write succeeds.

    Reported-by: syzbot+0937935b993956ba28ab@syzkaller.appspotmail.com
    CC: stable@vger.kernel.org
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7a965da79f2d22601f329cbfce588386b0847544
Author: Jan Kara <jack@suse.cz>
Date:   Fri Dec 16 12:37:51 2022 +0100

    udf: Do not bother merging very long extents

    commit 53cafe1d6d8ef9f93318e5bfccc0d24f27d41ced upstream.

    When merging very long extents we try to push as much length as possible
    to the first extent. However this is unnecessarily complicated and not
```

```
    really worth the trouble. Furthermore there was a bug in the logic
    resulting in corrupting extents in the file as syzbot reproducer shows.
    So just don't bother with the merging of extents that are too long
    together.

    CC: stable@vger.kernel.org
    Reported-by: syzbot+60f291a24acecb3c2bd5@syzkaller.appspotmail.com
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2b7da22a7842bd202fab11d3dc256099c7a2c2e0
Author: Jan Kara <jack@suse.cz>
Date:   Thu Dec 15 14:24:03 2022 +0100

    udf: Truncate added extents on failed expansion

    commit 70bfb3a8d661d4fdc742afc061b88a7f3fc9f500 upstream.

    When a file expansion failed because we didn't have enough space for
    indirect extents make sure we truncate extents created so far so that we
    don't leave extents beyond EOF.

    CC: stable@vger.kernel.org
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3d47428515684efbb319d5c3a422ae7609c45095
Author: Jeff Xu <jeffxu@google.com>
Date:   Sat Jan 14 02:03:06 2023 +0000

    selftests/landlock: Test ptrace as much as possible with Yama

    commit 8677e555f17f51321d0730b945aeb7d4b95f998f upstream.

    Update ptrace tests according to all potential Yama security policies.
    This is required to make such tests pass even if Yama is enabled.

    Tests are not skipped but they now check both Landlock and Yama boundary
    restrictions at run time to keep a maximum test coverage (i.e. positive
    and negative testing).

    Signed-off-by: Jeff Xu <jeffxu@google.com>
    Link: https://lore.kernel.org/r/20230114020306.1407195-2-jeffxu@google.com
    Cc: stable@vger.kernel.org
    [mic: Add curly braces around EXPECT_EQ() to make it build, and improve
    commit message]
    Co-developed-by: Mickaël Salaün <mic@digikod.net>
    Signed-off-by: Mickaël Salaün <mic@digikod.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3126dcbf40add0e1033f71c413356a13f0150414
Author: Jeff Xu <jeffxu@google.com>
Date:   Fri Jan 13 05:32:29 2023 +0000

    selftests/landlock: Skip overlayfs tests when not supported

    commit 366617a69e606109128365705 46f118006ebc7cb upstream.

    overlayfs may be disabled in the kernel configuration, causing related
    tests to fail.  Check that overlayfs is supported at runtime, so we can
    skip layout2_overlay.* accordingly.

    Signed-off-by: Jeff Xu <jeffxu@google.com>
    Reviewed-by: Guenter Roeck <groeck@chromium.org>
    Cc: stable@vger.kernel.org
```

```
    Link: https://lore.kernel.org/r/20230113053229.1281774-2-jeffxu@google.com
    [mic: Reword comments and constify variables]
    Signed-off-by: Mickaël Salaün <mic@digikod.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8929fd5fc5214114d1f092a27b9ac9258c40dbbb
Author: Andrew Morton <akpm@linux-foundation.org>
Date:   Sun Feb 26 12:31:11 2023 -0800

    fs/cramfs/inode.c: initialize file_ra_state

    commit 3e35102666f873a135d31a726ac1ec8af4905206 upstream.

    file_ra_state_init() assumes that the file_ra_state has been zeroed out.
    Fixes a KMSAN used-unintialized issue (at least).

    Fixes: cf948cbc35e80 ("cramfs: read_mapping_page() is synchronous")
    Reported-by: syzbot <syzbot+8ce7f8308d91e6b8bbe2@syzkaller.appspotmail.com>
      Link: https://lkml.kernel.org/r/0000000000008f74e905f56df987@google.com
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: Nicolas Pitre <nico@fluxnic.net>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c51d1a3e1d9ee85f41f90e50efd75b39b1a70e7d
Author: Heming Zhao via Ocfs2-devel <ocfs2-devel@oss.oracle.com>
Date:   Mon Feb 20 13:05:26 2023 +0800

    ocfs2: fix non-auto defrag path not working issue

    commit 236b9254f8d1edc273ad88b420aa85fbd84f492d upstream.

    This fixes three issues on move extents ioctl without auto defrag:

    a) In ocfs2_find_victim_alloc_group(), we have to convert bits to block
       first in case of global bitmap.

    b) In ocfs2_probe_alloc_group(), when finding enough bits in block
       group bitmap, we have to back off move_len to start pos as well,
       otherwise it may corrupt filesystem.

    c) In ocfs2_ioctl_move_extents(), set me_threshold both for non-auto
       and auto defrag paths.  Otherwise it will set move_max_hop to 0 and
       finally cause unexpectedly ENOSPC error.

    Currently there are no tools triggering the above issues since
    defragfs.ocfs2 enables auto defrag by default.  Tested with manually
    changing defragfs.ocfs2 to run non auto defrag path.

    Link: https://lkml.kernel.org/r/20230220050526.22020-1-heming.zhao@suse.com
    Signed-off-by: Heming Zhao <heming.zhao@suse.com>
    Reviewed-by: Joseph Qi <joseph.qi@linux.alibaba.com>
    Cc: Mark Fasheh <mark@fasheh.com>
    Cc: Joel Becker <jlbec@evilplan.org>
    Cc: Junxiao Bi <junxiao.bi@oracle.com>
    Cc: Changwei Ge <gechangwei@live.cn>
    Cc: Gang He <ghe@suse.com>
    Cc: Jun Piao <piaojun@huawei.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 590507ebabd33cd93324c04f9a5538309a5ba934
Author: Heming Zhao via Ocfs2-devel <ocfs2-devel@oss.oracle.com>
```

```
Date:   Fri Feb 17 08:37:17 2023 +0800

    ocfs2: fix defrag path triggering jbd2 ASSERT

    commit 60eed1e3d45045623e46944ebc7c42c30a4350f0 upstream.

    code path:

    ocfs2_ioctl_move_extents
     ocfs2_move_extents
      ocfs2_defrag_extent
       __ocfs2_move_extent
        + ocfs2_journal_access_di
        + ocfs2_split_extent  //sub-paths call jbd2_journal_restart
        + ocfs2_journal_dirty //crash by jbs2 ASSERT

    crash stacks:

    PID: 11297  TASK: ffff974a676dcd00  CPU: 67  COMMAND: "defragfs.ocfs2"
     #0 [ffffb25d8dad3900] machine_kexec at ffffffff8386fe01
     #1 [ffffb25d8dad3958] __crash_kexec at ffffffff8395959d
     #2 [ffffb25d8dad3a20] crash_kexec at ffffffff8395a45d
     #3 [ffffb25d8dad3a38] oops_end at ffffffff83836d3f
     #4 [ffffb25d8dad3a58] do_trap at ffffffff83833205
     #5 [ffffb25d8dad3aa0] do_invalid_op at ffffffff83833aa6
     #6 [ffffb25d8dad3ac0] invalid_op at ffffffff84200d18
        [exception RIP: jbd2_journal_dirty_metadata+0x2ba]
        RIP: ffffffffc09ca54a  RSP: ffffb25d8dad3b70  RFLAGS: 00010207
        RAX: 0000000000000000  RBX: ffff9706eedc5248  RCX: 0000000000000000
        RDX: 0000000000000001  RSI: ffff97337029ea28  RDI: ffff9706eedc5250
        RBP: ffff9703c3520200   R8: 000000000f46b0b2   R9: 0000000000000000
        R10: 0000000000000001  R11: 00000001000000fe  R12: ffff97337029ea28
        R13: 0000000000000000  R14: ffff9703de59bf60  R15: ffff9706eedc5250
        ORIG_RAX: ffffffffffffffff  CS: 0010  SS: 0018
     #7 [ffffb25d8dad3ba8] ocfs2_journal_dirty at ffffffffc137fb95 [ocfs2]
     #8 [ffffb25d8dad3be8] __ocfs2_move_extent at ffffffffc139a950 [ocfs2]
     #9 [ffffb25d8dad3c80] ocfs2_defrag_extent at ffffffffc139b2d2 [ocfs2]

    Analysis

    This bug has the same root cause of 'commit 7f27ec978b0e ("ocfs2: call
    ocfs2_journal_access_di() before ocfs2_journal_dirty() in
    ocfs2_write_end_nolock()")'.  For this bug, jbd2_journal_restart() is
    called by ocfs2_split_extent() during defragmenting.

    How to fix

    For ocfs2_split_extent() can handle journal operations totally by itself.
    Caller doesn't need to call journal access/dirty pair, and caller only
    needs to call journal start/stop pair.  The fix method is to remove
    journal access/dirty from __ocfs2_move_extent().

    The discussion for this patch:
    https://oss.oracle.com/pipermail/ocfs2-devel/2023-February/000647.html

    Link: https://lkml.kernel.org/r/20230217003717.32469-1-heming.zhao@suse.com
    Signed-off-by: Heming Zhao <heming.zhao@suse.com>
    Reviewed-by: Joseph Qi <joseph.qi@linux.alibaba.com>
    Cc: Mark Fasheh <mark@fasheh.com>
    Cc: Joel Becker <jlbec@evilplan.org>
    Cc: Junxiao Bi <junxiao.bi@oracle.com>
    Cc: Changwei Ge <gechangwei@live.cn>
    Cc: Gang He <ghe@suse.com>
    Cc: Jun Piao <piaojun@huawei.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 f7e28613856373b51906825d3a069edfbbbc9e50
Author: Jaegeuk Kim <jaegeuk@kernel.org>
Date:   Tue Feb 14 15:53:52 2023 -0800

    f2fs: Revert "f2fs: truncate blocks in batch in __complete_revoke_list()"

    commit c7dbc06688292db34c1bb9c715e29ac4935af994 upstream.

    We should not truncate replaced blocks, and were supposed to truncate the first
    part as well.

    This reverts commit 78a99fe6254cad4be310cd84af39f6c46b668c72.

    Cc: stable@vger.kernel.org
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eb52f13c6093ac761dbeaa459c810fc0253209fc
Author: Jaegeuk Kim <jaegeuk@kernel.org>
Date:   Sun Feb 5 19:13:39 2023 -0800

    f2fs: fix kernel crash due to null io->bio

    commit 267c159f9c7bcb7009dae16889b880c5ed8759a8 upstream.

    We should return when io->bio is null before doing anything. Otherwise, panic.

    BUG: kernel NULL pointer dereference, address: 0000000000000010
    RIP: 0010:__submit_merged_write_cond+0x164/0x240 [f2fs]
    Call Trace:
     <TASK>
     f2fs_submit_merged_write+0x1d/0x30 [f2fs]
     commit_checkpoint+0x110/0x1e0 [f2fs]
     f2fs_write_checkpoint+0x9f7/0xf00 [f2fs]
     ? __pfx_issue_checkpoint_thread+0x10/0x10 [f2fs]
     __checkpoint_and_complete_reqs+0x84/0x190 [f2fs]
     ? preempt_count_add+0x82/0xc0
     ? __pfx_issue_checkpoint_thread+0x10/0x10 [f2fs]
     issue_checkpoint_thread+0x4c/0xf0 [f2fs]
     ? __pfx_autoremove_wake_function+0x10/0x10
     kthread+0xff/0x130
     ? __pfx_kthread+0x10/0x10
     ret_from_fork+0x2c/0x50
     </TASK>

    Cc: stable@vger.kernel.org # v5.18+
    Fixes: 64bf0eef0171 ("f2fs: pass the bio operation to bio_alloc_bioset")
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3dd1cbfa11d6929ba04b0c055d4c7b5754372f2a
Author: Eric Biggers <ebiggers@google.com>
Date:   Thu Feb 2 17:02:39 2023 -0800

    f2fs: fix cgroup writeback accounting with fs-layer encryption

    commit 844545c51a5b2a524b22a2fe9d0b353b827d24b4 upstream.

    When writing a page from an encrypted file that is using
    filesystem-layer encryption (not inline encryption), f2fs encrypts the
```

```
    pagecache page into a bounce page, then writes the bounce page.

    It also passes the bounce page to wbc_account_cgroup_owner().  That's
    incorrect, because the bounce page is a newly allocated temporary page
    that doesn't have the memory cgroup of the original pagecache page.
    This makes wbc_account_cgroup_owner() not account the I/O to the owner
    of the pagecache page as it should.

    Fix this by always passing the pagecache page to
    wbc_account_cgroup_owner().

    Fixes: 578c647879f7 ("f2fs: implement cgroup writeback support")
    Cc: stable@vger.kernel.org
    Reported-by: Matthew Wilcox (Oracle) <willy@infradead.org>
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6b93f1e6fbef2939785977db39d0d57a674349a7
Author: Jaegeuk Kim <jaegeuk@kernel.org>
Date:   Mon Jan 30 15:20:09 2023 -0800

    f2fs: retry to update the inode page given data corruption

    commit 3aa51c61cb4a4dcb40df51ac61171e9ac5a35321 upstream.

    If the storage gives a corrupted node block due to short power failure and
    reset, f2fs stops the entire operations by setting the checkpoint failure flag.

    Let's give more chances to live by re-issuing IOs for a while in such critical
    path.

    Cc: stable@vger.kernel.org
    Suggested-by: Randall Huang <huangrandall@google.com>
    Suggested-by: Chao Yu <chao@kernel.org>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eebaecef0095bb8f493c03982da75c6e7bae1056
Author: Eric Biggers <ebiggers@google.com>
Date:   Sun Jan 22 23:04:14 2023 -0800

    f2fs: fix information leak in f2fs_move_inline_dirents()

    commit 9a5571cff4ffcfc24847df9fd545cc5799ac0ee5 upstream.

    When converting an inline directory to a regular one, f2fs is leaking
    uninitialized memory to disk because it doesn't initialize the entire
    directory block.  Fix this by zero-initializing the block.

    This bug was introduced by commit 4ec17d688d74 ("f2fs: avoid unneeded
    initializing when converting inline dentry"), which didn't consider the
    security implications of leaking uninitialized memory to disk.

    This was found by running xfstest generic/435 on a KMSAN-enabled kernel.

    Fixes: 4ec17d688d74 ("f2fs: avoid unneeded initializing when converting inline dentry")
    Cc: <stable@vger.kernel.org> # v4.3+
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 23a7f431531ee81373d13af2efa1da2f8159f806
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:37 2023 -0500

    fs: dlm: send FIN ack back in right cases

    commit 00908b3388255fc1d3782b744d07f327712f401f upstream.

    This patch moves to send a ack back for receiving a FIN message only
    when we are in valid states. In other cases and there might be a sender
    waiting for a ack we just let it timeout at the senders time and
    hopefully all other cleanups will remove the FIN message on their
    sending queue. As an example we should never send out an ACK being in
    LAST_ACK state or we cannot assume a working socket communication when
    we are in CLOSED state.

    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dcfde37c99a4f8325d379d54ff68eaa50cdc7b0f
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:36 2023 -0500

    fs: dlm: move sending fin message into state change handling

    commit a58496361802070996f9bd76e941d109c4a85ebd upstream.

    This patch moves the send fin handling, which should appear in a specific
    state change, into the state change handling while the per node
    state_lock is held. I experienced issues with other messages because
    we changed the state and a fin message was sent out in a different state.

    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5cdc5869b60b2a7ca8b543645b1c7edb0511a6e8
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:35 2023 -0500

    fs: dlm: don't set stop rx flag after node reset

    commit 15c63db8e86a72e0d5cfb9bf0cd1870e39a3e5fe upstream.

    Similar to the stop tx flag, the rx flag should warn about a dlm message
    being received at DLM_FIN state change, when we are assuming no other
    dlm application messages. If we receive a FIN message and we are in the
    state DLM_FIN_WAIT2 we call midcomms_node_reset() which puts the
    midcomms node into DLM_CLOSED state. Afterwards we should not set the
    DLM_NODE_FLAG_STOP_RX flag any more.  This patch changes the setting
    DLM_NODE_FLAG_STOP_RX in those state changes when we receive a FIN
    message and we assume there will be no other dlm application messages
    received until we hit DLM_CLOSED state.

    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit c7283d6461e944f609c7dca094e2a576aed1d353
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:34 2023 -0500

    fs: dlm: fix race setting stop tx flag

    commit 164272113b685927126c938b4a9cbd2075eb15ee upstream.

    This patch sets the stop tx flag before we commit the dlm message.
    This flag will report about unexpected transmissions after we
    send the DLM_FIN message out, which should be the last message sent.
    When we commit the dlm fin message, it could be that we already
    got an ack back and the CLOSED state change already happened.
    We should not set this flag when we are in CLOSED state. To avoid this
    race we simply set the tx flag before the state change can be in
    progress by moving it before dlm_midcomms_commit_mhandle().

    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 73a2bd498ab88518c0e1ddd49388cf1e9ba0e074
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:33 2023 -0500

    fs: dlm: be sure to call dlm_send_queue_flush()

    commit 7354fa4ef697191effedc2ae9a8293427708bbf5 upstream.

    If we release a midcomms node structure, there should be nothing left
    inside the dlm midcomms send queue. However, sometimes this is not true
    because I believe some DLM_FIN message was not acked... if we run
    into a shutdown timeout, then we should be sure there is no pending send
    dlm message inside this queue when releasing midcomms node structure.

    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2de9f9b686c71b4fa3663ae374f5f643c46a446
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:32 2023 -0500

    fs: dlm: fix use after free in midcomms commit

    commit 724b6bab0d75f1dc01fdfbf7fe8d4217a5cb90ba upstream.

    While working on processing dlm message in softirq context I experienced
    the following KASAN use-after-free warning:

    [  151.760477] ================================================================
    [  151.761803] BUG: KASAN: use-after-free in dlm_midcomms_commit_mhandle+0x19d/0x4b0
    [  151.763414] Read of size 4 at addr ffff88811a980c60 by task lock_torture/1347

    [  151.765284] CPU: 7 PID: 1347 Comm: lock_torture Not tainted 6.1.0-rc4+ #2828
    [  151.766778] Hardware name: Red Hat KVM/RHEL-AV, BIOS 1.16.0-
3.module+el8.7.0+16134+e5908aa2 04/01/2014
    [  151.768726] Call Trace:
    [  151.769277]  <TASK>
    [  151.769748]  dump_stack_lvl+0x5b/0x86
```

```
[  151.770556]  print_report+0x180/0x4c8
[  151.771378]  ? kasan_complete_mode_report_info+0x7c/0x1e0
[  151.772241]  ? dlm_midcomms_commit_mhandle+0x19d/0x4b0
[  151.773069]  kasan_report+0x93/0x1a0
[  151.773668]  ? dlm_midcomms_commit_mhandle+0x19d/0x4b0
[  151.774514]  __asan_load4+0x7e/0xa0
[  151.775089]  dlm_midcomms_commit_mhandle+0x19d/0x4b0
[  151.775890]  ? create_message.isra.29.constprop.64+0x57/0xc0
[  151.776770]  send_common+0x19f/0x1b0
[  151.777342]  ? remove_from_waiters+0x60/0x60
[  151.778017]  ? lock_downgrade+0x410/0x410
[  151.778648]  ? __this_cpu_preempt_check+0x13/0x20
[  151.779421]  ? rcu_lockdep_current_cpu_online+0x88/0xc0
[  151.780292]  _convert_lock+0x46/0x150
[  151.780893]  convert_lock+0x7b/0xc0
[  151.781459]  dlm_lock+0x3ac/0x580
[  151.781993]  ? 0xffffffffc0540000
[  151.782522]  ? torture_stop+0x120/0x120 [dlm_locktorture]
[  151.783379]  ? dlm_scan_rsbs+0xa70/0xa70
[  151.784003]  ? preempt_count_sub+0xd6/0x130
[  151.784661]  ? is_module_address+0x47/0x70
[  151.785309]  ? torture_stop+0x120/0x120 [dlm_locktorture]
[  151.786166]  ? 0xffffffffc0540000
[  151.786693]  ? lockdep_init_map_type+0xc3/0x360
[  151.787414]  ? 0xffffffffc0540000
[  151.787947]  torture_dlm_lock_sync.isra.3+0xe9/0x150 [dlm_locktorture]
[  151.789004]  ? torture_stop+0x120/0x120 [dlm_locktorture]
[  151.789858]  ? 0xffffffffc0540000
[  151.790392]  ? lock_torture_cleanup+0x20/0x20 [dlm_locktorture]
[  151.791347]  ? delay_tsc+0x94/0xc0
[  151.791898]  torture_ex_iter+0xc3/0xea [dlm_locktorture]
[  151.792735]  ? torture_start+0x30/0x30 [dlm_locktorture]
[  151.793606]  lock_torture+0x177/0x270 [dlm_locktorture]
[  151.794448]  ? torture_dlm_lock_sync.isra.3+0x150/0x150 [dlm_locktorture]
[  151.795539]  ? lock_torture_stats+0x80/0x80 [dlm_locktorture]
[  151.796476]  ? do_raw_spin_lock+0x11e/0x1e0
[  151.797152]  ? mark_held_locks+0x34/0xb0
[  151.797784]  ? _raw_spin_unlock_irqrestore+0x30/0x70
[  151.798581]  ? __kthread_parkme+0x79/0x110
[  151.799246]  ? trace_preempt_on+0x2a/0xf0
[  151.799902]  ? __kthread_parkme+0x79/0x110
[  151.800579]  ? preempt_count_sub+0xd6/0x130
[  151.801271]  ? __kasan_check_read+0x11/0x20
[  151.801963]  ? __kthread_parkme+0xec/0x110
[  151.802630]  ? lock_torture_stats+0x80/0x80 [dlm_locktorture]
[  151.803569]  kthread+0x192/0x1d0
[  151.804104]  ? kthread_complete_and_exit+0x30/0x30
[  151.804881]  ret_from_fork+0x1f/0x30
[  151.805480]  </TASK>

[  151.806111] Allocated by task 1347:
[  151.806681]  kasan_save_stack+0x26/0x50
[  151.807308]  kasan_set_track+0x25/0x30
[  151.807920]  kasan_save_alloc_info+0x1e/0x30
[  151.808609]  __kasan_slab_alloc+0x63/0x80
[  151.809263]  kmem_cache_alloc+0x1ad/0x830
[  151.809916]  dlm_allocate_mhandle+0x17/0x20
[  151.810590]  dlm_midcomms_get_mhandle+0x96/0x260
[  151.811344]  _create_message+0x95/0x180
[  151.811994]  create_message.isra.29.constprop.64+0x57/0xc0
[  151.812880]  send_common+0x129/0x1b0
[  151.813467]  _convert_lock+0x46/0x150
[  151.814074]  convert_lock+0x7b/0xc0
[  151.814648]  dlm_lock+0x3ac/0x580
[  151.815199]  torture_dlm_lock_sync.isra.3+0xe9/0x150 [dlm_locktorture]
```

```
     [  151.816258]  torture_ex_iter+0xc3/0xea [dlm_locktorture]
     [  151.817129]  lock_torture+0x177/0x270 [dlm_locktorture]
     [  151.817986]  kthread+0x192/0x1d0
     [  151.818518]  ret_from_fork+0x1f/0x30

     [  151.819369] Freed by task 1336:
     [  151.819890]  kasan_save_stack+0x26/0x50
     [  151.820514]  kasan_set_track+0x25/0x30
     [  151.821128]  kasan_save_free_info+0x2e/0x50
     [  151.821812]  __kasan_slab_free+0x107/0x1a0
     [  151.822483]  kmem_cache_free+0x204/0x5e0
     [  151.823152]  dlm_free_mhandle+0x18/0x20
     [  151.823781]  dlm_mhandle_release+0x2e/0x40
     [  151.824454]  rcu_core+0x583/0x1330
     [  151.825047]  rcu_core_si+0xe/0x20
     [  151.825594]  __do_softirq+0xf4/0x5c2

     [  151.826450] Last potentially related work creation:
     [  151.827238]  kasan_save_stack+0x26/0x50
     [  151.827870]  __kasan_record_aux_stack+0xa2/0xc0
     [  151.828609]  kasan_record_aux_stack_noalloc+0xb/0x20
     [  151.829415]  call_rcu+0x4c/0x760
     [  151.829954]  dlm_mhandle_delete+0x97/0xb0
     [  151.830718]  dlm_process_incoming_buffer+0x2fc/0xb30
     [  151.831524]  process_dlm_messages+0x16e/0x470
     [  151.832245]  process_one_work+0x505/0xa10
     [  151.832905]  worker_thread+0x67/0x650
     [  151.833507]  kthread+0x192/0x1d0
     [  151.834046]  ret_from_fork+0x1f/0x30

     [  151.834900] The buggy address belongs to the object at ffff88811a980c30
                     which belongs to the cache dlm_mhandle of size 88
     [  151.836894] The buggy address is located 48 bytes inside of
                     88-byte region [ffff88811a980c30, ffff88811a980c88)

     [  151.839007] The buggy address belongs to the physical page:
     [  151.839904] page:0000000076cf5d62 refcount:1 mapcount:0 mapping:0000000000000000
  index:0x0 pfn:0x11a980
     [  151.841378] flags: 0x8000000000000200(slab|zone=2)
     [  151.842141] raw: 8000000000000200 0000000000000000 dead000000000122 ffff8881089b43c0
     [  151.843401] raw: 0000000000000000 0000000000220022 00000001ffffffff 0000000000000000
     [  151.844640] page dumped because: kasan: bad access detected

     [  151.845822] Memory state around the buggy address:
     [  151.846602]  ffff88811a980b00: fb fb fb fb fc fc fc fc fa fb fb fb fb fb fb fb
     [  151.847761]  ffff88811a980b80: fb fb fb fc fc fc fc fa fb fb fb fb fb fb fb fb
     [  151.848921] >ffff88811a980c00: fb fb fc fc fc fc fa fb fb fb fb fb fb fb fb fb
     [  151.850076]                                                 ^
     [  151.851085]  ffff88811a980c80: fb fc fc fc fc fa fb fb fb fb fb fb fb fb fb fb
     [  151.852269]  ffff88811a980d00: fc fc fc fc fa fb fb fb fb fb fb fb fb fb fb fc
     [  151.853428] ==================================================================
     [  151.855618] Disabling lock debugging due to kernel taint

     It is accessing a mhandle in dlm_midcomms_commit_mhandle() and the mhandle
     was freed by a call_rcu() call in dlm_process_incoming_buffer(),
     dlm_mhandle_delete(). It looks like it was freed because an ack of
     this message was received. There is a short race between committing the
     dlm message to be transmitted and getting an ack back. If the ack is
     faster than returning from dlm_midcomms_commit_msg_3_2(), then we run
     into a use-after free because we still need to reference the mhandle when
     calling srcu_read_unlock().

     To avoid that, we don't allow that mhandle to be freed between
     dlm_midcomms_commit_msg_3_2() and srcu_read_unlock() by using rcu read
     lock. We can do that because mhandle is protected by rcu handling.
```

```
    Cc: stable@vger.kernel.org
    Fixes: 489d8e559c65 ("fs: dlm: add reliable connection if reconnect")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 714329e2d0f5efe9ff89e01ed111850bf224cc8e
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Jan 12 17:10:31 2023 -0500

    fs: dlm: start midcomms before scand

    commit aad633dc0cf90093998b1ae0ba9f19b5f1dab644 upstream.

    The scand kthread can send dlm messages out, especially dlm remove
    messages to free memory for unused rsb on other nodes. To send out dlm
    messages, midcomms must be initialized. This patch moves the midcomms
    start before scand is started.

    Cc: stable@vger.kernel.org
    Fixes: e7fd41792fc0 ("[DLM] The core of the DLM for GFS2/CLVM")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec03f6a9eb51b7d74ca14403de91c70cfa243521
Author: Yuezhang Mo <Yuezhang.Mo@sony.com>
Date:   Wed Jan 4 14:37:47 2023 +0800

    exfat: fix inode->i_blocks for non-512 byte sector size device

    commit 39c1ce8eafc0ff64fb9e28536ccc7df6a8e2999d upstream.

    inode->i_blocks is not real number of blocks, but 512 byte ones.

    Fixes: 98d917047e8b ("exfat: add file operations")
    Cc: stable@vger.kernel.org # v5.7+
    Reported-by: Wang Yugui <wangyugui@e16-tech.com>
    Tested-by: Wang Yugui <wangyugui@e16-tech.com>
    Signed-off-by: Yuezhang Mo <Yuezhang.Mo@sony.com>
    Reviewed-by: Andy Wu <Andy.Wu@sony.com>
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ea0f96fa694988edbfade443d832c4ba59727174
Author: Sungjong Seo <sj1557.seo@samsung.com>
Date:   Thu Dec 29 20:52:38 2022 +0900

    exfat: redefine DIR_DELETED as the bad cluster number

    commit bdaadfd343e3cba49ad0b009ff4b148dad0fa404 upstream.

    When a file or a directory is deleted, the hint for the cluster of
    its parent directory in its in-memory inode is set as DIR_DELETED.
    Therefore, DIR_DELETED must be one of invalid cluster numbers. According
    to the exFAT specification, a volume can have at most 2^32-11 clusters.
    However, DIR_DELETED is wrongly defined as 0xFFFF0321, which could be
    a valid cluster number. To fix it, let's redefine DIR_DELETED as
    0xFFFFFFF7, the bad cluster number.

    Fixes: 1acf1a564b60 ("exfat: add in-memory and on-disk structures and headers")
    Cc: stable@vger.kernel.org # v5.7+
    Reported-by: Yuezhang Mo <Yuezhang.Mo@sony.com>
    Signed-off-by: Sungjong Seo <sj1557.seo@samsung.com>
```

```
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 499c8e2882a1cc3eaa17bb7e4ae8dd51e9a79bf0
Author: Yuezhang Mo <Yuezhang.Mo@sony.com>
Date:   Thu Sep 22 14:43:47 2022 +0800

    exfat: fix unexpected EOF while reading dir

    commit 6cb5d1a16a51d080fbc1649a5144cbc5ca7d6f88 upstream.

    If the position is not aligned with the dentry size, the return
    value of readdir() will be NULL and errno is 0, which means the
    end of the directory stream is reached.

    If the position is aligned with dentry size, but there is no file
    or directory at the position, exfat_readdir() will continue to
    get dentry from the next dentry. So the dentry gotten by readdir()
    may not be at the position.

    After this commit, if the position is not aligned with the dentry
    size, round the position up to the dentry size and continue to get
    the dentry.

    Fixes: ca06197382bd ("exfat: add directory operations")
    Cc: stable@vger.kernel.org # v5.7+
    Reported-by: Wang Yugui <wangyugui@e16-tech.com>
    Signed-off-by: Yuezhang Mo <Yuezhang.Mo@sony.com>
    Reviewed-by: Andy Wu <Andy.Wu@sony.com>
    Reviewed-by: Aoyama Wataru <wataru.aoyama@sony.com>
    Reviewed-by: Sungjong Seo <sj1557.seo@samsung.com>
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9bb5e48c5270692934da8751db33ae0c80747701
Author: Yuezhang Mo <Yuezhang.Mo@sony.com>
Date:   Thu Oct 20 14:27:37 2022 +0800

    exfat: fix reporting fs error when reading dir beyond EOF

    commit 706fdcac002316893434d753be8cfb549fe1d40d upstream.

    Since seekdir() does not check whether the position is valid, the
    position may exceed the size of the directory. We found that for
    a directory with discontinuous clusters, if the position exceeds
    the size of the directory and the excess size is greater than or
    equal to the cluster size, exfat_readdir() will return -EIO,
    causing a file system error and making the file system unavailable.

    Reproduce this bug by:

    seekdir(dir, dir_size + cluster_size);
    dirent = readdir(dir);

    The following log will be printed if mount with 'errors=remount-ro'.

    [11166.712896] exFAT-fs (sdb1): error, invalid access to FAT (entry 0xffffffff)
    [11166.712905] exFAT-fs (sdb1): Filesystem has been set read-only

    Fixes: 1e5654de0f51 ("exfat: handle wrong stream entry size in exfat_readdir()")
    Cc: stable@vger.kernel.org # v5.7+
    Signed-off-by: Yuezhang Mo <Yuezhang.Mo@sony.com>
    Reviewed-by: Andy Wu <Andy.Wu@sony.com>
    Reviewed-by: Aoyama Wataru <wataru.aoyama@sony.com>
    Reviewed-by: Sungjong Seo <sj1557.seo@samsung.com>
```

```
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f348b373cebeaa91bf3552099026951189f14436
Author: Dongliang Mu <mudongliangabcd@gmail.com>
Date:   Sun Feb 26 20:49:47 2023 +0800

    fs: hfsplus: fix UAF issue in hfsplus_put_super

    commit 07db5e247ab5858439b14dd7cc1fe538b9efcf32 upstream.

    The current hfsplus_put_super first calls hfs_btree_close on
    sbi->ext_tree, then invokes iput on sbi->hidden_dir, resulting in an
    use-after-free issue in hfsplus_release_folio.

    As shown in hfsplus_fill_super, the error handling code also calls iput
    before hfs_btree_close.

    To fix this error, we move all iput calls before hfsplus_btree_close.

    Note that this patch is tested on Syzbot.

    Link: https://lkml.kernel.org/r/20230226124948.3175736-1-mudongliangabcd@gmail.com
    Reported-by: syzbot+57e3e98f7e3b80f64d56@syzkaller.appspotmail.com
    Tested-by: Dongliang Mu <mudongliangabcd@gmail.com>
    Signed-off-by: Dongliang Mu <mudongliangabcd@gmail.com>
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: Jens Axboe <axboe@kernel.dk>
    Cc: Muchun Song <songmuchun@bytedance.com>
    Cc: Roman Gushchin <roman.gushchin@linux.dev>
    Cc: "Theodore Ts'o" <tytso@mit.edu>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 38d72e6604b9f96dffcc0565090cc01622a37b2a
Author: Liu Shixin <liushixin2@huawei.com>
Date:   Mon Dec 12 10:16:27 2022 +0800

    hfs: fix missing hfs_bnode_get() in __hfs_bnode_create

    commit a9dc087fd3c484fd1ed18c5efb290efaaf44ce03 upstream.

    Syzbot found a kernel BUG in hfs_bnode_put():

     kernel BUG at fs/hfs/bnode.c:466!
     invalid opcode: 0000 [#1] PREEMPT SMP KASAN
     CPU: 0 PID: 3634 Comm: kworker/u4:5 Not tainted 6.1.0-rc7-syzkaller-00190-g97ee9d1c1696
#0
     Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
10/26/2022
     Workqueue: writeback wb_workfn (flush-7:0)
     RIP: 0010:hfs_bnode_put+0x46f/0x480 fs/hfs/bnode.c:466
     Code: 8a 80 ff e9 73 fe ff ff 89 d9 80 e1 07 80 c1 03 38 c1 0f 8c a0 fe ff ff 48 89 df
e8 db 8a 80 ff e9 93 fe ff ff e8 a1 68 2c ff <0f> 0b e8 9a 68 2c ff 0f 0b 0f 1f 84 00 00 00
00 00 55 41 57 41 56
     RSP: 0018:ffffc90003b4f258 EFLAGS: 00010293
     RAX: ffffffff825e318f RBX: 0000000000000000 RCX: ffff8880739dd7c0
     RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000000
     RBP: ffffc90003b4f430 R08: ffffffff825e2d9b R09: ffffed10045157d1
     R10: ffffed10045157d1 R11: 1ffff110045157d0 R12: ffff8880228abe80
     R13: ffff88807016c000 R14: dffffc0000000000 R15: ffff8880228abe00
     FS:  0000000000000000(0000) GS:ffff8880b9800000(0000) knlGS:0000000000000000
     CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
     CR2: 00007fa6ebe88718 CR3: 000000001e93d000 CR4: 00000000003506f0
```

```
 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
 Call Trace:
  <TASK>
  hfs_write_inode+0x1bc/0xb40
  write_inode fs/fs-writeback.c:1440 [inline]
  __writeback_single_inode+0x4d6/0x670 fs/fs-writeback.c:1652
  writeback_sb_inodes+0xb3b/0x18f0 fs/fs-writeback.c:1878
  __writeback_inodes_wb+0x125/0x420 fs/fs-writeback.c:1949
  wb_writeback+0x440/0x7b0 fs/fs-writeback.c:2054
  wb_check_start_all fs/fs-writeback.c:2176 [inline]
  wb_do_writeback fs/fs-writeback.c:2202 [inline]
  wb_workfn+0x827/0xef0 fs/fs-writeback.c:2235
  process_one_work+0x877/0xdb0 kernel/workqueue.c:2289
  worker_thread+0xb14/0x1330 kernel/workqueue.c:2436
  kthread+0x266/0x300 kernel/kthread.c:376
  ret_from_fork+0x1f/0x30 arch/x86/entry/entry_64.S:306
  </TASK>

The BUG_ON() is triggered at here:

/* Dispose of resources used by a node */
void hfs_bnode_put(struct hfs_bnode *node)
{
        if (node) {
                <skipped>
                BUG_ON(!atomic_read(&node->refcnt)); <- we have issue here!!!!
                <skipped>
        }
}

By tracing the refcnt, I found the node is created by hfs_bmap_alloc()
with refcnt 1.  Then the node is used by hfs_btree_write().  There is a
missing of hfs_bnode_get() after find the node.  The issue happened in
following path:

<alloc>
 hfs_bmap_alloc
   hfs_bnode_find
     __hfs_bnode_create   <- allocate a new node with refcnt 1.
   hfs_bnode_put          <- decrease the refcnt

<write>
 hfs_btree_write
   hfs_bnode_find
     __hfs_bnode_create
       hfs_bnode_findhash <- find the node without refcnt increased.
   hfs_bnode_put          <- trigger the BUG_ON() since refcnt is 0.

Link: https://lkml.kernel.org/r/20221212021627.3766829-1-liushixin2@huawei.com
Reported-by: syzbot+5b04b49a7ec7226c7426@syzkaller.appspotmail.com
Signed-off-by: Liu Shixin <liushixin2@huawei.com>
Cc: Fabio M. De Francesco <fmdefrancesco@gmail.com>
Cc: Viacheslav Dubeyko <slava@dubeyko.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 eaa2a19d24f1eb3d2b7bb83ed2c150e23505bd1a
Author: Jens Axboe <axboe@kernel.dk>
Date:   Mon Feb 6 08:20:46 2023 -0700

    io_uring: mark task TASK_RUNNING before handling resume/task work

    commit 2f2bb1ffc9983e227424d0787289da5483b0c74f upstream.
```

```
    Just like for task_work, set the task mode to TASK_RUNNING before doing
    any potential resume work. We're not holding any locks at this point,
    but we may have already set the task state to TASK_INTERRUPTIBLE in
    preparation for going to sleep waiting for events. Ensure that we set it
    back to TASK_RUNNING if we have work to process, to avoid warnings on
    calling blocking operations with !TASK_RUNNING.

    Fixes: b5d3ae202fbf ("io_uring: handle TIF_NOTIFY_RESUME when checking for task_work")
    Reported-by: kernel test robot <oliver.sang@intel.com>
    Link: https://lore.kernel.org/oe-lkp/202302062208.24d3e563-oliver.sang@intel.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 156aab6f9fc782d146ebccb5d78e033934113d6b
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Wed Jan 25 10:45:05 2023 +0100

    ARM: dts: exynos: correct HDMI phy compatible in Exynos4

    commit af1c89ddb74f170eccd5a57001d7317560b638ea upstream.

    The HDMI phy compatible was missing vendor prefix.

    Fixes: ed80d4cab772 ("ARM: dts: add hdmi related nodes for exynos4 SoCs")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230125094513.155063-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6f7cd8467c548e6cf3c60ad64aff619a1e3e2678
Author: Joel Fernandes (Google) <joel@joelfernandes.org>
Date:   Sun Jan 1 06:15:55 2023 +0000

    torture: Fix hang during kthread shutdown phase

    commit d52d3a2bf408ff86f3a79560b5cce80efb340239 upstream.

    During rcutorture shutdown, the rcu_torture_cleanup() function calls
    torture_cleanup_begin(), which sets the fullstop global variable to
    FULLSTOP_RMMOD. This causes the rcutorture threads for readers and
    fakewriters to exit all of their "while" loops and start shutting down.

    They then call torture_kthread_stopping(), which in turn waits for
    kthread_stop() to be called.  However, rcu_torture_cleanup() has
    not yet called kthread_stop() on those threads, and before it gets a
    chance to do so, multiple instances of torture_kthread_stopping() invoke
    schedule_timeout_interruptible(1) in a tight loop.  Tracing confirms that
    TIMER_SOFTIRQ can then continuously execute timer callbacks.  If that
    TIMER_SOFTIRQ preempts the task executing rcu_torture_cleanup(), that
    task might never invoke kthread_stop().

    This commit improves this situation by increasing the timeout passed to
    schedule_timeout_interruptible() from one jiffy to 1/20th of a second.
    This change prevents TIMER_SOFTIRQ from monopolizing its CPU, thus
    allowing rcu_torture_cleanup() to carry out the needed kthread_stop()
    invocations.  Testing has shown 100 runs of TREE07 passing reliably,
    as oppose to the tens-of-percent failure rates seen beforehand.

    Cc: Paul McKenney <paulmck@kernel.org>
    Cc: Frederic Weisbecker <fweisbec@gmail.com>
    Cc: Zhouyi Zhou <zhouzhouyi@gmail.com>
    Cc: <stable@vger.kernel.org> # 6.0.x
    Signed-off-by: Joel Fernandes (Google) <joel@joelfernandes.org>
    Tested-by: Zhouyi Zhou <zhouzhouyi@gmail.com>
```

```
    Reviewed-by: Davidlohr Bueso <dave@stgolabs.net>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 11d38f8a0c19763e34d2093b5ecb640e012cb2d2
Author: Hangyu Hua <hbh25y@gmail.com>
Date:   Fri Feb 17 22:29:34 2023 +0900

    ksmbd: fix possible memory leak in smb2_lock()

    commit d3ca9f7aeba793d74361d88a8800b2f205c9236b upstream.

    argv needs to be free when setup_async_work fails or when the current
    process is woken up.

    Fixes: e2f34481b24d ("cifsd: add server-side procedures for SMB3")
    Cc: stable@vger.kernel.org
    Signed-off-by: Hangyu Hua <hbh25y@gmail.com>
    Acked-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 237f4b71599d802ad87c8702a83465c6523b7527
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Sat Feb 11 00:27:34 2023 +0900

    ksmbd: do not allow the actual frame length to be smaller than the rfc1002 length

    commit fb533473d1595fe79ecb528fda1de33552b07178 upstream.

    ksmbd allowed the actual frame length to be smaller than the rfc1002
    length. If allowed, it is possible to allocates a large amount of memory
    that can be limited by credit management and can eventually cause memory
    exhaustion problem. This patch do not allow it except SMB2 Negotiate
    request which will be validated when message handling proceeds.
    Also, Allow a message that padded to 8byte boundary.

    Fixes: e2f34481b24d ("cifsd: add server-side procedures for SMB3")
    Cc: stable@vger.kernel.org
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit af2639fc998e79ccd781242193beab5ae395d697
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Wed Feb 8 09:50:46 2023 +0900

    ksmbd: fix wrong data area length for smb2 lock request

    commit 8f8c43b125882ac14372f8dca0c8e50a59e78d79 upstream.

    When turning debug mode on, The following error message from
    ksmbd_smb2_check_message() is coming.

    ksmbd: cli req padded more than expected. Length 112 not 88 for cmd:10 mid:14

    data area length calculation for smb2 lock request in smb2_get_data_area_len() is
    incorrect.

    Fixes: e2f34481b24d ("cifsd: add server-side procedures for SMB3")
    Cc: stable@vger.kernel.org
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit a513dd803d67b36ff81590c5c4cb2a42bc6e4429
Author: Waiman Long <longman@redhat.com>
Date:   Wed Jan 25 19:36:25 2023 -0500

    locking/rwsem: Prevent non-first waiter from spinning in down_write() slowpath

    commit b613c7f31476c44316bfac1af7cac714b7d6bef9 upstream.

    A non-first waiter can potentially spin in the for loop of
    rwsem_down_write_slowpath() without sleeping but fail to acquire the
    lock even if the rwsem is free if the following sequence happens:

      Non-first RT waiter    First waiter      Lock holder
      -------------------    ------------      -----------
      Acquire wait_lock
      rwsem_try_write_lock():
        Set handoff bit if RT or
          wait too long
        Set waiter->handoff_set
      Release wait_lock
                             Acquire wait_lock
                             Inherit waiter->handoff_set
                             Release wait_lock
                                               Clear owner
                                               Release lock
      if (waiter.handoff_set) {
        rwsem_spin_on_owner(();
        if (OWNER_NULL)
          goto trylock_again;
      }
      trylock_again:
      Acquire wait_lock
      rwsem_try_write_lock():
         if (first->handoff_set && (waiter != first))
            return false;
      Release wait_lock

    A non-first waiter cannot really acquire the rwsem even if it mistakenly
    believes that it can spin on OWNER_NULL value. If that waiter happens
    to be an RT task running on the same CPU as the first waiter, it can
    block the first waiter from acquiring the rwsem leading to live lock.
    Fix this problem by making sure that a non-first waiter cannot spin in
    the slowpath loop without sleeping.

    Fixes: d257cc8cb8d5 ("locking/rwsem: Make handoff bit handling more consistent")
    Signed-off-by: Waiman Long <longman@redhat.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Tested-by: Mukesh Ojha <quic_mojha@quicinc.com>
    Reviewed-by: Mukesh Ojha <quic_mojha@quicinc.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230126003628.365092-2-longman@redhat.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d534fa317606208ca00d6c7f49694159461c0e1a
Author: Qu Wenruo <wqu@suse.com>
Date:   Fri Jan 13 19:11:39 2023 +0800

    btrfs: sysfs: update fs features directory asynchronously

    commit b7625f461da6734a21c38ba6e7558538a116a2e3 upstream.

    [BUG]
    Since the introduction of per-fs feature sysfs interface
    (/sys/fs/btrfs/<UUID>/features/), the content of that directory is never
    updated.
```

```
    Thus for the following case, that directory will not show the new
    features like RAID56:

      # mkfs.btrfs -f $dev1 $dev2 $dev3
      # mount $dev1 $mnt
      # btrfs balance start -f -mconvert=raid5 $mnt
      # ls /sys/fs/btrfs/$uuid/features/
      extended_iref  free_space_tree  no_holes  skinny_metadata

    While after unmount and mount, we got the correct features:

      # umount $mnt
      # mount $dev1 $mnt
      # ls /sys/fs/btrfs/$uuid/features/
      extended_iref  free_space_tree  no_holes  raid56 skinny_metadata

    [CAUSE]
    Because we never really try to update the content of per-fs features/
    directory.

    We had an attempt to update the features directory dynamically in commit
    14e46e04958d ("btrfs: synchronize incompat feature bits with sysfs
    files"), but unfortunately it get reverted in commit e410e34fad91
    ("Revert "btrfs: synchronize incompat feature bits with sysfs files"").
    The problem in the original patch is, in the context of
    btrfs_create_chunk(), we can not afford to update the sysfs group.

    The exported but never utilized function, btrfs_sysfs_feature_update()
    is the leftover of such attempt.  As even if we go sysfs_update_group(),
    new files will need extra memory allocation, and we have no way to
    specify the sysfs update to go GFP_NOFS.

    [FIX]
    This patch will address the old problem by doing asynchronous sysfs
    update in the cleaner thread.

    This involves the following changes:

    - Make __btrfs_(set|clear)_fs_(incompat|compat_ro) helpers to set
      BTRFS_FS_FEATURE_CHANGED flag when needed

    - Update btrfs_sysfs_feature_update() to use sysfs_update_group()
      And drop unnecessary arguments.

    - Call btrfs_sysfs_feature_update() in cleaner_kthread
      If we have the BTRFS_FS_FEATURE_CHANGED flag set.

    - Wake up cleaner_kthread in btrfs_commit_transaction if we have
      BTRFS_FS_FEATURE_CHANGED flag

    By this, all the previously dangerous call sites like
    btrfs_create_chunk() need no new changes, as above helpers would
    have already set the BTRFS_FS_FEATURE_CHANGED flag.

    The real work happens at cleaner_kthread, thus we pay the cost of
    delaying the update to sysfs directory, but the delayed time should be
    small enough that end user can not distinguish though it might get
    delayed if the cleaner thread is busy with removing subvolumes or
    defrag.

    CC: stable@vger.kernel.org # 4.14+
    Reviewed-by: Anand Jain <anand.jain@oracle.com>
    Signed-off-by: Qu Wenruo <wqu@suse.com>
    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 ab71c266e4335d18ddee7a5669d072733c2c49fc
Author: Boris Burkov <boris@bur.io>
Date:   Thu Jan 12 16:05:11 2023 -0800

    btrfs: hold block group refcount during async discard

    commit 2b5463fcbdfb24e898916bcae2b1359042d26963 upstream.

    Async discard does not acquire the block group reference count while it
    holds a reference on the discard list. This is generally OK, as the
    paths which destroy block groups tend to try to synchronize on
    cancelling async discard work. However, relying on cancelling work
    requires careful analysis to be sure it is safe from races with
    unpinning scheduling more work.

    While I am unable to find a race with unpinning in the current code for
    either the unused bgs or relocation paths, I believe we have one in an
    older version of auto relocation in a Meta internal build. This suggests
    that this is in fact an error prone model, and could be fragile to
    future changes to these bg deletion paths.

    To make this ownership more clear, add a refcount for async discard. If
    work is queued for a block group, its refcount should be incremented,
    and when work is completed or canceled, it should be decremented.

    CC: stable@vger.kernel.org # 5.15+
    Signed-off-by: Boris Burkov <boris@bur.io>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8a3ebb7636002ca0866068701fcb11ee7dc105e1
Author: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
Date:   Tue Feb 14 09:50:17 2023 +0900

    scsi: mpi3mr: Remove unnecessary memcpy() to alltgt_info->dmi

    commit eeb270aee3e085411399f129fc14fa04bd6d83cf upstream.

    In the function mpi3mr_get_all_tgt_info(), devmap_info points to
    alltgt_info->dmi then there is no need to memcpy() data from devmap_info to
    alltgt_info->dmi. Remove the unnecessary memcpy(). This also allows to
    remove the local variable 'rval' and the goto label 'out'.

    Link: https://lore.kernel.org/r/20230214005019.1897251-3-shinichiro.kawasaki@wdc.com
    Cc: stable@vger.kernel.org
    Fixes: f5e6d5a34376 ("scsi: mpi3mr: Add support for driver commands")
    Signed-off-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Acked-by: Sathya Prakash Veerichetty <sathya.prakash@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2f3d3fa5b8ed7d3b147478f42b00b468eeb1ecd2
Author: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
Date:   Tue Feb 14 09:50:16 2023 +0900

    scsi: mpi3mr: Fix issues in mpi3mr_get_all_tgt_info()

    commit fb428a2005fc1260d18b989cc5199f281617f44d upstream.

    The function mpi3mr_get_all_tgt_info() has four issues:

    1) It calculates valid entry length in alltgt_info assuming the header part
```

```
        of the struct mpi3mr_device_map_info would equal to sizeof(u32).  The
        correct size is sizeof(u64).

    2) When it calculates the valid entry length kern_entrylen, it excludes one
       entry by subtracting 1 from num_devices.

    3) It copies num_device by calling memcpy(). Substitution is enough.

    4) It does not specify the calculated length to sg_copy_from_buffer().
       Instead, it specifies the payload length which is larger than the
       alltgt_info size. It causes "BUG: KASAN: slab-out-of-bounds".

    Fix the issues by using the correct header size, removing the subtraction
    from num_devices, replacing the memcpy() with substitution and specifying
    the correct length to sg_copy_from_buffer().

    Link: https://lore.kernel.org/r/20230214005019.1897251-2-shinichiro.kawasaki@wdc.com
    Cc: stable@vger.kernel.org
    Fixes: f5e6d5a34376 ("scsi: mpi3mr: Add support for driver commands")
    Signed-off-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Acked-by: Sathya Prakash Veerichetty <sathya.prakash@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 67989091e11a974003ddf2ec39bc613df8eadd83
Author: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
Date:   Tue Feb 14 09:50:19 2023 +0900

    scsi: mpi3mr: Fix missing mrioc->evtack_cmds initialization

    commit e39ea831ebad4ab15c4748cb62a397a8abcca36e upstream.

    Commit c1af985d27da ("scsi: mpi3mr: Add Event acknowledgment logic")
    introduced an array mrioc->evtack_cmds but initialization of the array
    elements was missed. They are just zero cleared. The function
    mpi3mr_complete_evt_ack() refers host_tag field of the elements. Due to the
    zero value of the host_tag field, the function calls clear_bit() for
    mrico->evtack_cmds_bitmap with wrong bit index. This results in memory
    access to invalid address and "BUG: KASAN: use-after-free". This BUG was
    observed at eHBA-9600 firmware update to version 8.3.1.0. To fix it, add
    the missing initialization of mrioc->evtack_cmds.

    Link: https://lore.kernel.org/r/20230214005019.1897251-5-shinichiro.kawasaki@wdc.com
    Cc: stable@vger.kernel.org
    Fixes: c1af985d27da ("scsi: mpi3mr: Add Event acknowledgment logic")
    Signed-off-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Reviewed-by: Damien Le Moal <damien.lemoal@opensource.wdc.com>
    Acked-by: Sathya Prakash Veerichetty <sathya.prakash@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 83f33885c5b046c8c6c1c7eb92baa2ed0d259702
Author: Ronnie Sahlberg <lsahlber@redhat.com>
Date:   Fri Feb 17 13:35:01 2023 +1000

    cifs: return a single-use cfid if we did not get a lease

    commit 8e843bf38f7be0766642a91523cfa65f2b021a8a upstream.

    If we did not get a lease we can still return a single use cfid to the caller.
    The cfid will not have has_lease set and will thus not be shared with any
    other concurrent users and will be freed immediately when the caller
    drops the handle.

    This avoids extra roundtrips for servers that do not support directory leases
```

```
    where they would first fail to get a cfid with a lease and then fallback
    to try a normal SMB2_open()

    Signed-off-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Bharath SM <bharathsm@microsoft.com>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@manguebit.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6391b89eeb44e175d8ae6a1a52a94e5393cdad9a
Author: Ronnie Sahlberg <lsahlber@redhat.com>
Date:   Fri Feb 17 13:35:00 2023 +1000

    cifs: Check the lease context if we actually got a lease

    commit 66d45ca1350a3bb8d5f4db8879ccad3ed492337a upstream.

    Some servers may return that we got a lease in rsp->OplockLevel
    but then in the lease context contradict this and say we got no lease
    at all.  Thus we need to check the context if we have a lease.
    Additionally, If we do not get a lease we need to make sure we close
    the handle before we return an error to the caller.

    Signed-off-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Bharath SM <bharathsm@microsoft.com>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@manguebit.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ef0cd88096f514f18b2f9e15c1703b5cc7b6cf5a
Author: Stefan Metzmacher <metze@samba.org>
Date:   Wed Feb 1 16:21:41 2023 +0100

    cifs: don't try to use rdma offload on encrypted connections

    commit 3891f6c7655a39065e44980f51ba46bb32be3133 upstream.

    The aim of using encryption on a connection is to keep
    the data confidential, so we must not use plaintext rdma offload
    for that data!

    It seems that current windows servers and ksmbd would allow
    this, but that's no reason to expose the users data in plaintext!
    And servers hopefully reject this in future.

    Note modern windows servers support signed or encrypted offload,
    see MS-SMB2 2.2.3.1.6 SMB2_RDMA_TRANSFORM_CAPABILITIES, but we don't
    support that yet.

    Signed-off-by: Stefan Metzmacher <metze@samba.org>
    Cc: Steve French <smfrench@gmail.com>
    Cc: Tom Talpey <tom@talpey.com>
    Cc: Long Li <longli@microsoft.com>
    Cc: Namjae Jeon <linkinjeon@kernel.org>
    Cc: David Howells <dhowells@redhat.com>
    Cc: linux-cifs@vger.kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 74ab77137c99438626f4eb8134d8e26204bb5b64
Author: Stefan Metzmacher <metze@samba.org>
Date:   Wed Feb 1 16:21:40 2023 +0100
```

```
    cifs: split out smb3_use_rdma_offload() helper

    commit a6559cc1d35d3eeafb0296aca347b2f745a28a74 upstream.

    We should have the logic to decide if we want rdma offload
    in a single spot in order to advance it in future.

    Signed-off-by: Stefan Metzmacher <metze@samba.org>
    Cc: Steve French <smfrench@gmail.com>
    Cc: Tom Talpey <tom@talpey.com>
    Cc: Long Li <longli@microsoft.com>
    Cc: Namjae Jeon <linkinjeon@kernel.org>
    Cc: David Howells <dhowells@redhat.com>
    Cc: linux-cifs@vger.kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7b4002f74e7440c032df3736f81f9a461ca63fc0
Author: Stefan Metzmacher <metze@samba.org>
Date:   Wed Feb 1 16:21:39 2023 +0100

    cifs: introduce cifs_io_parms in smb2_async_writev()

    commit d643a8a446fc46c06837d08a056f69da2ff16025 upstream.

    This will simplify the following changes and makes it easy to get
    in passed in from the caller in future.

    Signed-off-by: Stefan Metzmacher <metze@samba.org>
    Cc: Steve French <smfrench@gmail.com>
    Cc: Tom Talpey <tom@talpey.com>
    Cc: Long Li <longli@microsoft.com>
    Cc: Namjae Jeon <linkinjeon@kernel.org>
    Cc: David Howells <dhowells@redhat.com>
    Cc: linux-cifs@vger.kernel.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c53236ef6a546891be2f4562f8434be9b5fa7a6b
Author: Paulo Alcantara <pc@manguebit.com>
Date:   Thu Feb 16 15:33:22 2023 -0300

    cifs: fix mount on old smb servers

    commit d99e86ebde2d7b3a04190f8d14de5bf6814bf10f upstream.

    The client was sending rfc1002 session request packet with a wrong
    length field set, therefore failing to mount shares against old SMB
    servers over port 139.

    Fix this by calculating the correct length as specified in rfc1002.

    Fixes: d7173623bf0b ("cifs: use ALIGN() and round_up() macros")
    Cc: stable@vger.kernel.org
    Signed-off-by: Paulo Alcantara (SUSE) <pc@manguebit.com>
    Reviewed-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ad76ca6e9142963217fddb02b1aa60f2e9ff72f8
Author: Volker Lendecke <vl@samba.org>
Date:   Wed Jan 11 12:37:58 2023 +0100
```

```
    cifs: Fix uninitialized memory reads for oparms.mode

    commit de036dcaca65cf94bf7ff09c571c077f02bc92b4 upstream.

    Use a struct assignment with implicit member initialization

    Signed-off-by: Volker Lendecke <vl@samba.org>
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2c72ff0bff47b0df45e7eb5060f29529350f31f7
Author: Volker Lendecke <vl@samba.org>
Date:   Wed Jan 11 12:37:58 2023 +0100

    cifs: Fix uninitialized memory read in smb3_qfs_tcon()

    commit d447e794a37288ec7a080aa1b044a8d9deebbab7 upstream.

    oparms was not fully initialized

    Signed-off-by: Volker Lendecke <vl@samba.org>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Cc: stable@vger.kernel.org
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a4ee01374fe3ed4e968d6ac1d570f85aed59a09d
Author: Paulo Alcantara <pc@manguebit.com>
Date:   Tue Feb 28 19:01:54 2023 -0300

    cifs: improve checking of DFS links over STATUS_OBJECT_NAME_INVALID

    commit b9ee2e307c6b06384b6f9e393a9b8e048e8fc277 upstream.

    Do not map STATUS_OBJECT_NAME_INVALID to -EREMOTE under non-DFS
    shares, or 'nodfs' mounts or CONFIG_CIFS_DFS_UPCALL=n builds.
    Otherwise, in the slow path, get a referral to figure out whether it
    is an actual DFS link.

    This could be simply reproduced under a non-DFS share by running the
    following

      $ mount.cifs //srv/share /mnt -o ...
      $ cat /mnt/$(printf '\U110000')
      cat: '/mnt/'$'\364\220\200\200': Object is remote

    Fixes: c877ce47e137 ("cifs: reduce roundtrips on create/qinfo requests")
    CC: stable@vger.kernel.org # 6.2
    Signed-off-by: Paulo Alcantara (SUSE) <pc@manguebit.com>
    Reviewed-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d8682c469157ebe28243eb595880205403c81c95
Author: Nico Boehr <nrb@linux.ibm.com>
Date:   Fri Jan 27 15:05:32 2023 +0100

    KVM: s390: disable migration mode when dirty tracking is disabled

    commit f2d3155e2a6bac44d16f04415a321e8707d895c6 upstream.

    Migration mode is a VM attribute which enables tracking of changes in
    storage attributes (PGSTE). It assumes dirty tracking is enabled on all
```

```
    memslots to keep a dirty bitmap of pages with changed storage attributes.

    When enabling migration mode, we currently check that dirty tracking is
    enabled for all memslots. However, userspace can disable dirty tracking
    without disabling migration mode.

    Since migration mode is pointless with dirty tracking disabled, disable
    migration mode whenever userspace disables dirty tracking on any slot.

    Also update the documentation to clarify that dirty tracking must be
    enabled when enabling migration mode, which is already enforced by the
    code in kvm_s390_vm_start_migration().

    Also highlight in the documentation for KVM_S390_GET_CMMA_BITS that it
    can now fail with -EINVAL when dirty tracking is disabled while
    migration mode is on. Move all the error codes to a table so this stays
    readable.

    To disable migration mode, slots_lock should be held, which is taken
    in kvm_set_memory_region() and thus held in
    kvm_arch_prepare_memory_region().

    Restructure the prepare code a bit so all the sanity checking is done
    before disabling migration mode. This ensures migration mode isn't
    disabled when some sanity check fails.

    Cc: stable@vger.kernel.org
    Fixes: 190df4a212a7 ("KVM: s390: CMMA tracking, ESSA emulation, migration mode")
    Signed-off-by: Nico Boehr <nrb@linux.ibm.com>
    Reviewed-by: Janosch Frank <frankja@linux.ibm.com>
    Reviewed-by: Claudio Imbrenda <imbrenda@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230127140532.230651-2-nrb@linux.ibm.com
    Message-Id: <20230127140532.230651-2-nrb@linux.ibm.com>
    [frankja@linux.ibm.com: fixed commit message typo, moved api.rst error table upwards]
    Signed-off-by: Janosch Frank <frankja@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 141bdce1cb3c4829d965d1946965d70c867887df
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Wed Mar 1 17:58:06 2023 +0100

    s390/kprobes: fix current_kprobe never cleared after kprobes reenter

    commit cd57953936f2213dfaccce10d20f396956222c7d upstream.

    Recent test_kprobe_missed kprobes kunit test uncovers the following
    problem. Once kprobe is triggered from another kprobe (kprobe reenter),
    all future kprobes on this cpu are considered as kprobe reenter, thus
    pre_handler and post_handler are not being called and kprobes are counted
    as "missed".

    Commit b9599798f953 ("[S390] kprobes: activation and deactivation")
    introduced a simpler scheme for kprobes (de)activation and status
    tracking by using push_kprobe/pop_kprobe, which supposed to work for
    both initial kprobe entry as well as kprobe reentry and helps to avoid
    handling those two cases differently. The problem is that a sequence of
    calls in case of kprobes reenter:
    push_kprobe() <- NULL (current_kprobe)
    push_kprobe() <- kprobe1 (current_kprobe)
    pop_kprobe() -> kprobe1 (current_kprobe)
    pop_kprobe() -> kprobe1 (current_kprobe)
    leaves "kprobe1" as "current_kprobe" on this cpu, instead of setting it
    to NULL. In fact push_kprobe/pop_kprobe can only store a single state
    (there is just one prev_kprobe in kprobe_ctlblk). Which is a hack but
    sufficient, there is no need to have another prev_kprobe just to store
```

```
    NULL. To make a simple and backportable fix simply reset "prev_kprobe"
    when kprobe is poped from this "stack". No need to worry about
    "kprobe_status" in this case, because its value is only checked when
    current_kprobe != NULL.

    Cc: stable@vger.kernel.org
    Fixes: b9599798f953 ("[S390] kprobes: activation and deactivation")
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2110de0095e0deb4aaeba2af18d7ec0fc077da91
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Wed Mar 1 02:23:08 2023 +0100

    s390/kprobes: fix irq mask clobbering on kprobe reenter from post_handler

    commit 42e19e6f04984088b6f9f0507c4c89a8152d9730 upstream.

    Recent test_kprobe_missed kprobes kunit test uncovers the following error
    (reported when CONFIG_DEBUG_ATOMIC_SLEEP is enabled):

    BUG: sleeping function called from invalid context at kernel/locking/mutex.c:580
    in_atomic(): 0, irqs_disabled(): 1, non_block: 0, pid: 662, name: kunit_try_catch
    preempt_count: 0, expected: 0
    RCU nest depth: 0, expected: 0
    no locks held by kunit_try_catch/662.
    irq event stamp: 280
    hardirqs last  enabled at (279): [<00000003e60a3d42>] __do_pgm_check+0x17a/0x1c0
    hardirqs last disabled at (280): [<00000003e3bd774a>]
kprobe_exceptions_notify+0x27a/0x318
    softirqs last  enabled at (0): [<00000003e3c5c890>] copy_process+0x14a8/0x4c80
    softirqs last disabled at (0): [<0000000000000000>] 0x0
    CPU: 46 PID: 662 Comm: kunit_try_catch Tainted: G                 N 6.2.0-173644-
g44c18d77f0c0 #2
    Hardware name: IBM 3931 A01 704 (LPAR)
    Call Trace:
     [<00000003e60a3a00>] dump_stack_lvl+0x120/0x198
     [<00000003e3d02e82>] __might_resched+0x60a/0x668
     [<00000003e60b9908>] __mutex_lock+0xc0/0x14e0
     [<00000003e60bad5a>] mutex_lock_nested+0x32/0x40
     [<00000003e3f7b460>] unregister_kprobe+0x30/0xd8
     [<00000003e51b2602>] test_kprobe_missed+0xf2/0x268
     [<00000003e51b5406>] kunit_try_run_case+0x10e/0x290
     [<00000003e51b7dfa>] kunit_generic_run_threadfn_adapter+0x62/0xb8
     [<00000003e3ce30f8>] kthread+0x2d0/0x398
     [<00000003e3b96afa>] __ret_from_fork+0x8a/0xe8
     [<00000003e60ccada>] ret_from_fork+0xa/0x40

    The reason for this error report is that kprobes handling code failed
    to restore irqs.

    The problem is that when kprobe is triggered from another kprobe
    post_handler current sequence of enable_singlestep / disable_singlestep
    is the following:
    enable_singlestep  <- original kprobe (saves kprobe_saved_imask)
    enable_singlestep  <- kprobe triggered from post_handler (clobbers kprobe_saved_imask)
    disable_singlestep <- kprobe triggered from post_handler (restores kprobe_saved_imask)
    disable_singlestep <- original kprobe (restores wrong clobbered kprobe_saved_imask)

    There is just one kprobe_ctlblk per cpu and both calls saves and
    loads irq mask to kprobe_saved_imask. To fix the problem simply move
    resume_execution (which calls disable_singlestep) before calling
    post_handler. This also fixes the problem that post_handler is called
```

```
    with pt_regs which were not yet adjusted after single-stepping.

    Cc: stable@vger.kernel.org
    Fixes: 4ba069b802c2 ("[S390] add kprobes support.")
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit be1b2f19b77fa317861264622187b872bc30e295
Author: Sven Schnelle <svens@linux.ibm.com>
Date:   Sun Jan 29 16:39:19 2023 +0100

    s390/ipl: add loadparm parameter to eckd ipl/reipl data

    commit 6bb361d5d8eb1dbc9e0b190eeee27a2ac4d1119f upstream.

    commit 87fd22e0ae92 ("s390/ipl: add eckd support") missed to add the
    loadparm attribute to the new eckd ipl/reipl data.

    Fixes: 87fd22e0ae92 ("s390/ipl: add eckd support")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Sven Schnelle <svens@linux.ibm.com>
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit facf2a129f44432dbe883f8ebe83ec955452aa03
Author: Sven Schnelle <svens@linux.ibm.com>
Date:   Sun Jan 29 16:36:03 2023 +0100

    s390/ipl: add DEFINE_GENERIC_LOADPARM()

    commit c676aac66f5b2b03a1090bc6b1891486255f7159 upstream.

    In the current code each reipl type implements its own pair of loadparm
    show/store functions. Add a macro to deduplicate the code a bit.

    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sven Schnelle <svens@linux.ibm.com>
    Fixes: 87fd22e0ae92 ("s390/ipl: add eckd support")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 45301870452fc423707da40f1578b3b40fb84c32
Author: Ilya Leoshkevich <iii@linux.ibm.com>
Date:   Mon Jan 23 22:50:32 2023 +0100

    s390: discard .interp section

    commit e9c9cb90e76ffaabcc7ca8f275d9e82195fd6367 upstream.

    When debugging vmlinux with QEMU + GDB, the following GDB error may
    occur:

        (gdb) c
        Continuing.
        Warning:
        Cannot insert breakpoint -1.
        Cannot access memory at address 0xfffffffffff95c0

        Command aborted.
        (gdb)
```

```
    The reason is that, when .interp section is present, GDB tries to
    locate the file specified in it in memory and put a number of
    breakpoints there (see enable_break() function in gdb/solib-svr4.c).
    Sometimes GDB finds a bogus location that matches its heuristics,
    fails to set a breakpoint and stops. This makes further debugging
    impossible.

    The .interp section contains misleading information anyway (vmlinux
    does not need ld.so), so fix by discarding it.

    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c9d079a5b44c43412dceb2abaf70ee556263e9d0
Author: Gerald Schaefer <gerald.schaefer@linux.ibm.com>
Date:   Mon Feb 27 20:03:00 2023 +0100

    s390/extmem: return correct segment type in __segment_load()

    commit 8c42dd78df148c90e48efff204cce38743906a79 upstream.

    Commit f05f62d04271f ("s390/vmem: get rid of memory segment list")
    reshuffled the call to vmem_add_mapping() in __segment_load(), which now
    overwrites rc after it was set to contain the segment type code.

    As result, __segment_load() will now always return 0 on success, which
    corresponds to the segment type code SEG_TYPE_SW, i.e. a writeable
    segment. This results in a kernel crash when loading a read-only segment
    as dcssblk block device, and trying to write to it.

    Instead of reshuffling code again, make sure to return the segment type
    on success, and also describe this rather delicate and unexpected logic
    in the function comment. Also initialize new segtype variable with
    invalid value, to prevent possible future confusion.

    Fixes: f05f62d04271 ("s390/vmem: get rid of memory segment list")
    Cc: <stable@vger.kernel.org> # 5.9+
    Signed-off-by: Gerald Schaefer <gerald.schaefer@linux.ibm.com>
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 10fb2e16ee6ffaf1716b9e90d007e6b300bfa457
Author: Joseph Qi <joseph.qi@linux.alibaba.com>
Date:   Tue Feb 28 12:54:59 2023 +0800

    io_uring: fix fget leak when fs don't support nowait buffered read

    commit 54aa7f2330b82884f4a1afce0220add6e8312f8b upstream.

    Heming reported a BUG when using io_uring doing link-cp on ocfs2. [1]

    Do the following steps can reproduce this BUG:
    mount -t ocfs2 /dev/vdc /mnt/ocfs2
    cp testfile /mnt/ocfs2/
    ./link-cp /mnt/ocfs2/testfile /mnt/ocfs2/testfile.1
    umount /mnt/ocfs2

    Then umount will fail, and it outputs:
    umount: /mnt/ocfs2: target is busy.

    While tracing umount, it blames mnt_get_count() not return as expected.
    Do a deep investigation for fget()/fput() on related code flow, I've
```

```
    finally found that fget() leaks since ocfs2 doesn't support nowait
    buffered read.

    io_issue_sqe
    |-io_assign_file  // do fget() first
      |-io_read
      |-io_iter_do_read
        |-ocfs2_file_read_iter  // return -EOPNOTSUPP
      |-kiocb_done
        |-io_rw_done
          |-__io_complete_rw_common  // set REQ_F_REISSUE
        |-io_resubmit_prep
          |-io_req_prep_async  // override req->file, leak happens

    This was introduced by commit a196c78b5443 in v5.18. Fix it by don't
    re-assign req->file if it has already been assigned.

    [1] https://lore.kernel.org/ocfs2-devel/ab580a75-91c8-d68a-3455-
40361be1bfa8@linux.alibaba.com/T/#t

    Fixes: a196c78b5443 ("io_uring: assign non-fixed early for async work")
    Cc: <stable@vger.kernel.org>
    Reported-by: Heming Zhao <heming.zhao@suse.com>
    Signed-off-by: Joseph Qi <joseph.qi@linux.alibaba.com>
    Cc: Xiaoguang Wang <xiaoguang.wang@linux.alibaba.com>
    Link: https://lore.kernel.org/r/20230228045459.13524-1-joseph.qi@linux.alibaba.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6f1e7441712234b0bed52904e2c7d4cca3e4d2b4
Author: Jens Axboe <axboe@kernel.dk>
Date:   Sat Feb 25 12:53:53 2023 -0700

    io_uring/poll: allow some retries for poll triggering spuriously

    commit c16bda37594f83147b167d381d54c010024efecf upstream.

    If we get woken spuriously when polling and fail the operation with
    -EAGAIN again, then we generally only allow polling again if data
    had been transferred at some point. This is indicated with
    REQ_F_PARTIAL_IO. However, if the spurious poll triggers when the socket
    was originally empty, then we haven't transferred data yet and we will
    fail the poll re-arm. This either punts the socket to io-wq if it's
    blocking, or it fails the request with -EAGAIN if not. Neither condition
    is desirable, as the former will slow things down, while the latter
    will make the application confused.

    We want to ensure that a repeated poll trigger doesn't lead to infinite
    work making no progress, that's what the REQ_F_PARTIAL_IO check was
    for. But it doesn't protect against a loop post the first receive, and
    it's unnecessarily strict if we started out with an empty socket.

    Add a somewhat random retry count, just to put an upper limit on the
    potential number of retries that will be done. This should be high enough
    that we won't really hit it in practice, unless something needs to be
    aborted anyway.

    Cc: stable@vger.kernel.org # v5.10+
    Link: https://github.com/axboe/liburing/issues/364
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9b5ed6b4e6c3ad9b2ca8d5f6594a6ab4a2cd9531
Author: David Lamparter <equinox@diac24.net>
Date:   Fri Feb 24 16:01:24 2023 +0100
```

```
    io_uring: remove MSG_NOSIGNAL from recvmsg

    commit 7605c43d67face310b4b87dee1a28bc0c8cd8c0f upstream.

    MSG_NOSIGNAL is not applicable for the receiving side, SIGPIPE is
    generated when trying to write to a "broken pipe".  AF_PACKET's
    packet_recvmsg() does enforce this, giving back EINVAL when MSG_NOSIGNAL
    is set - making it unuseable in io_uring's recvmsg.

    Remove MSG_NOSIGNAL from io_recvmsg_prep().

    Cc: stable@vger.kernel.org # v5.10+
    Signed-off-by: David Lamparter <equinox@diac24.net>
    Cc: Eric Dumazet <edumazet@google.com>
    Cc: Jens Axboe <axboe@kernel.dk>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Link: https://lore.kernel.org/r/20230224150123.128346-1-equinox@diac24.net
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e559fb5d87eb5dce2d414aca400f343e1cac2761
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Wed Feb 22 14:36:48 2023 +0000

    io_uring/rsrc: disallow multi-source reg buffers

    commit edd478269640b360c6f301f2baa04abdda563ef3 upstream.

    If two or more mappings go back to back to each other they can be passed
    into io_uring to be registered as a single registered buffer. That would
    even work if mappings came from different sources, e.g. it's possible to
    mix in this way anon pages and pages from shmem or hugetlb. That is not
    a problem but it'd rather be less prone if we forbid such mixing.

    Cc: <stable@vger.kernel.org>
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb94f3c10a26a9213633da6aa03d4b4a484ae495
Author: Jens Axboe <axboe@kernel.dk>
Date:   Fri Jan 27 09:50:31 2023 -0700

    io_uring: add reschedule point to handle_tw_list()

    commit f58680085478dd292435727210122960d38e8014 upstream.

    If CONFIG_PREEMPT_NONE is set and the task_work chains are long, we
    could be running into issues blocking others for too long. Add a
    reschedule check in handle_tw_list(), and flush the ctx if we need to
    reschedule.

    Cc: stable@vger.kernel.org # 5.10+
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 300c2f021ee837689ffb1f5d0ac760967276e315
Author: Jens Axboe <axboe@kernel.dk>
Date:   Fri Jan 27 09:28:13 2023 -0700

    io_uring: add a conditional reschedule to the IOPOLL cancelation loop

    commit fcc926bb857949dbfa51a7d95f3f5ebc657f198c upstream.
```

```
    If the kernel is configured with CONFIG_PREEMPT_NONE, we could be
    sitting in a tight loop reaping events but not giving them a chance to
    finish. This results in a trace ala:

    rcu: INFO: rcu_sched self-detected stall on CPU
    rcu:    2-...!: (5249 ticks this GP) idle=935c/1/0x4000000000000000 softirq=4265/4274
 fqs=1
            (t=5251 jiffies g=465 q=4135 ncpus=4)
    rcu: rcu_sched kthread starved for 5249 jiffies! g465 f0x0 RCU_GP_WAIT_FQS(5) -
 >state=0x0 ->cpu=0
    rcu:    Unless rcu_sched kthread gets sufficient CPU time, OOM is now expected behavior.
    rcu: RCU grace-period kthread stack dump:
    task:rcu_sched       state:R  running task     stack:0     pid:12    ppid:2
 flags:0x00000008
    Call trace:
     __switch_to+0xb0/0xc8
     __schedule+0x43c/0x520
     schedule+0x4c/0x98
     schedule_timeout+0xbc/0xdc
     rcu_gp_fqs_loop+0x308/0x344
     rcu_gp_kthread+0xd8/0xf0
     kthread+0xb8/0xc8
     ret_from_fork+0x10/0x20
    rcu: Stack dump where RCU GP kthread last ran:
    Task dump for CPU 0:
    task:kworker/u8:10   state:R  running task     stack:0     pid:89    ppid:2
 flags:0x0000000a
    Workqueue: events_unbound io_ring_exit_work
    Call trace:
     __switch_to+0xb0/0xc8
     0xffff0000c8fefd28
    CPU: 2 PID: 95 Comm: kworker/u8:13 Not tainted 6.2.0-rc5-00042-g40316e337c80-dirty #2759
    Hardware name: linux,dummy-virt (DT)
    Workqueue: events_unbound io_ring_exit_work
    pstate: 61400005 (nZCv daif +PAN -UAO -TCO +DIT -SSBS BTYPE=--)
    pc : io_do_iopoll+0x344/0x360
    lr : io_do_iopoll+0xb8/0x360
    sp : ffff800009bebc60
    x29: ffff800009bebc60 x28: 0000000000000000 x27: 0000000000000000
    x26: ffff0000c0f67d48 x25: ffff0000c0f67840 x24: ffff800008950024
    x23: 0000000000000001 x22: 0000000000000000 x21: ffff0000c27d3200
    x20: ffff0000c0f67840 x19: ffff0000c0f67800 x18: 0000000000000000
    x17: 0000000000000000 x16: 0000000000000000 x15: 0000000000000000
    x14: 0000000000000001 x13: 0000000000000001 x12: 0000000000000000
    x11: 0000000000000179 x10: 0000000000000870 x9 : ffff800009bebd60
    x8 : ffff0000c27d3ad0 x7 : fefefefefefefeff x6 : 0000646e756f626e
    x5 : ffff0000c0f67840 x4 : 0000000000000000 x3 : ffff0000c2398000
    x2 : 0000000000000000 x1 : 0000000000000000 x0 : 0000000000000000
    Call trace:
     io_do_iopoll+0x344/0x360
     io_uring_try_cancel_requests+0x21c/0x334
     io_ring_exit_work+0x90/0x40c
     process_one_work+0x1a4/0x254
     worker_thread+0x1ec/0x258
     kthread+0xb8/0xc8
     ret_from_fork+0x10/0x20

    Add a cond_resched() in the cancelation IOPOLL loop to fix this.

    Cc: stable@vger.kernel.org # 5.10+
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit ff43f73a647da02e3e8dca24fb2a1b164ec540d3
 Author: Jens Axboe <axboe@kernel.dk>
```

```
Date:   Tue Jan 24 08:24:25 2023 -0700

    io_uring: handle TIF_NOTIFY_RESUME when checking for task_work

    commit b5d3ae202fbfe055aa2a8ae8524531ee1dcab717 upstream.

    If TIF_NOTIFY_RESUME is set, then we need to call resume_user_mode_work()
    for PF_IO_WORKER threads. They never return to usermode, hence never get
    a chance to process any items that are marked by this flag. Most notably
    this includes the final put of files, but also any throttling markers set
    by block cgroups.

    Cc: stable@vger.kernel.org # 5.10+
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 371248291304840fd8c1e6fa9e14b1c175085dcf
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Mon Jan 23 14:37:13 2023 +0000

    io_uring: use user visible tail in io_uring_poll()

    commit c10bb64684813a326174c3eebcafb3ee5af52ca3 upstream.

    We return POLLIN from io_uring_poll() depending on whether there are
    CQEs for the userspace, and so we should use the user visible tail
    pointer instead of a transient cached value.

    Cc: stable@vger.kernel.org
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/228ffcbf30ba98856f66ffdb9a6a60ead1dd96c0.1674484266.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2ad1065ffedce367947b9935158a91370600ea3c
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Jan 5 11:05:11 2023 -0800

    io_uring: Replace 0-length array with flexible array

    commit 36632d062975a9ff4410c90dd6d37922b68d0920 upstream.

    Zero-length arrays are deprecated[1]. Replace struct io_uring_buf_ring's
    "bufs" with a flexible array member. (How is the size of this array
    verified?) Detected with GCC 13, using -fstrict-flex-arrays=3:

    In function 'io_ring_buffer_select',
        inlined from 'io_buffer_select' at io_uring/kbuf.c:183:10:
    io_uring/kbuf.c:141:23: warning: array subscript 255 is outside the bounds of an
interior zero-length array 'struct io_uring_buf[0]' [-Wzero-length-bounds]
      141 |                 buf = &br->bufs[head];
          |                       ^~~~~~~~~~~~~~~
    In file included from include/linux/io_uring.h:7,
                     from io_uring/kbuf.c:10:
    include/uapi/linux/io_uring.h: In function 'io_buffer_select':
    include/uapi/linux/io_uring.h:628:41: note: while referencing 'bufs'
      628 |                 struct io_uring_buf     bufs[0];
          |                                         ^~~~

    [1] https://www.kernel.org/doc/html/latest/process/deprecated.html#zero-length-and-one-
element-arrays

    Fixes: c7fb19428d67 ("io_uring: add support for ring mapped supplied buffers")
```

```
    Cc: Jens Axboe <axboe@kernel.dk>
    Cc: Pavel Begunkov <asml.silence@gmail.com>
    Cc: "Gustavo A. R. Silva" <gustavoars@kernel.org>
    Cc: stable@vger.kernel.org
    Cc: io-uring@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Link: https://lore.kernel.org/r/20230105190507.gonna.131-kees@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4e26ea5f808148db119e0c20ae4e10a13294e977
Author: Corey Minyard <cminyard@mvista.com>
Date:   Wed Jan 25 10:34:47 2023 -0600

    ipmi:ssif: Add a timer between request retries

    commit 00bb7e763ec9f384cb382455cb6ba5588b5375cf upstream.

    The IPMI spec has a time (T6) specified between request retries.  Add
    the handling for that.

    Reported by: Tony Camuso <tcamuso@redhat.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Corey Minyard <cminyard@mvista.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 17e3c7f2b576bcca0a973a8f506f9e598d4fb704
Author: Corey Minyard <cminyard@mvista.com>
Date:   Wed Jan 25 10:13:13 2023 -0600

    ipmi_ssif: Rename idle state and check

    commit 8230831c43a328c2be6d28c65d3f77e14c59986b upstream.

    Rename the SSIF_IDLE() to IS_SSIF_IDLE(), since that is more clear, and
    rename SSIF_NORMAL to SSIF_IDLE, since that's more accurate.

    Cc: stable@vger.kernel.org
    Signed-off-by: Corey Minyard <cminyard@mvista.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5bf2e8077e20932c4372b331a1a7db7bc1422752
Author: Corey Minyard <cminyard@mvista.com>
Date:   Wed Jan 25 10:11:06 2023 -0600

    ipmi:ssif: resend_msg() cannot fail

    commit 95767ed78a181d5404202627499f9cde56053b96 upstream.

    The resend_msg() function cannot fail, but there was error handling
    around using it.  Rework the handling of the error, and fix the out of
    retries debug reporting that was wrong around this, too.

    Cc: stable@vger.kernel.org
    Signed-off-by: Corey Minyard <cminyard@mvista.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 292e9b1f4f8a8e08651d5fac587214f65eba0421
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sun Feb 5 11:04:01 2023 +0100

    ipmi: ipmb: Fix the MODULE_PARM_DESC associated to 'retry_time_ms'

    commit befb28f2676a65a5a4cc4626ae224461d8785af6 upstream.
```

```
    'This should be 'retry_time_ms' instead of 'max_retries'.

    Fixes: 63c4eb347164 ("ipmi:ipmb: Add initial support for IPMI over IPMB")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Message-Id:
<0d8670cff2c656e99a832a249e77dc90578f67de.1675591429.git.christophe.jaillet@wanadoo.fr>
    Cc: stable@vger.kernel.org
    Signed-off-by: Corey Minyard <cminyard@mvista.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c61cff02942c9884e7c59789416aaf8c864be227
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Thu Feb 2 16:54:27 2023 +0100

    rtc: pm8xxx: fix set-alarm race

    commit c88db0eff9722fc2b6c4d172a50471d20e08ecc6 upstream.

    Make sure to disable the alarm before updating the four alarm time
    registers to avoid spurious alarms during the update.

    Note that the disable needs to be done outside of the ctrl_reg_lock
    section to prevent a racing alarm interrupt from disabling the newly set
    alarm when the lock is released.

    Fixes: 9a9a54ad7aa2 ("drivers/rtc: add support for Qualcomm PMIC8xxx RTC")
    Cc: stable@vger.kernel.org      # 3.1
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Reviewed-by: David Collins <quic_collinsd@quicinc.com>
    Link: https://lore.kernel.org/r/20230202155448.6715-2-johan+linaro@kernel.org
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0510d5e654d05053ed0e6309a9b42043ac9903ab
Author: Jens Axboe <axboe@kernel.dk>
Date:   Fri Feb 24 10:01:19 2023 -0700

    block: be a bit more careful in checking for NULL bdev while polling

    commit 310726c33ad76cebdee312dbfafc12c1b44bf977 upstream.

    Wei reports a crash with an application using polled IO:

    PGD 14265e067 P4D 14265e067 PUD 47ec50067 PMD 0
    Oops: 0000 [#1] SMP
    CPU: 0 PID: 21915 Comm: iocore_0 Kdump: loaded Tainted: G S                5.12.0-
0_fbk12_clang_7346_g1bb6f2e7058f #1
    Hardware name: Wiwynn Delta Lake MP T8/Delta Lake-Class2, BIOS Y3DLM08 04/10/2022
    RIP: 0010:bio_poll+0x25/0x200
    Code: 0f 1f 44 00 00 0f 1f 44 00 00 55 41 57 41 56 41 55 41 54 53 48 83 ec 28 65 48 8b
04 25 28 00 00 00 48 89 44 24 20 48 8b 47 08 <48> 8b 80 70 02 00 00 4c 8b 70 50 8b 6f 34 31
db 83 fd ff 75 25 65
    RSP: 0018:ffffc90005fafdf8 EFLAGS: 00010292
    RAX: 0000000000000000 RBX: 0000000000000000 RCX: 74b43cd65dd66600
    RDX: 0000000000000003 RSI: ffffc90005fafe78 RDI: ffff8884b614e140
    RBP: ffff88849964df78 R08: 0000000000000000 R09: 0000000000000008
    R10: 0000000000000000 R11: 0000000000000000 R12: ffff88849964df00
    R13: ffffc90005fafe78 R14: ffff888137d3c378 R15: 0000000000000001
    FS:  00007fd195000640(0000) GS:ffff88903f400000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 0000000000000270 CR3: 0000000466121001 CR4: 00000000007706f0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    PKRU: 55555554
    Call Trace:
```

```
      iocb_bio_iopoll+0x1d/0x30
      io_do_iopoll+0xac/0x250
      __se_sys_io_uring_enter+0x3c5/0x5a0
      ? __x64_sys_write+0x89/0xd0
      do_syscall_64+0x2d/0x40
      entry_SYSCALL_64_after_hwframe+0x44/0xae
     RIP: 0033:0x94f225d
     Code: 24 cc 00 00 00 41 8b 84 24 d0 00 00 00 c1 e0 04 83 e0 10 41 09 c2 8b 33 8b 53 04
 4c 8b 43 18 4c 63 4b 0c b8 aa 01 00 00 0f 05 <85> c0 0f 88 85 00 00 00 29 03 45 84 f6 0f 84
 88 00 00 00 41 f6 c7
     RSP: 002b:00007fd194ffcd88 EFLAGS: 00000202 ORIG_RAX: 00000000000001aa
     RAX: ffffffffffffffda RBX: 00007fd194ffcdc0 RCX: 00000000094f225d
     RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000007
     RBP: 00007fd194ffcdb0 R08: 0000000000000000 R09: 0000000000000008
     R10: 0000000000000001 R11: 0000000000000202 R12: 00007fd269d68030
     R13: 0000000000000000 R14: 0000000000000001 R15: 0000000000000000

     which is due to bio->bi_bdev being NULL. This can happen if we have two
     tasks doing polled IO, and task B ends up completing IO from task A if
     they are sharing a poll queue. If task B completes the IO and puts the
     bio into our cache, then it can allocate that bio again before task A
     is done polling for it. As that would necessitate a preempt between the
     two tasks, it's enough to just be a bit more careful in checking for
     whether or not bio->bi_bdev is NULL.

     Reported-and-tested-by: Wei Zhang <wzhang@meta.com>
     Cc: stable@vger.kernel.org
     Fixes: be4d234d7aeb ("bio: add allocation cache abstraction")
     Reviewed-by: Keith Busch <kbusch@kernel.org>
     Signed-off-by: Jens Axboe <axboe@kernel.dk>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 6b6676da6b2a275dd4c7c6a72ec4e684ffdefdc7
 Author: Jens Axboe <axboe@kernel.dk>
 Date:   Fri Feb 24 09:59:44 2023 -0700

     block: clear bio->bi_bdev when putting a bio back in the cache

     commit 11eb695feb636fa5211067189cad25ac073e7fe5 upstream.

     This isn't strictly needed in terms of correctness, but it does allow
     polling to know if the bio has been put already by a different task
     and hence avoid polling something that we don't need to.

     Cc: stable@vger.kernel.org
     Fixes: be4d234d7aeb ("bio: add allocation cache abstraction")
     Reviewed-by: Keith Busch <kbusch@kernel.org>
     Signed-off-by: Jens Axboe <axboe@kernel.dk>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 3dbd3d1b33cea7d062983ebb990b28d55985ebb6
 Author: Jens Axboe <axboe@kernel.dk>
 Date:   Mon Jan 16 08:55:53 2023 -0700

     block: don't allow multiple bios for IOCB_NOWAIT issue

     commit 67d59247d4b52c917e373f05a807027756ab216f upstream.

     If we're doing a large IO request which needs to be split into multiple
     bios for issue, then we can run into the same situation as the below
     marked commit fixes - parts will complete just fine, one or more parts
     will fail to allocate a request. This will result in a partially
     completed read or write request, where the caller gets EAGAIN even though
     parts of the IO completed just fine.
```

```
    Do the same for large bios as we do for splits - fail a NOWAIT request
    with EAGAIN. This isn't technically fixing an issue in the below marked
    patch, but for stable purposes, we should have either none of them or
    both.

    This depends on: 613b14884b85 ("block: handle bio_split_to_limits() NULL return")

    Cc: stable@vger.kernel.org # 5.15+
    Fixes: 9cea62b2cbab ("block: don't allow splitting of a REQ_NOWAIT bio")
    Link: https://github.com/axboe/liburing/issues/766
    Reported-and-tested-by: Michael Kelley <mikelley@microsoft.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 22689a071d62286aed6b87a646f2d6ea8fdaf964
Author: Alper Nebi Yasak <alpernebiyasak@gmail.com>
Date:   Sun Jan 22 22:04:31 2023 +0300

    firmware: coreboot: framebuffer: Ignore reserved pixel color bits

    commit e6acaf25cba14661211bb72181c35dd13b24f5b3 upstream.

    The coreboot framebuffer doesn't support transparency, its 'reserved'
    bit field is merely padding for byte/word alignment of pixel colors [1].
    When trying to match the framebuffer to a simplefb format, the kernel
    driver unnecessarily requires the format's transparency bit field to
    exactly match this padding, even if the former is zero-width.

    Due to a coreboot bug [2] (fixed upstream), some boards misreport the
    reserved field's size as equal to its position (0x18 for both on a
    'Lick' Chromebook), and the driver fails to probe where it would have
    otherwise worked fine with e.g. the a8r8g8b8 or x8r8g8b8 formats.

    Remove the transparency comparison with reserved bits. When the
    bits-per-pixel and other color components match, transparency will
    already be in a subset of the reserved field. Not forcing it to match
    reserved bits allows the driver to work on the boards which misreport
    the reserved field. It also enables using simplefb formats that don't
    have transparency bits, although this doesn't currently happen due to
    format support and ordering in linux/platform_data/simplefb.h.

    [1]
https://review.coreboot.org/plugins/gitiles/coreboot/+/4.19/src/commonlib/include/commonlib/
coreboot_tables.h#255
    [2]
https://review.coreboot.org/plugins/gitiles/coreboot/+/4.13/src/drivers/intel/fsp2_0/graphic
s.c#82

    Signed-off-by: Alper Nebi Yasak <alpernebiyasak@gmail.com>
    Link: https://lore.kernel.org/r/20230122190433.195941-1-alpernebiyasak@gmail.com
    Cc: Salvatore Bonaccorso <carnil@debian.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 411590b87f0e5d483f6571c0bcb477cfad855c51
Author: Jun ASAKA <JunASAKA@zzy040330.moe>
Date:   Sat Dec 17 11:06:59 2022 +0800

    wifi: rtl8xxxu: fixing transmisison failure for rtl8192eu

    commit c6015bf3ff1ffb3caa27eb913797438a0fc634a0 upstream.

    Fixing transmission failure which results in
    "authentication with ... timed out". This can be
    fixed by disable the REG_TXPAUSE.
```

```
    Signed-off-by: Jun ASAKA <JunASAKA@zzy040330.moe>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221217030659.12577-1-JunASAKA@zzy040330.moe
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 75c00130e3a6b5cb037ed99e607b277f45186ab1
Author: Saravana Kannan <saravanak@google.com>
Date:   Fri Feb 24 22:41:47 2023 -0800

    driver core: fw_devlink: Avoid spurious error message

    [ Upstream commit 6309872413f14f3d58c13ae4dc85b1a7004b4193 ]

    fw_devlink can sometimes try to create a device link with the consumer
    and supplier as the same device. These attempts will fail (correctly),
    but are harmless. So, avoid printing an error for these cases. Also, add
    more detail to the error message.

    Fixes: 3fb16866b51d ("driver core: fw_devlink: Make cycle detection more robust")
    Reported-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Reported-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Link: https://lore.kernel.org/r/20230225064148.274376-1-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7bea49d937a87784e8809b0bb03122f18159d753
Author: Asahi Lina <lina@asahilina.net>
Date:   Mon Feb 27 18:04:21 2023 +0900

    drm/shmem-helper: Revert accidental non-GPL export

    [ Upstream commit 047a754558d640eaa080fce3b22ca9f3d4e04626 ]

    The referenced commit added a wrapper for drm_gem_shmem_get_pages_sgt(),
    but in the process it accidentally changed the export type from GPL to
    non-GPL. Switch it back to GPL.

    Reported-by: Dmitry Osipenko <dmitry.osipenko@collabora.com>
    Fixes: ddddedaa0db9 ("drm/shmem-helper: Fix locking for drm_gem_shmem_get_pages_sgt()")
    Signed-off-by: Asahi Lina <lina@asahilina.net>
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230227-shmem-export-fix-v1-1-
8880b2c25e81@asahilina.net
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e7e072cfa50839aefdf1b94e23298b264cbc8c3f
Author: Matt Roper <matthew.d.roper@intel.com>
Date:   Wed Jan 25 15:41:58 2023 -0800

    drm/i915/mtl: Correct implementation of Wa_18018781329

    [ Upstream commit eda94a6e6a4f2d3d1574ff4f2bd4b9f844504f71 ]

    Workaround Wa_18018781329 has applied to several recent Xe_HP-based
    platforms.  However there are some extra gotchas to implementing this
    properly for MTL that we need to take into account:

     * Due to the separation of media and render/compute into separate GTs,
       this workaround needs to be implemented on each GT, not just the
       primary GT.  Since each class of register only exists on one of the
       two GTs, we should program the appropriate registers on each GT.

     * As with past Xe_HP platforms, the registers on the primary GT (Xe_LPG
```

```
        IP) are multicast/replicated registers and should be handled with the
        MCR-aware functions.  However the registers on the media GT (Xe_LPM+
        IP) are regular singleton registers and should _not_ use MCR
        handling.  We need to create separate register definitions for the
        Xe_HP multicast form and the Xe_LPM+ singleton form and use each in
        the appropriate place.

      * Starting with MTL, workarounds documented by the hardware teams are
        technically associated with IP versions/steppings rather than
        top-level platforms.  That means we should take care to check the
        media IP version rather than the graphics IP version when deciding
        whether the workaround is needed on the Xe_LPM+ media GT (in this
        case the workaround applies to both IPs and the stepping bounds are
        identical, but we should still write the code appropriately to set a
        proper precedent for future workaround implementations).

      * It's worth noting that the GSC register and the CCS register are
        defined with the same MMIO offset (0xCF30).  Since the CCS is only
        relevant to the primary GT and the GSC is only relevant to the media
        GT there isn't actually a clash here (the media GT automatically adds
        the additional 0x380000 GSI offset).  However there's currently a
        glitch in the bspec where the CCS register doesn't show up at all and
        the GSC register is listed as existing on both GTs.  That's a known
        documentation problem for several registers with shared GSC/CCS
        offsets; rest assured that the CCS register really does still exist.

    Cc: Gustavo Sousa <gustavo.sousa@intel.com>
    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Fixes: 41bb543f5598 ("drm/i915/mtl: Add initial gt workarounds")
    Reviewed-by: Gustavo Sousa <gustavo.sousa@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230125234159.3015385-2-
matthew.d.roper@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa71316593cdbb907bf7209ba2136478cf8efb2b
Author: Paulo Alcantara <pc@cjr.nz>
Date:   Mon Jan 30 20:33:29 2023 -0300

    cifs: prevent data race in smb2_reconnect()

    [ Upstream commit 3c0070f54b3128de498c2dd9934a21f0dd867111 ]

    Make sure to get an up-to-date TCP_Server_Info::nr_targets value prior
    to waiting the server to be reconnected in smb2_reconnect().  It is
    set in cifs_tcp_ses_needs_reconnect() and protected by
    TCP_Server_Info::srv_lock.

    Signed-off-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b91486b092d8592be6166034480392db81bffaf
Author: Jeff Layton <jlayton@kernel.org>
Date:   Fri Jan 27 07:09:33 2023 -0500

    nfsd: don't hand out delegation on setuid files being opened for write

    [ Upstream commit 826b67e6376c2a788e3a62c4860dcd79500a27d5 ]

    We had a bug report that xfstest generic/355 was failing on NFSv4.0.
    This test sets various combinations of setuid/setgid modes and tests
    whether DIO writes will cause them to be stripped.

    What I found was that the server did properly strip those bits, but
    the client didn't notice because it held a delegation that was not
```

```
    recalled. The recall didn't occur because the client itself was the
    one generating the activity and we avoid recalls in that case.

    Clearing setuid bits is an "implicit" activity. The client didn't
    specifically request that we do that, so we need the server to issue a
    CB_RECALL, or avoid the situation entirely by not issuing a delegation.

    The easiest fix here is to simply not give out a delegation if the file
    is being opened for write, and the mode has the setuid and/or setgid bit
    set. Note that there is a potential race between the mode and lease
    being set, so we test for this condition both before and after setting
    the lease.

    This patch fixes generic/355, generic/683 and generic/684 for me. (Note
    that 355 fails only on v4.0, and 683 and 684 require NFSv4.2 to run and
    fail).

    Reported-by: Boyang Xue <bxue@redhat.com>
    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 048d4245ac4c58a9bea6c33ac8306a9c7672dbf9
Author: Jeff Layton <jlayton@kernel.org>
Date:   Tue Jan 17 14:38:30 2023 -0500

    nfsd: zero out pointers after putting nfsd_files on COPY setup error

    [ Upstream commit 1f0001d43d0c0ac2a19a34a914f6595ad97cbc1d ]

    At first, I thought this might be a source of nfsd_file overputs, but
    the current callers seem to avoid an extra put when nfsd4_verify_copy
    returns an error.

    Still, it's "bad form" to leave the pointers filled out when we don't
    have a reference to them anymore, and that might lead to bugs later.
    Zero them out as a defensive coding measure.

    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 31966a7b2d573b5e5ba1c6151251ea9e19bb60d8
Author: Mike Snitzer <snitzer@kernel.org>
Date:   Thu Feb 16 15:31:08 2023 -0500

    dm cache: add cond_resched() to various workqueue loops

    [ Upstream commit 76227f6dc805e9e960128bcc6276647361e0827c ]

    Otherwise on resource constrained systems these workqueues may be too
    greedy.

    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dcb25ac0350ff77b4f8c684e09eb33512c9efdda
Author: Mike Snitzer <snitzer@kernel.org>
Date:   Thu Feb 16 15:29:44 2023 -0500

    dm thin: add cond_resched() to various workqueue loops

    [ Upstream commit e4f80303c2353952e6e980b23914e4214487f2a6 ]

    Otherwise on resource constrained systems these workqueues may be too
```

```
    greedy.

    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63a412dbb5e8fba5bba20f39dedf1a4744e139be
Author: Aurabindo Pillai <aurabindo.pillai@amd.com>
Date:   Wed Feb 15 12:51:20 2023 -0500

    drm/amd/display: disable SubVP + DRR to prevent underflow

    [ Upstream commit 80c6d6804f31451848a3956a70c2bcb1f07cfcb0 ]

    [Why&How]
    Temporarily disable SubVP+DRR since Xorg has an architectural limitation
    where freesync will not work in a multi monitor configuration. SubVP+DRR
    requires that freesync be working.

    Whether OS has variable refresh setting enabled or not, the state on
    the crtc remains same unless an application requests VRR. Due to this,
    there is no way to know whether freesync will actually work or not
    while we are on the desktop from the kernel's perspective.

    If userspace does not have a limitation with multi-display freesync (for
    example wayland), then this feature can be enabled by adding a
    dcfeaturemask option to amdgpu on the kernel cmdline like:

    amdgpu.dcfeaturemask=0x200

    Signed-off-by: Aurabindo Pillai <aurabindo.pillai@amd.com>
    Reviewed-by: Rodrigo Siqueira <Rodrigo.Siqueira@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14dec567b02d409251b018b8d60df02a18d04070
Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
Date:   Thu Feb 2 15:15:51 2023 -0500

    drm/amd/display: Disable HUBP/DPP PG on DCN314 for now

    [ Upstream commit b7c67f72408b11b922f23f06c7df0f6743a2e89d ]

    [Why]
    The DMCUB implementation required to workaround corruption is
    not currently stable and may cause intermittent corruption or hangs.

    [How]
    Disable PG until the sequence is stable.

    Reviewed-by: Hansen Dsouza <hansen.dsouza@amd.com>
    Acked-by: Qingqing Zhuo <qingqing.zhuo@amd.com>
    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1982d5f08fd60f1b9c901d882ea81d969b2242d
Author: Darrell Kavanagh <darrell.kavanagh@gmail.com>
Date:   Tue Feb 14 16:46:59 2023 +0000

    drm: panel-orientation-quirks: Add quirk for Lenovo IdeaPad Duet 3 10IGL5

    [ Upstream commit 38b2d8efd03d2e56431b611e3523f0158306451d ]

    Another Lenovo convertable where the panel is installed landscape but is
```

```
    reported to the kernel as portrait.

    Signed-off-by: Darrell Kavanagh <darrell.kavanagh@gmail.com>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230214164659.3583-1-
darrell.kavanagh@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d54c791543389e2150d0294c4031eee6fd106fbd
Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
Date:   Fri Jan 27 10:03:45 2023 -0500

    drm/amd/display: Enable P-state validation checks for DCN314

    [ Upstream commit 37d184b548db0f64d4a878960b2c6988b38a3e7e ]

    [Why]
    To align with DCN31 behavior. This helps avoid p-state hangs in
    the case where underflow does occur.

    [How]
    Flip the bit to true.

    Reviewed-by: Hansen Dsouza <hansen.dsouza@amd.com>
    Acked-by: Qingqing Zhuo <qingqing.zhuo@amd.com>
    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04deaae5132f5054180bd512c9db287a87fb8157
Author: Bastien Nocera <hadess@hadess.net>
Date:   Wed Jan 25 13:17:22 2023 +0100

    HID: logitech-hidpp: Don't restart communication if not necessary

    [ Upstream commit 498ba20690357691103de0f766960355247c78be ]

    Don't stop and restart communication with the device unless we need to
    modify the connect flags used because of a device quirk.

    Signed-off-by: Bastien Nocera <hadess@hadess.net>
    Link: https://lore.kernel.org/r/20230125121723.3122-1-hadess@hadess.net
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ffd2cd644e0f1eea01339831bac4b1054e8817c
Author: Mason Zhang <Mason.Zhang@mediatek.com>
Date:   Fri Dec 16 11:25:33 2022 +0800

    scsi: ufs: core: Fix device management cmd timeout flow

    [ Upstream commit 36822124f9de200cedc2f42516301b50d386a6cd ]

    In the UFS error handling flow, the host will send a device management cmd
    (NOP OUT) to the device for link recovery. If this cmd times out and
    clearing the doorbell fails, ufshcd_wait_for_dev_cmd() will do nothing and
    return. hba->dev_cmd.complete struct is not set to NULL.

    When this happens, if cmd has been completed by device, then we will call
    complete() in __ufshcd_transfer_req_compl(). Because the complete struct is
    allocated on the stack, the following crash will occur:

      ipanic_die+0x24/0x38 [mrdump]
```

```
        die+0x344/0x748
        arm64_notify_die+0x44/0x104
        do_debug_exception+0x104/0x1e0
        el1_dbg+0x38/0x54
        el1_sync_handler+0x40/0x88
        el1_sync+0x8c/0x140
        queued_spin_lock_slowpath+0x2e4/0x3c0
        __ufshcd_transfer_req_compl+0x3b0/0x1164
        ufshcd_trc_handler+0x15c/0x308
        ufshcd_host_reset_and_restore+0x54/0x260
        ufshcd_reset_and_restore+0x28c/0x57c
        ufshcd_err_handler+0xeb8/0x1b6c
        process_one_work+0x288/0x964
        worker_thread+0x4bc/0xc7c
        kthread+0x15c/0x264
        ret_from_fork+0x10/0x30

    Link: https://lore.kernel.org/r/20221216032532.1280-1-mason.zhang@mediatek.com
    Signed-off-by: Mason Zhang <Mason.Zhang@mediatek.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 995424f59ab52fb432b26ccb3abced63745ea041
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Feb 2 15:10:09 2023 +0100

    scsi: snic: Fix memory leak with using debugfs_lookup()

    [ Upstream commit ad0e4e2fab928477f74d742e6e77d79245d3d3e7 ]

    When calling debugfs_lookup() the result must have dput() called on it,
    otherwise the memory will leak over time.  To make things simpler, just
    call debugfs_lookup_and_remove() instead which handles all of the logic at
    once.

    Link: https://lore.kernel.org/r/20230202141009.2290380-1-gregkh@linuxfoundation.org
    Cc: Karan Tilak Kumar <kartilak@cisco.com>
    Cc: Sesidhar Baddela <sebaddel@cisco.com>
    Cc: "James E.J. Bottomley" <jejb@linux.ibm.com>
    Cc: "Martin K. Petersen" <martin.petersen@oracle.com>
    Cc: linux-scsi@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1ba3306c3d45d6921d6eed34a05685fa43968a0
Author: Wesley Chalmers <Wesley.Chalmers@amd.com>
Date:   Fri Nov 4 15:19:22 2022 -0400

    drm/amd/display: Do not commit pipe when updating DRR

    [ Upstream commit 8f0d304d21b351d65e8c434c5399a40231876ba1 ]

    [WHY]
    DRR and Pipe cannot be updated on
    the same frame, or else underflow will
    occur.

    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Qingqing Zhuo <qingqing.zhuo@amd.com>
    Signed-off-by: Wesley Chalmers <Wesley.Chalmers@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit e5a6db7937a092b061a6b28e0d74edd646c7ea0a
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Fri Feb 3 15:27:14 2023 +0200

    pinctrl: at91: use devm_kasprintf() to avoid potential leaks

    [ Upstream commit 1c4e5c470a56f7f7c649c0c70e603abc1eab15c4 ]

    Use devm_kasprintf() instead of kasprintf() to avoid any potential
    leaks. At the moment drivers have no remove functionality thus
    there is no need for fixes tag.

    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20230203132714.1931596-1-claudiu.beznea@microchip.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e4aacc79e1d4241610d15c63ff11a881fbc7a0dc
Author: Denis Pauk <pauk.denis@gmail.com>
Date:   Wed Jan 11 23:22:41 2023 +0200

    hwmon: (nct6775) B650/B660/X670 ASUS boards support

    [ Upstream commit e2e09989ccc21ad428d6393450add78584b143bd ]

    Boards such as:
      "EX-B660M-V5 PRO D4",
      "PRIME B650-PLUS",
      "PRIME B650M-A",
      "PRIME B650M-A AX",
      "PRIME B650M-A II",
      "PRIME B650M-A WIFI",
      "PRIME B650M-A WIFI II",
      "PRIME B660M-A D4",
      "PRIME B660M-A WIFI D4",
      "PRIME X670-P",
      "PRIME X670-P WIFI",
      "PRIME X670E-PRO WIFI",
      "Pro B660M-C-D4",
      "ProArt B660-CREATOR D4",
      "ProArt X670E-CREATOR WIFI",
      "ROG CROSSHAIR X670E EXTREME",
      "ROG CROSSHAIR X670E GENE",
      "ROG CROSSHAIR X670E HERO",
      "ROG MAXIMUS XIII EXTREME GLACIAL",
      "ROG MAXIMUS Z690 EXTREME",
      "ROG MAXIMUS Z690 EXTREME GLACIAL",
      "ROG STRIX B650-A GAMING WIFI",
      "ROG STRIX B650E-E GAMING WIFI",
      "ROG STRIX B650E-F GAMING WIFI",
      "ROG STRIX B650E-I GAMING WIFI",
      "ROG STRIX B660-A GAMING WIFI D4",
      "ROG STRIX B660-F GAMING WIFI",
      "ROG STRIX B660-G GAMING WIFI",
      "ROG STRIX B660-I GAMING WIFI",
      "ROG STRIX X670E-A GAMING WIFI",
      "ROG STRIX X670E-E GAMING WIFI",
      "ROG STRIX X670E-F GAMING WIFI",
      "ROG STRIX X670E-I GAMING WIFI",
      "ROG STRIX Z590-A GAMING WIFI II",
      "ROG STRIX Z690-A GAMING WIFI D4",
      "TUF GAMING B650-PLUS",
      "TUF GAMING B650-PLUS WIFI",
      "TUF GAMING B650M-PLUS",
```

```
        "TUF GAMING B650M-PLUS WIFI",
        "TUF GAMING B660M-PLUS WIFI",
        "TUF GAMING X670E-PLUS",
        "TUF GAMING X670E-PLUS WIFI",
        "TUF GAMING Z590-PLUS WIFI",
    have got a NCT6799D chip, but by default there's no use of it
    because of resource conflict with WMI method.

    This commit adds such boards to the monitoring list with new ACPI device
    UID.

    BugLink: https://bugzilla.kernel.org/show_bug.cgi?id=204807
    Signed-off-by: Denis Pauk <pauk.denis@gmail.com>
    Co-developed-by: Ahmad Khalifa <ahmad@khalifa.ws>
    Signed-off-by: Ahmad Khalifa <ahmad@khalifa.ws>
    Tested-by: Jeroen Beerstra <jeroen@beerstra.org>
    Tested-by: Slawomir Stepien <sst@poczta.fm>
    Link: https://lore.kernel.org/r/20230111212241.7456-2-pauk.denis@gmail.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edb8e4bda141f4efa4716a3338b1fd9766ac1e85
Author: Denis Pauk <pauk.denis@gmail.com>
Date:   Wed Jan 11 23:22:40 2023 +0200

    hwmon: (nct6775) Directly call ASUS ACPI WMI method

    [ Upstream commit c3b3747d02f571da2543e719066a50dd966989d8 ]

    New ASUS B650/B660/X670 boards firmware have not exposed WMI monitoring
    GUID  and entrypoint method WMBD could be implemented for different device
    UID.

    Implement the direct call to entrypoint method for monitoring the device
    UID of B550/X570 boards.

    BugLink: https://bugzilla.kernel.org/show_bug.cgi?id=204807
    Signed-off-by: Denis Pauk <pauk.denis@gmail.com>
    Co-developed-by: Ahmad Khalifa <ahmad@khalifa.ws>
    Signed-off-by: Ahmad Khalifa <ahmad@khalifa.ws>
    Link: https://lore.kernel.org/r/20230111212241.7456-1-pauk.denis@gmail.com
    [groeck: Fix multi-line formatting]
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c57a8d14d7880521150ee801d53a0a64fdffd9c8
Author: Robin Murphy <robin.murphy@arm.com>
Date:   Tue Jan 3 12:46:20 2023 +0100

    hwmon: (coretemp) Simplify platform device handling

    [ Upstream commit 6d03bbff456befeccdd4d663177c4d6c75d0c4ff ]

    Coretemp's platform driver is unconventional. All the real work is done
    globally by the initcall and CPU hotplug notifiers, while the "driver"
    effectively just wraps an allocation and the registration of the hwmon
    interface in a long-winded round-trip through the driver core.  The whole
    logic of dynamically creating and destroying platform devices to bring
    the interfaces up and down is error prone, since it assumes
    platform_device_add() will synchronously bind the driver and set drvdata
    before it returns, thus results in a NULL dereference if drivers_autoprobe
    is turned off for the platform bus. Furthermore, the unusual approach of
    doing that from within a CPU hotplug notifier, already commented in the
    code that it deadlocks suspend, also causes lockdep issues for other
    drivers or subsystems which may want to legitimately register a CPU
```

```
    hotplug notifier from a platform bus notifier.

    All of these issues can be solved by ripping this unusual behaviour out
    completely, simply tying the platform devices to the lifetime of the
    module itself, and directly managing the hwmon interfaces from the
    hotplug notifiers. There is a slight user-visible change in that
    /sys/bus/platform/drivers/coretemp will no longer appear, and
    /sys/devices/platform/coretemp.n will remain present if package n is
    hotplugged off, but hwmon users should really only be looking for the
    presence of the hwmon interfaces, whose behaviour remains unchanged.

    Link: https://lore.kernel.org/lkml/20220922101036.87457-1-
janusz.krzysztofik@linux.intel.com/
    Link: https://gitlab.freedesktop.org/drm/intel/issues/6641
    Signed-off-by: Robin Murphy <robin.murphy@arm.com>
    Signed-off-by: Janusz Krzysztofik <janusz.krzysztofik@linux.intel.com>
    Link: https://lore.kernel.org/r/20230103114620.15319-1-
janusz.krzysztofik@linux.intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9fc32dad3cdba18669c71893f3e6d96905b39b3f
Author: Andreas Gruenbacher <agruenba@redhat.com>
Date:   Tue Jan 31 15:06:53 2023 +0100

    gfs2: Improve gfs2_make_fs_rw error handling

    [ Upstream commit b66f723bb552ad59c2acb5d45ea45c890f84498b ]

    In gfs2_make_fs_rw(), make sure to call gfs2_consist() to report an
    inconsistency and mark the filesystem as withdrawn when
    gfs2_find_jhead() fails.

    At the end of gfs2_make_fs_rw(), when we discover that the filesystem
    has been withdrawn, make sure we report an error.  This also replaces
    the gfs2_withdrawn() check after gfs2_find_jhead().

    Reported-by: Tetsuo Handa <penguin-kernel@i-love.sakura.ne.jp>
    Cc: syzbot+f51cb4b9afbd87ec06f2@syzkaller.appspotmail.com
    Signed-off-by: Andreas Gruenbacher <agruenba@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6853d56dba56d1c24db403ff3885c71e18d572c4
Author: Vladimir Stempen <vladimir.stempen@amd.com>
Date:   Tue Jan 17 14:14:42 2023 -0500

    drm/amd/display: fix FCLK pstate change underflow

    [ Upstream commit 972243f973eb0821084e5833d5f7f4ed025f42da ]

    [Why]
    Currently we set FCLK p-state change
    watermark calculated based on dummy
    p-state latency when UCLK p-state is
    not supported

    [How]
    Calculate FCLK p-state change watermark
    based on on FCLK pstate change latency
    in case UCLK p-state is not supported

    Reviewed-by: Nevenko Stupar <Nevenko.Stupar@amd.com>
    Acked-by: Alex Hung <alex.hung@amd.com>
    Signed-off-by: Vladimir Stempen <vladimir.stempen@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
```

```
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1ec5ede1aec776cc91686a96321f61a0c41bc641
Author: Vitaly Prosyak <vitaly.prosyak@amd.com>
Date:   Fri Jan 20 13:19:50 2023 -0500

    Revert "drm/amdgpu: TA unload messages are not actually sent to psp when amdgpu is
uninstalled"

    [ Upstream commit 39934d3ed5725c5e3570ed1b67f612f1ea60ce03 ]

    This reverts commit fac53471d0ea9693d314aa2df08d62b2e7e3a0f8.
    The following change: move the drm_dev_unplug call after
    amdgpu_driver_unload_kms in amdgpu_pci_remove. The reason is
    the following: amdgpu_pci_remove calls drm_dev_unregister
    and it should be called first to ensure userspace can't access the
    device instance anymore. If we call drm_dev_unplug after
    amdgpu_driver_unload_kms then we observe IGT PCI software unplug
    test failure (kernel hung) for all ASICs. This is how this
    regression was found.

    After this revert, the following commands do work not, but it would
    be fixed in the next commit:
     - sudo modprobe -r amdgpu
     - sudo modprobe amdgpu

    Signed-off-by: Vitaly Prosyak <vitaly.prosyak@amd.com>
    Reviewed-by Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb97499a97b42f0a32cde911aa92190cccfbbebe
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Jan 27 16:53:58 2023 -0800

    regulator: s5m8767: Bounds check id indexing into arrays

    [ Upstream commit e314e15a0b58f9d051c00b25951073bcdae61953 ]

    The compiler has no way to know if "id" is within the array bounds of
    the regulators array. Add a check for this and a build-time check that
    the regulators and reg_voltage_map arrays are sized the same. Seen with
    GCC 13:

    ../drivers/regulator/s5m8767.c: In function 's5m8767_pmic_probe':
    ../drivers/regulator/s5m8767.c:936:35: warning: array subscript [0, 36] is outside array
bounds of 'struct regulator_desc[37]' [-Warray-bounds=]
      936 |                         regulators[id].vsel_reg =
          |                         ~~~~~~~~~~^~~~

    Cc: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Cc: Liam Girdwood <lgirdwood@gmail.com>
    Cc: Mark Brown <broonie@kernel.org>
    Cc: linux-samsung-soc@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230128005358.never.313-kees@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f96ad1b9086194f86e66620c5c5c376eb7527f7617
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Jan 27 14:52:07 2023 -0800
```

```
    regulator: max77802: Bounds check regulator id against opmode

    [ Upstream commit 4fd8bcec5fd7c0d586206fa2f42bd67b06cdaa7e ]

    Explicitly bounds-check the id before accessing the opmode array. Seen
    with GCC 13:

    ../drivers/regulator/max77802-regulator.c: In function 'max77802_enable':
    ../drivers/regulator/max77802-regulator.c:217:29: warning: array subscript [0, 41] is
 outside array bounds of 'unsigned int[42]' [-Warray-bounds=]
      217 |         if (max77802->opmode[id] == MAX77802_OFF_PWRREQ)
          |             ~~~~~~~~~~~~~~~~^~~~
    ../drivers/regulator/max77802-regulator.c:62:22: note: while referencing 'opmode'
       62 |         unsigned int opmode[MAX77802_REG_MAX];
          |                      ^~~~~~

    Cc: Javier Martinez Canillas <javier@dowhile0.org>
    Cc: Liam Girdwood <lgirdwood@gmail.com>
    Cc: Mark Brown <broonie@kernel.org>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Acked-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://lore.kernel.org/r/20230127225203.never.864-kees@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 466b8849a4eb9013aa27cd90b2fd486e6b16134f
 Author: Kees Cook <keescook@chromium.org>
 Date:   Fri Jan 27 14:41:29 2023 -0800

    ASoC: kirkwood: Iterate over array indexes instead of using pointer math

    [ Upstream commit b3bcedc0402fcdc5c8624c433562d9d1882749d8 ]

    Walking the dram->cs array was seen as accesses beyond the first array
    item by the compiler. Instead, use the array index directly. This allows
    for run-time bounds checking under CONFIG_UBSAN_BOUNDS as well. Seen
    with GCC 13 with -fstrict-flex-arrays:

    ../sound/soc/kirkwood/kirkwood-dma.c: In function
    'kirkwood_dma_conf_mbus_windows.constprop':
    ../sound/soc/kirkwood/kirkwood-dma.c:90:24: warning: array subscript 0 is outside array
 bounds of 'const struct mbus_dram_window[0]' [-Warray-bounds=]
       90 |                 if ((cs->base & 0xffff0000) < (dma & 0xffff0000)) {
          |                      ~~^~~~~

    Cc: Liam Girdwood <lgirdwood@gmail.com>
    Cc: Mark Brown <broonie@kernel.org>
    Cc: Jaroslav Kysela <perex@perex.cz>
    Cc: Takashi Iwai <tiwai@suse.com>
    Cc: alsa-devel@alsa-project.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20230127224128.never.410-kees@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 37a3eb6054d17676ce2a0bb5dd1fbf7733ecfa7d
 Author: 강신형 <s47.kang@samsung.com>
 Date:   Thu Dec 8 14:18:18 2022 +0900

    ASoC: soc-compress: Reposition and add pcm_mutex

    [ Upstream commit aa9ff6a4955fdba02b54fbc4386db876603703b7 ]

    If panic_on_warn is set and compress stream(DPCM) is started,
    then kernel panic occurred because card->pcm_mutex isn't held appropriately.
```

```
In the following functions, warning were issued at this line
"snd_soc_dpcm_mutex_assert_held".

static int dpcm_be_connect(struct snd_soc_pcm_runtime *fe,
                struct snd_soc_pcm_runtime *be, int stream)
{
        ...
        snd_soc_dpcm_mutex_assert_held(fe);
        ...
}

void dpcm_be_disconnect(struct snd_soc_pcm_runtime *fe, int stream)
{
        ...
        snd_soc_dpcm_mutex_assert_held(fe);
        ...
}

void snd_soc_runtime_action(struct snd_soc_pcm_runtime *rtd,
                            int stream, int action)
{
        ...
        snd_soc_dpcm_mutex_assert_held(rtd);
        ...
}

int dpcm_dapm_stream_event(struct snd_soc_pcm_runtime *fe, int dir,
        int event)
{
        ...
        snd_soc_dpcm_mutex_assert_held(fe);
        ...
}

These functions are called by soc_compr_set_params_fe, soc_compr_open_fe
and soc_compr_free_fe
without pcm_mutex locking. And this is call stack.

[  414.527841][ T2179] pc : dpcm_process_paths+0x5a4/0x750
[  414.527848][ T2179] lr : dpcm_process_paths+0x37c/0x750
[  414.527945][ T2179] Call trace:
[  414.527949][ T2179]  dpcm_process_paths+0x5a4/0x750
[  414.527955][ T2179]  soc_compr_open_fe+0xb0/0x2cc
[  414.527972][ T2179]  snd_compr_open+0x180/0x248
[  414.527981][ T2179]  snd_open+0x15c/0x194
[  414.528003][ T2179]  chrdev_open+0x1b0/0x220
[  414.528023][ T2179]  do_dentry_open+0x30c/0x594
[  414.528045][ T2179]  vfs_open+0x34/0x44
[  414.528053][ T2179]  path_openat+0x914/0xb08
[  414.528062][ T2179]  do_filp_open+0xc0/0x170
[  414.528068][ T2179]  do_sys_openat2+0x94/0x18c
[  414.528076][ T2179]  __arm64_sys_openat+0x78/0xa4
[  414.528084][ T2179]  invoke_syscall+0x48/0x10c
[  414.528094][ T2179]  el0_svc_common+0xbc/0x104
[  414.528099][ T2179]  do_el0_svc+0x34/0xd8
[  414.528103][ T2179]  el0_svc+0x34/0xc4
[  414.528125][ T2179]  el0t_64_sync_handler+0x8c/0xfc
[  414.528133][ T2179]  el0t_64_sync+0x1a0/0x1a4
[  414.528142][ T2179] Kernel panic - not syncing: panic_on_warn set ...

So, I reposition and add pcm_mutex to resolve lockdep error.

Signed-off-by: Shinhyung Kang <s47.kang@samsung.com>
Link: https://lore.kernel.org/r/016401d90ac4$7b6848c0$7238da40$@samsung.com
Signed-off-by: Mark Brown <broonie@kernel.org>
```

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

commit 6f1c10536d2b21eb2c5dbf3ca818093de5b42e3a
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Wed Jan 25 11:14:11 2023 +0100

    drm/msm/dpu: Add DSC hardware blocks to register snapshot

    [ Upstream commit a7efe60e36b9c0e966d7f82ac90a89b591d984e9 ]

    Add missing DSC hardware block register ranges to the snapshot utility
    to include them in dmesg (on MSM_DISP_SNAPSHOT_DUMP_IN_CONSOLE) and the
    kms debugfs file.

    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/520175/
    Link: https://lore.kernel.org/r/20230125101412.216924-1-marijn.suijten@somainline.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9eba635cfc915c8e1ad022d101520f865ced27ec
Author: Jakob Koschel <jkl820.git@gmail.com>
Date:   Fri Jan 20 00:23:20 2023 +0100

    docs/scripts/gdb: add necessary make scripts_gdb step

    [ Upstream commit 6b219431037bf98c9efd49716aea9b68440477a3 ]

    In order to debug the kernel successfully with gdb you need to run
    'make scripts_gdb' nowadays.

    This was changed with the following commit:

    Commit 67274c083438340ad16c ("scripts/gdb: delay generation of gdb
    constants.py")

    In order to have a complete guide for beginners this remark
    should be added to the offial documentation.

    Signed-off-by: Jakob Koschel <jkl820.git@gmail.com>
    Link: https://lore.kernel.org/r/20230112-documentation-gdb-v2-1-292785c43dc9@gmail.com
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e5a6618bfef0194c4bb51c759f531546e0cef360
Author: farah kassabri <fkassabri@habana.ai>
Date:   Tue Jan 10 12:29:55 2023 +0200

    habanalabs: fix bug in timestamps registration code

    [ Upstream commit ac5af9900f82b7034de7c9eb1d70d030ba325607 ]

    Protect re-using the same timestamp buffer record before actually
    adding it to the to interrupt wait list.
    Mark ts buff offset as in use in the spinlock protection area of the
    interrupt wait list to avoid getting in the re-use section in
    ts_buff_get_kernel_ts_record before adding the node to the list.
    this scenario might happen when multiple threads are racing on
    same offset and one thread could set data in the ts buff in
    ts_buff_get_kernel_ts_record then the other thread takes over
    and get to ts_buff_get_kernel_ts_record and we will try
    to re-use the same ts buff offset then we will try to
    delete a non existing node from the list.
```

```
    Signed-off-by: farah kassabri <fkassabri@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b08669fc7dacc78789c1814a6c41b4ee59598d5a
Author: Moti Haimovski <mhaimovski@habana.ai>
Date:   Thu Dec 29 12:44:09 2022 +0200

    habanalabs: extend fatal messages to contain PCI info

    [ Upstream commit 2a0a839b6a28f7c4c528bb75b740c7f38ef79a37 ]

    This commit attaches the PCI device address to driver fatal messages
    in order to ease debugging in multi-device setups.

    Signed-off-by: Moti Haimovski <mhaimovski@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 159361617e4c004f3c7dd4e75c40d63026168a70
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Wed Jan 25 21:04:06 2023 +0100

    drm/client: Test for connectors before sending hotplug event

    [ Upstream commit c2bb3be64eb7182285846123219230375af61abd ]

    Test for connectors in the client code and remove a similar test
    from the generic fbdev emulation. Do nothing if the test fails.
    Not having connectors indicates a driver bug.

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230125200415.14123-2-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6dce564255330721987c75b7c36060a8d3e6bc84
Author: Roman Li <roman.li@amd.com>
Date:   Thu Jan 5 17:44:10 2023 -0500

    drm/amd/display: Set hvm_enabled flag for S/G mode

    [ Upstream commit 40e9f3f067bc6fb47b878f8ba0a9cc7b93abbf49 ]

    [Why]
    After enabling S/G on dcn314 a screen corruption may be observed.
    HostVM flag should be set in S/G mode to be included in DML calculations.

    [How]
    In S/G mode gpu_vm_support flag is set.
    Use its value to init is_hvm_enabled.

    Reviewed-by: Nicholas Kazlauskas <Nicholas.Kazlauskas@amd.com>
    Acked-by: Alan Liu <HaoPing.Liu@amd.com>
    Signed-off-by: Roman Li <roman.li@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d099b0f127c1c1a05db320bc1d67be5ca4edd60d
Author: Wayne Lin <Wayne.Lin@amd.com>
```

```
Date:   Fri Dec 23 09:57:02 2022 +0800

    drm/drm_print: correct format problem

    [ Upstream commit d987150b539271b0394f24c1c648d2846662adb4 ]

    [why & how]
    __drm_dbg() parameter set format is wrong and not aligned with the
    format under CONFIG_DRM_USE_DYNAMIC_DEBUG is on. Fix it.

    Signed-off-by: Wayne Lin <Wayne.Lin@amd.com>
    Signed-off-by: Harry Wentland <harry.wentland@amd.com>
    Acked-by: Harry Wentland <harry.wentland@amd.com>
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20eda27847c24611bad27bd2f5ef90326b7a2fed
Author: Tomi Valkeinen <tomi.valkeinen+renesas@ideasonboard.com>
Date:   Mon Jan 23 12:47:41 2023 +0200

    drm: rcar-du: Fix setting a reserved bit in DPLLCR

    [ Upstream commit 5fbc2f3b91d27e12b61494704876409957 0cbb55 ]

    On H3 ES1.x two bits in DPLLCR are used to select the DU input dot clock
    source. These are bits 20 and 21 for DU2, and bits 22 and 23 for DU1. On
    non-ES1.x, only the higher bits are used (bits 21 and 23), and the lower
    bits are reserved and should be set to 0.

    The current code always sets the lower bits, even on non-ES1.x.

    For both DU1 and DU2, on all SoC versions, when writing zeroes to those
    bits the input clock is DCLKIN, and thus there's no difference between
    ES1.x and non-ES1.x.

    For DU1, writing 0b10 to the bits (or only writing the higher bit)
    results in using PLL0 as the input clock, so in this case there's also
    no difference between ES1.x and non-ES1.x.

    However, for DU2, writing 0b10 to the bits results in using PLL0 as the
    input clock on ES1.x, whereas on non-ES1.x it results in using PLL1. On
    ES1.x you need to write 0b11 to select PLL1.

    The current code always writes 0b11 to PLCS0 field to select PLL1 on all
    SoC versions, which works but causes an illegal (in the sense of not
    allowed by the documentation) write to a reserved bit field.

    To remove the illegal bit write on PLSC0 we need to handle the input dot
    clock selection differently for ES1.x and non-ES1.x.

    Add a new quirk, RCAR_DU_QUIRK_H3_ES1_PLL, for this. This way we can
    always set the bit 21 on PLSC0 when choosing the PLL as the source
    clock, and additionally set the bit 20 when on ES1.x.

    Signed-off-by: Tomi Valkeinen <tomi.valkeinen+renesas@ideasonboard.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c51dab5e5237b3fa9c37f6b6ef9d587269b3ce4a
Author: Tomi Valkeinen <tomi.valkeinen+renesas@ideasonboard.com>
Date:   Mon Jan 23 12:47:40 2023 +0200

    drm: rcar-du: Add quirk for H3 ES1.x pclk workaround
```

```
    [ Upstream commit 4f548bc48a2b4c4e54eecfddb6f7d24aa1b98768 ]

    rcar_du_crtc.c does a soc_device_match() in
    rcar_du_crtc_set_display_timing() to find out if the SoC is H3 ES1.x, and
    if so, apply a workaround.

    We will need another H3 ES1.x check in the following patch, so rather than
    adding more soc_device_match() calls, let's add a rcar_du_device_info
    entry for the ES1, and a quirk flag,
    RCAR_DU_QUIRK_H3_ES1_PCLK_STABILITY, for the workaround.

    Signed-off-by: Tomi Valkeinen <tomi.valkeinen+renesas@ideasonboard.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 759ea5677c362fb1e3edc667260ba9f409dc931d
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Jan 10 10:16:51 2023 +0800

    drm/msm/dsi: Add missing check for alloc_ordered_workqueue

    [ Upstream commit 115906ca7b535afb1fe7b5406c566ccd3873f82b ]

    Add check for the return value of alloc_ordered_workqueue as it may return
    NULL pointer and cause NULL pointer dereference.

    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/517646/
    Link: https://lore.kernel.org/r/20230110021651.12770-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5478c800166a09346277c03f6133e8c3476e50f9
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Mon Dec 26 13:54:54 2022 +0100

    HID: uclogic: Add support for XP-PEN Deco Pro MW

    [ Upstream commit 9266a88156d1fbb8e50d6eeff7bac44ad4eaecc2 ]

    The XP-PEN Deco Pro MW is a UGEE v2 device with a frame with 8 buttons,
    a bitmap dial and a mouse. Its pen has 2 buttons, supports tilt and
    pressure.

    It can be connected using a USB cable or, to use it in wireless mode,
    using a USB Bluetooth dongle. When it is connected in wireless mode the
    device battery is used to power it.

    All the pieces to support it are already in place. Add its ID and
    quirks in order to support the device.

    Link: https://github.com/DIGImend/digimend-kernel-drivers/issues/622
    Tested-by: Andreas Grosse <andig.mail@t-online.de>
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0924f4bcb6d5c8e37918f1b9ec7b5010a056cb63
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Mon Dec 26 13:54:53 2022 +0100

    HID: uclogic: Add support for XP-PEN Deco Pro SW
```

```
    [ Upstream commit 7744ca571af55b794595cff2da9d51a26904998f ]

    The XP-PEN Deco Pro SW is a UGEE v2 device with a frame with 8 buttons,
    a bitmap dial and a mouse; however, the UCLOGIC_MOUSE_FRAME_QUIRK is
    required because it reports an incorrect frame type. Its pen has 2
    buttons, supports tilt and pressure.

    It can be connected using a USB cable or, to use it in wireless mode,
    using a USB Bluetooth dongle. When it is connected in wireless mode the
    device battery is used to power it.

    All the pieces to support it are already in place. Add its ID and
    quirks in order to support the device.

    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 96e0f2b11395670e4ddcad401612beeb0a382311
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Mon Dec 26 13:54:50 2022 +0100

    HID: uclogic: Add battery quirk

    [ Upstream commit f60c377f52de37f8705c5fc6d57737fdaf309ff9 ]

    Some UGEE v2 tablets have a wireless version with an internal battery
    and their firmware is able to report their battery level.

    However, there was not found a field on their descriptor indicating
    whether the tablet has battery or not, making it mandatory to classify
    such devices through the UCLOGIC_BATTERY_QUIRK quirk.

    Tested-by: Mia Kanashi <chad@redpilled.dev>
    Tested-by: Andreas Grosse <andig.mail@t-online.de>
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d329bacd8baff4f73cdfa9ad3aab5b96fbd4c015
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Mon Dec 26 13:54:49 2022 +0100

    HID: uclogic: Add frame type quirk

    [ Upstream commit 14b71e6ad8ca59dd734c7fa9676f3d60bddee2a9 ]

    The report descriptor used to get information about UGEE v2 devices is
    incorrect in the XP-PEN Deco Pro SW. It indicates that the device frame
    is of type UCLOGIC_PARAMS_FRAME_BUTTONS but the device has a frame of
    type UCLOGIC_PARAMS_FRAME_MOUSE.

    Here is the original report descriptor:

      0x0e 0x03 0xc8 0xb3 0x34 0x65 0x08 0x00 0xff 0x1f 0xd8 0x13 0x00 0x00
                                        ^ This byte should be 2

    Add a quirk to be able to fix the reported frame type.

    Tested-by: Mia Kanashi <chad@redpilled.dev>
    Tested-by: Andreas Grosse <andig.mail@t-online.de>
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 24aaf6603600d6d1159973c809ea2737664b28c4
Author: Brandon Syu <Brandon.Syu@amd.com>
Date:   Thu Dec 22 16:43:37 2022 +0800

    drm/amd/display: fix mapping to non-allocated address

    [ Upstream commit 9190d4a263264eabf715f5fc1827da45e3fdc247 ]

    [Why]
    There is an issue mapping non-allocated location of memory.
    It would allocate gpio registers from an array out of bounds.

    [How]
    Patch correct numbers of bounds for using.

    Tested-by: Daniel Wheeler <Daniel.Wheeler@amd.com>
    Reviewed-by: Martin Leung <Martin.Leung@amd.com>
    Acked-by: Rodrigo Siqueira <Rodrigo.Siqueira@amd.com>
    Signed-off-by: Brandon Syu <Brandon.Syu@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1bdea8ee92a6abc650b2189fd5c53f36859baecb
Author: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
Date:   Wed Nov 30 10:50:46 2022 +0800

    drm: amd: display: Fix memory leakage

    [ Upstream commit 6b8701be1f66064ca72733c5f6e13748cdbf8397 ]

    This commit fixes memory leakage in dc_construct_ctx() function.

    Signed-off-by: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 586fda3d939e9c4edacd9d8cdcc180c8b22f9afb
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Wed Jan 11 15:52:31 2023 -0600

    drm/amd: Avoid ASSERT for some message failures

    [ Upstream commit 3e5019ee67760cd61b2a5fd605e1289c2f92d983 ]

    On DCN314 when resuming from s0i3 an ASSERT is shown indicating that
    `VBIOSSMC_MSG_SetHardMinDcfclkByFreq` returned `VBIOSSMC_Result_Failed`.

    This isn't a driver bug; it's a BIOS/configuration bug. To make this
    easier to triage, add an explicit warning when this issue happens.

    This matches the behavior utilized for failures with
    `VBIOSSMC_MSG_TransferTableDram2Smu` configuration.

    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1300f91c97c1a615d6214add6372091e6715766e
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Mon Dec 19 17:05:00 2022 +0100

    Revert "fbcon: don't lose the console font across generic->chip driver switch"
```

```
    [ Upstream commit 12d5796d55f9fd9e4b621003127c99e176665064 ]

    This reverts commit ae1287865f5361fa138d4d3b1b6277908b54eac9.

    Always free the console font when deinitializing the framebuffer
    console. Subsequent framebuffer consoles will then use the default
    font. Rely on userspace to load any user-configured font for these
    consoles.

    Commit ae1287865f53 ("fbcon: don't lose the console font across
    generic->chip driver switch") was introduced to work around losing
    the font during graphics-device handover. [1][2] It kept a dangling
    pointer with the font data between loading the two consoles, which is
    fairly adventurous hack. It also never covered cases when the other
    consoles, such as VGA text mode, where involved.

    The problem has meanwhile been solved in userspace. Systemd comes
    with a udev rule that re-installs the configured font when a console
    comes up. [3] So the kernel workaround can be removed.

    This also removes one of the two special cases triggered by setting
    FBINFO_MISC_FIRMWARE in an fbdev driver.

    Tested during device handover from efifb and simpledrm to radeon. Udev
    reloads the configured console font for the new driver's terminal.

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://bugzilla.redhat.com/show_bug.cgi?id=892340 # 1
    Link: https://bugzilla.redhat.com/show_bug.cgi?id=1074624 # 2
    Link: https://cgit.freedesktop.org/systemd/systemd/tree/src/vconsole/90-
vconsole.rules.in?h=v222 # 3
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221219160516.23436-3-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8becb97918f04bb177bc9c4e00c2bdb302e00944
Author: Justin Tee <justin.tee@broadcom.com>
Date:   Mon Jan 9 15:33:13 2023 -0800

    scsi: lpfc: Fix use-after-free KFENCE violation during sysfs firmware write

    [ Upstream commit 21681b81b9ae548c5dae7ae00d931197a27f480c ]

    During the sysfs firmware write process, a use-after-free read warning is
    logged from the lpfc_wr_object() routine:

      BUG: KFENCE: use-after-free read in lpfc_wr_object+0x235/0x310 [lpfc]
      Use-after-free read at 0x0000000000cf164d (in kfence-#111):
      lpfc_wr_object+0x235/0x310 [lpfc]
      lpfc_write_firmware.cold+0x206/0x30d [lpfc]
      lpfc_sli4_request_firmware_update+0xa6/0x100 [lpfc]
      lpfc_request_firmware_upgrade_store+0x66/0xb0 [lpfc]
      kernfs_fop_write_iter+0x121/0x1b0
      new_sync_write+0x11c/0x1b0
      vfs_write+0x1ef/0x280
      ksys_write+0x5f/0xe0
      do_syscall_64+0x59/0x90
      entry_SYSCALL_64_after_hwframe+0x63/0xcd

    The driver accessed wr_object pointer data, which was initialized into
    mailbox payload memory, after the mailbox object was released back to the
    mailbox pool.

    Fix by moving the mailbox free calls to the end of the routine ensuring
```

```
      that we don't reference internal mailbox memory after release.

      Signed-off-by: Justin Tee <justin.tee@broadcom.com>
      Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 81ea4f89b8380c6116726b6da19fa39488298009
  Author: Philip Yang <Philip.Yang@amd.com>
  Date:   Mon Jan 9 18:08:17 2023 -0500

      drm/amdkfd: Page aligned memory reserve size

      [ Upstream commit 0c2dece8fb541ab07b68c3312a1065fa9c927a81 ]

      Use page aligned size to reserve memory usage because page aligned TTM
      BO size is used to unreserve memory usage, otherwise no page aligned
      size causes memory usage accounting unbalanced.

      Change vram_used definition type to int64_t to be able to trigger
      WARN_ONCE(adev && adev->kfd.vram_used < 0, "..."), to help debug the
      accounting issue with warning and backtrace.

      Signed-off-by: Philip Yang <Philip.Yang@amd.com>
      Reviewed-by: Felix Kuehling <Felix.Kuehling@amd.com>
      Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 383781beaf0ed13c2f1ebefd47136792b2f66079
  Author: Mario Limonciello <mario.limonciello@amd.com>
  Date:   Tue Jan 3 16:03:41 2023 -0600

      drm/amd: Avoid BUG() for case of SRIOV missing IP version

      [ Upstream commit 93fec4f8c158584065134b4d45e875499bf517c8 ]

      No need to crash the kernel.  AMDGPU will now fail to probe.

      Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
      Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
      Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
      Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit e2791f2f4d1d804e45fa91b14295c326b64c65f1
  Author: Liwei Song <liwei.song@windriver.com>
  Date:   Fri Jan 6 17:47:29 2023 +0800

      drm/radeon: free iio for atombios when driver shutdown

      [ Upstream commit 4773fadedca918faec443daaca5e4ea1c0ced144 ]

      Fix below kmemleak when unload radeon driver:

      unreferenced object 0xffff9f8608ede200 (size 512):
        comm "systemd-udevd", pid 326, jiffies 4294682822 (age 716.338s)
        hex dump (first 32 bytes):
          00 00 00 00 c4 aa ec aa 14 ab 00 00 00 00 00 00  ................
          00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
        backtrace:
          [<0000000062fadebe>] kmem_cache_alloc_trace+0x2f1/0x500
          [<00000000b6883cea>] atom_parse+0x117/0x230 [radeon]
          [<00000000158c23fd>] radeon_atombios_init+0xab/0x170 [radeon]
          [<00000000683f672e>] si_init+0x57/0x750 [radeon]
          [<00000000566cc31f>] radeon_device_init+0x559/0x9c0 [radeon]
          [<0000000046efabb3>] radeon_driver_load_kms+0xc1/0x1a0 [radeon]
```

```
          [<00000000b5155064>] drm_dev_register+0xdd/0x1d0
          [<0000000045fec835>] radeon_pci_probe+0xbd/0x100 [radeon]
          [<00000000e69ecca3>] pci_device_probe+0xe1/0x160
          [<0000000019484b76>] really_probe.part.0+0xc1/0x2c0
          [<000000003f2649da>] __driver_probe_device+0x96/0x130
          [<00000000231c5bb1>] driver_probe_device+0x24/0xf0
          [<0000000000a42377>] __driver_attach+0x77/0x190
          [<00000000d7574da6>] bus_for_each_dev+0x7f/0xd0
          [<00000000633166d2>] driver_attach+0x1e/0x30
          [<00000000313b05b8>] bus_add_driver+0x12c/0x1e0

      iio was allocated in atom_index_iio() called by atom_parse(),
      but it doesn't got released when the dirver is shutdown.
      Fix this kmemleak by free it in radeon_atombios_fini().

      Signed-off-by: Liwei Song <liwei.song@windriver.com>
      Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 0c0da18568a23f5993cef8a9c845f868b354b52a
  Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
  Date:   Tue Dec 6 16:10:22 2022 -0500

      drm/amd/display: Defer DIG FIFO disable after VID stream enable

      [ Upstream commit 2d90a1c054831338d57b39aec4d273cf3e867590 ]

      [Why]
      On some monitors we see a brief flash of corruption during the
      monitor disable sequence caused by FIFO being disabled in the middle
      of an active DP stream.

      [How]
      Wait until DP vid stream is disabled before turning off the FIFO.

      The FIFO reset on DP unblank should take care of clearing any FIFO
      error, if any.

      Acked-by: Aurabindo Pillai <aurabindo.pillai@amd.com>
      Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
      Reviewed-by: Syed Hassan <Syed.Hassan@amd.com>
      Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
      Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 2c24102fc36be944d0a07bf5f3b8418f346fc6ab
  Author: Carlo Caione <ccaione@baylibre.com>
  Date:   Mon Dec 19 10:02:38 2022 +0100

      drm/tiny: ili9486: Do not assume 8-bit only SPI controllers

      [ Upstream commit 77772e607522daa61f3af74df018559db75c43d6 ]

      The pixel data for the ILI9486 is always 16-bits wide and it must be
      sent over the SPI bus. When the controller is only able to deal with
      8-bit transfers, this 16-bits data needs to be swapped before the
      sending to account for the big endian bus, this is on the contrary not
      needed when the SPI controller already supports 16-bits transfers.

      The decision about swapping the pixel data or not is taken in the MIPI
      DBI code by probing the controller capabilities: if the controller only
      suppors 8-bit transfers the data is swapped, otherwise it is not.

      This swapping/non-swapping is relying on the assumption that when the
      controller does support 16-bit transactions then the data is sent
```

```
    unswapped in 16-bits-per-word over SPI.

    The problem with the ILI9486 driver is that it is forcing 8-bit
    transactions also for controllers supporting 16-bits, violating the
    assumption and corrupting the pixel data.

    Align the driver to what is done in the MIPI DBI code by adjusting the
    transfer size to the maximum allowed by the SPI controller.

    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Carlo Caione <ccaione@baylibre.com>
    Reviewed-by: Kamlesh Gurudasani <kamlesh.gurudasani@gmail.com>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221116-s905x_spi_ili9486-v4-2-
f86b4463b9e4@baylibre.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 811b20319cd841801178939a45216d9896d3a5eb
Author: Jingyuan Liang <jingyliang@chromium.org>
Date:   Tue Dec 13 22:53:30 2022 +0000

    HID: Add Mapping for System Microphone Mute

    [ Upstream commit 2d60f9f4f26785a00273cb81fe60eff129ebd449 ]

    HUTRR110 added a new usage code for a key that is supposed to
    mute/unmute microphone system-wide.

    Map the new usage code(0x01 0xa9) to keycode KEY_MICMUTE.
    Additionally hid-debug is adjusted to recognize this keycode as well.

    Signed-off-by: Jingyuan Liang <jingyliang@chromium.org>
    Reviewed-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fbf3b9f47d782ca50eeb023f614c9e58b1d927bf
Author: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
Date:   Fri Sep 16 11:22:05 2022 +0300

    drm/omap: dsi: Fix excessive stack usage

    [ Upstream commit cfca78971b9233aef0891507a98fba62046d4542 ]

    dsi_dump_dsi_irqs(), a function used for debugfs prints, has a large
    struct in its frame, which can result in:

    drivers/gpu/drm/omapdrm/dss/dsi.c:1126:1: warning: the frame size of 1060 bytes is
larger than 1024 bytes [-Wframe-larger-than=]

    As the performance of the function is of no concern, let's allocate the
    struct with kmalloc instead.

    Compile-tested only.

    Signed-off-by: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220916082206.167427-1-
tomi.valkeinen@ideasonboard.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00b655fa96b4e941351cc4bf5ca755a65ae94a8e
Author: Roman Li <roman.li@amd.com>
Date:   Thu Dec 1 09:06:42 2022 -0500
```

```
    drm/amd/display: Fix potential null-deref in dm_resume

    [ Upstream commit 7a7175a2cd84b7874bebbf8e59f134557a34161b ]

    [Why]
    Fixing smatch error:
    dm_resume() error: we previously assumed 'aconnector->dc_link' could be null

    [How]
    Check if dc_link null at the beginning of the loop,
    so further checks can be dropped.

    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>

    Reviewed-by: Wayne Lin <Wayne.Lin@amd.com>
    Acked-by: Jasdeep Dhillon <jdhillon@amd.com>
    Signed-off-by: Roman Li <roman.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d5273f0f903f79e7d35e75e4282f548c760a243
Author: Ian Chen <ian.chen@amd.com>
Date:   Mon Nov 28 16:17:34 2022 +0800

    drm/amd/display: Revert Reduce delay when sink device not able to ACK 00340h write

    [ Upstream commit 639f6ad6df7f47db48b59956b469a6917a136afb ]

    [WHY]
    It causes regression AMD source will not write DPCD 340.

    Reviewed-by: Wayne Lin <Wayne.Lin@amd.com>
    Acked-by: Jasdeep Dhillon <jdhillon@amd.com>
    Signed-off-by: Ian Chen <ian.chen@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffa70a73b46c0616fa395c0f86b16233bcc7dd4d
Author: Dillon Varone <Dillon.Varone@amd.com>
Date:   Fri Dec 2 22:56:57 2022 -0500

    drm/amd/display: Reduce expected sdp bandwidth for dcn321

    [ Upstream commit 6b81090d6d4cc0fd818c9ec9dbb6906f921ad396 ]

    [Description]
    Modify soc BB to reduce expected sdp bandwidth and align with measurements to
    fix underflow issues.

    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Jasdeep Dhillon <jdhillon@amd.com>
    Signed-off-by: Dillon Varone <Dillon.Varone@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e5a1bdfefc9e7ce5bfa6a77c7e47ffd179474062
Author: Allen Ballway <ballway@chromium.org>
Date:   Wed Nov 30 17:08:22 2022 +0000

    drm: panel-orientation-quirks: Add quirk for DynaBook K50

    [ Upstream commit a3caf7ea0c3d5872ed0f2c51f5476aee0c47a73a ]
```

```
    Like the ASUS T100HAN for which there is already a quirk,
    the DynaBook K50 has a 800x1280 portrait screen mounted
    in the tablet part of a landscape oriented 2-in-1.
    Update the quirk to be more generic and apply to this device.

    Signed-off-by: Allen Ballway <ballway@chromium.org>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20221130170811.1.Iee9a494547541dade9eeee9521cc
8b811e76a8a0@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03eb529cd84fde8c864e07861d5723b0138e49ab
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Nov 27 19:15:39 2022 +0100

    drm: panel-orientation-quirks: Add quirk for Lenovo Yoga Tab 3 X90F

    [ Upstream commit 8a238d7f7eea7592e0764bc3b9e79e7c6354b04c ]

    The Lenovo Yoga Tab 3 X90F has a portrait 1600x2560 LCD used in
    landscape mode, add a quirk for this.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221127181539.104223-1-
hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1d18236eff5339082e27ab56e6b2c474b4b5da0
Author: Eric Dumazet <edumazet@google.com>
Date:   Fri Feb 17 18:24:54 2023 +0000

    scm: add user copy checks to put_cmsg()

    [ Upstream commit 5f1eb1ff58ea122e24adf0bc940f268ed2227462 ]

    This is a followup of commit 2558b8039d05 ("net: use a bounce
    buffer for copying skb->mark")

    x86 and powerpc define user_access_begin, meaning
    that they are not able to perform user copy checks
    when using user_write_access_begin() / unsafe_copy_to_user()
    and friends [1]

    Instead of waiting bugs to trigger on other arches,
    add a check_object_size() in put_cmsg() to make sure
    that new code tested on x86 with CONFIG_HARDENED_USERCOPY=y
    will perform more security checks.

    [1] We can not generically call check_object_size() from
    unsafe_copy_to_user() because UACCESS is enabled at this point.

    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Acked-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 46ecdc0f0986b8491342466a9028b7db345b1834
Author: Moshe Shemesh <moshe@nvidia.com>
Date:   Tue Feb 14 18:38:06 2023 +0200

    devlink: Fix TP_STRUCT_entry in trace of devlink health report
```

```
    [ Upstream commit d0ab772c1f1558af84f3293a52e9e886e08e0754 ]

    Fix a bug in trace point definition for devlink health report, as
    TP_STRUCT_entry of reporter_name should get reporter_name and not msg.

    Note no fixes tag as this is a harmless bug as both reporter_name and
    msg are strings and TP_fast_assign for this entry is correct.

    Signed-off-by: Moshe Shemesh <moshe@nvidia.com>
    Reviewed-by: Jiri Pirko <jiri@nvidia.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 53d766c968953bb6c48a3129708d190d174d05d1
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Feb 13 19:38:58 2023 +0100

    s390/kfence: fix page fault reporting

    [ Upstream commit d9c2cf67b9cfd643ba85d51bc865a89a92e4f979 ]

    Baoquan He reported lots of KFENCE reports when /proc/kcore is read,
    e.g. with crash or even simpler with dd:

     BUG: KFENCE: invalid read in copy_from_kernel_nofault+0x5e/0x120
     Invalid read at 0x00000000f4f5149f:
      copy_from_kernel_nofault+0x5e/0x120
      read_kcore+0x6b2/0x870
      proc_reg_read+0x9a/0xf0
      vfs_read+0x94/0x270
      ksys_read+0x70/0x100
      __do_syscall+0x1d0/0x200
      system_call+0x82/0xb0

    The reason for this is that read_kcore() simply reads memory that might
    have been unmapped by KFENCE with copy_from_kernel_nofault(). Any fault due
    to pages being unmapped by KFENCE would be handled gracefully by the fault
    handler (exception table fixup).

    However the s390 fault handler first reports the fault, and only afterwards
    would perform the exception table fixup. Most architectures have this in
    reversed order, which also avoids the false positive KFENCE reports when an
    unmapped page is accessed.

    Therefore change the s390 fault handler so it handles exception table
    fixups before KFENCE page faults are reported.

    Reported-by: Baoquan He <bhe@redhat.com>
    Tested-by: Baoquan He <bhe@redhat.com>
    Acked-by: Alexander Potapenko <glider@google.com>
    Link: https://lore.kernel.org/r/20230213183858.1473681-1-hca@linux.ibm.com
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ca6d6e730cd42b0fb3853dcc22f15127a9a32c72
Author: Michael Kelley <mikelley@microsoft.com>
Date:   Sun Feb 12 21:08:01 2023 -0800

    hv_netvsc: Check status in SEND_RNDIS_PKT completion message

    [ Upstream commit dca5161f9bd052e9e73be90716ffd57e8762c697 ]

    Completion responses to SEND_RNDIS_PKT messages are currently processed
    regardless of the status in the response, so that resources associated
```

```
    with the request are freed.  While this is appropriate, code bugs that
    cause sending a malformed message, or errors on the Hyper-V host, go
    undetected. Fix this by checking the status and outputting a rate-limited
    message if there is an error.

    Signed-off-by: Michael Kelley <mikelley@microsoft.com>
    Reviewed-by: Haiyang Zhang <haiyangz@microsoft.com>
    Link: https://lore.kernel.org/r/1676264881-48928-1-git-send-email-mikelley@microsoft.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee9732c0650d339c5e50a0106e6f34873b9e9a50
Author: Zong-Zhe Yang <kevin_yang@realtek.com>
Date:   Thu Jan 19 14:35:28 2023 +0800

    wifi: rtw89: debug: avoid invalid access on RTW89_DBG_SEL_MAC_30

    [ Upstream commit c074da21dd346e0cfef5d08b0715078d7aea7f8d ]

    Only 8852C chip has valid pages on RTW89_DBG_SEL_MAC_30. To other chips,
    this section is an address hole. It will lead to crash if trying to access
    this section on chips except for 8852C. So, we avoid that.

    Signed-off-by: Zong-Zhe Yang <kevin_yang@realtek.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230119063529.61563-2-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e2eba59b0b356582b478b5d37943821849786dc1
Author: Moises Cardona <moisesmcardona@gmail.com>
Date:   Sun Dec 25 14:07:13 2022 -0500

    Bluetooth: btusb: Add VID:PID 13d3:3529 for Realtek RTL8821CE

    [ Upstream commit 1eec3b95b5ce7fb2cdd273ac4f8b24b1ed6776a1 ]

    This patch adds VID:PID 13d3:3529 to the btusb.c file.

    This VID:PID is found in the Realtek RTL8821CE module
    (M.2 module AW-CB304NF on an ASUS E210MA laptop)

    Output of /sys/kernel/debug/usb/devices:

    T:  Bus=01 Lev=01 Prnt=01 Port=07 Cnt=02 Dev#=  3 Spd=12   MxCh= 0
    D:  Ver= 1.10 Cls=e0(wlcon) Sub=01 Prot=01 MxPS=64 #Cfgs=  1
    P:  Vendor=13d3 ProdID=3529 Rev= 1.10
    S:  Manufacturer=Realtek
    S:  Product=Bluetooth Radio
    C:* #Ifs= 2 Cfg#= 1 Atr=e0 MxPwr=500mA
    I:* If#= 0 Alt= 0 #EPs= 3 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=81(I) Atr=03(Int.) MxPS=  16 Ivl=1ms
    E:  Ad=02(O) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    E:  Ad=82(I) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    I:* If#= 1 Alt= 0 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=03(O) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    I:  If#= 1 Alt= 1 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=03(O) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    I:  If#= 1 Alt= 2 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    I:  If#= 1 Alt= 3 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  25 Ivl=1ms
```

```
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  25 Ivl=1ms
      I:  If#= 1 Alt= 4 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  33 Ivl=1ms
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  33 Ivl=1ms
      I:  If#= 1 Alt= 5 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  49 Ivl=1ms
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  49 Ivl=1ms

      Signed-off-by: Moises Cardona <moisesmcardona@gmail.com>
      Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 8b51d437fc586393f7323d8525781f11a8e21cb0
  Author: Mario Limonciello <mario.limonciello@amd.com>
  Date:   Wed Dec 14 20:18:54 2022 -0600

      Bluetooth: btusb: Add new PID/VID 0489:e0f2 for MT7921

      [ Upstream commit 83458a5f272b303479e7d2f451600817a7350b6b ]

      This bluetooth device is found in a combo WLAN/BT card
      for a MediaTek 7921e.

      The device information:

      T:  Bus=01 Lev=01 Prnt=01 Port=02 Cnt=01 Dev#=  2 Spd=480  MxCh= 0
      D:  Ver= 2.10 Cls=ef(misc ) Sub=02 Prot=01 MxPS=64 #Cfgs=  1
      P:  Vendor=0489 ProdID=e0f2 Rev= 1.00
      S:  Manufacturer=MediaTek Inc.
      S:  Product=Wireless_Device
      S:  SerialNumber=000000000
      C:* #Ifs= 3 Cfg#= 1 Atr=e0 MxPwr=100mA
      A:  FirstIf#= 0 IfCount= 3 Cls=e0(wlcon) Sub=01 Prot=01
      I:* If#= 0 Alt= 0 #EPs= 3 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=81(I) Atr=03(Int.) MxPS=  16 Ivl=125us
      E:  Ad=82(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
      E:  Ad=02(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms
      I:* If#= 1 Alt= 0 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=   0 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=   0 Ivl=1ms
      I:  If#= 1 Alt= 1 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=   9 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=   9 Ivl=1ms
      I:  If#= 1 Alt= 2 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  17 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  17 Ivl=1ms
      I:  If#= 1 Alt= 3 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  25 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  25 Ivl=1ms
      I:  If#= 1 Alt= 4 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  33 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  33 Ivl=1ms
      I:  If#= 1 Alt= 5 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  49 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  49 Ivl=1ms
      I:  If#= 1 Alt= 6 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
      E:  Ad=83(I) Atr=01(Isoc) MxPS=  63 Ivl=1ms
      E:  Ad=03(O) Atr=01(Isoc) MxPS=  63 Ivl=1ms
      I:* If#= 2 Alt= 0 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=(none)
      E:  Ad=8a(I) Atr=03(Int.) MxPS=  64 Ivl=125us
      E:  Ad=0a(O) Atr=03(Int.) MxPS=  64 Ivl=125us
      I:  If#= 2 Alt= 1 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=(none)
      E:  Ad=8a(I) Atr=03(Int.) MxPS= 512 Ivl=125us
      E:  Ad=0a(O) Atr=03(Int.) MxPS= 512 Ivl=125us
```

```
    Cc: Sean Wang <sean.wang@mediatek.com>
    Cc: Anson Tsao <anson.tsao@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e261de03d5a2b1504abfc8dc3075fb38b1ecb243
Author: Marcel Holtmann <marcel@holtmann.org>
Date:   Fri Dec 16 21:12:47 2022 +0100

    Bluetooth: Fix issue with Actions Semi ATS2851 based devices

    [ Upstream commit 5d44f76fab0839799b19cbc88d13575da968dc08 ]

    Their devices claim to support the erroneous data reporting, but don't
    actually support the required commands. So blacklist them and add a
    quirk.

      < HCI Command: Read Default Erroneous Data Reporting (0x03|0x005a) plen 0
      > HCI Event: Command Status (0x0f) plen 4
            Read Default Erroneous Data Reporting (0x03|0x005a) ncmd 1
              Status: Unknown HCI Command (0x01)

    T:  Bus=02 Lev=02 Prnt=08 Port=02 Cnt=01 Dev#= 10 Spd=12   MxCh= 0
    D:  Ver= 2.00 Cls=e0(wlcon) Sub=01 Prot=01 MxPS=64 #Cfgs=  1
    P:  Vendor=10d7 ProdID=b012 Rev=88.91
    S:  Manufacturer=Actions
    S:  Product=general adapter
    S:  SerialNumber=ACTIONS1234
    C:* #Ifs= 2 Cfg#= 1 Atr=c0 MxPwr=100mA
    I:* If#= 0 Alt= 0 #EPs= 3 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=81(I) Atr=03(Int.) MxPS=  64 Ivl=1ms
    E:  Ad=02(O) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    E:  Ad=82(I) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    I:* If#= 1 Alt= 0 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    I:  If#= 1 Alt= 1 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    I:  If#= 1 Alt= 2 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    I:  If#= 1 Alt= 3 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=  25 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  25 Ivl=1ms
    I:  If#= 1 Alt= 4 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=  33 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  33 Ivl=1ms
    I:  If#= 1 Alt= 5 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=01(O) Atr=01(Isoc) MxPS=  49 Ivl=1ms
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  49 Ivl=1ms

    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 30fee10192e1239478a0987bc7ee445d5e980d46
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Feb 2 16:15:15 2023 +0100

    PM: EM: fix memory leak with using debugfs_lookup()

    [ Upstream commit a0e8c13ccd6a9a636d27353da62c2410c4eca337 ]
```

```
    When calling debugfs_lookup() the result must have dput() called on it,
    otherwise the memory will leak over time.  To make things simpler, just
    call debugfs_lookup_and_remove() instead which handles all of the logic
    at once.

    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 543d7113c37206ed7dae7bfb0b7e50955077770e
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Feb 2 15:15:45 2023 +0100

    PM: domains: fix memory leak with using debugfs_lookup()

    [ Upstream commit 0b6200e1e9f53dabdc30d0f6c51af9a5f664d32b ]

    When calling debugfs_lookup() the result must have dput() called on it,
    otherwise the memory will leak over time.  To make things simpler, just
    call debugfs_lookup_and_remove() instead which handles all of the logic
    at once.

    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Reviewed-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b588b42d077ce93c98704b41003bcec6a564b738
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Feb 2 16:12:14 2023 +0100

    time/debug: Fix memory leak with using debugfs_lookup()

    [ Upstream commit 5b268d8abaec6cbd4bd70d062e769098d96670aa ]

    When calling debugfs_lookup() the result must have dput() called on it,
    otherwise the memory will leak over time.  To make things simpler, just
    call debugfs_lookup_and_remove() instead which handles all of the logic at
    once.

    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lore.kernel.org/r/20230202151214.2306822-1-gregkh@linuxfoundation.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fc60c4f12d8a056f20d8f4d0086a36c68ffa9fdc
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Feb 6 14:49:40 2023 +0100

    s390/idle: mark arch_cpu_idle() noinstr

    [ Upstream commit a9cbc1b471d291c865907542394f1c483b93a811 ]

    linux-next commit ("cpuidle: tracing: Warn about !rcu_is_watching()")
    adds a new warning which hits on s390's arch_cpu_idle() function:

    RCU not on for: arch_cpu_idle+0x0/0x28
    WARNING: CPU: 2 PID: 0 at include/linux/trace_recursion.h:162
arch_ftrace_ops_list_func+0x24c/0x258
    Modules linked in:
    CPU: 2 PID: 0 Comm: swapper/2 Not tainted 6.2.0-rc6-next-20230202 #4
    Hardware name: IBM 8561 T01 703 (z/VM 7.3.0)
    Krnl PSW : 0404d00180000000 00000000002b55c0 (arch_ftrace_ops_list_func+0x250/0x258)
               R:0 T:1 IO:0 EX:0 Key:0 M:1 W:0 P:0 AS:3 CC:1 PM:0 RI:0 EA:3
    Krnl GPRS: c0000000ffffbfff 0000000080000002 0000000000000026 0000000000000000
```

```
               0000037ffffe3a28 0000037ffffe3a20 0000000000000000 0000000000000000
               0000000000000000 0000000000f4acf6 00000000001044f0 0000037ffffe3cb0
               0000000000000000 0000000000000000 00000000002b55bc 0000037ffffe3bb8
    Krnl Code: 00000000002b55b0: c02000840051        larl    %r2,0000000001335652
               00000000002b55b6: c0e5fff512d1        brasl   %r14,0000000000157b58
              #00000000002b55bc: af000000            mc      0,0
              >00000000002b55c0: a7f4ffe7            brc     15,00000000002b558e
               00000000002b55c4: 0707                bcr     0,%r7
               00000000002b55c6: 0707                bcr     0,%r7
               00000000002b55c8: eb6ff0480024        stmg    %r6,%r15,72(%r15)
               00000000002b55ce: b90400ef            lgr     %r14,%r15
    Call Trace:
     [<00000000002b55c0>] arch_ftrace_ops_list_func+0x250/0x258
    ([<00000000002b55bc>] arch_ftrace_ops_list_func+0x24c/0x258)
     [<0000000000f5f0fc>] ftrace_common+0x1c/0x20
     [<00000000001044f6>] arch_cpu_idle+0x6/0x28
     [<0000000000f4acf6>] default_idle_call+0x76/0x128
     [<00000000001cc374>] do_idle+0xf4/0x1b0
     [<00000000001cc6ce>] cpu_startup_entry+0x36/0x40
     [<0000000000119d00>] smp_start_secondary+0x140/0x150
     [<0000000000f5d2ae>] restart_int_handler+0x6e/0x90

    Mark arch_cpu_idle() noinstr like all other architectures with
    CONFIG_ARCH_WANTS_NO_INSTR (should) have it to fix this.

    Reviewed-by: Sven Schnelle <svens@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ff7fc0ed62b4955f2327b8fa53c3319b7d54b79
Author: Kees Cook <keescook@chromium.org>
Date:   Tue Jan 31 17:37:59 2023 -0800

    uaccess: Add minimum bounds check on kernel buffer size

    [ Upstream commit 04ffde1319a715bd0550ded3580d4ea3bc003776 ]

    While there is logic about the difference between ksize and usize,
    copy_struct_from_user() didn't check the size of the destination buffer
    (when it was known) against ksize. Add this check so there is an upper
    bounds check on the possible memset() call, otherwise lower bounds
    checks made by callers will trigger bounds warnings under -Warray-bounds.
    Seen under GCC 13:

    In function 'copy_struct_from_user',
        inlined from 'iommufd_fops_ioctl' at
    ../drivers/iommu/iommufd/main.c:333:8:
    ../include/linux/fortify-string.h:59:33: warning: '__builtin_memset' offset [57,
4294967294] is out of the bounds [0, 56] of object 'buf' with type 'union ucmd_buffer' [-
Warray-bounds=]
       59 | #define __underlying_memset     __builtin_memset
          |                                 ^
    ../include/linux/fortify-string.h:453:9: note: in expansion of macro
'__underlying_memset'
      453 |         __underlying_memset(p, c, __fortify_size); \
          |         ^~~~~~~~~~~~~~~~~~~
    ../include/linux/fortify-string.h:461:25: note: in expansion of macro
'__fortify_memset_chk'
      461 | #define memset(p, c, s) __fortify_memset_chk(p, c, s, \
          |                         ^~~~~~~~~~~~~~~~~~~~
    ../include/linux/uaccess.h:334:17: note: in expansion of macro 'memset'
      334 |                 memset(dst + size, 0, rest);
          |                 ^~~~~~
    ../drivers/iommu/iommufd/main.c: In function 'iommufd_fops_ioctl':
    ../drivers/iommu/iommufd/main.c:311:27: note: 'buf' declared here
```

```
      311 |         union ucmd_buffer buf;
          |                           ^~~
```

    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Rasmus Villemoes <linux@rasmusvillemoes.dk>
    Cc: Arnd Bergmann <arnd@arndb.de>
    Cc: Dinh Nguyen <dinguyen@kernel.org>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Andrew Morton <akpm@linux-foundation.org>
    Cc: Geert Uytterhoeven <geert@linux-m68k.org>
    Cc: Alexander Potapenko <glider@google.com>
    Acked-by: Aleksa Sarai <cyphar@cyphar.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/lkml/20230203193523.never.667-kees@kernel.org/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1015949696c099de175d8e5af786d129657c8fa
Author: Kees Cook <keescook@chromium.org>
Date: Fri Jan 27 14:39:21 2023 -0800

    coda: Avoid partial allocation of sig_inputArgs

    [Upstream commit 48df133578c70185a95a49390d42df1996ddba2a]

    GCC does not like having a partially allocated object, since it cannot
    reason about it for bounds checking when it is passed to other code.
    Instead, fully allocate sig_inputArgs. (Alternatively, sig_inputArgs
    should be defined as a struct coda_in_hdr, if it is actually not using
    any other part of the union.) Seen under GCC 13:

    ../fs/coda/upcall.c: In function 'coda_upcall':
    ../fs/coda/upcall.c:801:22: warning: array subscript 'union inputArgs[0]' is partly
outside array bounds of 'unsigned char[20]' [-Warray-bounds=]
      801 | sig_inputArgs->ih.opcode = CODA_SIGNAL;
          | ^~

    Cc: Jan Harkes <jaharkes@cs.cmu.edu>
    Cc: coda@cs.cmu.edu
    Cc: codalist@coda.cs.cmu.edu
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20230127223921.never.882-kees@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fea8f26b5649c401ddb82d36a51514bb9bb48bfc
Author: Shay Drory <shayd@nvidia.com>
Date: Wed Jan 11 13:34:02 2023 +0200

    net/mlx5: fw_tracer: Fix debug print

    [Upstream commit 988c2352273997a242f15c4fc3711773515006a2]

    The debug message specify tdsn, but takes as an argument the
    tmsn. The correct argument is tmsn, hence, fix the print.

    Signed-off-by: Shay Drory <shayd@nvidia.com>
    Reviewed-by: Moshe Shemesh <moshe@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 46652b0d0ac4aecd266f645cddaf6941faa4b630
Author: Hans de Goede <hdegoede@redhat.com>
Date: Thu Feb 2 13:44:49 2023 +0100

    ACPI: video: Fix Lenovo Ideapad Z570 DMI match

```
    [ Upstream commit 2d11eae42d52a131f06061015e49dc0f085c5bfc ]

    Multiple Ideapad Z570 variants need acpi_backlight=native to force native
    use on these pre Windows 8 machines since acpi_video backlight control
    does not work here.

    The original DMI quirk matches on a product_name of "102434U" but other
    variants may have different product_name-s such as e.g. "1024D9U".

    Move to checking product_version instead as is more or less standard for
    Lenovo DMI quirks for similar reasons.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 44e8c429f81c126f0531f2000d8656e349962ff6
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun Jan 1 12:47:57 2023 +0100

    wifi: mt76: dma: free rx_head in mt76_dma_rx_cleanup

    [ Upstream commit 1b88b47e898edef0e56e3a2f4e49f052a136153d ]

    Free rx_head skb in mt76_dma_rx_cleanup routine in order to avoid
    possible memory leak at module unload.

    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7663d70eb85ac882d106d5ece17f19bd693f8a2e
Author: Armin Wolf <W_Armin@gmx.de>
Date:   Thu Jan 26 20:40:17 2023 +0100

    platform/x86: dell-ddv: Add support for interface version 3

    [ Upstream commit 3e899fec5dfce37701d49d656954a825275bf867 ]

    While trying to solve a bugreport on bugzilla, i learned that
    some devices (for example the Dell XPS 17 9710) provide a more
    recent DDV WMI interface (version 3).
    Since the new interface version just adds an additional method,
    no code changes are necessary apart from whitelisting the version.

    Signed-off-by: Armin Wolf <W_Armin@gmx.de>
    Link: https://lore.kernel.org/r/20230126194021.381092-2-W_Armin@gmx.de
    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 f396f2c21238750c75d6254656b0cc41e30312ff
Author: Zhang Rui <rui.zhang@intel.com>
Date:   Wed Aug 24 15:44:42 2022 +0800

    tools/power/x86/intel-speed-select: Add Emerald Rapid quirk

    [ Upstream commit 61f9fdcdcd01f9a996b6db4e7092fcdfe8414ad5 ]

    Need memory frequency quirk as Sapphire Rapids in Emerald Rapids.
    So add Emerald Rapids CPU model check in is_spr_platform().

    Signed-off-by: Zhang Rui <rui.zhang@intel.com>
    [srinivas.pandruvada@linux.intel.com: Subject, changelog and code edits]
    Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
```

```
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9aa116f4bc30f076c15d90bff66b875e41d53a2f
Author: Sam James <sam@gentoo.org>
Date:   Wed Feb 1 23:00:09 2023 +0000

    gcc-plugins: drop -std=gnu++11 to fix GCC 13 build

    [ Upstream commit 5a6b64adc18d9adfb497a529ff004d59b6df151f ]

    The latest GCC 13 snapshot (13.0.1 20230129) gives the following:
    ```
    cc1: error: cannot load plugin ./scripts/gcc-plugins/randomize_layout_plugin.so
     :./scripts/gcc-plugins/randomize_layout_plugin.so: undefined symbol: tree_code_type
    ```

    This ends up being because of https://gcc.gnu.org/git/gitweb.cgi?
p=gcc.git;h=b0241ce6e37031
    upstream in GCC which changes the visibility of some types used by the kernel's
    plugin infrastructure like tree_code_type.

    After discussion with the GCC folks, we found that the kernel needs to be building
    plugins with the same flags used to build GCC - and GCC defaults to gnu++17
    right now. The minimum GCC version needed to build the kernel is GCC 5.1
    and GCC 5.1 already defaults to gnu++14 anyway, so just drop the flag, as
    all GCCs that could be used to build GCC already default to an acceptable
    version which was >= the version we forced via flags until now.

    Bug: https://gcc.gnu.org/bugzilla/show_bug.cgi?id=108634
    Signed-off-by: Sam James <sam@gentoo.org>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20230201230009.2252783-1-sam@gentoo.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd4faace51e41a82a8c0770ee0cc26088f9d9d06
Author: Oliver Hartkopp <socketcan@hartkopp.net>
Date:   Wed Jan 4 21:18:44 2023 +0100

    can: isotp: check CAN address family in isotp_bind()

    [ Upstream commit c6adf659a8ba85913e16a571d5a9bcd17d3d1234 ]

    Add missing check to block non-AF_CAN binds.

    Syzbot created some code which matched the right sockaddr struct size
    but used AF_XDP (0x2C) instead of AF_CAN (0x1D) in the address family
    field:

    bind$xdp(r2, &(0x7f0000000540)={0x2c, 0x0, r4, 0x0, r2}, 0x10)
                                       ^^^^
    This has no funtional impact but the userspace should be notified about
    the wrong address family field content.

    Link: https://syzkaller.appspot.com/text?tag=CrashLog&x=11ff9d8c480000
    Reported-by: syzbot+5aed6c3aaba661f5b917@syzkaller.appspotmail.com
    Signed-off-by: Oliver Hartkopp <socketcan@hartkopp.net>
    Link: https://lore.kernel.org/all/20230104201844.13168-1-socketcan@hartkopp.net
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 30bd1235e7703d9b190b9a22e035fec6bf1455db
Author: Alok Tiwari <alok.a.tiwari@oracle.com>
Date:   Tue Jan 17 07:45:49 2023 -0800
```

```
    netfilter: nf_tables: NULL pointer dereference in nf_tables_updobj()

    [ Upstream commit dac7f50a45216d652887fb92d6cd3b7ca7f006ea ]

    static analyzer detect null pointer dereference case for 'type'
    function __nft_obj_type_get() can return NULL value which require to handle
    if type is NULL pointer return -ENOENT.

    This is a theoretical issue, since an existing object has a type, but
    better add this failsafe check.

    Signed-off-by: Alok Tiwari <alok.a.tiwari@oracle.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d0365248bed9597b0d2e4d598e8df7aee3c6c2dd
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Tue Jan 24 18:08:38 2023 +0100

    s390/mm,ptdump: avoid Kasan vs Memcpy Real markers swapping

    [ Upstream commit 05178996e1a77e2a4664536e6d101a086a905034 ]

    ---[ Real Memory Copy Area Start ]---
    0x001bfffffffff000-0x001c000000000000         4K PTE I
    ---[ Kasan Shadow Start ]---
    ---[ Real Memory Copy Area End ]---
    0x001c000000000000-0x001c000200000000         8G PMD RW NX
    ...
    ---[ Kasan Shadow End ]---

    ptdump does a stable sort of markers. Move kasan markers after
    memcpy real to avoid swapping.

    Reviewed-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a873e7734ea3f48f67e08e9f251d73c260d181f
Author: Michael Schmitz <schmitzmic@gmail.com>
Date:   Thu Jan 12 16:55:27 2023 +1300

    m68k: Check syscall_trace_enter() return code

    [ Upstream commit 2ca8a1de4437f21562e57f9ac123914747a8e7a1 ]

    Check return code of syscall_trace_enter(), and skip syscall
    if -1. Return code will be left at what had been set by
    ptrace or seccomp (in regs->d0).

    No regression seen in testing with strace on ARAnyM.

    Signed-off-by: Michael Schmitz <schmitzmic@gmail.com>
    Reviewed-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Link: https://lore.kernel.org/r/20230112035529.13521-2-schmitzmic@gmail.com
    Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 411317d2a4a7d6049d8efeef0d32ae43f8baefce
Author: Florian Fainelli <f.fainelli@gmail.com>
Date:   Thu Jan 26 16:08:19 2023 -0800

    net: bcmgenet: Add a check for oversized packets
```

```
    [ Upstream commit 5c0862c2c962052ed5055220a00ac1cefb92fbcd ]

    Occasionnaly we may get oversized packets from the hardware which
    exceed the nomimal 2KiB buffer size we allocate SKBs with. Add an early
    check which drops the packet to avoid invoking skb_over_panic() and move
    on to processing the next packet.

    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be64d4b114f9ef90e4a37769e2eb0643097dcddd
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Jan 5 20:19:48 2023 -0800

    crypto: hisilicon: Wipe entire pool on error

    [ Upstream commit aa85923a954e7704bc9d3847dabeb8540aa98d13 ]

    To work around a Clang __builtin_object_size bug that shows up under
    CONFIG_FORTIFY_SOURCE and UBSAN_BOUNDS, move the per-loop-iteration
    mem_block wipe into a single wipe of the entire pool structure after
    the loop.

    Reported-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://github.com/ClangBuiltLinux/linux/issues/1780
    Cc: Weili Qian <qianweili@huawei.com>
    Cc: Zhou Wang <wangzhou1@hisilicon.com>
    Cc: Herbert Xu <herbert@gondor.apana.org.au>
    Cc: "David S. Miller" <davem@davemloft.net>
    Cc: linux-crypto@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Tested-by: Nathan Chancellor <nathan@kernel.org> # build
    Link: https://lore.kernel.org/r/20230106041945.never.831-kees@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68acead2f03cc6f357f2906be8c967f6150f2513
Author: Feng Tang <feng.tang@intel.com>
Date:   Tue Dec 20 16:25:12 2022 +0800

    clocksource: Suspend the watchdog temporarily when high read latency detected

    [ Upstream commit b7082cdfc464bf9231300605d03eebf943dda307 ]

    Bugs have been reported on 8 sockets x86 machines in which the TSC was
    wrongly disabled when the system is under heavy workload.

     [ 818.380354] clocksource: timekeeping watchdog on CPU336: hpet wd-wd read-back delay
of 1203520ns
     [ 818.436160] clocksource: wd-tsc-wd read-back delay of 181880ns, clock-skew test
skipped!
     [ 819.402962] clocksource: timekeeping watchdog on CPU338: hpet wd-wd read-back delay
of 324000ns
     [ 819.448036] clocksource: wd-tsc-wd read-back delay of 337240ns, clock-skew test
skipped!
     [ 819.880863] clocksource: timekeeping watchdog on CPU339: hpet read-back delay of
150280ns, attempt 3, marking unstable
     [ 819.936243] tsc: Marking TSC unstable due to clocksource watchdog
     [ 820.068173] TSC found unstable after boot, most likely due to broken BIOS. Use
'tsc=unstable'.
     [ 820.092382] sched_clock: Marking unstable (818769414384, 1195404998)
     [ 820.643627] clocksource: Checking clocksource tsc synchronization from CPU 267 to
CPUs 0,4,25,70,126,430,557,564.
     [ 821.067990] clocksource: Switched to clocksource hpet
```

```
    This can be reproduced by running memory intensive 'stream' tests,
    or some of the stress-ng subcases such as 'ioport'.

    The reason for these issues is the when system is under heavy load, the
    read latency of the clocksources can be very high.  Even lightweight TSC
    reads can show high latencies, and latencies are much worse for external
    clocksources such as HPET or the APIC PM timer.  These latencies can
    result in false-positive clocksource-unstable determinations.

    These issues were initially reported by a customer running on a production
    system, and this problem was reproduced on several generations of Xeon
    servers, especially when running the stress-ng test.  These Xeon servers
    were not production systems, but they did have the latest steppings
    and firmware.

    Given that the clocksource watchdog is a continual diagnostic check with
    frequency of twice a second, there is no need to rush it when the system
    is under heavy load.  Therefore, when high clocksource read latencies
    are detected, suspend the watchdog timer for 5 minutes.

    Signed-off-by: Feng Tang <feng.tang@intel.com>
    Acked-by: Waiman Long <longman@redhat.com>
    Cc: John Stultz <jstultz@google.com>
    Cc: Thomas Gleixner <tglx@linutronix.de>
    Cc: Stephen Boyd <sboyd@kernel.org>
    Cc: Feng Tang <feng.tang@intel.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63fb588b143968e5b55a0df69c64f5e399367393
Author: Tim Zimmermann <tim@linux4.de>
Date:   Sat Jan 7 20:25:13 2023 +0100

    thermal: intel: intel_pch: Add support for Wellsburg PCH

    [ Upstream commit 40dc1929089fc844ea06d9f8bdb6211ed4517c2e ]

    Add the PCI ID for the Wellsburg C610 series chipset PCH.

    The driver can read the temperature from the Wellsburg PCH with only
    the PCI ID added and no other modifications.

    Signed-off-by: Tim Zimmermann <tim@linux4.de>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8ff9e3f8a9252fbd61ff42cb8f48fc87418aa561
Author: Dave Thaler <dthaler@microsoft.com>
Date:   Tue Jan 24 00:12:18 2023 +0000

    bpf, docs: Fix modulo zero, division by zero, overflow, and underflow

    [ Upstream commit 0eb9d19e2201068260e439a5c96dc85f9f3722a2 ]

    Fix modulo zero, division by zero, overflow, and underflow. Also clarify how
    a negative immediate value is used in unsigned division.

    Signed-off-by: Dave Thaler <dthaler@microsoft.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230124001218.827-1-dthaler1968@googlemail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4948cdce42f6e5eb4e87a852d27ef186d1b1fb9c
Author: Mark Rutland <mark.rutland@arm.com>
Date:   Mon Jan 23 13:45:58 2023 +0000
```

```
    ACPI: Don't build ACPICA with '-Os'

    [ Upstream commit 8f9e0a52810dd83406c768972d022c37e7a18f1f ]

    The ACPICA code has been built with '-Os' since the beginning of git
    history, though there's no explanatory comment as to why.

    This is unfortunate as GCC drops the alignment specificed by
    '-falign-functions=N' when '-Os' is used, as reported in GCC bug 88345:

      https://gcc.gnu.org/bugzilla/show_bug.cgi?id=88345

    This prevents CONFIG_FUNCTION_ALIGNMENT and
    CONFIG_DEBUG_FORCE_FUNCTION_ALIGN_64B from having their expected effect
    on the ACPICA code. This is doubly unfortunate as in subsequent patches
    arm64 will depend upon CONFIG_FUNCTION_ALIGNMENT for its ftrace
    implementation.

    Drop the '-Os' flag when building the ACPICA code. With this removed,
    the code builds cleanly and works correctly in testing so far.

    I've tested this by selecting CONFIG_DEBUG_FORCE_FUNCTION_ALIGN_64B=y,
    building and booting a kernel using ACPI, and looking for misaligned
    text symbols:

    * arm64:

      Before, v6.2-rc3:
        # uname -rm
        6.2.0-rc3 aarch64
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        5009

      Before, v6.2-rc3 + fixed __cold:
        # uname -rm
        6.2.0-rc3-00001-g2a2bedf8bfa9 aarch64
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        919

      After:
        # uname -rm
        6.2.0-rc3-00002-g267bddc38572 aarch64
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        323
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | grep acpi | wc -l
        0

    * x86_64:

      Before, v6.2-rc3:
        # uname -rm
        6.2.0-rc3 x86_64
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        11537

      Before, v6.2-rc3 + fixed __cold:
        # uname -rm
        6.2.0-rc3-00001-g2a2bedf8bfa9 x86_64
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        2805

      After:
        # uname -rm
        6.2.0-rc3-00002-g267bddc38572 x86_64
```

```
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
        1357
        # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | grep acpi | wc -l
        0
```

With the patch applied, the remaining unaligned text labels are a
combination of static call trampolines and labels in assembly, which can
be dealt with in subsequent patches.

Signed-off-by: Mark Rutland <mark.rutland@arm.com>
Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Cc: Florent Revest <revest@chromium.org>
Cc: Len Brown <lenb@kernel.org>
Cc: Masami Hiramatsu <mhiramat@kernel.org>
Cc: Peter Zijlstra <peterz@infradead.org>
Cc: Robert Moore <robert.moore@intel.com>
Cc: Steven Rostedt <rostedt@goodmis.org>
Cc: Will Deacon <will@kernel.org>
Cc: linux-acpi@vger.kernel.org
Link: https://lore.kernel.org/r/20230123134603.1064407-4-mark.rutland@arm.com
Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76751e14d179119739c548cc95ffbaa6ff52785d
Author: Mark Rutland <mark.rutland@arm.com>
Date: Mon Jan 23 13:45:57 2023 +0000

Compiler attributes: GCC cold function alignment workarounds

[Upstream commit c27cd083cfb9d392f304657ed00fcde1136704e7]

Contemporary versions of GCC (e.g. GCC 12.2.0) drop the alignment
specified by '-falign-functions=N' for functions marked with the
__cold__ attribute, and potentially for callees of __cold__ functions as
these may be implicitly marked as __cold__ by the compiler. LLVM appears
to respect '-falign-functions=N' in such cases.

This has been reported to GCC in bug 88345:

  https://gcc.gnu.org/bugzilla/show_bug.cgi?id=88345

... which also covers alignment being dropped when '-Os' is used, which
will be addressed in a separate patch.

Currently, use of '-falign-functions=N' is limited to
CONFIG_FUNCTION_ALIGNMENT, which is largely used for performance and/or
analysis reasons (e.g. with CONFIG_DEBUG_FORCE_FUNCTION_ALIGN_64B), but
isn't necessary for correct functionality. However, this dropped
alignment isn't great for the performance and/or analysis cases.

Subsequent patches will use CONFIG_FUNCTION_ALIGNMENT as part of arm64's
ftrace implementation, which will require all instrumented functions to
be aligned to at least 8-bytes.

This patch works around the dropped alignment by avoiding the use of the
__cold__ attribute when CONFIG_FUNCTION_ALIGNMENT is non-zero, and by
specifically aligning abort(), which GCC implicitly marks as __cold__.
As the __cold macro is now dependent upon config options (which is
against the policy described at the top of compiler_attributes.h), it is
moved into compiler_types.h.

I've tested this by building and booting a kernel configured with
defconfig + CONFIG_EXPERT=y + CONFIG_DEBUG_FORCE_FUNCTION_ALIGN_64B=y,
and looking for misaligned text symbols in /proc/kallsyms:

```
      * arm64:

        Before:
          # uname -rm
          6.2.0-rc3 aarch64
          # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
          5009

        After:
          # uname -rm
          6.2.0-rc3-00001-g2a2bedf8bfa9 aarch64
          # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
          919

      * x86_64:

        Before:
          # uname -rm
          6.2.0-rc3 x86_64
          # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
          11537

        After:
          # uname -rm
          6.2.0-rc3-00001-g2a2bedf8bfa9 x86_64
          # grep ' [Tt] ' /proc/kallsyms | grep -iv '[048c]0 [Tt] ' | wc -l
          2805

      There's clearly a substantial reduction in the number of misaligned
      symbols. From manual inspection, the remaining unaligned text labels are
      a combination of ACPICA functions (due to the use of '-Os'), static call
      trampolines, and non-function labels in assembly, which will be dealt
      with in subsequent patches.

      Signed-off-by: Mark Rutland <mark.rutland@arm.com>
      Cc: Florent Revest <revest@chromium.org>
      Cc: Masami Hiramatsu <mhiramat@kernel.org>
      Cc: Peter Zijlstra <peterz@infradead.org>
      Cc: Steven Rostedt <rostedt@goodmis.org>
      Cc: Will Deacon <will@kernel.org>
      Cc: Miguel Ojeda <ojeda@kernel.org>
      Cc: Nick Desaulniers <ndesaulniers@google.com>
      Link: https://lore.kernel.org/r/20230123134603.1064407-3-mark.rutland@arm.com
      Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 99d0ce822433c3d4d958f6b46df1dcfdad7aa6ee
  Author: Jesse Brandeburg <jesse.brandeburg@intel.com>
  Date:   Tue Dec 13 16:01:31 2022 -0800

      ice: add missing checks for PF vsi type

      [ Upstream commit 6a8d013e904ad9a66706fcc926ec9993bed7d190 ]

      There were a few places we had missed checking the VSI type to make sure
      it was definitely a PF VSI, before calling setup functions intended only
      for the PF VSI.

      This doesn't fix any explicit bugs but cleans up the code in a few
      places and removes one explicit != vsi->type check that can be
      superseded by this code (it's a super set)

      Signed-off-by: Jesse Brandeburg <jesse.brandeburg@intel.com>
      Tested-by: Gurucharan G <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 5bcb0a51c89e2950c80f014ba314e5de58d79920
Author: Siddaraju DH <siddaraju.dh@intel.com>
Date:   Tue Nov 15 15:11:35 2022 +0530

    ice: restrict PTP HW clock freq adjustments to 100, 000, 000 PPB

    [ Upstream commit 8aa4318c3a122b8670bc09af142de3872ca63b88 ]

    The PHY provides only 39b timestamp. With current timing
    implementation, we discard lower 7b, leaving 32b timestamp.
    The driver reconstructs the full 64b timestamp by correlating the
    32b timestamp with cached_time for performance. The reconstruction
    algorithm does both forward & backward interpolation.

    The 32b timeval has overflow duration of 2^32 counts ~= 4.23 second.
    Due to interpolation in both direction, its now ~= 2.125 second
    IIRC, going with at least half a duration, the cached_time is updated
    with periodic thread of 1 second (worst-case) periodicity.

    But the 1 second periodicity is based on System-timer.
    With PPB adjustments, if the 1588 timers increments at say
    double the rate, (2s in-place of 1s), the Nyquist rate/half duration
    sampling/update of cached_time with 1 second periodic thread will
    lead to incorrect interpolations.

    Hence we should restrict the PPB adjustments to at least half duration
    of cached_time update which translates to 500,000,000 PPB.

    Since the periodicity of the cached-time system thread can vary,
    it is good to have some buffer time and considering practicality of
    PPB adjustments, limiting the max_adj to 100,000,000.

    Signed-off-by: Siddaraju DH <siddaraju.dh@intel.com>
    Tested-by: Gurucharan G <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 2bed027b3d9d471578b91fd5680f042011eab273
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sat Jan 14 13:11:41 2023 +0000

    inet: fix fast path in __inet_hash_connect()

    [ Upstream commit 21cbd90a6fab7123905386985e3e4a80236b8714 ]

    __inet_hash_connect() has a fast path taken if sk_head(&tb->owners) is
    equal to the sk parameter.
    sk_head() returns the hlist_entry() with respect to the sk_node field.
    However entries in the tb->owners list are inserted with respect to the
    sk_bind_node field with sk_add_bind_node().
    Thus the check would never pass and the fast path never execute.

    This fast path has never been executed or tested as this bug seems
    to be present since commit 1da177e4c3f4 ("Linux-2.6.12-rc2"), thus
    remove it to reduce code complexity.

    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Reviewed-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Link: https://lore.kernel.org/r/20230112-inet_hash_connect_bind_head-v3-1-
b591fd212b93@diag.uniroma1.it
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit d0db59e2f718d1e2f1d2a2d8092168fdd2f3add0
Author: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
Date:   Thu Dec 29 18:29:06 2022 +0900

    wifi: mt7601u: fix an integer underflow

    [ Upstream commit 803f3176c5df3b5582c27ea690f204abb60b19b9 ]

    Fix an integer underflow that leads to a null pointer dereference in
    'mt7601u_rx_skb_from_seg()'. The variable 'dma_len' in the URB packet
    could be manipulated, which could trigger an integer underflow of
    'seg_len' in 'mt7601u_rx_process_seg()'. This underflow subsequently
    causes the 'bad_frame' checks in 'mt7601u_rx_skb_from_seg()' to be
    bypassed, eventually leading to a dereference of the pointer 'p', which
    is a null pointer.

    Ensure that 'dma_len' is greater than 'min_seg_len'.

    Found by a modified version of syzkaller.

    KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f]
    CPU: 0 PID: 12 Comm: ksoftirqd/0 Tainted: G        W  O      5.14.0+
    #139
    Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS
    rel-1.12.1-0-ga5cab58e9a3f-prebuilt.qemu.org 04/01/2014
    RIP: 0010:skb_add_rx_frag+0x143/0x370
    Code: e2 07 83 c2 03 38 ca 7c 08 84 c9 0f 85 86 01 00 00 4c 8d 7d 08 44
    89 68 08 48 b8 00 00 00 00 00 fc ff df 4c 89 fa 48 c1 ea 03 <80> 3c 02
    00 0f 85 cd 01 00 00 48 8b 45 08 a8 01 0f 85 3d 01 00 00
    RSP: 0018:ffffc900000cfc90 EFLAGS: 00010202
    RAX: dffffc0000000000 RBX: ffff888115520dc0 RCX: 0000000000000000
    RDX: 0000000000000001 RSI: ffff8881118430c0 RDI: ffff8881118430f8
    RBP: 0000000000000000 R08: 0000000000000e09 R09: 0000000000000010
    R10: ffff888111843017 R11: ffffed1022308602 R12: 0000000000000000
    R13: 0000000000000e09 R14: 0000000000000010 R15: 0000000000000008
    FS:  0000000000000000(0000) GS:ffff88811a800000(0000)
    knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 000000004035af40 CR3: 00000001157f2000 CR4: 0000000000750ef0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    PKRU: 55555554
    Call Trace:
     mt7601u_rx_tasklet+0xc73/0x1270
     ? mt7601u_submit_rx_buf.isra.0+0x510/0x510
     ? tasklet_action_common.isra.0+0x79/0x2f0
     tasklet_action_common.isra.0+0x206/0x2f0
     __do_softirq+0x1b5/0x880
     ? tasklet_unlock+0x30/0x30
     run_ksoftirqd+0x26/0x50
     smpboot_thread_fn+0x34f/0x7d0
     ? smpboot_register_percpu_thread+0x370/0x370
     kthread+0x3a1/0x480
     ? set_kthread_struct+0x120/0x120
     ret_from_fork+0x1f/0x30
    Modules linked in: 88XXau(O) 88x2bu(O)
    ---[ end trace 57f34f93b4da0f9b ]---
    RIP: 0010:skb_add_rx_frag+0x143/0x370
    Code: e2 07 83 c2 03 38 ca 7c 08 84 c9 0f 85 86 01 00 00 4c 8d 7d 08 44
    89 68 08 48 b8 00 00 00 00 00 fc ff df 4c 89 fa 48 c1 ea 03 <80> 3c 02
    00 0f 85 cd 01 00 00 48 8b 45 08 a8 01 0f 85 3d 01 00 00
    RSP: 0018:ffffc900000cfc90 EFLAGS: 00010202
    RAX: dffffc0000000000 RBX: ffff888115520dc0 RCX: 0000000000000000
    RDX: 0000000000000001 RSI: ffff8881118430c0 RDI: ffff8881118430f8
```

```
      RBP: 0000000000000000 R08: 0000000000000e09 R09: 0000000000000010
      R10: ffff888111843017 R11: ffffed1022308602 R12: 0000000000000000
      R13: 0000000000000e09 R14: 0000000000000010 R15: 0000000000000008
      FS:  0000000000000000(0000) GS:ffff88811a800000(0000)
      knlGS:0000000000000000
      CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
      CR2: 000000004035af40 CR3: 00000001157f2000 CR4: 0000000000750ef0
      DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
      DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
      PKRU: 55555554

      Signed-off-by: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
      Acked-by: Jakub Kicinski <kuba@kernel.org>
      Signed-off-by: Kalle Valo <kvalo@kernel.org>
      Link: https://lore.kernel.org/r/20221229092906.2328282-1-jisoo.jang@yonsei.ac.kr
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 81ccbdffc9b040175dee8b9a21c32bc50dc0e364
  Author: Zong-Zhe Yang <kevin_yang@realtek.com>
  Date:   Fri Jan 13 17:06:31 2023 +0800

      wifi: rtw89: fix assignation of TX BD RAM table

      [ Upstream commit 7f495de6ae7d31f098970fb45a038c9f69b1bf75 ]

      TX BD's RAM table describes how HW allocates usable buffer section
      for each TX channel at fetch time. The total RAM size for TX BD is
      chip-dependent. For 8852BE, it has only half size (32) for TX channels
      of single band. Original table arrange total size (64) for dual band.
      It will overflow on 8852BE circuit and cause section conflicts between
      different TX channels.

      So, we do the changes below.
      * add another table for single band chip and export both kind of tables
      * point to the expected one in rtw89_pci_info by chip

      Signed-off-by: Zong-Zhe Yang <kevin_yang@realtek.com>
      Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
      Signed-off-by: Kalle Valo <kvalo@kernel.org>
      Link: https://lore.kernel.org/r/20230113090632.60957-4-pkshih@realtek.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c02f733024d70105f22de8dd0a1252a0350cd516
  Author: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
  Date:   Fri Dec 30 16:51:39 2022 +0900

      wifi: brcmfmac: ensure CLM version is null-terminated to prevent stack-out-of-bounds

      [ Upstream commit 660145d708be52f946a82e5b633c020f58f996de ]

      Fix a stack-out-of-bounds read in brcmfmac that occurs
      when 'buf' that is not null-terminated is passed as an argument of
      strreplace() in brcmf_c_preinit_dcmds(). This buffer is filled with
      a CLM version string by memcpy() in brcmf_fil_iovar_data_get().
      Ensure buf is null-terminated.

      Found by a modified version of syzkaller.

      [   33.004414][ T1896] brcmfmac: brcmf_c_process_clm_blob: no clm_blob available
  (err=-2), device may have limited channels available
      [   33.013486][ T1896] brcmfmac: brcmf_c_preinit_dcmds: Firmware: BCM43236/3 wl0: Nov 30
  2011 17:33:42 version 5.90.188.22
      [   33.021554][ T1896]
  ================================================================
      [   33.022379][ T1896] BUG: KASAN: stack-out-of-bounds in strreplace+0xf2/0x110
```

```
     [   33.023122][ T1896] Read of size 1 at addr ffffc90001d6efc8 by task kworker/0:2/1896
     [   33.023852][ T1896]
     [   33.024096][ T1896] CPU: 0 PID: 1896 Comm: kworker/0:2 Tainted: G           O
  5.14.0+ #132
     [   33.024927][ T1896] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-
  1.12.1-0-ga5cab58e9a3f-prebuilt.qemu.org 04/01/2014
     [   33.026065][ T1896] Workqueue: usb_hub_wq hub_event
     [   33.026581][ T1896] Call Trace:
     [   33.026896][ T1896]  dump_stack_lvl+0x57/0x7d
     [   33.027372][ T1896]  print_address_description.constprop.0.cold+0xf/0x334
     [   33.028037][ T1896]  ? strreplace+0xf2/0x110
     [   33.028403][ T1896]  ? strreplace+0xf2/0x110
     [   33.028807][ T1896]  kasan_report.cold+0x83/0xdf
     [   33.029283][ T1896]  ? strreplace+0xf2/0x110
     [   33.029666][ T1896]  strreplace+0xf2/0x110
     [   33.029966][ T1896]  brcmf_c_preinit_dcmds+0xab1/0xc40
     [   33.030351][ T1896]  ? brcmf_c_set_joinpref_default+0x100/0x100
     [   33.030787][ T1896]  ? rcu_read_lock_sched_held+0xa1/0xd0
     [   33.031223][ T1896]  ? rcu_read_lock_bh_held+0xb0/0xb0
     [   33.031661][ T1896]  ? lock_acquire+0x19d/0x4e0
     [   33.032091][ T1896]  ? find_held_lock+0x2d/0x110
     [   33.032605][ T1896]  ? brcmf_usb_deq+0x1a7/0x260
     [   33.033087][ T1896]  ? brcmf_usb_rx_fill_all+0x5a/0xf0
     [   33.033582][ T1896]  brcmf_attach+0x246/0xd40
     [   33.034022][ T1896]  ? wiphy_new_nm+0x1476/0x1d50
     [   33.034383][ T1896]  ? kmemdup+0x30/0x40
     [   33.034722][ T1896]  brcmf_usb_probe+0x12de/0x1690
     [   33.035223][ T1896]  ? brcmf_usbdev_qinit.constprop.0+0x470/0x470
     [   33.035833][ T1896]  usb_probe_interface+0x25f/0x710
     [   33.036315][ T1896]  really_probe+0x1be/0xa90
     [   33.036656][ T1896]  __driver_probe_device+0x2ab/0x460
     [   33.037026][ T1896]  ? usb_match_id.part.0+0x88/0xc0
     [   33.037383][ T1896]  driver_probe_device+0x49/0x120
     [   33.037790][ T1896]  __device_attach_driver+0x18a/0x250
     [   33.038300][ T1896]  ? driver_allows_async_probing+0x120/0x120
     [   33.038986][ T1896]  bus_for_each_drv+0x123/0x1a0
     [   33.039906][ T1896]  ? bus_rescan_devices+0x20/0x20
     [   33.041412][ T1896]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
     [   33.041861][ T1896]  ? trace_hardirqs_on+0x1c/0x120
     [   33.042330][ T1896]  __device_attach+0x207/0x330
     [   33.042664][ T1896]  ? device_bind_driver+0xb0/0xb0
     [   33.043026][ T1896]  ? kobject_uevent_env+0x230/0x12c0
     [   33.043515][ T1896]  bus_probe_device+0x1a2/0x260
     [   33.043914][ T1896]  device_add+0xa61/0x1ce0
     [   33.044227][ T1896]  ? __mutex_unlock_slowpath+0xe7/0x660
     [   33.044891][ T1896]  ? __fw_devlink_link_to_suppliers+0x550/0x550
     [   33.045531][ T1896]  usb_set_configuration+0x984/0x1770
     [   33.046051][ T1896]  ? kernfs_create_link+0x175/0x230
     [   33.046548][ T1896]  usb_generic_driver_probe+0x69/0x90
     [   33.046931][ T1896]  usb_probe_device+0x9c/0x220
     [   33.047434][ T1896]  really_probe+0x1be/0xa90
     [   33.047760][ T1896]  __driver_probe_device+0x2ab/0x460
     [   33.048134][ T1896]  driver_probe_device+0x49/0x120
     [   33.048516][ T1896]  __device_attach_driver+0x18a/0x250
     [   33.048910][ T1896]  ? driver_allows_async_probing+0x120/0x120
     [   33.049437][ T1896]  bus_for_each_drv+0x123/0x1a0
     [   33.049814][ T1896]  ? bus_rescan_devices+0x20/0x20
     [   33.050164][ T1896]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
     [   33.050579][ T1896]  ? trace_hardirqs_on+0x1c/0x120
     [   33.050936][ T1896]  __device_attach+0x207/0x330
     [   33.051399][ T1896]  ? device_bind_driver+0xb0/0xb0
     [   33.051888][ T1896]  ? kobject_uevent_env+0x230/0x12c0
     [   33.052314][ T1896]  bus_probe_device+0x1a2/0x260
     [   33.052688][ T1896]  device_add+0xa61/0x1ce0
     [   33.053121][ T1896]  ? __fw_devlink_link_to_suppliers+0x550/0x550
```

```
     [   33.053568][ T1896]  usb_new_device.cold+0x463/0xf66
     [   33.053953][ T1896]  ? hub_disconnect+0x400/0x400
     [   33.054313][ T1896]  ? rwlock_bug.part.0+0x90/0x90
     [   33.054661][ T1896]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
     [   33.055094][ T1896]  hub_event+0x10d5/0x3330
     [   33.055530][ T1896]  ? hub_port_debounce+0x280/0x280
     [   33.055934][ T1896]  ? __lock_acquire+0x1671/0x5790
     [   33.056387][ T1896]  ? wq_calc_node_cpumask+0x170/0x2a0
     [   33.056924][ T1896]  ? lock_release+0x640/0x640
     [   33.057383][ T1896]  ? rcu_read_lock_sched_held+0xa1/0xd0
     [   33.057916][ T1896]  ? rcu_read_lock_bh_held+0xb0/0xb0
     [   33.058402][ T1896]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
     [   33.059019][ T1896]  process_one_work+0x873/0x13e0
     [   33.059488][ T1896]  ? lock_release+0x640/0x640
     [   33.059932][ T1896]  ? pwq_dec_nr_in_flight+0x320/0x320
     [   33.060446][ T1896]  ? rwlock_bug.part.0+0x90/0x90
     [   33.060898][ T1896]  worker_thread+0x8b/0xd10
     [   33.061348][ T1896]  ? __kthread_parkme+0xd9/0x1d0
     [   33.061810][ T1896]  ? process_one_work+0x13e0/0x13e0
     [   33.062288][ T1896]  kthread+0x379/0x450
     [   33.062660][ T1896]  ? _raw_spin_unlock_irq+0x24/0x30
     [   33.063148][ T1896]  ? set_kthread_struct+0x100/0x100
     [   33.063606][ T1896]  ret_from_fork+0x1f/0x30
     [   33.064070][ T1896]
     [   33.064313][ T1896]
     [   33.064545][ T1896] addr ffffc90001d6efc8 is located in stack of task
  kworker/0:2/1896 at offset 512 in frame:
     [   33.065478][ T1896]  brcmf_c_preinit_dcmds+0x0/0xc40
     [   33.065973][ T1896]
     [   33.066191][ T1896] this frame has 4 objects:
     [   33.066614][ T1896]  [48, 56) 'ptr'
     [   33.066618][ T1896]  [80, 148) 'revinfo'
     [   33.066957][ T1896]  [192, 210) 'eventmask'
     [   33.067338][ T1896]  [256, 512) 'buf'
     [   33.067742][ T1896]
     [   33.068304][ T1896] Memory state around the buggy address:
     [   33.068838][ T1896]  ffffc90001d6ee80: f2 00 00 02 f2 f2 f2 f2 f2 00 00 00 00 00 00
  00
     [   33.069545][ T1896]  ffffc90001d6ef00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00
     [   33.070626][ T1896] >ffffc90001d6ef80: 00 00 00 00 00 00 00 00 f3 f3 f3 f3 f3 f3
  f3
     [   33.072052][ T1896]                                                 ^
     [   33.073043][ T1896]  ffffc90001d6f000: f3 00 00 00 00 00 00 00 00 00 00 00 00 00
  00
     [   33.074230][ T1896]  ffffc90001d6f080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00
     [   33.074914][ T1896]
  ==================================================================
     [   33.075713][ T1896] Disabling lock debugging due to kernel taint

     Reviewed-by: Arend van Spriel<arend.vanspriel@broadcom.com>
     Signed-off-by: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
     Signed-off-by: Kalle Valo <kvalo@kernel.org>
     Link: https://lore.kernel.org/r/20221230075139.56591-1-jisoo.jang@yonsei.ac.kr
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 70e37d8bf243dea2b94b89d2889598b4c53530e0
  Author: Holger Hoffstätte <holger@applied-asynchrony.com>
  Date:   Fri Jan 13 16:40:23 2023 +0100

     bpftool: Always disable stack protection for BPF objects

     [ Upstream commit 878625e1c7a10dfbb1fdaaaae2c4d2a58fbce627 ]
```

```
    When the clang toolchain has stack protection enabled in order to be
    consistent with gcc - which just happens to be the case on Gentoo -
    the bpftool build fails:

      [...]
      clang \
            -I. \
            -I/tmp/portage/dev-util/bpftool-6.0.12/work/linux-6.0/tools/include/uapi/ \
            -I/tmp/portage/dev-util/bpftool-6.0.12/work/linux-
6.0/tools/bpf/bpftool/bootstrap/libbpf/include \
            -g -O2 -Wall -target bpf -c skeleton/pid_iter.bpf.c -o pid_iter.bpf.o
      clang \
            -I. \
            -I/tmp/portage/dev-util/bpftool-6.0.12/work/linux-6.0/tools/include/uapi/ \
            -I/tmp/portage/dev-util/bpftool-6.0.12/work/linux-
6.0/tools/bpf/bpftool/bootstrap/libbpf/include \
            -g -O2 -Wall -target bpf -c skeleton/profiler.bpf.c -o profiler.bpf.o
      skeleton/profiler.bpf.c:40:14: error: A call to built-in function '__stack_chk_fail'
is not supported.
      int BPF_PROG(fentry_XXX)
                   ^
      skeleton/profiler.bpf.c:94:14: error: A call to built-in function '__stack_chk_fail'
is not supported.
      int BPF_PROG(fexit_XXX)
                   ^
      2 errors generated.
      [...]

    Since stack-protector makes no sense for the BPF bits just unconditionally
    disable it.

    Bug: https://bugs.gentoo.org/890638
    Signed-off-by: Holger Hoffstätte <holger@applied-asynchrony.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Quentin Monnet <quentin@isovalent.com>
    Link: https://lore.kernel.org/bpf/74cd9d2e-6052-312a-241e-2b514a75c92c@applied-
asynchrony.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 135da293756f7979c564b01d36a22604320ef1ce
Author: Breno Leitao <leitao@debian.org>
Date:   Mon Nov 28 07:31:48 2022 -0800

    x86/bugs: Reset speculation control settings on init

    [ Upstream commit 0125acda7d76b943ca55811df40ed6ec0ecf670f ]

    Currently, x86_spec_ctrl_base is read at boot time and speculative bits
    are set if Kconfig items are enabled. For example, IBRS is enabled if
    CONFIG_CPU_IBRS_ENTRY is configured, etc. These MSR bits are not cleared
    if the mitigations are disabled.

    This is a problem when kexec-ing a kernel that has the mitigation
    disabled from a kernel that has the mitigation enabled. In this case,
    the MSR bits are not cleared during the new kernel boot. As a result,
    this might have some performance degradation that is hard to pinpoint.

    This problem does not happen if the machine is (hard) rebooted because
    the bit will be cleared by default.

      [ bp: Massage. ]

    Suggested-by: Pawan Gupta <pawan.kumar.gupta@linux.intel.com>
    Signed-off-by: Breno Leitao <leitao@debian.org>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
```

```
    Link: https://lore.kernel.org/r/20221128153148.1129350-1-leitao@debian.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cbd2ca88035b11070b6e88d48228df6289dbbe79
Author: Jann Horn <jannh@google.com>
Date:   Thu Jan 5 14:44:03 2023 +0100

    timers: Prevent union confusion from unexpected restart_syscall()

    [ Upstream commit 9f76d59173d9d146e96c66886b671c1915a5c5e5 ]

    The nanosleep syscalls use the restart_block mechanism, with a quirk:
    The `type` and `rmtp`/`compat_rmtp` fields are set up unconditionally on
    syscall entry, while the rest of the restart_block is only set up in the
    unlikely case that the syscall is actually interrupted by a signal (or
    pseudo-signal) that doesn't have a signal handler.

    If the restart_block was set up by a previous syscall (futex(...,
    FUTEX_WAIT, ...) or poll()) and hasn't been invalidated somehow since then,
    this will clobber some of the union fields used by futex_wait_restart() and
    do_restart_poll().

    If userspace afterwards wrongly calls the restart_syscall syscall,
    futex_wait_restart()/do_restart_poll() will read struct fields that have
    been clobbered.

    This doesn't actually lead to anything particularly interesting because
    none of the union fields contain trusted kernel data, and
    futex(..., FUTEX_WAIT, ...) and poll() aren't syscalls where it makes much
    sense to apply seccomp filters to their arguments.

    So the current consequences are just of the "if userspace does bad stuff,
    it can damage itself, and that's not a problem" flavor.

    But still, it seems like a hazard for future developers, so invalidate the
    restart_block when partly setting it up in the nanosleep syscalls.

    Signed-off-by: Jann Horn <jannh@google.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lore.kernel.org/r/20230105134403.754986-1-jannh@google.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eddbff59aa516b45bea71c4de5ef0a7560d38877
Author: Yang Li <yang.lee@linux.alibaba.com>
Date:   Fri Jan 6 08:59:51 2023 +0800

    thermal: intel: Fix unsigned comparison with less than zero

    [ Upstream commit e7fcfe67f9f410736b758969477b17ea285e8e6c ]

    The return value from the call to intel_tcc_get_tjmax() is int, which can
    be a negative error code. However, the return value is being assigned to
    an u32 variable 'tj_max', so making 'tj_max' an int.

    Eliminate the following warning:
    ./drivers/thermal/intel/intel_soc_dts_iosf.c:394:5-11: WARNING: Unsigned expression
compared with zero: tj_max < 0

    Link: https://bugzilla.openanolis.cn/show_bug.cgi?id=3637
    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Signed-off-by: Yang Li <yang.lee@linux.alibaba.com>
    Acked-by: Zhang Rui <rui.zhang@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8ec481e85e05665d133d679b7d36e4adcb4fb836
Author: Kalle Valo <quic_kvalo@quicinc.com>
Date:   Thu Dec 22 19:15:59 2022 +0200

    wifi: ath11k: debugfs: fix to work with multiple PCI devices

    [ Upstream commit 323d91d4684d238f6bc3693fed93caf795378fe0 ]

    ath11k fails to load if there are multiple ath11k PCI devices with same name:

     ath11k_pci 0000:01:00.0: Hardware name qcn9074 hw1.0
     debugfs: Directory 'ath11k' with parent '/' already present!
     ath11k_pci 0000:01:00.0: failed to create ath11k debugfs
     ath11k_pci 0000:01:00.0: failed to create soc core: -17
     ath11k_pci 0000:01:00.0: failed to init core: -17
     ath11k_pci: probe of 0000:01:00.0 failed with error -17

    Fix this by creating a directory for each ath11k device using schema
    <bus>-<devname>, for example "pci-0000:06:00.0". This directory created under
    the top-level ath11k directory, for example /sys/kernel/debug/ath11k.

    The reference to the toplevel ath11k directory is not stored anymore within ath11k,
instead
    it's retrieved using debugfs_lookup(). If the directory does not exist it will
    be created. After the last directory from the ath11k directory is removed, for
    example when doing rmmod ath11k, the empty ath11k directory is left in place,
    it's a minor cosmetic issue anyway.

    Here's an example hierarchy with one WCN6855:

    ath11k
    `-- pci-0000:06:00.0
        |-- mac0
        |   |-- dfs_block_radar_events
        |   |-- dfs_simulate_radar
        |   |-- ext_rx_stats
        |   |-- ext_tx_stats
        |   |-- fw_dbglog_config
        |   |-- fw_stats
        |   |   |-- beacon_stats
        |   |   |-- pdev_stats
        |   |   `-- vdev_stats
        |   |-- htt_stats
        |   |-- htt_stats_reset
        |   |-- htt_stats_type
        |   `-- pktlog_filter
        |-- simulate_fw_crash
        `-- soc_dp_stats

    I didn't have a test setup where I could connect multiple ath11k devices to the
    same the host, so I have only tested this with one device.

    Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-03125-QCAHSPSWPL_V1_V2_SILICONZ_LITE-3.6510.9
    Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-01208-QCAHKSWPL_SILICONZ-1
    Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.5.0.1-01208-QCAHKSWPL_SILICONZ-1

    Tested-by: Robert Marko <robert.marko@sartura.hr>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20221220121231.20120-1-kvalo@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0de2180ba3599437d4f42c71eb9ed90a965875e2
Author: Zqiang <qiang1.zhang@intel.com>
Date:   Sat Dec 3 10:25:03 2022 +0800
```

rcu-tasks: Handle queue-shrink/callback-enqueue race condition

[Upstream commit a4fcfbee8f6274f9b3f9a71dd5b03e6772ce33f3]

The rcu_tasks_need_gpcb() determines whether or not: (1) There are callbacks needing another grace period, (2) There are callbacks ready to be invoked, and (3) It would be a good time to shrink back down to a single-CPU callback list. This third case is interesting because some other CPU might be adding new callbacks, which might suddenly make this a very bad time to be shrinking.

This is currently handled by requiring call_rcu_tasks_generic() to enqueue callbacks under the protection of rcu_read_lock() and requiring rcu_tasks_need_gpcb() to wait for an RCU grace period to elapse before finalizing the transition. This works well in practice.

Unfortunately, the current code assumes that a grace period whose end is detected by the poll_state_synchronize_rcu() in the second "if" condition actually ended before the earlier code counted the callbacks queued on CPUs other than CPU 0 (local variable "ncbsnz"). Given the current code, it is possible that a long-delayed call_rcu_tasks_generic() invocation will queue a callback on a non-zero CPU after these CPUs have had their callbacks counted and zero has been stored to ncbsnz. Such a callback would trigger the WARN_ON_ONCE() in the second "if" statement.

To see this, consider the following sequence of events:

o CPU 0 invokes rcu_tasks_one_gp(), and counts fewer than rcu_task_collapse_lim callbacks. It sees at least one callback queued on some other CPU, thus setting ncbsnz to a non-zero value.

o CPU 1 invokes call_rcu_tasks_generic() and loads 42 from ->percpu_enqueue_lim. It therefore decides to enqueue its callback onto CPU 1's callback list, but is delayed.

o CPU 0 sees the rcu_task_cb_adjust is non-zero and that the number of callbacks does not exceed rcu_task_collapse_lim. It therefore checks percpu_enqueue_lim, and sees that its value is greater than the value one. CPU 0 therefore starts the shift back to a single callback list. It sets ->percpu_enqueue_lim to 1, but CPU 1 has already read the old value of 42. It also gets a grace-period state value from get_state_synchronize_rcu().

o CPU 0 sees that ncbsnz is non-zero in its second "if" statement, so it declines to finalize the shrink operation.

o CPU 0 again invokes rcu_tasks_one_gp(), and counts fewer than rcu_task_collapse_lim callbacks. It also sees that there are no callback queued on any other CPU, and thus sets ncbsnz to zero.

o CPU 1 resumes execution and enqueues its callback onto its own list. This invalidates the value of ncbsnz.

o CPU 0 sees the rcu_task_cb_adjust is non-zero and that the number of callbacks does not exceed rcu_task_collapse_lim. It therefore checks percpu_enqueue_lim, but sees that its value is already unity. It therefore does not get a new grace-period state value.

o CPU 0 sees that rcu_task_cb_adjust is non-zero, ncbsnz is zero, and that poll_state_synchronize_rcu() says that the grace period has completed. it therefore finalizes the shrink operation, setting ->percpu_dequeue_lim to the value one.

o CPU 0 does a debug check, scanning the other CPUs' callback lists.

```
             It sees that CPU 1's list has a callback, so it (rightly)
             triggers the WARN_ON_ONCE().  After all, the new value of
             ->percpu_dequeue_lim says to not bother looking at CPU 1's
             callback list, which means that this callback will never be
             invoked.  This can result in hangs and maybe even OOMs.

     Based on long experience with rcutorture, this is an extremely
     low-probability race condition, but it really can happen, especially in
     preemptible kernels or within guest OSes.

     This commit therefore checks for completion of the grace period
     before counting callbacks.  With this change, in the above failure
     scenario CPU 0 would know not to prematurely end the shrink operation
     because the grace period would not have completed before the count
     operation started.

     [ paulmck: Adjust grace-period end rather than adding RCU reader. ]
     [ paulmck: Avoid spurious WARN_ON_ONCE() with ->percpu_dequeue_lim check. ]

     Signed-off-by: Zqiang <qiang1.zhang@intel.com>
     Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 5c12121eecd7946e68d499309ecc65c5601ab441
  Author: Zqiang <qiang1.zhang@intel.com>
  Date:   Thu Dec 1 07:45:33 2022 +0800

     rcu-tasks: Make rude RCU-Tasks work well with CPU hotplug

     [ Upstream commit ea5c8987fef20a8cca07e428aa28bc64649c5104 ]

     The synchronize_rcu_tasks_rude() function invokes rcu_tasks_rude_wait_gp()
     to wait one rude RCU-tasks grace period.  The rcu_tasks_rude_wait_gp()
     function in turn checks if there is only a single online CPU.  If so, it
     will immediately return, because a call to synchronize_rcu_tasks_rude()
     is by definition a grace period on a single-CPU system.  (We could
     have blocked!)

     Unfortunately, this check uses num_online_cpus() without synchronization,
     which can result in too-short grace periods.  To see this, consider the
     following scenario:

             CPU0                                    CPU1 (going offline)
                                                migration/1 task:
                                             cpu_stopper_thread
                                              -> take_cpu_down
                                                 -> _cpu_disable
                                                  (dec __num_online_cpus)
                                                 ->cpuhp_invoke_callback
                                                       preempt_disable
                                                       access old_data0
                task1
      del old_data0                                    .....
      synchronize_rcu_tasks_rude()
      task1 schedule out
      ....
      task2 schedule in
      rcu_tasks_rude_wait_gp()
          ->__num_online_cpus == 1
            ->return
      ....
      task1 schedule in
      ->free old_data0
                                                       preempt_enable
```

```
    When CPU1 decrements __num_online_cpus, its value becomes 1.  However,
    CPU1 has not finished going offline, and will take one last trip through
    the scheduler and the idle loop before it actually stops executing
    instructions.  Because synchronize_rcu_tasks_rude() is mostly used for
    tracing, and because both the scheduler and the idle loop can be traced,
    this means that CPU0's prematurely ended grace period might disrupt the
    tracing on CPU1.  Given that this disruption might include CPU1 executing
    instructions in memory that was just now freed (and maybe reallocated),
    this is a matter of some concern.

    This commit therefore removes that problematic single-CPU check from the
    rcu_tasks_rude_wait_gp() function.  This dispenses with the single-CPU
    optimization, but there is no evidence indicating that this optimization
    is important.  In addition, synchronize_rcu_tasks_generic() contains a
    similar optimization (albeit only for early boot), which also splats.
    (As in exactly why are you invoking synchronize_rcu_tasks_rude() so
    early in boot, anyway???)

    It is OK for the synchronize_rcu_tasks_rude() function's check to be
    unsynchronized because the only times that this check can evaluate to
    true is when there is only a single CPU running with preemption
    disabled.

    While in the area, this commit also fixes a minor bug in which a
    call to synchronize_rcu_tasks_rude() would instead be attributed to
    synchronize_rcu_tasks().

    [ paulmck: Add "synchronize_" prefix and "()" suffix. ]

    Signed-off-by: Zqiang <qiang1.zhang@intel.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7c0bc03fa44942fe0fdc5ac52cda6e11529c0ea
Author: Pingfan Liu <kernelfans@gmail.com>
Date:   Mon Oct 31 09:52:37 2022 +0800

    srcu: Delegate work to the boot cpu if using SRCU_SIZE_SMALL

    [ Upstream commit 7f24626d6dd844bfc6d1f492d214d29c86d02550 ]

    Commit 994f706872e6 ("srcu: Make Tree SRCU able to operate without
    snp_node array") assumes that cpu 0 is always online.  However, there
    really are situations when some other CPU is the boot CPU, for example,
    when booting a kdump kernel with the maxcpus=1 boot parameter.

    On PowerPC, the kdump kernel can hang as follows:
    ...
    [    1.740036] systemd[1]: Hostname set to <xyz.com>
    [  243.686240] INFO: task systemd:1 blocked for more than 122 seconds.
    [  243.686264]       Not tainted 6.1.0-rc1 #1
    [  243.686272] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
    [  243.686281] task:systemd         state:D stack:0     pid:1     ppid:0
flags:0x00042000
    [  243.686296] Call Trace:
    [  243.686301] [c000000016657640] [c000000016657670] 0xc000000016657670 (unreliable)
    [  243.686317] [c000000016657830] [c00000001001dec0] __switch_to+0x130/0x220
    [  243.686333] [c000000016657890] [c000000010f607b8] __schedule+0x1f8/0x580
    [  243.686347] [c000000016657940] [c000000010f60bb4] schedule+0x74/0x140
    [  243.686361] [c0000000166579b0] [c000000010f699b8] schedule_timeout+0x168/0x1c0
    [  243.686374] [c000000016657a80] [c000000010f61de8] __wait_for_common+0x148/0x360
    [  243.686387] [c000000016657b20] [c000000010176bb0] __flush_work.isra.0+0x1c0/0x3d0
    [  243.686401] [c000000016657bb0] [c0000000105f2768]
fsnotify_wait_marks_destroyed+0x28/0x40
    [  243.686415] [c000000016657bd0] [c0000000105f21b8] fsnotify_destroy_group+0x68/0x160
```

```
     [  243.686428] [c000000016657c40] [c0000000105f6500] inotify_release+0x30/0xa0
     [  243.686440] [c000000016657cb0] [c0000000105751a8] __fput+0xc8/0x350
     [  243.686452] [c000000016657d00] [c00000001017d524] task_work_run+0xe4/0x170
     [  243.686464] [c000000016657d50] [c000000010020e94] do_notify_resume+0x134/0x140
     [  243.686478] [c000000016657d80] [c00000001002eb18]
 interrupt_exit_user_prepare_main+0x198/0x270
     [  243.686493] [c000000016657de0] [c00000001002ec60] syscall_exit_prepare+0x70/0x180
     [  243.686505] [c000000016657e10] [c00000001000bf7c]
 system_call_vectored_common+0xfc/0x280
     [  243.686520] --- interrupt: 3000 at 0x7fffa47d5ba4
     [  243.686528] NIP:  00007fffa47d5ba4 LR: 0000000000000000 CTR: 0000000000000000
     [  243.686538] REGS: c000000016657e80 TRAP: 3000   Not tainted  (6.1.0-rc1)
     [  243.686548] MSR:  800000000000d033 <SF,EE,PR,ME,IR,DR,RI,LE>  CR: 42044440  XER:
 00000000
     [  243.686572] IRQMASK: 0
     [  243.686572] GPR00: 0000000000000006 00007ffffa606710 00007fffa48e7200
 0000000000000000
     [  243.686572] GPR04: 0000000000000002 000000000000000a 0000000000000000
 0000000000000001
     [  243.686572] GPR08: 000001000c172dd0 0000000000000000 0000000000000000
 0000000000000000
     [  243.686572] GPR12: 0000000000000000 00007fffa4ff4bc0 0000000000000000
 0000000000000000
     [  243.686572] GPR16: 0000000000000000 0000000000000000 0000000000000000
 0000000000000000
     [  243.686572] GPR20: 0000000132dfdc50 000000000000000e 0000000000189375
 0000000000000000
     [  243.686572] GPR24: 00007ffffa606ae0 0000000000000005 000001000c185490
 000001000c172570
     [  243.686572] GPR28: 000001000c172990 000001000c184850 000001000c172e00
 00007fffa4fedd98
     [  243.686683] NIP [00007fffa47d5ba4] 0x7fffa47d5ba4
     [  243.686691] LR [0000000000000000] 0x0
     [  243.686698] --- interrupt: 3000
     [  243.686708] INFO: task kworker/u16:1:24 blocked for more than 122 seconds.
     [  243.686717]       Not tainted 6.1.0-rc1 #1
     [  243.686724] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
     [  243.686733] task:kworker/u16:1   state:D stack:0     pid:24    ppid:2
 flags:0x00000800
     [  243.686747] Workqueue: events_unbound fsnotify_mark_destroy_workfn
     [  243.686758] Call Trace:
     [  243.686762] [c0000000166736e0] [c00000004fd91000] 0xc00000004fd91000 (unreliable)
     [  243.686775] [c0000000166738d0] [c00000001001dec0] __switch_to+0x130/0x220
     [  243.686788] [c000000016673930] [c000000010f607b8] __schedule+0x1f8/0x580
     [  243.686801] [c0000000166739e0] [c000000010f60bb4] schedule+0x74/0x140
     [  243.686814] [c000000016673a50] [c000000010f699b8] schedule_timeout+0x168/0x1c0
     [  243.686827] [c000000016673b20] [c000000010f61de8] __wait_for_common+0x148/0x360
     [  243.686840] [c000000016673bc0] [c000000010210840] __synchronize_srcu.part.0+0xa0/0xe0
     [  243.686855] [c000000016673c30] [c0000000105f2c64]
 fsnotify_mark_destroy_workfn+0xc4/0x1a0
     [  243.686868] [c000000016673ca0] [c000000010174ea8] process_one_work+0x2a8/0x570
     [  243.686882] [c000000016673d40] [c000000010175208] worker_thread+0x98/0x5e0
     [  243.686895] [c000000016673dc0] [c0000000101828d4] kthread+0x124/0x130
     [  243.686908] [c000000016673e10] [c00000001000cd40] ret_from_kernel_thread+0x5c/0x64
     [  366.566274] INFO: task systemd:1 blocked for more than 245 seconds.
     [  366.566298]       Not tainted 6.1.0-rc1 #1
     [  366.566305] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
     [  366.566314] task:systemd         state:D stack:0     pid:1     ppid:0
 flags:0x00042000
     [  366.566329] Call Trace:
     ...

    The above splat occurs because PowerPC really does use maxcpus=1
    instead of nr_cpus=1 in the kernel command line.  Consequently, the
    (quite possibly non-zero) kdump CPU is the only online CPU in the kdump
```

```
    kernel.  SRCU unconditionally queues a sdp->work on cpu 0, for which no
    worker thread has been created, so sdp->work will be never executed and
    __synchronize_srcu() will never be completed.

    This commit therefore replaces CPU ID 0 with get_boot_cpu_id() in key
    places in Tree SRCU.  Since the CPU indicated by get_boot_cpu_id()
    is guaranteed to be online, this avoids the above splat.

    Signed-off-by: Pingfan Liu <kernelfans@gmail.com>
    Cc: "Paul E. McKenney" <paulmck@kernel.org>
    Cc: Lai Jiangshan <jiangshanlai@gmail.com>
    Cc: Josh Triplett <josh@joshtriplett.org>
    Cc: Steven Rostedt <rostedt@goodmis.org>
    Cc: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    To: rcu@vger.kernel.org
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f713285cb58db821e7a52d2821d8993c702061c1
Author: Paul E. McKenney <paulmck@kernel.org>
Date:   Fri Dec 16 15:55:48 2022 -0800

    rcu: Suppress smp_processor_id() complaint in synchronize_rcu_expedited_wait()

    [ Upstream commit 2d7f00b2f01301d6e41fd4a28030dab0442265be ]

    The normal grace period's RCU CPU stall warnings are invoked from the
    scheduling-clock interrupt handler, and can thus invoke smp_processor_id()
    with impunity, which allows them to directly invoke dump_cpu_task().
    In contrast, the expedited grace period's RCU CPU stall warnings are
    invoked from process context, which causes the dump_cpu_task() function's
    calls to smp_processor_id() to complain bitterly in debug kernels.

    This commit therefore causes synchronize_rcu_expedited_wait() to disable
    preemption around its call to dump_cpu_task().

    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 639bdc733fc40d263cca30a262193ed893e9e5c3
Author: Paul E. McKenney <paulmck@kernel.org>
Date:   Wed Dec 14 11:41:44 2022 -0800

    rcu: Make RCU_LOCKDEP_WARN() avoid early lockdep checks

    [ Upstream commit 0cae5ded535c3a80aed94f119bbd4ee3ae284a65 ]

    Currently, RCU_LOCKDEP_WARN() checks the condition before checking
    to see if lockdep is still enabled.  This is necessary to avoid the
    false-positive splats fixed by commit 3066820034b5dd ("rcu: Reject
    RCU_LOCKDEP_WARN() false positives").  However, the current state can
    result in false-positive splats during early boot before lockdep is fully
    initialized.  This commit therefore checks debug_lockdep_rcu_enabled()
    both before and after checking the condition, thus avoiding both sets
    of false-positive error reports.

    Reported-by: Steven Rostedt <rostedt@goodmis.org>
    Reported-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Reported-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Reviewed-by: Mathieu Desnoyers <mathieu.desnoyers@efficios.com>
    Cc: Boqun Feng <boqun.feng@gmail.com>
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: Thomas Gleixner <tglx@linutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit ba166e0ebdde3dfa833f0a3edaf2b2934d4a87f7
Author: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
Date:   Tue Nov 15 13:34:58 2022 +0900

    wifi: brcmfmac: Fix potential stack-out-of-bounds in brcmf_c_preinit_dcmds()

    [ Upstream commit 0a06cadcc2a0044e4a117cc0e61436fc3a0dad69 ]

    This patch fixes a stack-out-of-bounds read in brcmfmac that occurs
    when 'buf' that is not null-terminated is passed as an argument of
    strsep() in brcmf_c_preinit_dcmds(). This buffer is filled with a firmware
    version string by memcpy() in brcmf_fil_iovar_data_get().
    The patch ensures buf is null-terminated.

    Found by a modified version of syzkaller.

    [   47.569679][ T1897] brcmfmac: brcmf_fw_alloc_request: using brcm/brcmfmac43236b for
chip BCM43236/3
    [   47.582839][ T1897] brcmfmac: brcmf_c_process_clm_blob: no clm_blob available
(err=-2), device may have limited channels available
    [   47.601565][ T1897]
================================================================
    [   47.602574][ T1897] BUG: KASAN: stack-out-of-bounds in strsep+0x1b2/0x1f0
    [   47.603447][ T1897] Read of size 1 at addr ffffc90001f6f000 by task kworker/0:2/1897
    [   47.604336][ T1897]
    [   47.604621][ T1897] CPU: 0 PID: 1897 Comm: kworker/0:2 Tainted: G           O
5.14.0+ #131
    [   47.605617][ T1897] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-
1.12.1-0-ga5cab58e9a3f-prebuilt.qemu.org 04/01/2014
    [   47.606907][ T1897] Workqueue: usb_hub_wq hub_event
    [   47.607453][ T1897] Call Trace:
    [   47.607801][ T1897]  dump_stack_lvl+0x8e/0xd1
    [   47.608295][ T1897]  print_address_description.constprop.0.cold+0xf/0x334
    [   47.609009][ T1897]  ? strsep+0x1b2/0x1f0
    [   47.609434][ T1897]  ? strsep+0x1b2/0x1f0
    [   47.609863][ T1897]  kasan_report.cold+0x83/0xdf
    [   47.610366][ T1897]  ? strsep+0x1b2/0x1f0
    [   47.610882][ T1897]  strsep+0x1b2/0x1f0
    [   47.611300][ T1897]  ? brcmf_fil_iovar_data_get+0x3a/0xf0
    [   47.611883][ T1897]  brcmf_c_preinit_dcmds+0x995/0xc40
    [   47.612434][ T1897]  ? brcmf_c_set_joinpref_default+0x100/0x100
    [   47.613078][ T1897]  ? rcu_read_lock_sched_held+0xa1/0xd0
    [   47.613662][ T1897]  ? rcu_read_lock_bh_held+0xb0/0xb0
    [   47.614208][ T1897]  ? lock_acquire+0x19d/0x4e0
    [   47.614704][ T1897]  ? find_held_lock+0x2d/0x110
    [   47.615236][ T1897]  ? brcmf_usb_deq+0x1a7/0x260
    [   47.615741][ T1897]  ? brcmf_usb_rx_fill_all+0x5a/0xf0
    [   47.616288][ T1897]  brcmf_attach+0x246/0xd40
    [   47.616758][ T1897]  ? wiphy_new_nm+0x1703/0x1dd0
    [   47.617280][ T1897]  ? kmemdup+0x43/0x50
    [   47.617720][ T1897]  brcmf_usb_probe+0x12de/0x1690
    [   47.618244][ T1897]  ? brcmf_usbdev_qinit.constprop.0+0x470/0x470
    [   47.618901][ T1897]  usb_probe_interface+0x2aa/0x760
    [   47.619429][ T1897]  ? usb_probe_device+0x250/0x250
    [   47.619950][ T1897]  really_probe+0x205/0xb70
    [   47.620435][ T1897]  ? driver_allows_async_probing+0x130/0x130
    [   47.621048][ T1897]  __driver_probe_device+0x311/0x4b0
    [   47.621595][ T1897]  ? driver_allows_async_probing+0x130/0x130
    [   47.622209][ T1897]  driver_probe_device+0x4e/0x150
    [   47.622739][ T1897]  __device_attach_driver+0x1cc/0x2a0
    [   47.623287][ T1897]  bus_for_each_drv+0x156/0x1d0
    [   47.623796][ T1897]  ? bus_rescan_devices+0x30/0x30
    [   47.624309][ T1897]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
    [   47.624907][ T1897]  ? trace_hardirqs_on+0x46/0x160
```

```
[   47.625437][ T1897]  __device_attach+0x23f/0x3a0
[   47.625924][ T1897]  ? device_bind_driver+0xd0/0xd0
[   47.626433][ T1897]  ? kobject_uevent_env+0x287/0x14b0
[   47.627057][ T1897]  bus_probe_device+0x1da/0x290
[   47.627557][ T1897]  device_add+0xb7b/0x1eb0
[   47.628027][ T1897]  ? wait_for_completion+0x290/0x290
[   47.628593][ T1897]  ? __fw_devlink_link_to_suppliers+0x5a0/0x5a0
[   47.629249][ T1897]  usb_set_configuration+0xf59/0x16f0
[   47.629829][ T1897]  usb_generic_driver_probe+0x82/0xa0
[   47.630385][ T1897]  usb_probe_device+0xbb/0x250
[   47.630927][ T1897]  ? usb_suspend+0x590/0x590
[   47.631397][ T1897]  really_probe+0x205/0xb70
[   47.631855][ T1897]  ? driver_allows_async_probing+0x130/0x130
[   47.632469][ T1897]  __driver_probe_device+0x311/0x4b0
[   47.633002][ T1897]  ? usb_generic_driver_match+0x75/0x90
[   47.633573][ T1897]  ? driver_allows_async_probing+0x130/0x130
[   47.634170][ T1897]  driver_probe_device+0x4e/0x150
[   47.634703][ T1897]  __device_attach_driver+0x1cc/0x2a0
[   47.635248][ T1897]  bus_for_each_drv+0x156/0x1d0
[   47.635748][ T1897]  ? bus_rescan_devices+0x30/0x30
[   47.636271][ T1897]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
[   47.636881][ T1897]  ? trace_hardirqs_on+0x46/0x160
[   47.637396][ T1897]  __device_attach+0x23f/0x3a0
[   47.637904][ T1897]  ? device_bind_driver+0xd0/0xd0
[   47.638426][ T1897]  ? kobject_uevent_env+0x287/0x14b0
[   47.638985][ T1897]  bus_probe_device+0x1da/0x290
[   47.639512][ T1897]  device_add+0xb7b/0x1eb0
[   47.639977][ T1897]  ? __fw_devlink_link_to_suppliers+0x5a0/0x5a0
[   47.640612][ T1897]  ? kfree+0x14a/0x6b0
[   47.641055][ T1897]  ? __usb_get_extra_descriptor+0x116/0x160
[   47.641679][ T1897]  usb_new_device.cold+0x49c/0x1029
[   47.642245][ T1897]  ? hub_disconnect+0x450/0x450
[   47.642756][ T1897]  ? rwlock_bug.part.0+0x90/0x90
[   47.643273][ T1897]  ? _raw_spin_unlock_irq+0x24/0x30
[   47.643822][ T1897]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
[   47.644445][ T1897]  hub_event+0x1c98/0x3950
[   47.644939][ T1897]  ? hub_port_debounce+0x2e0/0x2e0
[   47.645467][ T1897]  ? check_irq_usage+0x861/0xf20
[   47.645975][ T1897]  ? drain_workqueue+0x280/0x360
[   47.646506][ T1897]  ? lock_release+0x640/0x640
[   47.646994][ T1897]  ? rcu_read_lock_sched_held+0xa1/0xd0
[   47.647572][ T1897]  ? rcu_read_lock_bh_held+0xb0/0xb0
[   47.648111][ T1897]  ? lockdep_hardirqs_on_prepare+0x273/0x3e0
[   47.648735][ T1897]  process_one_work+0x92b/0x1460
[   47.649262][ T1897]  ? pwq_dec_nr_in_flight+0x330/0x330
[   47.649816][ T1897]  ? rwlock_bug.part.0+0x90/0x90
[   47.650336][ T1897]  worker_thread+0x95/0xe00
[   47.650830][ T1897]  ? __kthread_parkme+0x115/0x1e0
[   47.651361][ T1897]  ? process_one_work+0x1460/0x1460
[   47.651904][ T1897]  kthread+0x3a1/0x480
[   47.652329][ T1897]  ? set_kthread_struct+0x120/0x120
[   47.652878][ T1897]  ret_from_fork+0x1f/0x30
[   47.653370][ T1897]
[   47.653608][ T1897]
[   47.653848][ T1897] addr ffffc90001f6f000 is located in stack of task
 kworker/0:2/1897 at offset 512 in frame:
[   47.654891][ T1897]  brcmf_c_preinit_dcmds+0x0/0xc40
[   47.655442][ T1897]
[   47.655690][ T1897] this frame has 4 objects:
[   47.656151][ T1897]  [48, 56) 'ptr'
[   47.656159][ T1897]  [80, 148) 'revinfo'
[   47.656534][ T1897]  [192, 210) 'eventmask'
[   47.656953][ T1897]  [256, 512) 'buf'
[   47.657410][ T1897]
[   47.658035][ T1897] Memory state around the buggy address:
```

```
    [   47.658743][ T1897]  ffffc90001f6ef00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 00
    [   47.659577][ T1897]  ffffc90001f6ef80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 00
    [   47.660394][ T1897] >ffffc90001f6f000: f3 f3 f3 f3 f3 f3 f3 f3 00 00 00 00 00 00 00 00
 00
    [   47.661199][ T1897]                       ^
    [   47.661625][ T1897]  ffffc90001f6f080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 00
    [   47.662455][ T1897]  ffffc90001f6f100: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 f1
 f1
    [   47.663318][ T1897]
 ==================================================================
    [   47.664147][ T1897] Disabling lock debugging due to kernel taint

    Reported-by: Dokyung Song <dokyungs@yonsei.ac.kr>
    Reported-by: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
    Reported-by: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
    Signed-off-by: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221115043458.37562-1-jisoo.jang@yonsei.ac.kr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 9089c3080a98f1452335e08b8014a28003a211ce
 Author: Nagarajan Maran <quic_nmaran@quicinc.com>
 Date:   Tue Nov 29 19:55:32 2022 +0530

    wifi: ath11k: fix monitor mode bringup crash

    [ Upstream commit 950b43f8bd8a4d476d2da6d2a083a89bcd3c90d7 ]

    When the interface is brought up in monitor mode, it leads
    to NULL pointer dereference crash. This crash happens when
    the packet type is extracted for a SKB. This extraction
    which is present in the received msdu delivery path,is
    not needed for the monitor ring packets since they are
    all RAW packets. Hence appending the flags with
    "RX_FLAG_ONLY_MONITOR" to skip that extraction.

    Observed calltrace:

    Unable to handle kernel NULL pointer dereference at virtual address
    0000000000000064
    Mem abort info:
      ESR = 0x0000000096000004
      EC = 0x25: DABT (current EL), IL = 32 bits
      SET = 0, FnV = 0
      EA = 0, S1PTW = 0
      FSC = 0x04: level 0 translation fault
    Data abort info:
      ISV = 0, ISS = 0x00000004
      CM = 0, WnR = 0
    user pgtable: 4k pages, 48-bit VAs, pgdp=0000000048517000
    [0000000000000064] pgd=0000000000000000, p4d=0000000000000000
    Internal error: Oops: 0000000096000004 [#1] PREEMPT SMP
    Modules linked in: ath11k_pci ath11k qmi_helpers
    CPU: 2 PID: 1781 Comm: napi/-271 Not tainted
    6.1.0-rc5-wt-ath-656295-gef907406320c-dirty #6
    Hardware name: Qualcomm Technologies, Inc. IPQ8074/AP-HK10-C2 (DT)
    pstate: 60000005 (nZCv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
    pc : ath11k_hw_qcn9074_rx_desc_get_decap_type+0x34/0x60 [ath11k]
    lr : ath11k_hw_qcn9074_rx_desc_get_decap_type+0x5c/0x60 [ath11k]
    sp : ffff80000ef5bb10
    x29: ffff80000ef5bb10 x28: 0000000000000000 x27: ffff000007baafa0
    x26: ffff000014a91ed0 x25: 0000000000000000 x24: 0000000000000000
```

```
      x23: ffff800002b77378 x22: ffff000014a91ec0 x21: ffff000006c8d600
      x20: 0000000000000000 x19: ffff800002b77740 x18: 0000000000000006
      x17: 736564203634343a x16: 656e694c20657079 x15: 0000000000000143
      x14: 00000000ffffffea x13: ffff80000ef5b8b8 x12: ffff80000ef5b8c8
      x11: ffff80000a591d30 x10: ffff80000a579d40 x9 : c0000000ffffefff
      x8 : 0000000000000003 x7 : 0000000000017fe8 x6 : ffff80000a579ce8
      x5 : 0000000000000000 x4 : 0000000000000000 x3 : 0000000000000000
      x2 : 3a35ec12ed7f8900 x1 : 0000000000000000 x0 : 0000000000000052
      Call trace:
       ath11k_hw_qcn9074_rx_desc_get_decap_type+0x34/0x60 [ath11k]
       ath11k_dp_rx_deliver_msdu.isra.42+0xa4/0x3d0 [ath11k]
       ath11k_dp_rx_mon_deliver.isra.43+0x2f8/0x458 [ath11k]
       ath11k_dp_rx_process_mon_rings+0x310/0x4c0 [ath11k]
       ath11k_dp_service_srng+0x234/0x338 [ath11k]
       ath11k_pcic_ext_grp_napi_poll+0x30/0xb8 [ath11k]
       __napi_poll+0x5c/0x190
       napi_threaded_poll+0xf0/0x118
       kthread+0xf4/0x110
       ret_from_fork+0x10/0x20

      Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.7.0.1-01744-QCAHKSWPL_SILICONZ-1
      Reported-by: Florian Schmidt <florian@fls.name>
      Link: https://bugzilla.kernel.org/show_bug.cgi?id=216573
      Signed-off-by: Nagarajan Maran <quic_nmaran@quicinc.com>
      Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
      Link: https://lore.kernel.org/r/20221129142532.23421-1-quic_nmaran@quicinc.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit de15e8bbd9eb26fe94a06d0ec7be82dc490eb729
  Author: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
  Date:   Mon Dec 5 10:43:08 2022 +0900

      wifi: ath9k: Fix use-after-free in ath9k_hif_usb_disconnect()

      [ Upstream commit f099c5c9e2ba08a379bd354a82e05ef839ae29ac ]

      This patch fixes a use-after-free in ath9k that occurs in
      ath9k_hif_usb_disconnect() when ath9k_destroy_wmi() is trying to access
      'drv_priv' that has already been freed by ieee80211_free_hw(), called by
      ath9k_htc_hw_deinit(). The patch moves ath9k_destroy_wmi() before
      ieee80211_free_hw(). Note that urbs from the driver should be killed
      before freeing 'wmi' with ath9k_destroy_wmi() as their callbacks will
      access 'wmi'.

      Found by a modified version of syzkaller.

      ==================================================================
      BUG: KASAN: use-after-free in ath9k_destroy_wmi+0x38/0x40
      Read of size 8 at addr ffff8881069132a0 by task kworker/0:1/7

      CPU: 0 PID: 7 Comm: kworker/0:1 Tainted: G O 5.14.0+ #131
      Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.12.1-0-ga5cab58e9a3f-
  prebuilt.qemu.org 04/01/2014
      Workqueue: usb_hub_wq hub_event
      Call Trace:
       dump_stack_lvl+0x8e/0xd1
       print_address_description.constprop.0.cold+0x93/0x334
       ? ath9k_destroy_wmi+0x38/0x40
       ? ath9k_destroy_wmi+0x38/0x40
       kasan_report.cold+0x83/0xdf
       ? ath9k_destroy_wmi+0x38/0x40
       ath9k_destroy_wmi+0x38/0x40
       ath9k_hif_usb_disconnect+0x329/0x3f0
       ? ath9k_hif_usb_suspend+0x120/0x120
       ? usb_disable_interface+0xfc/0x180
```

```
       usb_unbind_interface+0x19b/0x7e0
       ? usb_autoresume_device+0x50/0x50
       device_release_driver_internal+0x44d/0x520
       bus_remove_device+0x2e5/0x5a0
       device_del+0x5b2/0xe30
       ? __device_link_del+0x370/0x370
       ? usb_remove_ep_devs+0x43/0x80
       ? remove_intf_ep_devs+0x112/0x1a0
       usb_disable_device+0x1e3/0x5a0
       usb_disconnect+0x267/0x870
       hub_event+0x168d/0x3950
       ? rcu_read_lock_sched_held+0xa1/0xd0
       ? hub_port_debounce+0x2e0/0x2e0
       ? check_irq_usage+0x860/0xf20
       ? drain_workqueue+0x281/0x360
       ? lock_release+0x640/0x640
       ? rcu_read_lock_sched_held+0xa1/0xd0
       ? rcu_read_lock_bh_held+0xb0/0xb0
       ? lockdep_hardirqs_on_prepare+0x273/0x3e0
       process_one_work+0x92b/0x1460
       ? pwq_dec_nr_in_flight+0x330/0x330
       ? rwlock_bug.part.0+0x90/0x90
       worker_thread+0x95/0xe00
       ? __kthread_parkme+0x115/0x1e0
       ? process_one_work+0x1460/0x1460
       kthread+0x3a1/0x480
       ? set_kthread_struct+0x120/0x120
       ret_from_fork+0x1f/0x30

     The buggy address belongs to the page:
     page:ffffea00041a44c0 refcount:0 mapcount:0 mapping:0000000000000000 index:0x0
  pfn:0x106913
     flags: 0x200000000000000(node=0|zone=2)
     raw: 0200000000000000 0000000000000000 dead000000000122 0000000000000000
     raw: 0000000000000000 0000000000000000 00000000ffffffff 0000000000000000
     page dumped because: kasan: bad access detected
     page_owner tracks the page as freed
     page last allocated via order 3, migratetype Unmovable, gfp_mask
  0x40dc0(GFP_KERNEL|__GFP_COMP|__GFP_ZERO), pid 7, ts 38347963444, free_ts 41399957635
       prep_new_page+0x1aa/0x240
       get_page_from_freelist+0x159a/0x27c0
       __alloc_pages+0x2da/0x6a0
       alloc_pages+0xec/0x1e0
       kmalloc_order+0x39/0xf0
       kmalloc_order_trace+0x19/0x120
       __kmalloc+0x308/0x390
       wiphy_new_nm+0x6f5/0x1dd0
       ieee80211_alloc_hw_nm+0x36d/0x2230
       ath9k_htc_probe_device+0x9d/0x1e10
       ath9k_htc_hw_init+0x34/0x50
       ath9k_hif_usb_firmware_cb+0x25f/0x4e0
       request_firmware_work_func+0x131/0x240
       process_one_work+0x92b/0x1460
       worker_thread+0x95/0xe00
       kthread+0x3a1/0x480
     page last free stack trace:
       free_pcp_prepare+0x3d3/0x7f0
       free_unref_page+0x1e/0x3d0
       device_release+0xa4/0x240
       kobject_put+0x186/0x4c0
       put_device+0x20/0x30
       ath9k_htc_disconnect_device+0x1cf/0x2c0
       ath9k_htc_hw_deinit+0x26/0x30
       ath9k_hif_usb_disconnect+0x2d9/0x3f0
       usb_unbind_interface+0x19b/0x7e0
```

```
       device_release_driver_internal+0x44d/0x520
       bus_remove_device+0x2e5/0x5a0
       device_del+0x5b2/0xe30
       usb_disable_device+0x1e3/0x5a0
       usb_disconnect+0x267/0x870
       hub_event+0x168d/0x3950
       process_one_work+0x92b/0x1460

      Memory state around the buggy address:
       ffff888106913180: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
       ffff888106913200: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
      >ffff888106913280: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
                         ^
       ffff888106913300: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
       ffff888106913380: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
      ==================================================================

      Reported-by: Dokyung Song <dokyungs@yonsei.ac.kr>
      Reported-by: Jisoo Jang <jisoo.jang@yonsei.ac.kr>
      Reported-by: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
      Signed-off-by: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
      Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
      Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
      Link: https://lore.kernel.org/r/20221205014308.1617597-1-linuxlovemin@yonsei.ac.kr
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c2cb786500e47b9fd869227f5c9f51dbc0735370
  Author: Kan Liang <kan.liang@linux.intel.com>
  Date:   Fri Feb 10 11:02:38 2023 -0800

      perf/x86/intel/uncore: Add Meteor Lake support

      [ Upstream commit c828441f21ddc819a28b5723a72e3c840e9de1c6 ]

      The uncore subsystem for Meteor Lake is similar to the previous Alder
      Lake. The main difference is that MTL provides PMU support for different
      tiles, while ADL only provides PMU support for the whole package. On
      ADL, there are CBOX, ARB, and clockbox uncore PMON units. On MTL, they
      are split into CBOX/HAC_CBOX, ARB/HAC_ARB, and cncu/sncu which provides
      a fixed counter for clockticks. Also, new MSR addresses are introduced
      on MTL.

      The IMC uncore PMON is the same as Alder Lake. Add new PCIIDs of IMC for
      Meteor Lake.

      Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
      Signed-off-by: Ingo Molnar <mingo@kernel.org>
      Link: https://lore.kernel.org/r/20230210190238.1726237-1-kan.liang@linux.intel.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 28a2cd87dd09fea2c804d07155a7f011755641aa
  Author: Peter Zijlstra <peterz@infradead.org>
  Date:   Thu Jan 26 16:08:31 2023 +0100

      cpuidle: lib/bug: Disable rcu_is_watching() during WARN/BUG

      [ Upstream commit 5a5d7e9badd2cb8065db171961bd30bd3595e4b6 ]

      In order to avoid WARN/BUG from generating nested or even recursive
      warnings, force rcu_is_watching() true during
      WARN/lockdep_rcu_suspicious().

      Notably things like unwinding the stack can trigger rcu_dereference()
      warnings, which then triggers more unwinding which then triggers more
      warnings etc..
```

```
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link: https://lore.kernel.org/r/20230126151323.408156109@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8d0e32f66a1b74f25a532c5e9c3f5040f7b5accb
Author: Mark Rutland <mark.rutland@arm.com>
Date:   Thu Jan 26 16:08:30 2023 +0100

    cpuidle: drivers: firmware: psci: Dont instrument suspend code

    [ Upstream commit 393e2ea30aec634b37004d401863428e120d5e1b ]

    The PSCI suspend code is currently instrumentable, which is not safe as
    instrumentation (e.g. ftrace) may try to make use of RCU during idle
    periods when RCU is not watching.

    To fix this we need to ensure that psci_suspend_finisher() and anything
    it calls are not instrumented. We can do this fairly simply by marking
    psci_suspend_finisher() and the psci*_cpu_suspend() functions as
    noinstr, and the underlying helper functions as __always_inline.

    When CONFIG_DEBUG_VIRTUAL=y, __pa_symbol() can expand to an out-of-line
    instrumented function, so we must use __pa_symbol_nodebug() within
    psci_suspend_finisher().

    The raw SMCCC invocation functions are written in assembly, and are not
    subject to compiler instrumentation.

    Signed-off-by: Mark Rutland <mark.rutland@arm.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link: https://lore.kernel.org/r/20230126151323.349423061@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bcb553fb19350960f262574ca355817d93a60d71
Author: Jens Axboe <axboe@kernel.dk>
Date:   Tue Jan 24 08:23:20 2023 -0700

    x86/fpu: Don't set TIF_NEED_FPU_LOAD for PF_IO_WORKER threads

    [ Upstream commit cb3ea4b7671b7cfbac3ee609976b790aebd0bbda ]

    We don't set it on PF_KTHREAD threads as they never return to userspace,
    and PF_IO_WORKER threads are identical in that regard. As they keep
    running in the kernel until they die, skip setting the FPU flag on them.

    More of a cosmetic thing that was found while debugging and
    issue and pondering why the FPU flag is set on these threads.

    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Acked-by: Peter Zijlstra <peterz@infradead.org>
    Link: https://lore.kernel.org/r/560c844c-f128-555b-40c6-31baff27537f@kernel.dk
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3fb1b2ca5d80863cf3fccc616131002db1f2317b
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Thu Jan 12 20:43:33 2023 +0100

    cpuidle, intel_idle: Fix CPUIDLE_FLAG_INIT_XSTATE

    [ Upstream commit 821ad23d0eaff73ef599ece39ecc77482df20a8c ]
```

```
    Fix instrumentation bugs objtool found:

      vmlinux.o: warning: objtool: intel_idle_s2idle+0xd5: call to fpu_idle_fpregs() leaves
.noinstr.text section
      vmlinux.o: warning: objtool: intel_idle_xstate+0x11: call to fpu_idle_fpregs() leaves
.noinstr.text section
      vmlinux.o: warning: objtool: fpu_idle_fpregs+0x9: call to xfeatures_in_use() leaves
.noinstr.text section

    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Tested-by: Tony Lindgren <tony@atomide.com>
    Tested-by: Ulf Hansson <ulf.hansson@linaro.org>
    Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Acked-by: Frederic Weisbecker <frederic@kernel.org>
    Link: https://lore.kernel.org/r/20230112195540.494977795@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03fd93cda063978611a2c4365fb97a400512ba24
Author: Michael Grzeschik <m.grzeschik@pengutronix.de>
Date:   Sun Oct 23 23:56:49 2022 +0200

    arm64: zynqmp: Enable hs termination flag for USB dwc3 controller

    [ Upstream commit 32405e532d358a2f9d4befae928b9883c8597616 ]

    Since we need to support legacy phys with the dwc3 controller,
    we enable this quirk on the zynqmp platforms.

    Signed-off-by: Michael Grzeschik <m.grzeschik@pengutronix.de>
    Link: https://lore.kernel.org/r/20221023215649.221726-1-m.grzeschik@pengutronix.de
    Signed-off-by: Michal Simek <michal.simek@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03d7193114bcebcde395c5531d0d9bbb015e697c
Author: Qu Wenruo <wqu@suse.com>
Date:   Sat Dec 17 10:34:29 2022 +0800

    btrfs: scrub: improve tree block error reporting

    [ Upstream commit 28232909ba43561887508a6ef46d7f33a648f375 ]

    [BUG]
    When debugging a scrub related metadata error, it turns out that our
    metadata error reporting is not ideal.

    The only 3 error messages are:

    - BTRFS error (device dm-2): bdev /dev/mapper/test-scratch1 errs: wr 0, rd 0, flush 0,
corrupt 0, gen 1
      Showing we have metadata generation mismatch errors.

    - BTRFS error (device dm-2): unable to fixup (regular) error at logical 7110656 on dev
/dev/mapper/test-scratch1
      Showing which tree blocks are corrupted.

    - BTRFS warning (device dm-2): checksum/header error at logical 24772608 on dev
/dev/mapper/test-scratch2, physical 3801088: metadata node (level 1) in tree 5
      Showing which physical range the corrupted metadata is at.

    We have to combine the above 3 to know we have a corrupted metadata with
    generation mismatch.

    And this is already the better case, if we have other problems, like
    fsid mismatch, we can not even know the cause.
```

```
    [CAUSE]
    The problem is caused by the fact that, scrub_checksum_tree_block()
    never outputs any error message.

    It just return two bits for scrub: sblock->header_error, and
    sblock->generation_error.

    And later we report error in scrub_print_warning(), but unfortunately we
    only have two bits, there is not really much thing we can done to print
    any detailed errors.

    [FIX]
    This patch will do the following to enhance the error reporting of
    metadata scrub:

    - Add extra warning (ratelimited) for every error we hit
      This can help us to distinguish the different types of errors.
      Some errors can help us to know what's going wrong immediately,
      like bytenr mismatch.

    - Re-order the checks
      Currently we check bytenr first, then immediately generation.
      This can lead to false generation mismatch reports, while the fsid
      mismatches.

    Here is the new output for the bug I'm debugging (we forgot to
    writeback tree blocks for commit roots):

     BTRFS warning (device dm-2): tree block 24117248 mirror 1 has bad fsid, has b77cd862-
f150-4c71-90ec-7baf0544d83f want 17df6abf-23cd-445f-b350-5b3e40bfd2fc
     BTRFS warning (device dm-2): tree block 24117248 mirror 0 has bad fsid, has b77cd862-
f150-4c71-90ec-7baf0544d83f want 17df6abf-23cd-445f-b350-5b3e40bfd2fc

    Now we can immediately know it's some tree blocks didn't even get written
    back, other than the original confusing generation mismatch.

    Signed-off-by: Qu Wenruo <wqu@suse.com>
    Reviewed-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5286b72fb425291af5f4ca7285d73c16a08d8691
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Feb 2 15:19:56 2023 +0100

    trace/blktrace: fix memory leak with using debugfs_lookup()

    [ Upstream commit 83e8864fee26f63a7435e941b7c36a20fd6fe93e ]

    When calling debugfs_lookup() the result must have dput() called on it,
    otherwise the memory will leak over time.  To make things simpler, just
    call debugfs_lookup_and_remove() instead which handles all of the logic
    at once.

    Cc: Jens Axboe <axboe@kernel.dk>
    Cc: Steven Rostedt <rostedt@goodmis.org>
    Cc: Masami Hiramatsu <mhiramat@kernel.org>
    Cc: linux-block@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Cc: linux-trace-kernel@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Link: https://lore.kernel.org/r/20230202141956.2299521-1-gregkh@linuxfoundation.org
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
```

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

commit bfe46d2efe46c5c952f982e2ca94fe2ec5e58e2a
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Thu Jan 19 19:03:50 2023 +0800

    blk-cgroup: synchronize pd_free_fn() from blkg_free_workfn() and
blkcg_deactivate_policy()

    [ Upstream commit f1c006f1c6850c14040f8337753a63119bba39b9 ]

    Currently parent pd can be freed before child pd:

    t1: remove cgroup C1
    blkcg_destroy_blkgs
     blkg_destroy
      list_del_init(&blkg->q_node)
      // remove blkg from queue list
      percpu_ref_kill(&blkg->refcnt)
       blkg_release
        call_rcu

    t2: from t1
    __blkg_release
     blkg_free
      schedule_work
                            t4: deactivate policy
                            blkcg_deactivate_policy
                             pd_free_fn
                             // parent of C1 is freed first
    t3: from t2
     blkg_free_workfn
      pd_free_fn

    If policy(for example, ioc_timer_fn() from iocost) access parent pd from
    child pd after pd_offline_fn(), then UAF can be triggered.

    Fix the problem by delaying 'list_del_init(&blkg->q_node)' from
    blkg_destroy() to blkg_free_workfn(), and using a new disk level mutex to
    synchronize blkg_free_workfn() and blkcg_deactivate_policy().

    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230119110350.2287325-4-yukuai1@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57a425badc05c2e87e9f25713e5c3c0298e4202c
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Thu Jan 19 19:03:48 2023 +0800

    blk-cgroup: dropping parent refcount after pd_free_fn() is done

    [ Upstream commit c7241babf0855d8a6180cd1743ff0ec34de40b4e ]

    Some cgroup policies will access parent pd through child pd even
    after pd_offline_fn() is done. If pd_free_fn() for parent is called
    before child, then UAF can be triggered. Hence it's better to guarantee
    the order of pd_free_fn().

    Currently refcount of parent blkg is dropped in __blkg_release(), which
    is before pd_free_fn() is called in blkg_free_work_fn() while
    blkg_free_work_fn() is called asynchronously.
```

```
    This patch make sure pd_free_fn() called from removing cgroup is ordered
    by delaying dropping parent refcount after calling pd_free_fn() for
    child.

    BTW, pd_free_fn() will also be called from blkcg_deactivate_policy()
    from deleting device, and following patches will guarantee the order.

    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230119110350.2287325-2-yukuai1@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b96d7b4a9745fbd0c8384608ceb1f50415e862fa
Author: Li Nan <linan122@huawei.com>
Date:   Tue Jan 17 15:08:05 2023 +0800

    blk-iocost: fix divide by 0 error in calc_lcoefs()

    [ Upstream commit 984af1e66b4126cf145153661cc24c213e2ec231 ]

    echo max of u64 to cost.model can cause divide by 0 error.

      # echo 8:0 rbps=18446744073709551615 > /sys/fs/cgroup/io.cost.model

      divide error: 0000 [#1] PREEMPT SMP
      RIP: 0010:calc_lcoefs+0x4c/0xc0
      Call Trace:
       <TASK>
       ioc_refresh_params+0x2b3/0x4f0
       ioc_cost_model_write+0x3cb/0x4c0
       ? _copy_from_iter+0x6d/0x6c0
       ? kernfs_fop_write_iter+0xfc/0x270
       cgroup_file_write+0xa0/0x200
       kernfs_fop_write_iter+0x17d/0x270
       vfs_write+0x414/0x620
       ksys_write+0x73/0x160
       __x64_sys_write+0x1e/0x30
       do_syscall_64+0x35/0x80
       entry_SYSCALL_64_after_hwframe+0x63/0xcd

    calc_lcoefs() uses the input value of cost.model in DIV_ROUND_UP_ULL,
    overflow would happen if bps plus IOC_PAGE_SIZE is greater than
    ULLONG_MAX, it can cause divide by 0 error.

    Fix the problem by setting basecost

    Signed-off-by: Li Nan <linan122@huawei.com>
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Link: https://lore.kernel.org/r/20230117070806.3857142-5-yukuai1@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88c61805a8d09acaa9759894e9a0d0ed6728db9a
Author: Jann Horn <jannh@google.com>
Date:   Mon Jan 16 20:14:25 2023 +0100

    fs: Use CHECK_DATA_CORRUPTION() when kernel bugs are detected

    [ Upstream commit 47d586913f2abec4d240bae33417f537fda987ec ]

    Currently, filp_close() and generic_shutdown_super() use printk() to log
    messages when bugs are detected. This is problematic because infrastructure
```

```
    like syzkaller has no idea that this message indicates a bug.
    In addition, some people explicitly want their kernels to BUG() when kernel
    data corruption has been detected (CONFIG_BUG_ON_DATA_CORRUPTION).
    And finally, when generic_shutdown_super() detects remaining inodes on a
    system without CONFIG_BUG_ON_DATA_CORRUPTION, it would be nice if later
    accesses to a busy inode would at least crash somewhat cleanly rather than
    walking through freed memory.

    To address all three, use CHECK_DATA_CORRUPTION() when kernel bugs are
    detected.

    Signed-off-by: Jann Horn <jannh@google.com>
    Reviewed-by: Christian Brauner (Microsoft) <brauner@kernel.org>
    Reviewed-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Christian Brauner (Microsoft) <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29961ee63dd676cc67f7c00f76faa21e11f0d7c6
Author: Markuss Broks <markuss.broks@gmail.com>
Date:   Sat Jan 21 22:18:42 2023 +0200

    ARM: dts: exynos: Use Exynos5420 compatible for the MIPI video phy

    [ Upstream commit 5d5aa219a790d61cad2c38e1aa32058f16ad2f0b ]

    For some reason, the driver adding support for Exynos5420 MIPI phy
    back in 2016 wasn't used on Exynos5420, which caused a kernel panic.
    Add the proper compatible for it.

    Signed-off-by: Markuss Broks <markuss.broks@gmail.com>
    Link: https://lore.kernel.org/r/20230121201844.46872-2-markuss.broks@gmail.com
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 439076fb6553a60973e0d36617c3961ecca0095f
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Fri Jan 20 11:18:20 2023 +1000

    exit: Detect and fix irq disabled state in oops

    [ Upstream commit 001c28e57187570e4b5aa4492c7a957fb6d65d7b ]

    If a task oopses with irqs disabled, this can cause various cascading
    problems in the oops path such as sleep-from-invalid warnings, and
    potentially worse.

    Since commit 0258b5fd7c712 ("coredump: Limit coredumps to a single
    thread group"), the unconditional irq enable in coredump_task_exit()
    will "fix" the irq state to be enabled early in do_exit(), so currently
    this may not be triggerable, but that is coincidental and fragile.

    Detect and fix the irqs_disabled() condition in the oops path before
    calling do_exit(), similarly to the way in_atomic() is handled.

    Reported-by: Michael Ellerman <mpe@ellerman.id.au>
    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Acked-by: "Eric W. Biederman" <ebiederm@xmission.com>
    Link: https://lore.kernel.org/lkml/20221004094401.708299-1-npiggin@gmail.com/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 655bed11a5445d3d8c239c2b6545f694e3a93596
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Thu Jan 12 20:44:05 2023 +0100
```

```
    context_tracking: Fix noinstr vs KASAN

    [ Upstream commit 0e26e1de0032779e43929174339429c16307a299 ]

    Low level noinstr context-tracking code is calling out to instrumented
    code on KASAN:

      vmlinux.o: warning: objtool: __ct_user_enter+0x72: call to __kasan_check_write()
leaves .noinstr.text section
      vmlinux.o: warning: objtool: __ct_user_exit+0x47: call to __kasan_check_write() leaves
.noinstr.text section

    Use even lower level atomic methods to avoid the instrumentation.

    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link: https://lore.kernel.org/r/20230112195542.458034262@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ed4492fe412d6b8c7723ce77894b4d346390bb2
Author: Jan Kara <jack@suse.cz>
Date:   Thu Sep 29 16:34:45 2022 +0200

    udf: Define EFSCORRUPTED error code

    [ Upstream commit 3d2d7e61553dbcc8ba45201d8ae4f383742c8202 ]

    Similarly to other filesystems define EFSCORRUPTED error code for
    reporting internal filesystem corruption.

    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd62fd3a49c4ed6da4703801521852d94cdc2197
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Sat Dec 10 21:03:53 2022 +0100

    arm64: dts: qcom: msm8996: Add additional A2NoC clocks

    [ Upstream commit 67fb53745e0b38275fa0b422b6a3c6c1c028c9a2 ]

    On eMMC devices, the UFS clocks aren't started in the bootloader (or well,
    at least it should not be, as that would just leak power..), which results
    in platform reboots when trying to access the unclocked UFS hardware,
    which unfortunately happens on each and every boot, as interconnect calls
    sync_state and goes over each and every path.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Tested-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org> #db820c
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221210200353.418391-6-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 049875b76660bbdc4873a915afb294f954eb7320
Author: Liang He <windhl@126.com>
Date:   Tue Jun 28 19:29:39 2022 +0800

    ARM: OMAP2+: omap4-common: Fix refcount leak bug

    [ Upstream commit 7c32919a378782c95c72bc028b5c30dfe8c11f82 ]

    In omap4_sram_init(), of_find_compatible_node() will return a node
    pointer with refcount incremented. We should use of_node_put() when
    it is not used anymore.
```

```
    Signed-off-by: Liang He <windhl@126.com>
    Message-Id: <20220628112939.160737-1-windhl@126.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1cd17d9a8a02b96a99a218b0b09faf92356793c
Author: Bjorn Andersson <quic_bjorande@quicinc.com>
Date:   Mon Jan 9 14:39:31 2023 -0800

    rpmsg: glink: Release driver_override

    [ Upstream commit fb80ef67e8ff6a00d3faad4cb348dafdb8eccfd8 ]

    Upon termination of the rpmsg_device, driver_override needs to be freed
    to avoid leaking the potentially assigned string.

    Fixes: 42cd402b8fd4 ("rpmsg: Fix kfree() of static memory on setting driver_override")
    Fixes: 39e47767ec9b ("rpmsg: Add driver_override device attribute for rpmsg_device")
    Reviewed-by: Chris Lew <quic_clew@quicinc.com>
    Signed-off-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230109223931.1706429-1-quic_bjorande@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed6c1ac6f9e939052fc5842a15e7aa4709ca636c
Author: Bjorn Andersson <quic_bjorande@quicinc.com>
Date:   Tue Feb 14 15:42:31 2023 -0800

    rpmsg: glink: Avoid infinite loop on intent for missing channel

    [ Upstream commit 3e74ec2f39362bffbd42854acbb67c7f4cb808f9 ]

    In the event that an intent advertisement arrives on an unknown channel
    the fifo is not advanced, resulting in the same message being handled
    over and over.

    Fixes: dacbb35e930f ("rpmsg: glink: Receive and store the remote intent buffers")
    Signed-off-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Reviewed-by: Chris Lew <quic_clew@quicinc.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230214234231.2069751-1-quic_bjorande@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce33877a70020d4283b1a874f965534f69a8ec8c
Author: Tasos Sahanidis <tasos@tasossah.com>
Date:   Thu Jan 26 12:00:59 2023 +0100

    media: saa7134: Use video_unregister_device for radio_dev

    [ Upstream commit bc7635c6435c77a0c168e2cc6535740adfaff4e4 ]

    The radio device doesn't use vb2, thus calling vb2_video_unregister_device()
    which results in the following warning being printed on module unload.

    WARNING: CPU: 1 PID: 215963 at drivers/media/common/videobuf2/videobuf2-v4l2.c:1236
vb2_video_unregister_device+0xc6/0xe0 [videobuf2_v4l2]

    Fixes: 11788d9b7e91 ("media: media/pci: use vb2_video_unregister_device()")
    Signed-off-by: Tasos Sahanidis <tasos@tasossah.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19aadf0eb70edae7180285dbb9bfa237d1ddb34d
```

```
Author: Duoming Zhou <duoming@zju.edu.cn>
Date:   Mon Jan 23 03:04:38 2023 +0100

    media: usb: siano: Fix use after free bugs caused by do_submit_urb

    [ Upstream commit ebad8e731c1c06adf04621d6fd327b860c0861b5 ]

    There are UAF bugs caused by do_submit_urb(). One of the KASan reports
    is shown below:

    [   36.403605] BUG: KASAN: use-after-free in worker_thread+0x4a2/0x890
    [   36.406105] Read of size 8 at addr ffff8880059600e8 by task kworker/0:2/49
    [   36.408316]
    [   36.408867] CPU: 0 PID: 49 Comm: kworker/0:2 Not tainted 6.2.0-rc3-15798-
g5a41237ad1d4-dir8
    [   36.411696] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-
g15584
    [   36.416157] Workqueue:  0x0 (events)
    [   36.417654] Call Trace:
    [   36.418546]  <TASK>
    [   36.419320]  dump_stack_lvl+0x96/0xd0
    [   36.420522]  print_address_description+0x75/0x350
    [   36.421992]  print_report+0x11b/0x250
    [   36.423174]  ? _raw_spin_lock_irqsave+0x87/0xd0
    [   36.424806]  ? __virt_addr_valid+0xcf/0x170
    [   36.426069]  ? worker_thread+0x4a2/0x890
    [   36.427355]  kasan_report+0x131/0x160
    [   36.428556]  ? worker_thread+0x4a2/0x890
    [   36.430053]  worker_thread+0x4a2/0x890
    [   36.431297]  ? worker_clr_flags+0x90/0x90
    [   36.432479]  kthread+0x166/0x190
    [   36.433493]  ? kthread_blkcg+0x50/0x50
    [   36.434669]  ret_from_fork+0x22/0x30
    [   36.435923]  </TASK>
    [   36.436684]
    [   36.437215] Allocated by task 24:
    [   36.438289]  kasan_set_track+0x50/0x80
    [   36.439436]  __kasan_kmalloc+0x89/0xa0
    [   36.440566]  smsusb_probe+0x374/0xc90
    [   36.441920]  usb_probe_interface+0x2d1/0x4c0
    [   36.443253]  really_probe+0x1d5/0x580
    [   36.444539]  __driver_probe_device+0xe3/0x130
    [   36.446085]  driver_probe_device+0x49/0x220
    [   36.447423]  __device_attach_driver+0x19e/0x1b0
    [   36.448931]  bus_for_each_drv+0xcb/0x110
    [   36.450217]  __device_attach+0x132/0x1f0
    [   36.451470]  bus_probe_device+0x59/0xf0
    [   36.452563]  device_add+0x4ec/0x7b0
    [   36.453830]  usb_set_configuration+0xc63/0xe10
    [   36.455230]  usb_generic_driver_probe+0x3b/0x80
    [   36.456166] printk: console [ttyGS0] disabled
    [   36.456569]  usb_probe_device+0x90/0x110
    [   36.459523]  really_probe+0x1d5/0x580
    [   36.461027]  __driver_probe_device+0xe3/0x130
    [   36.462465]  driver_probe_device+0x49/0x220
    [   36.463847]  __device_attach_driver+0x19e/0x1b0
    [   36.465229]  bus_for_each_drv+0xcb/0x110
    [   36.466466]  __device_attach+0x132/0x1f0
    [   36.467799]  bus_probe_device+0x59/0xf0
    [   36.469010]  device_add+0x4ec/0x7b0
    [   36.470125]  usb_new_device+0x863/0xa00
    [   36.471374]  hub_event+0x18c7/0x2220
    [   36.472746]  process_one_work+0x34c/0x5b0
    [   36.474041]  worker_thread+0x4b7/0x890
    [   36.475216]  kthread+0x166/0x190
```

```
[   36.476267]  ret_from_fork+0x22/0x30
[   36.477447]
[   36.478160] Freed by task 24:
[   36.479239]  kasan_set_track+0x50/0x80
[   36.480512]  kasan_save_free_info+0x2b/0x40
[   36.481808]  ____kasan_slab_free+0x122/0x1a0
[   36.483173]  __kmem_cache_free+0xc4/0x200
[   36.484563]  smsusb_term_device+0xcd/0xf0
[   36.485896]  smsusb_probe+0xc85/0xc90
[   36.486976]  usb_probe_interface+0x2d1/0x4c0
[   36.488303]  really_probe+0x1d5/0x580
[   36.489498]  __driver_probe_device+0xe3/0x130
[   36.491140]  driver_probe_device+0x49/0x220
[   36.492475]  __device_attach_driver+0x19e/0x1b0
[   36.493988]  bus_for_each_drv+0xcb/0x110
[   36.495171]  __device_attach+0x132/0x1f0
[   36.496617]  bus_probe_device+0x59/0xf0
[   36.497875]  device_add+0x4ec/0x7b0
[   36.498972]  usb_set_configuration+0xc63/0xe10
[   36.500264]  usb_generic_driver_probe+0x3b/0x80
[   36.501740]  usb_probe_device+0x90/0x110
[   36.503084]  really_probe+0x1d5/0x580
[   36.504241]  __driver_probe_device+0xe3/0x130
[   36.505548]  driver_probe_device+0x49/0x220
[   36.506766]  __device_attach_driver+0x19e/0x1b0
[   36.508368]  bus_for_each_drv+0xcb/0x110
[   36.509646]  __device_attach+0x132/0x1f0
[   36.510911]  bus_probe_device+0x59/0xf0
[   36.512103]  device_add+0x4ec/0x7b0
[   36.513215]  usb_new_device+0x863/0xa00
[   36.514736]  hub_event+0x18c7/0x2220
[   36.516130]  process_one_work+0x34c/0x5b0
[   36.517396]  worker_thread+0x4b7/0x890
[   36.518591]  kthread+0x166/0x190
[   36.519599]  ret_from_fork+0x22/0x30
[   36.520851]
[   36.521405] Last potentially related work creation:
[   36.523143]  kasan_save_stack+0x3f/0x60
[   36.524275]  kasan_record_aux_stack_noalloc+0x9d/0xb0
[   36.525831]  insert_work+0x25/0x130
[   36.527039]  __queue_work+0x4d4/0x620
[   36.528236]  queue_work_on+0x72/0xb0
[   36.529344]  __usb_hcd_giveback_urb+0x13f/0x1b0
[   36.530819]  dummy_timer+0x350/0x1a40
[   36.532149]  call_timer_fn+0x2c/0x190
[   36.533567]  expire_timers+0x69/0x1f0
[   36.534736]  __run_timers+0x289/0x2d0
[   36.535841]  run_timer_softirq+0x2d/0x60
[   36.537110]  __do_softirq+0x116/0x380
[   36.538377]
[   36.538950] Second to last potentially related work creation:
[   36.540855]  kasan_save_stack+0x3f/0x60
[   36.542084]  kasan_record_aux_stack_noalloc+0x9d/0xb0
[   36.543592]  insert_work+0x25/0x130
[   36.544891]  __queue_work+0x4d4/0x620
[   36.546168]  queue_work_on+0x72/0xb0
[   36.547328]  __usb_hcd_giveback_urb+0x13f/0x1b0
[   36.548805]  dummy_timer+0x350/0x1a40
[   36.550116]  call_timer_fn+0x2c/0x190
[   36.551570]  expire_timers+0x69/0x1f0
[   36.552762]  __run_timers+0x289/0x2d0
[   36.553916]  run_timer_softirq+0x2d/0x60
[   36.555118]  __do_softirq+0x116/0x380
[   36.556239]
[   36.556807] The buggy address belongs to the object at ffff888005960000
```

```
    [   36.556807]  which belongs to the cache kmalloc-4k of size 4096
    [   36.560652] The buggy address is located 232 bytes inside of
    [   36.560652]  4096-byte region [ffff888005960000, ffff888005961000)
    [   36.564791]
    [   36.565355] The buggy address belongs to the physical page:
    [   36.567212] page:000000004f0a0731 refcount:1 mapcount:0 mapping:0000000000000000
index:0x00
    [   36.570534] head:000000004f0a0731 order:3 compound_mapcount:0 subpages_mapcount:0
compound0
    [   36.573717] flags: 0x100000000010200(slab|head|node=0|zone=1)
    [   36.575481] raw: 0100000000010200 ffff888001042140 dead000000000122 0000000000000000
    [   36.577842] raw: 0000000000000000 0000000000040004 00000001ffffffff 0000000000000000
    [   36.580175] page dumped because: kasan: bad access detected
    [   36.581994]
    [   36.582548] Memory state around the buggy address:
    [   36.583983]  ffff88800595ff80: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
    [   36.586240]  ffff888005960000: fa fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
    [   36.588884] >ffff888005960080: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
    [   36.591071]                                                 ^
    [   36.593295]  ffff888005960100: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
    [   36.595705]  ffff888005960180: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
    [   36.598026] ==================================================================
    [   36.600224] Disabling lock debugging due to kernel taint
    [   36.602681] general protection fault, probably for non-canonical address
0x43600a000000060I
    [   36.607129] CPU: 0 PID: 49 Comm: kworker/0:2 Tainted: G    B             6.2.0-rc3-
15798-8
    [   36.611115] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-
g15584
    [   36.615026] Workqueue: events do_submit_urb
    [   36.616290] RIP: 0010:_raw_spin_lock_irqsave+0x8a/0xd0
    [   36.618107] Code: 24 00 00 00 00 48 89 df be 04 00 00 00 e8 9e b5 c6 fe 48 89 ef be
04 00 5
    [   36.623522] RSP: 0018:ffff888004b6fcf0 EFLAGS: 00010046
    [   36.625072] RAX: 0000000000000000 RBX: 043600a000000060 RCX: ffffffff9fc0e0d7
    [   36.627206] RDX: 0000000000000000 RSI: dfffc0000000000 RDI: ffff888004b6fcf0
    [   36.629813] RBP: ffff888004b6fcf0 R08: dfffc0000000000 R09: ffffed100096df9f
    [   36.631974] R10: dfffe9100096dfa0 R11: 1ffff1100096df9e R12: ffff888005960020
    [   36.634285] R13: ffff8880059600f0 R14: 0000000000000246 R15: 0000000000000001
    [   36.636438] FS:  0000000000000000(0000) GS:ffff88806d600000(0000)
knlGS:0000000000000000
    [   36.639092] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [   36.640951] CR2: 00007f07476819a3 CR3: 0000000004a34000 CR4: 00000000000006f0
    [   36.643411] Call Trace:
    [   36.644215]  <TASK>
    [   36.644902]  smscore_getbuffer+0x3e/0x1e0
    [   36.646147]  do_submit_urb+0x4f/0x190
    [   36.647449]  process_one_work+0x34c/0x5b0
    [   36.648777]  worker_thread+0x4b7/0x890
    [   36.649984]  ? worker_clr_flags+0x90/0x90
    [   36.651166]  kthread+0x166/0x190
    [   36.652151]  ? kthread_blkcg+0x50/0x50
    [   36.653547]  ret_from_fork+0x22/0x30
    [   36.655051]  </TASK>
    [   36.655733] Modules linked in:
    [   36.656787] ---[ end trace 0000000000000000 ]---
    [   36.658328] RIP: 0010:_raw_spin_lock_irqsave+0x8a/0xd0
    [   36.660045] Code: 24 00 00 00 00 48 89 df be 04 00 00 00 e8 9e b5 c6 fe 48 89 ef be
04 00 5
    [   36.665730] RSP: 0018:ffff888004b6fcf0 EFLAGS: 00010046
    [   36.667448] RAX: 0000000000000000 RBX: 043600a000000060 RCX: ffffffff9fc0e0d7
    [   36.669675] RDX: 0000000000000000 RSI: dfffc0000000000 RDI: ffff888004b6fcf0
    [   36.672645] RBP: ffff888004b6fcf0 R08: dfffc0000000000 R09: ffffed100096df9f
    [   36.674921] R10: dfffe9100096dfa0 R11: 1ffff1100096df9e R12: ffff888005960020
    [   36.677034] R13: ffff8880059600f0 R14: 0000000000000246 R15: 0000000000000001
```

```
    [   36.679184] FS:  0000000000000000(0000) GS:ffff88806d600000(0000)
knlGS:0000000000000000
    [   36.681655] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [   36.683383] CR2: 00007f07476819a3 CR3: 0000000004a34000 CR4: 00000000000006f0
    [   36.685733] Kernel panic - not syncing: Fatal exception
    [   36.688585] Kernel Offset: 0x1d400000 from 0xffffffff81000000 (relocation range:
0xffffffff)
    [   36.692199] ---[ end Kernel panic - not syncing: Fatal exception ]---

    When the siano device is plugged in, it may call the following functions
    to initialize the device.

    smsusb_probe()-->smsusb_init_device()-->smscore_start_device().

    When smscore_start_device() gets failed, the function smsusb_term_device()
    will be called and smsusb_device_t will be deallocated. Although we use
    usb_kill_urb() in smsusb_stop_streaming() to cancel transfer requests
    and wait for them to finish, the worker threads that are scheduled by
    smsusb_onresponse() may be still running. As a result, the UAF bugs
    could happen.

    We add cancel_work_sync() in smsusb_stop_streaming() in order that the
    worker threads could finish before the smsusb_device_t is deallocated.

    Fixes: dd47fbd40e6e ("[media] smsusb: don't sleep while atomic")
    Signed-off-by: Duoming Zhou <duoming@zju.edu.cn>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f84631a03795123dfce516bae23f0319051ea5a4
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Thu Jan 26 14:03:51 2023 +0100

    media: i2c: ov7670: 0 instead of -EINVAL was returned

    [ Upstream commit 6a4c664539e6de9b32b65ddcf767ec1bcc1d7f8a ]

    If the media bus is unsupported, then return -EINVAL. Instead it
    returned 'ret' which happened to be 0.

    This fixes a smatch warning:

    ov7670.c:1843 ov7670_parse_dt() warn: missing error code? 'ret'

    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Fixes: 01b8444828fc ("media: v4l2: i2c: ov7670: Implement OF mbus configuration")
    Acked-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac058bd815d197cfac6701ce750191f440e01f97
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Wed Dec 28 23:11:47 2022 +0100

    media: atomisp: Only set default_run_mode on first open of a stream/asd

    [ Upstream commit 60ec70a71a9f9975a5d2dd4a7d97c20da0e41976 ]

    Calling v4l2_ctrl_s_ctrl(asd->run_mode, pipe->default_run_mode) when
    the stream is already active (through another /dev/video# node) causes
    the stream to stop.

    Move the call to set the default run-mode so that it is only done
    on the first open of one of the 4 /dev/video# nodes of one of
```

```
    the 2 streams (atomisp-sub-devices / asd-s).

    Fixes: 2c45e343c581 ("media: atomisp: set per-device's default mode")
    Reviewed-by: Andy Shevchenko <andy@kernel.org>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97443d473fadc9868ca8d7613250b78162e9e0d4
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Wed Jan 4 09:21:36 2023 +0100

    media: atomisp: fix videobuf2 Kconfig depenendency

    [ Upstream commit be94be1b7fc7e51f9ccef20a0ef76583587275f3 ]

    The recent conversion missed the Kconfig bit, so it can now
    end up in a link error on randconfig builds:

    ld.lld: error: undefined symbol: vb2_vmalloc_memops
    >>> referenced by atomisp_fops.c
    >>>               drivers/staging/media/atomisp/pci/atomisp_fops.o:(atomisp_open) in
archive vmlinux.a

    Link: https://lore.kernel.org/r/20230104082212.3770415-1-arnd@kernel.org

    Fixes: cb48ae89be3b ("media: atomisp: Convert to videobuf2")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Tested-by: Andy Shevchenko <andy@kernel.org>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 182ea492aae5b64067277e60a4ea5995c4628555
Author: Duoming Zhou <duoming@zju.edu.cn>
Date:   Tue Jan 24 08:55:33 2023 +0100

    media: rc: Fix use-after-free bugs caused by ene_tx_irqsim()

    [ Upstream commit 29b0589a865b6f66d141d79b2dd1373e4e50fe17 ]

    When the ene device is detaching, function ene_remove() will
    be called. But there is no function to cancel tx_sim_timer
    in ene_remove(), the timer handler ene_tx_irqsim() could race
    with ene_remove(). As a result, the UAF bugs could happen,
    the process is shown below.

        (cleanup routine)          |        (timer routine)
                                   | mod_timer(&dev->tx_sim_timer, ..)
    ene_remove()                   | (wait a time)
                                   | ene_tx_irqsim()
                                   |   dev->hw_lock //USE
                                   |   ene_tx_sample(dev) //USE

    Fix by adding del_timer_sync(&dev->tx_sim_timer) in ene_remove(),
    The tx_sim_timer could stop before ene device is deallocated.

    What's more, The rc_unregister_device() and del_timer_sync()
    should be called first in ene_remove() and the deallocated
    functions such as free_irq(), release_region() and so on
    should be called behind them. Because the rc_unregister_device()
    is well synchronized. Otherwise, race conditions may happen. The
    situations that may lead to race conditions are shown below.

    Firstly, the rx receiver is disabled with ene_rx_disable()
```

```
    before rc_unregister_device() in ene_remove(), which means it
    can be enabled again if a process opens /dev/lirc0 between
    ene_rx_disable() and rc_unregister_device().

    Secondly, the irqaction descriptor is freed by free_irq()
    before the rc device is unregistered, which means irqaction
    descriptor may be accessed again after it is deallocated.

    Thirdly, the timer can call ene_tx_sample() that can write
    to the io ports, which means the io ports could be accessed
    again after they are deallocated by release_region().

    Therefore, the rc_unregister_device() and del_timer_sync()
    should be called first in ene_remove().

    Suggested by: Sean Young <sean@mess.org>

    Fixes: 9ea53b74df9c ("V4L/DVB: STAGING: remove lirc_ene0100 driver")
    Signed-off-by: Duoming Zhou <duoming@zju.edu.cn>
    Signed-off-by: Sean Young <sean@mess.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bfe402cfb032d49ba88c4c560eaf6527392240dd
Author: Dong Chuanjian <chuanjian@nfschina.com>
Date:   Tue Dec 27 03:36:25 2022 +0100

    media: drivers/media/v4l2-core/v4l2-h264 : add detection of null pointers

    [ Upstream commit be3ae7cf4326e95bb1d5413b63baabc26f4a1324 ]

    When the pointer variable is judged to be null, null is returned
    directly.

    [hverkuil: fix two checkpatch warnings]

    Signed-off-by: Dong Chuanjian <chuanjian@nfschina.com>
    Acked-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Fixes: d3f756ad629b ("media: v4l2: Trace calculated p/b0/b1 initial reflist")
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a32be7ff63f311af09846c061bbd5779c5f277f
Author: Ming Qian <ming.qian@nxp.com>
Date:   Fri Jan 13 06:25:51 2023 +0100

    media: amphion: correct the unspecified color space

    [ Upstream commit 809060c8a357e020010dd8f797a5efd3c5432b13 ]

    in the E.2.1 of Rec. ITU-T H.264 (06/2019),
    0 of colour primaries is reserved, and 2 is unspecified.
    driver can map V4L2_COLORSPACE_LAST to 0,
    and map V4L2_COLORSPACE_DEFAULT to 2.

    v4l2_xfer_func and v4l2_ycbcr_encoding are similar case.

    Fixes: 3cd084519c6f ("media: amphion: add vpu v4l2 m2m support")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9534ddc03058a987feb8673b295d9beb09354ae3
```

```
Author: Ming Qian <ming.qian@nxp.com>
Date:   Thu Jan 12 10:47:02 2023 +0100

    media: imx-jpeg: Apply clk_bulk api instead of operating specific clk

    [ Upstream commit 61fe43dc9f454bc3caa99dbdd8f5fa3ba813981a ]

    using the api of clk_bulk can simplify the code.
    and the clock of the jpeg codec may be changed,
    the clk_bulk api can be compatible with the future change.

    Fixes: 4c2e5156d9fa ("media: imx-jpeg: Add pm-runtime support for imx-jpeg")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47cdece7b029cbbc98a5559274f496e008400511
Author: Nicolas Dufresne <nicolas.dufresne@collabora.com>
Date:   Fri Dec 23 19:16:47 2022 +0100

    media: hantro: Fix JPEG encoder ENUM_FRMSIZE on RK3399

    [ Upstream commit 29bd426764dee14a09e37700406f4a5920825fcc ]

    Since 79c987de8b354, enumerating framesize on format set with "MODE_NONE"
    (any raw formats) is reporting an invalid frmsize.

      Size: Stepwise 0x0 - 0x0 with step 0/0

    Before this change, the driver would return EINVAL, which is also invalid
    but worked in GStreamer. The original intent was not to implement it, hence
    the -ENOTTY return in this change. While drivers should implement
    ENUM_FRMSIZE for all formats and queues, this change is limited in scope to
    fix the regression.

    This fixes taking picture in Gnome Cheese software, or any software using
    GSteamer to encode JPEG with hardware acceleration.

    Fixes: 79c987de8b35 ("media: hantro: Use post processor scaling capacities")
    Reported-by: Robert Mader <robert.mader@collabora.com>
    Signed-off-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Reviewed-by: Benjamin Gaignard <benjamin.gaignard@collabora.com>
    Reviewed-by: Ezequiel Garcia <ezequiel@vanguardiasur.com.ar>
    Tested-by: Robert Mader <robert.mader@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3037ef82d60a38e80e2d98d2c3346e63b18d0c37
Author: Ming Qian <ming.qian@nxp.com>
Date:   Fri Dec 16 10:08:44 2022 +0100

    media: v4l2-jpeg: ignore the unknown APP14 marker

    [ Upstream commit 251c0ea6efd3c3ea0f8a55fdd96c749a98639bd3 ]

    The legal identifier of APP14 is "Adobe\0",
    but sometimes it may be
    "This is an unknown APP marker . Compliant decoders must ignore it."
    In this case, just ignore it.
    It won't affect the decode result.

    Fixes: b8035f7988a8 ("media: Add parsing for APP14 data segment in jpeg helpers")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
```

```
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc3f86a51595128f3ffb7e1cdf11093004cc2b3b
Author: Ming Qian <ming.qian@nxp.com>
Date:   Fri Dec 16 09:30:33 2022 +0100

    media: v4l2-jpeg: correct the skip count in jpeg_parse_app14_data

    [ Upstream commit 41959c4f973b837a12061b84d3a436fc64c73a30 ]

    The curr pointer has advanced 14 bytes in jpeg_parse_app14_data.
    1. jpeg_get_word_be(stream), it goes forward 2 bytes.
    2. jpeg_skip(stream, 11), it goes forward 11 bytes.
    3. jpeg_get_byte(stream), it goes forward 1 bytes.

    so the remain bytes of this segment should be (lp - 2 - 11 - 1),
    but not (lp - 2 - 11).

    if driver skip 1 extra bytes, the following parsing may go wrong.

    Fixes: b8035f7988a8 ("media: Add parsing for APP14 data segment in jpeg helpers")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 34026913597e2ea65eb5700d6d358cfe674efa42
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Thu Dec 15 17:40:08 2022 +0100

    media: platform: mtk-mdp3: fix Kconfig dependencies

    [ Upstream commit e3f7feb6d89311f369dd4ad903ea62e45328cdbe ]

    The new mdp3 driver uses 'select' to force-enable a couple of drivers
    it depends on. This is error-prone and likely to cause dependency
    loops as well as warnings like:

    WARNING: unmet direct dependencies detected for VIDEO_MEDIATEK_VPU
      Depends on [n]: MEDIA_SUPPORT [=m] && MEDIA_PLATFORM_SUPPORT [=y] &&
MEDIA_PLATFORM_DRIVERS [=y] && V4L_MEM2MEM_DRIVERS [=n] && VIDEO_DEV [=m] && (ARCH_MEDIATEK
[=y] || COMPILE_TEST [=y])
      Selected by [m]:
      - VIDEO_MEDIATEK_MDP3 [=m] && MEDIA_SUPPORT [=m] && MEDIA_PLATFORM_SUPPORT [=y] &&
MEDIA_PLATFORM_DRIVERS [=y] && (MTK_IOMMU [=m] || COMPILE_TEST [=y]) && VIDEO_DEV [=m] &&
(ARCH_MEDIATEK [=y] || COMPILE_TEST [=y]) && HAS_DMA [=y] && REMOTEPROC [=y]

    This specific warning was already addressed in a previous patch,
    but there are similar unnecessary 'select' statements, so turn those
    into 'depends on'. This also means the dependency on ARCH_MEDIATEK
    is redundant and can be dropped.

    Fixes: 61890ccaefaf ("media: platform: mtk-mdp3: add MediaTek MDP3 driver")
    Fixes: 9195a860ef0a ("media: platform: mtk-mdp3: remove unused VIDEO_MEDIATEK_VPU
config")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e057d12daed6986cb5850e63dadbd6980c86cc0c
Author: Arnd Bergmann <arnd@arndb.de>
```

```
Date:   Thu Dec 15 17:28:46 2022 +0100

    media: camss: csiphy-3ph: avoid undefined behavior

    [ Upstream commit 05fb9ace34b8645cb76f7e3a21b5c7b754329cae ]

    Marking a case of the switch statement as unreachable means the
    compiler treats it as undefined behavior, which is then caught by
    an objtool warning:

    drivers/media/platform/qcom/camss/camss-csiphy-3ph-1-0.o: warning: objtool:
csiphy_lanes_enable() falls through to next function csiphy_lanes_disable()

    Instead of simply continuing execution at a random place of the
    driver, print a warning and return from to the caller, which
    makes it possible to understand what happens and avoids the
    warning.

    Fixes: 53655d2a0ff2 ("media: camss: csiphy-3ph: add support for SM8250 CSI DPHY")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76ae28181c43b21c2af94b9798f4563c0af745ba
Author: Qiheng Lin <linqiheng@huawei.com>
Date:   Fri Dec 2 11:18:36 2022 +0100

    media: platform: mtk-mdp3: Fix return value check in mdp_probe()

    [ Upstream commit 1963689bed4d500236938d90c91cdd5e63c1eb28 ]

    In case of error, the function mtk_mutex_get()
    returns ERR_PTR() and never returns NULL. The NULL test in the
    return value check should be replaced with IS_ERR().
    And also fix the err_free_mutex case.

    Fixes: 61890ccaefaf ("media: platform: mtk-mdp3: add MediaTek MDP3 driver")
    Signed-off-by: Qiheng Lin <linqiheng@huawei.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec50986a1c23cf36aa74bc321a573d806da6f787
Author: Jai Luthra <j-luthra@ti.com>
Date:   Tue Jan 17 09:16:23 2023 +0100

    media: i2c: imx219: Fix binning for RAW8 capture

    [ Upstream commit ef86447e775fb1f2ced00d4c7fff2c0a1c63f165 ]

    2x2 binning works fine for RAW10 capture, but for RAW8 1232p mode it
    leads to corrupted frames [1][2].

    Using the special 2x2 analog binning mode fixes the issue, but causes
    artefacts for RAW10 1232p capture. So here we choose the binning mode
    depending upon the frame format selected.

    As both binning modes work fine for 480p RAW8 and RAW10 capture, it can
    share the same code path as 1232p for selecting binning mode.

    [1] https://forums.raspberrypi.com/viewtopic.php?t=332103
    [2] https://github.com/raspberrypi/libcamera-apps/issues/281
```

```
    Fixes: 22da1d56e982 ("media: i2c: imx219: Add support for RAW8 bit bayer format")
    Signed-off-by: Jai Luthra <j-luthra@ti.com>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6547093e3011773e83fb2590747a2c19a8f17b27
Author: Adam Ford <aford173@gmail.com>
Date:   Tue Dec 20 13:07:53 2022 +0100

    media: i2c: imx219: Split common registers from mode tables

    [ Upstream commit 8508455961d5a9e8907bcfd8dcd58f19d9b6ce47 ]

    There are four modes, and each mode has a table of registers.
    Some of the registers are common to all modes, so create new
    tables for these common registers to reduce duplicate code.

    Signed-off-by: Adam Ford <aford173@gmail.com>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Stable-dep-of: ef86447e775f ("media: i2c: imx219: Fix binning for RAW8 capture")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c86d760c1c6855a6131e78d0ddacc48c79324ac3
Author: Yuan Can <yuancan@huawei.com>
Date:   Thu Dec 8 09:06:25 2022 +0100

    media: i2c: ov772x: Fix memleak in ov772x_probe()

    [ Upstream commit 7485edb2b6ca5960205c0a49bedfd09bba30e521 ]

    A memory leak was reported when testing ov772x with bpf mock device:

    AssertionError: unreferenced object 0xffff888109afa7a8 (size 8):
      comm "python3", pid 279, jiffies 4294805921 (age 20.681s)
      hex dump (first 8 bytes):
        80 22 88 15 81 88 ff ff                          ."......
      backtrace:
        [<000000009990b438>] __kmalloc_node+0x44/0x1b0
        [<000000009e32f7d7>] kvmalloc_node+0x34/0x180
        [<00000000faf48134>] v4l2_ctrl_handler_init_class+0x11d/0x180 [videodev]
        [<00000000da376937>] ov772x_probe+0x1c3/0x68c [ov772x]
        [<000000003f0d225e>] i2c_device_probe+0x28d/0x680
        [<00000000e0b6db89>] really_probe+0x17c/0x3f0
        [<000000001b19fcee>] __driver_probe_device+0xe3/0x170
        [<0000000048370519>] driver_probe_device+0x49/0x120
        [<000000005ead07a0>] __device_attach_driver+0xf7/0x150
        [<0000000043f452b8>] bus_for_each_drv+0x114/0x180
        [<00000000358e5596>] __device_attach+0x1e5/0x2d0
        [<0000000043f83c5d>] bus_probe_device+0x126/0x140
        [<00000000ee0f3046>] device_add+0x810/0x1130
        [<00000000e0278184>] i2c_new_client_device+0x359/0x4f0
        [<0000000070baf34f>] of_i2c_register_device+0xf1/0x110
        [<00000000a9f2159d>] of_i2c_notify+0x100/0x160
    unreferenced object 0xffff888119825c00 (size 256):
      comm "python3", pid 279, jiffies 4294805921 (age 20.681s)
      hex dump (first 32 bytes):
        00 b4 a5 17 81 88 ff ff 00 5e 82 19 81 88 ff ff  .........^......
        10 5c 82 19 81 88 ff ff 10 5c 82 19 81 88 ff ff  .\.......\......
      backtrace:
        [<000000009990b438>] __kmalloc_node+0x44/0x1b0
```

```
          [<000000009e32f7d7>] kvmalloc_node+0x34/0x180
          [<0000000073d88e0b>] v4l2_ctrl_new.cold+0x19b/0x86f [videodev]
          [<00000000b1f576fb>] v4l2_ctrl_new_std+0x16f/0x210 [videodev]
          [<00000000caf7ac99>] ov772x_probe+0x1fa/0x68c [ov772x]
          [<000000003f0d225e>] i2c_device_probe+0x28d/0x680
          [<00000000e0b6db89>] really_probe+0x17c/0x3f0
          [<000000001b19fcee>] __driver_probe_device+0xe3/0x170
          [<0000000048370519>] driver_probe_device+0x49/0x120
          [<000000005ead07a0>] __device_attach_driver+0xf7/0x150
          [<0000000043f452b8>] bus_for_each_drv+0x114/0x180
          [<00000000358e5596>] __device_attach+0x1e5/0x2d0
          [<0000000043f83c5d>] bus_probe_device+0x126/0x140
          [<00000000ee0f3046>] device_add+0x810/0x1130
          [<00000000e0278184>] i2c_new_client_device+0x359/0x4f0
          [<0000000070baf34f>] of_i2c_register_device+0xf1/0x110

      The reason is that if priv->hdl.error is set, ov772x_probe() jumps to the
      error_mutex_destroy without doing v4l2_ctrl_handler_free(), and all
      resources allocated in v4l2_ctrl_handler_init() and v4l2_ctrl_new_std()
      are leaked.

      Fixes: 1112babde214 ("media: i2c: Copy ov772x soc_camera sensor driver")
      Signed-off-by: Yuan Can <yuancan@huawei.com>
      Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
      Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
      Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c955f7f62c2664882a5207e5cf05c386be494141
  Author: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
  Date:   Mon Dec 12 14:25:04 2022 +0100

      media: mc: Get media_device directly from pad

      [ Upstream commit a967a3a788028f541e4db54beabcebc3648997db ]

      Various functions access the media_device from a pad by going through
      the entity the pad belongs to. Remove the level of indirection and get
      the media_device from the pad directly.

      Fixes: 9e3576a1ae2b ("media: mc: convert pipeline funcs to take media_pad")
      Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
      Reviewed-by: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
      Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
      Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 243553acabe139fff628cf3fd63887cb97d66dc7
  Author: Jai Luthra <j-luthra@ti.com>
  Date:   Tue Jan 3 13:27:36 2023 +0100

      media: ov5640: Handle delays when no reset_gpio set

      [ Upstream commit d7ff69139908842adf824be4f50c7e9ac5886c04 ]

      Some module manufacturers [1][2] don't expose the RESETB and PWDN pins
      of the sensor directly through the 15-pin FFC connector. Instead wiring
      ~PWDN gpio to the sensor pins with appropriate delays.

      In such cases, reset_gpio will not be available to the driver, but it
      will still be toggled when the sensor is powered on, and thus we should
      still honor the wait time of >= 5ms + 1ms + 20ms (see figure 2-3 in [3])
      before attempting any i/o operations over SCCB.

      Also, rename the function to ov5640_powerup_sequence to better match the
```

```
    datasheet (section 2.7).

    [1] https://digilent.com/reference/_media/reference/add-ons/pcam-5c/pcam_5c_sch.pdf
    [2] https://www.alinx.com/public/upload/file/AN5641_User_Manual.pdf
    [3] https://cdn.sparkfun.com/datasheets/Sensors/LightImaging/OV5640_datasheet.pdf

    Fixes: 19a81c1426c1 ("[media] add Omnivision OV5640 sensor driver")
    Reported-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Jai Luthra <j-luthra@ti.com>
    Reviewed-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c263c44035bd4ec614c6fb6dac21c27b3c13181
Author: Jai Luthra <j-luthra@ti.com>
Date:   Tue Jan 3 13:27:35 2023 +0100

    media: ov5640: Fix soft reset sequence and timings

    [ Upstream commit decea0a98b7ac04536c7d659f74783e8d67a06c0 ]

    Move the register-based reset out of the init_setting[] and into the
    powerup_sequence function. The sensor is power cycled and reset using
    the gpio pins so the soft reset is not always necessary.

    This also ensures that soft reset honors the timing sequence
    from the datasheet [1].

    [1] https://cdn.sparkfun.com/datasheets/Sensors/LightImaging/OV5640_datasheet.pdf

    Fixes: 19a81c1426c1 ("[media] add Omnivision OV5640 sensor driver")
    Reported-by: Nishanth Menon <nm@ti.com>
    Suggested-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Signed-off-by: Jai Luthra <j-luthra@ti.com>
    Reviewed-by: Jacopo Mondi <jacopo.mondi@ideasonaboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d67fc82ff6d0c606d0f47afc547185b794377f0a
Author: Marco Felsch <m.felsch@pengutronix.de>
Date:   Fri Dec 16 11:35:45 2022 +0100

    media: i2c: tc358746: fix possible endianness issue

    [ Upstream commit 5ad2e46030ad97de7fdbdaf63bb1af45c7caf3dd ]

    Using the u64 v4l2_dbg_register.val directly can lead to unexpected
    results depending on machine endianness. Fix this by using a local
    variable which is assigned afterwards. Since tc358746_read() will init
    the val variable to 0 we can assing it without checking the return value
    first.

    Addresses-Coverity-ID: 1527256 ("Integer handling issues")

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Fixes: 80a21da36051 ("media: tc358746: add Toshiba TC358746 Parallel to CSI-2 bridge
driver")
    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04ec166f798a48fda4e14cddb00cd540fbc843a6
```

```
Author: Marco Felsch <m.felsch@pengutronix.de>
Date:   Fri Dec 16 11:35:44 2022 +0100

    media: i2c: tc358746: fix ignoring read error in g_register callback

    [ Upstream commit 9d33802c8bcf96c4099ffea4f392afa52897e556 ]

    Currently we ignore the return value of tc358746_read() and return
    alawys return 0 which is wrong. Fix this by returning the actual return
    value of the read operation which is either 0 on success or an error
    value.

    Addresses-Coverity-ID: 1527254 ("Error handling issues")

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Fixes: 80a21da36051 ("media: tc358746: add Toshiba TC358746 Parallel to CSI-2 bridge
driver")
    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3905ae4718c3335993eb84bdaa6ef0935ccab1e
Author: Marco Felsch <m.felsch@pengutronix.de>
Date:   Fri Dec 16 11:35:43 2022 +0100

    media: i2c: tc358746: fix missing return assignment

    [ Upstream commit 0605081142070a41de8f1deb8fdaeb8677e97741 ]

    It was intended to return an error if tc358746_update_bits() call fail.
    Fix this by storing the return code.

    Addresses-Coverity-ID: 1527252 ("Control flow issues")

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Fixes: 80a21da36051 ("media: tc358746: add Toshiba TC358746 Parallel to CSI-2 bridge
driver")
    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a36a6be694df87d019663863b922913947b42af
Author: Shang XiaoJing <shangxiaojing@huawei.com>
Date:   Thu Dec 8 08:59:38 2022 +0100

    media: ov5675: Fix memleak in ov5675_init_controls()

    [ Upstream commit dd74ed6c213003533e3abf4c204374ef01d86978 ]

    There is a kmemleak when testing the media/i2c/ov5675.c with bpf mock
    device:

    AssertionError: unreferenced object 0xffff888107362160 (size 16):
      comm "python3", pid 277, jiffies 4294832798 (age 20.722s)
      hex dump (first 16 bytes):
        00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
      backtrace:
        [<00000000abe7d67c>] __kmalloc_node+0x44/0x1b0
        [<000000008a725aac>] kvmalloc_node+0x34/0x180
        [<000000009a53cd11>] v4l2_ctrl_handler_init_class+0x11d/0x180
    [videodev]
        [<0000000055b46db0>] ov5675_probe+0x38b/0x897 [ov5675]
        [<00000000153d886c>] i2c_device_probe+0x28d/0x680
```

```
          [<000000004afb7e8f>] really_probe+0x17c/0x3f0
          [<00000000ff2f18e4>] __driver_probe_device+0xe3/0x170
          [<000000000a001029>] driver_probe_device+0x49/0x120
          [<00000000e39743c7>] __device_attach_driver+0xf7/0x150
          [<00000000d32fd070>] bus_for_each_drv+0x114/0x180
          [<000000009083ac41>] __device_attach+0x1e5/0x2d0
          [<0000000015b4a830>] bus_probe_device+0x126/0x140
          [<000000007813deaf>] device_add+0x810/0x1130
          [<000000007becb867>] i2c_new_client_device+0x386/0x540
          [<000000007f9cf4b4>] of_i2c_register_device+0xf1/0x110
          [<00000000ebfdd032>] of_i2c_notify+0xfc/0x1f0

      ov5675_init_controls() won't clean all the allocated resources in fail
      path, which may causes the memleaks. Add v4l2_ctrl_handler_free() to
      prevent memleak.

      Fixes: bf27502b1f3b ("media: ov5675: Add support for OV5675 sensor")
      Signed-off-by: Shang XiaoJing <shangxiaojing@huawei.com>
      Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
      Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 7c405ee63447f14eefcfe12a18aa749abbd596ea
  Author: Shang XiaoJing <shangxiaojing@huawei.com>
  Date:   Thu Dec 8 08:59:37 2022 +0100

      media: ov2740: Fix memleak in ov2740_init_controls()

      [ Upstream commit 2d899592ed7829d0d5140853bac4d58742a6b8af ]

      There is a kmemleak when testing the media/i2c/ov2740.c with bpf mock
      device:

      unreferenced object 0xffff8881090e19e0 (size 16):
        comm "51-i2c-ov2740", pid 278, jiffies 4294781584 (age 23.613s)
        hex dump (first 16 bytes):
          00 f3 7c 0b 81 88 ff ff 80 75 6a 09 81 88 ff ff  ..|......uj.....
        backtrace:
          [<000000004e9fad8f>] __kmalloc_node+0x44/0x1b0
          [<0000000039c802f4>] kvmalloc_node+0x34/0x180
          [<000000009b8b5c63>] v4l2_ctrl_handler_init_class+0x11d/0x180
      [videodev]
          [<0000000038644056>] ov2740_probe+0x37d/0x84f [ov2740]
          [<0000000092489f59>] i2c_device_probe+0x28d/0x680
          [<000000001038babe>] really_probe+0x17c/0x3f0
          [<0000000098c7af1c>] __driver_probe_device+0xe3/0x170
          [<00000000e1b3dc24>] device_driver_attach+0x34/0x80
          [<000000005a04a34d>] bind_store+0x10b/0x1a0
          [<00000000ce25d4f2>] drv_attr_store+0x49/0x70
          [<000000007d9f4e9a>] sysfs_kf_write+0x8c/0xb0
          [<00000000be6cff0f>] kernfs_fop_write_iter+0x216/0x2e0
          [<0000000031ddb40a>] vfs_write+0x658/0x810
          [<0000000041beecdd>] ksys_write+0xd6/0x1b0
          [<0000000023755840>] do_syscall_64+0x38/0x90
          [<00000000b2cc2da2>] entry_SYSCALL_64_after_hwframe+0x63/0xcd

      ov2740_init_controls() won't clean all the allocated resources in fail
      path, which may causes the memleaks. Add v4l2_ctrl_handler_free() to
      prevent memleak.

      Fixes: 866edc895171 ("media: i2c: Add ov2740 image sensor driver")
      Signed-off-by: Shang XiaoJing <shangxiaojing@huawei.com>
      Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
      Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5e31213fa017c20ccc989033a5f4a626473aa2ca
Author: Shang XiaoJing <shangxiaojing@huawei.com>
Date:   Tue Dec 6 14:05:55 2022 +0100

    media: max9286: Fix memleak in max9286_v4l2_register()

    [ Upstream commit 8636c5fc7658c7c6299fb8b352d24ea4b9ba99e2 ]

    There is a kmemleak when testing the media/i2c/max9286.c with bpf mock
    device:

    kmemleak: 5 new suspected memory leaks (see /sys/kernel/debug/kmemleak)

    unreferenced object 0xffff88810defc400 (size 256):
      comm "python3", pid 278, jiffies 4294737563 (age 31.978s)
      hex dump (first 32 bytes):
        28 06 a7 0a 81 88 ff ff 00 fe 22 12 81 88 ff ff  (.........".....
        10 c4 ef 0d 81 88 ff ff 10 c4 ef 0d 81 88 ff ff  ................
      backtrace:
        [<00000000191de6a7>] __kmalloc_node+0x44/0x1b0
        [<000000002f4912b7>] kvmalloc_node+0x34/0x180
        [<0000000057dc4cae>] v4l2_ctrl_new+0x325/0x10f0 [videodev]
        [<0000000026030272>] v4l2_ctrl_new_std+0x16f/0x210 [videodev]
        [<00000000f0d9ea2f>] max9286_probe+0x76e/0xbff [max9286]
        [<00000000ea8f6455>] i2c_device_probe+0x28d/0x680
        [<0000000087529af3>] really_probe+0x17c/0x3f0
        [<00000000b08be526>] __driver_probe_device+0xe3/0x170
        [<000000004382edea>] driver_probe_device+0x49/0x120
        [<000000007bde528a>] __device_attach_driver+0xf7/0x150
        [<000000009f9c6ab4>] bus_for_each_drv+0x114/0x180
        [<00000000c8aaf588>] __device_attach+0x1e5/0x2d0
        [<0000000041cc06b9>] bus_probe_device+0x126/0x140
        [<000000002309860d>] device_add+0x810/0x1130
        [<000000002827bf98>] i2c_new_client_device+0x359/0x4f0
        [<00000000593bdc85>] of_i2c_register_device+0xf1/0x110

    max9286_v4l2_register() calls v4l2_ctrl_new_std(), but won't free the
    created v412_ctrl when fwnode_graph_get_endpoint_by_id() failed, which
    causes the memleak. Call v4l2_ctrl_handler_free() to free the v412_ctrl.

    Fixes: 66d8c9d2422d ("media: i2c: Add MAX9286 driver")
    Signed-off-by: Shang XiaoJing <shangxiaojing@huawei.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6b9fec0d900a3fdf7a70ae40d455570dfcb5fe0a
Author: Bastian Germann <bage@linutronix.de>
Date:   Wed May 26 01:01:37 2021 +0200

    builddeb: clean generated package content

    [ Upstream commit c9f9cf2560e40b62015c6c4a04be60f55ce5240e ]

    For each binary Debian package, a directory with the package name is
    created in the debian directory. Correct the generated file matches in the
    package's clean target, which were renamed without adjusting the target.

    Fixes: 1694e94e4f46 ("builddeb: match temporary directory name to the package name")
    Signed-off-by: Bastian Germann <bage@linutronix.de>
    Signed-off-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 05bd9185a7602dd0fcfde2888cf08f3df9f3a5a1
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed Jan 11 20:05:07 2023 -0700

    s390/vdso: Drop '-shared' from KBUILD_CFLAGS_64

    [ Upstream commit fd8589dce8107e2ce62e92f76089654462dd67b4 ]

    When clang's -Qunused-arguments is dropped from KBUILD_CPPFLAGS, it
    points out that there is a linking phase flag added to CFLAGS, which
    will only be used for compiling

      clang-16: error: argument unused during compilation: '-shared' [-Werror,-Wunused-
command-line-argument]

    '-shared' is already present in ldflags-y so it can just be dropped.

    Fixes: 2b2a25845d53 ("s390/vdso: Use $(LD) instead of $(CC) to link vDSO")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Acked-by: Heiko Carstens <hca@linux.ibm.com>
    Reviewed-by: Sven Schnelle <svens@linux.ibm.com>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Anders Roxell <anders.roxell@linaro.org>
    Signed-off-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e49b665f5122b47496ba6c2bcedc1c88e578343c
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed Jan 11 20:05:02 2023 -0700

    powerpc: Remove linker flag from KBUILD_AFLAGS

    [ Upstream commit 31f48f16264bc70962fb3e7ec62da64d0a2ba04a ]

    When clang's -Qunused-arguments is dropped from KBUILD_CPPFLAGS, it
    points out that KBUILD_AFLAGS contains a linker flag, which will be
    unused:

      clang: error: -Wl,-a32: 'linker' input unused [-Werror,-Wunused-command-line-argument]

    This was likely supposed to be '-Wa,-a$(BITS)'. However, this change is
    unnecessary, as all supported versions of clang and gcc will pass '-a64'
    or '-a32' to GNU as based on the value of '-m'; the behavior of the
    latest stable release of the oldest supported major version of each
    compiler is shown below and each compiler's latest release exhibits the
    same behavior (GCC 12.2.0 and Clang 15.0.6).

      $ powerpc64-linux-gcc --version | head -1
      powerpc64-linux-gcc (GCC) 5.5.0

      $ powerpc64-linux-gcc -m64 -### -x assembler-with-cpp -c -o /dev/null /dev/null &|
grep 'as '
      .../as -a64 -mppc64 -many -mbig -o /dev/null /tmp/cctwuBzZ.s

      $ powerpc64-linux-gcc -m32 -### -x assembler-with-cpp -c -o /dev/null /dev/null &|
grep 'as '
      .../as -a32 -mppc -many -mbig -o /dev/null /tmp/ccaZP4mF.sg

      $ clang --version | head -1
      Ubuntu clang version 11.1.0-++20211011094159+1fdec59bffc1-1~exp1~20211011214622.5

      $ clang --target=powerpc64-linux-gnu -fno-integrated-as -m64 -### \
        -x assembler-with-cpp -c -o /dev/null /dev/null &| grep gnu-as
       "/usr/bin/powerpc64-linux-gnu-as" "-a64" "-mppc64" "-many" "-o" "/dev/null"
"/tmp/null-80267c.s"
```

```
      $ clang --target=powerpc64-linux-gnu -fno-integrated-as -m64 -### \
        -x assembler-with-cpp -c -o /dev/null /dev/null &| grep gnu-as
       "/usr/bin/powerpc64-linux-gnu-as" "-a32" "-mppc" "-many" "-o" "/dev/null" "/tmp/null-
ab8f8d.s"

    Remove this flag altogether to avoid future issues.

    Fixes: 1421dc6d4829 ("powerpc/kbuild: Use flags variables rather than overriding
LD/CC/AS")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Anders Roxell <anders.roxell@linaro.org>
    Acked-by: Michael Ellerman <mpe@ellerman.id.au>
    Signed-off-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16f7e93d9aae33ec2e73cefd403c3c49c1381511
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Jan 4 10:39:21 2023 +0100

    media: imx: imx7-media-csi: fix missing clk_disable_unprepare() in imx7_csi_init()

    [ Upstream commit cea606d9e996a77eed57fc60709e0728341450e3 ]

    Add missing clk_disable_unprepare(), if imx7_csi_dma_setup() fails
    in imx7_csi_init().

    Fixes: ff43ca911978 ("media: imx: imx7-media-csi: Move CSI configuration before source
start")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Reviewed-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e616bb87cf60cdba579297c0971556345e5b0eed
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Jan 4 09:55:37 2023 +0100

    media: platform: ti: Add missing check for devm_regulator_get

    [ Upstream commit da8e05f84a11c3cc3b0ba0a3c62d20e358002d99 ]

    Add check for the return value of devm_regulator_get since it may return
    error pointer.

    Fixes: 448de7e7850b ("[media] omap3isp: OMAP3 ISP core")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bdc64b8de48059cc78ca3dd73c59274eb47224be
Author: Gaosheng Cui <cuigaosheng1@huawei.com>
Date:   Tue Nov 29 12:01:59 2022 +0100

    media: ti: cal: fix possible memory leak in cal_ctx_create()

    [ Upstream commit 7acd650a0484d92985a0d6d867d980c6dd019885 ]

    The memory of ctx is allocated in cal_ctx_create(), but it will
    not be freed when cal_ctx_v4l2_init() fails, so add kfree() when
    cal_ctx_v4l2_init() fails to fix it.
```

```
    Fixes: d68a94e98a89 ("media: ti-vpe: cal: Split video device initialization and
registration")
    Signed-off-by: Gaosheng Cui <cuigaosheng1@huawei.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e36e77b104f1c13b9993d8227d2f628e3c259099
Author: Sibi Sankar <quic_sibis@quicinc.com>
Date:   Tue Jan 17 14:28:35 2023 +0530

    remoteproc: qcom_q6v5_mss: Use a carveout to authenticate modem headers

    [ Upstream commit 57f72170a2b2a362c35bb9407fc844eac5afdec1 ]

    Any access to the dynamically allocated metadata region by the application
    processor after assigning it to the remote Q6 will result in a XPU
    violation. Fix this by replacing the dynamically allocated memory region
    with a no-map carveout and unmap the modem metadata memory region before
    passing control to the remote Q6.

    Reported-and-tested-by: Amit Pundir <amit.pundir@linaro.org>
    Fixes: 6c5a9dc2481b ("remoteproc: qcom: Make secure world call for mem ownership
switch")
    Signed-off-by: Sibi Sankar <quic_sibis@quicinc.com>
    Reviewed-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230117085840.32356-7-quic_sibis@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57551c3ff8c4d5197fb18b8e8189df1e31bc41ef
Author: Christoph Hellwig <hch@lst.de>
Date:   Tue Jan 17 14:28:34 2023 +0530

    Revert "remoteproc: qcom_q6v5_mss: map/unmap metadata region before/after use"

    [ Upstream commit a899d542b687c9b04ccbd9eefabc829ba5fef791 ]

    This reverts commit fc156629b23a21181e473e60341e3a78af25a1d4.

    This commit manages to do three API violations at once:

     - dereference the return value of dma_alloc_attrs with the
       DMA_ATTR_NO_KERNEL_MAPPING mapping, which is clearly forbidden and
       will do the wrong thing on various dma mapping implementations.  The
       fact that dma-direct uses a struct page as a cookie is an undocumented
       implementation detail
     - include dma-map-ops.h and use pgprot_dmacoherent despite a clear
       comment documenting that this is not acceptable
     - use of the VM_DMA_COHERENT for something that is not the dma-mapping
       code
     - use of VM_FLUSH_RESET_PERMS for vmap, while it is only supported for
       vmalloc

    Acked-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sibi Sankar <quic_sibis@quicinc.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230117085840.32356-6-quic_sibis@quicinc.com
    Stable-dep-of: 57f72170a2b2 ("remoteproc: qcom_q6v5_mss: Use a carveout to authenticate
modem headers")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0ef9594936d1f078e8599a1cf683b052df2bec00
```

```
Author: Patrick Kelsey <pat.kelsey@cornelisnetworks.com>
Date:   Thu Feb 16 11:56:28 2023 -0500

    IB/hfi1: Fix sdma.h tx->num_descs off-by-one errors

    [ Upstream commit fd8958efe8779d3db19c9124fce593ce681ac709 ]

    Fix three sources of error involving struct sdma_txreq.num_descs.

    When _extend_sdma_tx_descs() extends the descriptor array, it uses the
    value of tx->num_descs to determine how many existing entries from the
    tx's original, internal descriptor array to copy to the newly allocated
    one.  As this value was incremented before the call, the copy loop will
    access one entry past the internal descriptor array, copying its contents
    into the corresponding slot in the new array.

    If the call to _extend_sdma_tx_descs() fails, _pad_smda_tx_descs() then
    invokes __sdma_tx_clean() which uses the value of tx->num_desc to drive a
    loop that unmaps all descriptor entries in use.  As this value was
    incremented before the call, the unmap loop will invoke sdma_unmap_desc()
    on a descriptor entry whose contents consist of whatever random data was
    copied into it during (1), leading to cascading further calls into the
    kernel and driver using arbitrary data.

    _sdma_close_tx() was using tx->num_descs instead of tx->num_descs - 1.

    Fix all of the above by:
    - Only increment .num_descs after .descp is extended.
    - Use .num_descs - 1 instead of .num_descs for last .descp entry.

    Fixes: f4d26d81ad7f ("staging/rdma/hfi1: Add coalescing support for SDMA TX
descriptors")
    Link: https://lore.kernel.org/r/167656658879.2223096.10026561343022570690.stgit@awfm-
02.cornelisnetworks.com
    Signed-off-by: Brendan Cunningham <bcunningham@cornelisnetworks.com>
    Signed-off-by: Patrick Kelsey <pat.kelsey@cornelisnetworks.com>
    Signed-off-by: Dennis Dalessandro <dennis.dalessandro@cornelisnetworks.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aba23d0d4d523d8171fef02eb12a05182fa90f37
Author: Patrick Kelsey <pat.kelsey@cornelisnetworks.com>
Date:   Thu Feb 16 11:56:23 2023 -0500

    IB/hfi1: Fix math bugs in hfi1_can_pin_pages()

    [ Upstream commit a0d198f79a8d033bd46605b779859193649f1f99 ]

    Fix arithmetic and logic errors in hfi1_can_pin_pages() that  would allow
    hfi1 to attempt pinning pages in cases where it should not because of
    resource limits or lack of required capability.

    Fixes: 2c97ce4f3c29 ("IB/hfi1: Add pin query function")
    Link: https://lore.kernel.org/r/167656658362.2223096.10954762619837718026.stgit@awfm-
02.cornelisnetworks.com
    Signed-off-by: Brendan Cunningham <bcunningham@cornelisnetworks.com>
    Signed-off-by: Patrick Kelsey <pat.kelsey@cornelisnetworks.com>
    Signed-off-by: Dennis Dalessandro <dennis.dalessandro@cornelisnetworks.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c171386d3c3f8f998ac79c8377784f6274f202ae
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Tue Feb 14 01:10:54 2023 -0600
```

```
    RDMA/rxe: Fix missing memory barriers in rxe_queue.h

    [ Upstream commit a77a52385e9a761f896a88a4162e69fb7ccafe3f ]

    An earlier patch which introduced smp_load_acquire/smp_store_release
    into rxe_queue.h incorrectly assumed that surrounding spin-locks in
    rxe_verbs.c around queue updates for kernel ulps was sufficient to
    protect the passing of data through the queues between the ulp and
    the rxe tasklets. But this was incorrect. The typical sequence was

            ulp                             rxe requester tasklet
            ------------------------        --------------------
            spin_lock_irqsave()             wqe = queue_head(queue)
            if (!queue_full(q)) {           if (!wqe)
                    spin_unlock_irqrestore          return;
                    return -ENOMEM
            }                               <process wqe>
            wqe = queue_producer_addr(q)
            <fill in wqe>                   queue_advance_consumer(queue)
            queue_advance_producer(q)
            spin_unlock_irqrestore()

    queue_head() calls queue_empty() which calls smp_load_acquire()
    For user space apps queue_advance_producer() calls smp_store_release()
    so that there is a memory barrier between the producer and the
    consumer but for kernel ulps queue_advance_produce() just incremented
    the producer index because the lock function is a release function.
    But to work the barrier has to come between filling in the wqe and
    updating the producer index. This patch adds the missing barriers.
    It also changes the enum names for the ulp queue types to
            QUEUE_TYPE_FROM/TO_ULP instead of QUEUE_TYPE_TO/FROM_DRIVER
    which is very ambiguous. This bug is suspected as the cause of very
    rare lockups in a very high scale storage application. It is a bug
    in any case and should be corrected.

    Fixes: 0a67c46d2e99 ("RDMA/rxe: Protect user space index loads/stores")
    Link: https://lore.kernel.org/r/20230214071053.5395-1-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc5bd6d731d25e66b3a5a2ba1ca33e0ba971637e
Author: Long Li <longli@microsoft.com>
Date:   Wed Feb 15 16:32:02 2023 -0800

    RDMA/mana_ib: Fix a bug when the PF indicates more entries for registering memory on
first packet

    [ Upstream commit 89d42b8c85b4c67d310c5ccaf491acbf71a260c3 ]

    When registering memory in a large chunk that doesn't fit into a single PF
    message, the PF may return GDMA_STATUS_MORE_ENTRIES on the first message if
    there are more messages needed for registering more chunks.

    Fix the VF to make it process the correct return code.

    Fixes: 0266a177631d ("RDMA/mana_ib: Add a driver for Microsoft Azure Network Adapter")
    Link: https://lore.kernel.org/r/1676507522-21018-1-git-send-email-
longli@linuxonhyperv.com
    Signed-off-by: Long Li <longli@microsoft.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32c790560fdffd609943f6db55844ab889e9cef7
Author: Bob Pearson <rpearsonhpe@gmail.com>
```

```
Date:   Wed Feb 1 22:42:41 2023 -0600

    Subject: RDMA/rxe: Handle zero length rdma

    [ Upstream commit 5ff31dfcd6d23f9c1bd5dd1a2c648ba499659357 ]

    Currently the rxe driver does not handle all cases of zero length rdma
    operations correctly. The client does not have to provide an rkey for zero
    length RDMA read or write operations so the rkey provided may be invalid
    and should not be used to lookup an mr.

    This patch corrects the driver to ignore the provided rkey if the reth
    length is zero for read or write operations and make sure to set the mr to
    NULL. In read_reply() if length is zero rxe_recheck_mr() is not
    called. Warnings are added in the routines in rxe_mr.c to catch NULL MRs
    when the length is non-zero.

    Fixes: 8700e3e7c485 ("Soft RoCE driver")
    Link: https://lore.kernel.org/r/20230202044240.6304-1-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Reviewed-by: Daisuke Matsuda <matsuda-daisuke@fujitsu.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e443760b8b7b1e6723f4408afa056b2bc4fea12
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:37 2023 -0600

    RDMA/rxe: Replace rxe_map and rxe_phys_buf by xarray

    [ Upstream commit 592627ccbdff0ec6fff00fc761142a76db750dd4 ]

    Replace struct rxe-phys_buf and struct rxe_map by struct xarray
    in rxe_verbs.h. This allows using rcu locking on reads for
    the memory maps stored in each mr.

    This is based off of a sketch of a patch from Jason Gunthorpe in the
    link below. Some changes were needed to make this work. It applies
    cleanly to the current for-next and passes the pyverbs, perftest
    and the same blktests test cases which run today.

    Link: https://lore.kernel.org/r/20230119235936.19728-7-rpearsonhpe@gmail.com
    Link: https://lore.kernel.org/linux-rdma/Y3gvZr6%2FNCii9Avy@nvidia.com/
    Co-developed-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5264698f71e7d4250ba93a7532e09a8e463ede6f
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:36 2023 -0600

    RDMA/rxe: Cleanup page variables in rxe_mr.c

    [ Upstream commit 325a7eb85199ec9c5b5a7af812f43ea16b735569 ]

    Cleanup usage of mr->page_shift and mr->page_mask and introduce
    an extractor for mr->ibmr.page_size. Normal usage in the kernel
    has page_mask masking out offset in page rather than masking out
    the page number. The rxe driver had reversed that which was confusing.
    Implicitly there can be a per mr page_size which was not uniformly
    supported.

    Link: https://lore.kernel.org/r/20230119235936.19728-6-rpearsonhpe@gmail.com
```

```
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9e7cf79594b9fe035508911d6c3726691652371c
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:35 2023 -0600

    RDMA-rxe: Isolate mr code from atomic_write_reply()

    [ Upstream commit d8bdb0ebca086b5845d782e800ad2bf2a7eb4877 ]

    Isolate mr specific code from atomic_write_reply() in rxe_resp.c into
    a subroutine rxe_mr_do_atomic_write() in rxe_mr.c.
    Check length for atomic write operation.
    Make iova_to_vaddr() static.

    Link: https://lore.kernel.org/r/20230119235936.19728-5-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97bf601b73db4b8f1ec5b796e587eb4987e2e3a2
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:34 2023 -0600

    RDMA-rxe: Isolate mr code from atomic_reply()

    [ Upstream commit f04d5b3d916c61752ac2c2adea5dfe78f8e12f78 ]

    Isolate mr specific code from atomic_reply() in rxe_resp.c into
    a subroutine rxe_mr_do_atomic_op() in rxe_mr.c.
    Minor cleanups to rxe_check_range() and iova_to_vaddr().
    Move enum resp_state to rxe.h

    Link: https://lore.kernel.org/r/20230119235936.19728-4-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c54b725fb1a4d22a4daa2a5b08ac7545e9b5393a
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:33 2023 -0600

    RDMA/rxe: Move rxe_map_mr_sg to rxe_mr.c

    [ Upstream commit db4729a5251992ed535da09c0fcf9b590ac7fe6c ]

    Move rxe_map_mr_sg() to rxe_mr.c where it makes a little more sense.

    Link: https://lore.kernel.org/r/20230119235936.19728-3-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8ef37e094dde16d15ff7db8255720cd423a90750
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Thu Jan 19 17:59:32 2023 -0600

    RDMA/rxe: Cleanup mr_check_range
```

```
    [ Upstream commit ade58da2a73de1b65616e4b1080dc078d1ce0b5d ]

    Remove blank lines and replace EFAULT by EINVAL when an invalid
    mr type is used.

    Link: https://lore.kernel.org/r/20230119235936.19728-2-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: 5ff31dfcd6d2 ("Subject: RDMA/rxe: Handle zero length rdma")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b1cdbacebb77befaabdb95fe67641679f24f81dc
Author: Tina Zhang <tina.zhang@intel.com>
Date:   Thu Feb 16 21:08:16 2023 +0800

    iommu/vt-d: Allow to use flush-queue when first level is default

    [ Upstream commit 257ec290741924f8df678927d0dfecb1deebb9c5 ]

    Commit 29b32839725f ("iommu/vt-d: Do not use flush-queue when caching-mode
    is on") forced default domains to be strict mode as long as IOMMU
    caching-mode is flagged. The reason for doing this is that when vIOMMU
    uses VT-d caching mode to synchronize shadowing page tables, the strict
    mode shows better performance.

    However, this optimization is orthogonal to the first-level page table
    because the Intel VT-d architecture does not define the caching mode of
    the first-level page table. Refer to VT-d spec, section 6.1, "When the
    CM field is reported as Set, any software updates to remapping
    structures other than first-stage mapping (including updates to not-
    present entries or present entries whose programming resulted in
    translation faults) requires explicit invalidation of the caches."
    Exclude the first-level page table from this optimization.

    Generally using first-stage translation in vIOMMU implies nested
    translation enabled in the physical IOMMU. In this case the first-stage
    page table is wholly captured by the guest. The vIOMMU only needs to
    transfer the cache invalidations on vIOMMU to the physical IOMMU.
    Forcing the default domain to strict mode will cause more frequent
    cache invalidations, resulting in performance degradation. In a real
    performance benchmark test measured by iperf receive, the performance
    result on Sapphire Rapids 100Gb NIC shows:
    w/ this fix ~51 Gbits/s, w/o this fix ~39.3 Gbits/s.

    Theoretically a first-stage IOMMU page table can still be shadowed
    in absence of the caching mode, e.g. with host write-protecting guest
    IOMMU page table to synchronize changed PTEs with the physical
    IOMMU page table. In this case the shadowing overhead is decoupled
    from emulating IOTLB invalidation then the overhead of the latter part
    is solely decided by the frequency of IOTLB invalidations. Hence
    allowing guest default dma domain to be lazy can also benefit the
    overall performance by reducing the total VM-exit numbers.

    Fixes: 29b32839725f ("iommu/vt-d: Do not use flush-queue when caching-mode is on")
    Reported-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Suggested-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Signed-off-by: Tina Zhang <tina.zhang@intel.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Link: https://lore.kernel.org/r/20230214025618.2292889-1-tina.zhang@intel.com
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a405f3f5e71e966c367e4d98aac1c139fd30d0c5
Author: Lu Baolu <baolu.lu@linux.intel.com>
```

```
Date:   Thu Feb 16 21:08:13 2023 +0800

    iommu/vt-d: Fix error handling in sva enable/disable paths

    [ Upstream commit 60b1daa3b168fbc648ae2ad28a84759223e49e18 ]

    Roll back all previous actions in error paths of intel_iommu_enable_sva()
    and intel_iommu_disable_sva().

    Fixes: d5b9e4bfe0d8 ("iommu/vt-d: Report prq to io-pgfault framework")
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Link: https://lore.kernel.org/r/20230208051559.700109-1-baolu.lu@linux.intel.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5bba023b1241c7af5d40447503a68de282ad5190
Author: Eric Pilmore <epilmore@gigaio.com>
Date:   Thu Feb 9 23:51:43 2023 -0800

    dmaengine: ptdma: check for null desc before calling pt_cmd_callback

    [ Upstream commit 928469986171a6f763b34b039427f5667ba3fd50 ]

    Resolves a panic that can occur on AMD systems, typically during host
    shutdown, after the PTDMA driver had been exercised. The issue was
    the pt_issue_pending() function is mistakenly assuming that there will
    be at least one descriptor in the Submitted queue when the function
    is called. However, it is possible that both the Submitted and Issued
    queues could be empty, which could result in pt_cmd_callback() being
    mistakenly called with a NULL pointer.
    Ref: Bugzilla Bug 216856.

    Fixes: 6fa7e0e836e2 ("dmaengine: ptdma: fix concurrency issue with multiple dma
transfer")
    Signed-off-by: Eric Pilmore <epilmore@gigaio.com>
    Link: https://lore.kernel.org/r/20230210075142.58253-1-epilmore@gigaio.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8d0afce410ddc1e4b542b0fea3de9dffbebca04
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Jan 27 14:36:27 2023 -0800

    dmaengine: dw-axi-dmac: Do not dereference NULL structure

    [ Upstream commit be4d46edeee4b2459d2f53f37ada88bbfb634b6c ]

    If "vdesc" is NULL, it cannot be used with vd_to_axi_desc(). Leave
    "bytes" unchanged at 0. Seen under GCC 13 with -Warray-bounds:

    ../drivers/dma/dw-axi-dmac/dw-axi-dmac-platform.c: In function 'dma_chan_tx_status':
    ../drivers/dma/dw-axi-dmac/dw-axi-dmac-platform.c:329:46: warning: array subscript 0 is
outside array bounds of 'struct
    virt_dma_desc[46116860184273879]' [-Warray-bounds=]
      329 |                 bytes = vd_to_axi_desc(vdesc)->length;
          |                                        ^~

    Fixes: 8e55444da65c ("dmaengine: dw-axi-dmac: Support burst residue granularity")
    Cc: Eugeniy Paltsev <Eugeniy.Paltsev@synopsys.com>
    Cc: Vinod Koul <vkoul@kernel.org>
    Cc: dmaengine@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20230127223623.never.507-kees@kernel.org
```

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

commit 8bd5040bd43f2b5ba3c898b09a3197a0c7ace126
Author: Shravan Chippa <shravan.chippa@microchip.com>
Date:   Fri Jan 20 15:36:23 2023 +0530

    dmaengine: sf-pdma: pdma_desc memory leak fix

    [ Upstream commit b02e07015a5ac7bbc029da931ae17914b8ae0339 ]

    Commit b2cc5c465c2c ("dmaengine: sf-pdma: Add multithread support for a
    DMA channel") changed sf_pdma_prep_dma_memcpy() to unconditionally
    allocate a new sf_pdma_desc each time it is called.

    The driver previously recycled descs, by checking the in_use flag, only
    allocating additional descs if the existing one was in use. This logic
    was removed in commit b2cc5c465c2c ("dmaengine: sf-pdma: Add multithread
    support for a DMA channel"), but sf_pdma_free_desc() was not changed to
    handle the new behaviour.

    As a result, each time sf_pdma_prep_dma_memcpy() is called, the previous
    descriptor is leaked, over time leading to memory starvation:

      unreferenced object 0xffffffe008447300 (size 192):
      comm "irq/39-mchp_dsc", pid 343, jiffies 4294906910 (age 981.200s)
      hex dump (first 32 bytes):
        00 00 00 ff 00 00 00 00 b8 c1 00 00 00 00 00 00  ................
        00 00 70 08 10 00 00 00 00 00 00 00 c0 00 00 00 00  ..p.............
      backtrace:
        [<00000000064a04f4>] kmemleak_alloc+0x1e/0x28
        [<00000000018927a7>] kmem_cache_alloc+0x11e/0x178
        [<000000002aea8d16>] sf_pdma_prep_dma_memcpy+0x40/0x112

    Add the missing kfree() to sf_pdma_free_desc(), and remove the redundant
    in_use flag.

    Fixes: b2cc5c465c2c ("dmaengine: sf-pdma: Add multithread support for a DMA channel")
    Signed-off-by: Shravan Chippa <shravan.chippa@microchip.com>
    Reviewed-by: Conor Dooley <conor.dooley@microchip.com>
    Link: https://lore.kernel.org/r/20230120100623.3530634-1-shravan.chippa@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 697cd3829521b7c40468bc25f29c3810c6794c0e
Author: Vasant Hegde <vasant.hegde@amd.com>
Date:   Tue Feb 7 09:17:52 2023 +0000

    iommu/amd: Do not identity map v2 capable device when snp is enabled

    [ Upstream commit 18792e99ea2fea27c72eb1ecca1879e5e6be304d ]

    Flow:
      - Booted system with SNP enabled, memory encryption off and
        IOMMU DMA translation mode
      - AMD driver detects v2 capable device and amd_iommu_def_domain_type()
        returns identity mode
      - amd_iommu_domain_alloc() returns NULL an SNP is enabled
      - System will fail to register device

    On SNP enabled system, passthrough mode is not supported. IOMMU default
    domain is set to translation mode. We need to return zero from
    amd_iommu_def_domain_type() so that it allocates translation domain.

    Fixes: fb2accadaa94 ("iommu/amd: Introduce function to check and enable SNP")
```

```
    CC: Suravee Suthikulpanit <suravee.suthikulpanit@amd.com>
    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
    Link: https://lore.kernel.org/r/20230207091752.7656-1-vasant.hegde@amd.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8602fa38a63fc689eca29c8c586fac2d8a1f8739
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date:   Wed Feb 15 21:21:16 2023 -0400

    iommu: Fix error unwind in iommu_group_alloc()

    [ Upstream commit 4daa861174d56023c2068ddb03de0752f07fa199 ]

    If either iommu_group_grate_file() fails then the
    iommu_group is leaked.

    Destroy it on these error paths.

    Found by kselftest/iommu/iommufd_fail_nth

    Fixes: bc7d12b91bd3 ("iommu: Implement reserved_regions iommu-group sysfs file")
    Fixes: c52c72d3dee8 ("iommu: Add sysfs attribyte for domain type")
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Reviewed-by: Lu Baolu <baolu.lu@linux.intel.com>
    Link: https://lore.kernel.org/r/0-v1-8f616bee028d+8b-
iommu_group_alloc_leak_jgg@nvidia.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd55240e4364d64befcc575b0d33091881524f42
Author: Dan Carpenter <error27@gmail.com>
Date:   Tue Feb 14 18:43:38 2023 +0300

    iw_cxgb4: Fix potential NULL dereference in c4iw_fill_res_cm_id_entry()

    [ Upstream commit 4ca446b127c568b59cb8d9748b6f70499624bb18 ]

    This condition needs to match the previous "if (epcp->state == LISTEN) {"
    exactly to avoid a NULL dereference of either "listen_ep" or "ep". The
    problem is that "epcp" has been re-assigned so just testing
    "if (epcp->state == LISTEN) {" a second time is not sufficient.

    Fixes: 116aeb887371 ("iw_cxgb4: provide detailed provider-specific CM_ID information")
    Signed-off-by: Dan Carpenter <error27@gmail.com>
    Link: https://lore.kernel.org/r/Y+usKuWIKr4dimZh@kili
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 090745df5a195a5825e3f0f06b9fcc48975c3338
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Mon Oct 17 13:47:04 2022 +0200

    PCI: qcom: Fix host-init error handling

    [ Upstream commit 997e010de9134474dbfde52be03efd7d1bce902d ]

    Implement the new host_deinit() callback so that the PHY is powered off
    and regulators and clocks are disabled also on late host-init errors.

    Link: https://lore.kernel.org/r/20221017114705.8277-2-johan+linaro@kernel.org
    Fixes: 82a823833f4e ("PCI: qcom: Add Qualcomm PCIe controller driver")
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
```

```
    Reviewed-by: Manivannan Sadhasivam <mani@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c560aa03901e51b85c25b5052aba89e0c7927504
Author: Neill Kapron <nkapron@google.com>
Date:   Thu Jan 26 00:10:12 2023 +0000

    phy: rockchip-typec: fix tcphy_get_mode error case

    [ Upstream commit 4ca651df07183e29cdad7272255e23aec0169a1b ]

    The existing logic in tcphy_get_mode() can cause the phy to be
    incorrectly configured to USB UFP or DisplayPort mode when
    extcon_get_state returns an error code.

    extcon_get_state() can return 0, 1, or a negative error code.

    It is possible to get into the failing state with an extcon driver
    which does not support the extcon connector id specified as the
    second argument to extcon_get_state().

    tcphy_get_mode()
    ->extcon_get_state()
    -->find_cable_index_by_id()
    --->return -EINVAL;

    Fixes: e96be45cb84e ("phy: Add USB Type-C PHY driver for rk3399")
    Signed-off-by: Neill Kapron <nkapron@google.com>
    Reviewed-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20230126001013.3707873-1-nkapron@google.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e6f2714d93cdf977b6124a80af2cf0e14e2d407
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Feb 10 14:46:39 2023 +0100

    PCI: Fix dropping valid root bus resources with .end = zero

    [ Upstream commit 9d8ba74a181b1c81def21168795ed96cbe6f05ed ]

    On r8a7791/koelsch:

      kmemleak: 1 new suspected memory leaks (see /sys/kernel/debug/kmemleak)
      # cat /sys/kernel/debug/kmemleak
      unreferenced object 0xc3a34e00 (size 64):
        comm "swapper/0", pid 1, jiffies 4294937460 (age 199.080s)
        hex dump (first 32 bytes):
          b4 5d 81 f0 b4 5d 81 f0 c0 b0 a2 c3 00 00 00 00  .]...]..........
          00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
        backtrace:
          [<fe3aa979>] __kmalloc+0xf0/0x140
          [<34bd6bc0>] resource_list_create_entry+0x18/0x38
          [<767046bc>] pci_add_resource_offset+0x20/0x68
          [<b3f3edf2>] devm_of_pci_get_host_bridge_resources.constprop.0+0xb0/0x390

    When coalescing two resources for a contiguous aperture, the second
    resource is enlarged to cover the full contiguous range, while the first
    resource is marked invalid.  This invalidation is done by clearing the
    flags, start, and end members.

    When adding the initial resources to the bus later, invalid resources are
    skipped.  Unfortunately, the check for an invalid resource considers only
    the end member, causing false positives.
```

```
    E.g. on r8a7791/koelsch, root bus resource 0 ("bus 00") is skipped, and no
    longer registered with pci_bus_insert_busn_res() (causing the memory leak),
    nor printed:

       pci-rcar-gen2 ee090000.pci: host bridge /soc/pci@ee090000 ranges:
       pci-rcar-gen2 ee090000.pci:      MEM 0x00ee080000..0x00ee08ffff -> 0x00ee080000
       pci-rcar-gen2 ee090000.pci: PCI: revision 11
       pci-rcar-gen2 ee090000.pci: PCI host bridge to bus 0000:00
      -pci_bus 0000:00: root bus resource [bus 00]
       pci_bus 0000:00: root bus resource [mem 0xee080000-0xee08ffff]

    Fix this by only skipping resources where all of the flags, start, and end
    members are zero.

    Fixes: 7c3855c423b17f6c ("PCI: Coalesce host bridge contiguous apertures")
    Link:
https://lore.kernel.org/r/da0fcd5e86c74239be79c7cb03651c0fce31b515.1676036673.git.geert+rene
sas@glider.be
    Tested-by: Niklas Schnelle <schnelle@linux.ibm.com>
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 126eda513d348acf0a1ed95dc61ff6bf52f914a3
Author: Serge Semin <Sergey.Semin@baikalelectronics.ru>
Date:   Wed Jan 25 18:19:00 2023 +0300

    dmaengine: dw-edma: Fix readq_ch() return value truncation

    [ Upstream commit 5fdca4a995bcd4cf61bda40af154a730589dc524 ]

    Previously, readq_ch() did a 64-bit readq(), but truncated the result by
    storing it in the u32 "value".  Change "value" to u64 to avoid the
    truncation.

    Note: the method is currently unused, so the bug hasn't caused any problem
    so far.

    Fixes: 04e0a39fc10f ("dmaengine: dw-edma: Add writeq() and readq() for 64 bits
architectures")
    Signed-off-by: Serge Semin <Sergey.Semin@baikalelectronics.ru>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c8f3a4f90f366ad6cbd6d5b3c3f9f4522b5c5276
Author: Alexander Stein <alexander.stein@ew.tq-group.com>
Date:   Tue Feb 7 12:05:31 2023 +0100

    usb: host: fsl-mph-dr-of: reuse device_set_of_node_from_dev

    [ Upstream commit e2ffae3ed92a9f768902c1cf82642c3a09cd0345 ]

    This sets both of_node fields and takes a of_node reference as well.

    Fixes: bb160ee61c04 ("drivers/usb/host/ehci-fsl: Fix interrupt setup in host mode.")
    Signed-off-by: Alexander Stein <alexander.stein@ew.tq-group.com>
    Reviewed-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20230207110531.1060252-4-alexander.stein@ew.tq-group.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d6b7d87d4967e9f1045448dd0ce6e81a6181c84
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:42:04 2023 -0800
```

```
    mtd: mtdpart: Don't create platform device that'll never probe

    [ Upstream commit fb42378dcc7f247df56f0ecddfdae85487495fbc ]

    These "nvmem-cells" platform devices never get probed because there's no
    platform driver for it and it's never used anywhere else. So it's a
    waste of memory. These devices also cause fw_devlink to block nvmem
    consumers of "nvmem-cells" partition from probing because the supplier
    device never probes.

    So stop creating platform devices for nvmem-cells partitions to avoid
    wasting memory and to avoid blocking probing of consumers.

    Reported-by: Maxim Kiselev <bigunclemax@gmail.com>
    Fixes: bcdf0315a61a ("mtd: call of_platform_populate() for MTD partitions")
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Maksim Kiselev <bigunclemax@gmail.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-13-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8c62cf7378351edf249c558095b9481f5cd78759
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:42:00 2023 -0800

    driver core: fw_devlink: Make cycle detection more robust

    [ Upstream commit 3fb16866b51ded6c016b664caad53f8d4fd9dc56 ]

    fw_devlink could only detect a single and simple cycle because it relied
    mainly on device link cycle detection code that only checked for cycles
    between devices. The expectation was that the firmware wouldn't have
    complicated cycles and multiple cycles between devices. That expectation
    has been proven to be wrong.

    For example, fw_devlink could handle:

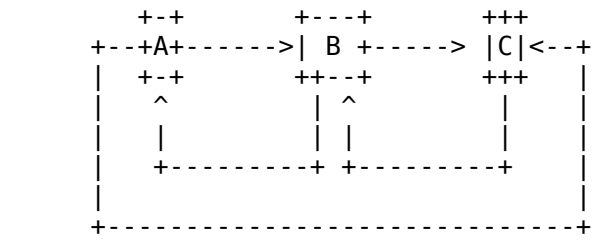
    +-+       +-+
    |A+------> |B+
    +-+       +++
     ^          |
     |          |
     +----------+

    But it couldn't handle even something as "simple" as:

     +---------------------+
     |                     |
     v                     |
    +-+       +-+       +++
    |A+------> |B+------> |C|
    +-+       +++       +-+
     ^          |
     |          |
     +----------+

    But firmware has even more complicated cycles like:

        +---------------------+
        |                     |
        v                     |
```



    And this is without including parent child dependencies or nodes in the
    cycle that are just firmware nodes that'll never have a struct device
    created for them.

    The proper way to treat these devices it to not force any probe ordering
    between them, while still enforce dependencies between node in the
    cycles (A, B and C) and their consumers.

    So this patch goes all out and just deals with all types of cycles. It
    does this by:

    1. Following dependencies across device links, parent-child and fwnode
       links.
    2. When it find cycles, it mark the device links and fwnode links as
       such instead of just deleting them or making the indistinguishable
       from proxy SYNC_STATE_ONLY device links.

    This way, when new nodes get added, we can immediately find and mark any
    new cycles whether the new node is a device or firmware node.

    Fixes: 2de9d8e0d2fe ("driver core: fw_devlink: Improve handling of cyclic dependencies")
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-9-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b5f2cd799be2d4a1a9dcd563b010ef5f79cc207
Author: Saravana Kannan <saravanak@google.com>
Date: Mon Feb 6 17:41:54 2023 -0800

    driver core: fw_devlink: Improve check for fwnode with no device/driver

    [Upstream commit 411c0d58ca6faa9bc4b9f5382118a31c7bb92a6f]

    fw_devlink shouldn't defer the probe of a device to wait on a supplier
    that'll never have a struct device or will never be probed by a driver.
    We currently check if a supplier falls into this category, but don't
    check its ancestors. We need to check the ancestors too because if the
    ancestor will never probe, then the supplier will never probe either.

    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-3-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Stable-dep-of: 3fb16866b51d ("driver core: fw_devlink: Make cycle detection more
robust")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

```
commit 1392204a981195b145bc424da22a1888b379666b
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:41:59 2023 -0800

    driver core: fw_devlink: Consolidate device link flag computation

    [ Upstream commit cd115c0409f283edde94bd5a9a42dc42bee0aba8 ]

    Consolidate the code that computes the flags to be used when creating a
    device link from a fwnode link.

    Fixes: 2de9d8e0d2fe ("driver core: fw_devlink: Improve handling of cyclic dependencies")
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-8-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3641c013963d4a8001d0546b4d335a47c3baf2bc
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:41:58 2023 -0800

    driver core: fw_devlink: Allow marking a fwnode link as being part of a cycle

    [ Upstream commit 6a6dfdf8b3ff337be5a447e9f4e71969f18370ad ]

    To improve detection and handling of dependency cycles, we need to be
    able to mark fwnode links as being part of cycles. fwnode links marked
    as being part of a cycle should not block their consumers from probing.

    Fixes: 2de9d8e0d2fe ("driver core: fw_devlink: Improve handling of cyclic dependencies")
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-7-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 083dcec8062801e0c0adb3b2ce2bbd7930b55f22
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:41:53 2023 -0800

    driver core: fw_devlink: Don't purge child fwnode's consumer links

    [ Upstream commit 3a2dbc510c437ca392516b0105bad8e7970e6614 ]

    When a device X is bound successfully to a driver, if it has a child
    firmware node Y that doesn't have a struct device created by then, we
    delete fwnode links where the child firmware node Y is the supplier. We
    did this to avoid blocking the consumers of the child firmware node Y
    from deferring probe indefinitely.

    While that a step in the right direction, it's better to make the
    consumers of the child firmware node Y to be consumers of the device X
    because device X is probably implementing whatever functionality is
    represented by child firmware node Y. By doing this, we capture the
    device dependencies more accurately and ensure better
```

```
    probe/suspend/resume ordering.

    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-2-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Stable-dep-of: 6a6dfdf8b3ff ("driver core: fw_devlink: Allow marking a fwnode link as
being part of a cycle")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6716dda4b9797c864397229713e7cba4c849a33
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Feb 6 17:41:57 2023 -0800

    driver core: fw_devlink: Add DL_FLAG_CYCLE support to device links

    [ Upstream commit 67cad5c67019c38126b749621665b6723d3ae7e6 ]

    fw_devlink uses DL_FLAG_SYNC_STATE_ONLY device link flag for two
    purposes:

    1. To allow a parent device to proxy its child device's dependency on a
       supplier so that the supplier doesn't get its sync_state() callback
       before the child device/consumer can be added and probed. In this
       usage scenario, we need to ignore cycles for ensure correctness of
       sync_state() callbacks.

    2. When there are dependency cycles in firmware, we don't know which of
       those dependencies are valid. So, we have to ignore them all wrt
       probe ordering while still making sure the sync_state() callbacks
       come correctly.

    However, when detecting dependency cycles, there can be multiple
    dependency cycles between two devices that we need to detect. For
    example:

    A -> B -> A and A -> C -> B -> A.

    To detect multiple cycles correct, we need to be able to differentiate
    DL_FLAG_SYNC_STATE_ONLY device links used for (1) vs (2) above.

    To allow this differentiation, add a DL_FLAG_CYCLE that can be use to
    mark use case (2). We can then use the DL_FLAG_CYCLE to decide which
    DL_FLAG_SYNC_STATE_ONLY device links to follow when looking for
    dependency cycles.

    Fixes: 2de9d8e0d2fe ("driver core: fw_devlink: Improve handling of cyclic dependencies")
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Colin Foster <colin.foster@in-advantage.com>
    Tested-by: Sudeep Holla <sudeep.holla@arm.com>
    Tested-by: Douglas Anderson <dianders@chromium.org>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Luca Weiss <luca.weiss@fairphone.com> # qcom/sm7225-fairphone-fp4
    Link: https://lore.kernel.org/r/20230207014207.1678715-6-saravanak@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 963875b0655197281775b0ea614aab8b6b3eb001
Author: Peng Fan <peng.fan@nxp.com>
Date:   Mon Feb 6 09:30:16 2023 +0800
```

```
    tty: serial: imx: disable Ageing Timer interrupt request irq

    [ Upstream commit ef25e16ea9674b713a68c3bda821556ce9901254 ]

    There maybe pending USR interrupt before requesting irq, however
    uart_add_one_port has not executed, so there will be kernel panic:
    [    0.795668] Unable to handle kernel NULL pointer dereference at virtual addre
    ss 0000000000000080
    [    0.802701] Mem abort info:
    [    0.805367]   ESR = 0x0000000096000004
    [    0.808950]   EC = 0x25: DABT (current EL), IL = 32 bits
    [    0.814033]   SET = 0, FnV = 0
    [    0.816950]   EA = 0, S1PTW = 0
    [    0.819950]   FSC = 0x04: level 0 translation fault
    [    0.824617] Data abort info:
    [    0.827367]   ISV = 0, ISS = 0x00000004
    [    0.831033]   CM = 0, WnR = 0
    [    0.833866] [0000000000000080] user address but active_mm is swapper
    [    0.839951] Internal error: Oops: 0000000096000004 [#1] PREEMPT SMP
    [    0.845953] Modules linked in:
    [    0.848869] CPU: 0 PID: 1 Comm: swapper/0 Not tainted 6.1.1+g56321e101aca #1
    [    0.855617] Hardware name: Freescale i.MX8MP EVK (DT)
    [    0.860452] pstate: 000000c5 (nzcv daIF -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
    [    0.867117] pc : __imx_uart_rxint.constprop.0+0x11c/0x2c0
    [    0.872283] lr : imx_uart_int+0xf8/0x1ec

    The issue only happends in the inmate linux when Jailhouse hypervisor
    enabled. The test procedure is:
    while true; do
            jailhouse enable imx8mp.cell
            jailhouse cell linux xxxx
            sleep 10
            jailhouse cell destroy 1
            jailhouse disable
            sleep 5
    done

    And during the upper test, press keys to the 2nd linux console.
    When `jailhouse cell destroy 1`, the 2nd linux has no chance to put
    the uart to a quiese state, so USR1/2 may has pending interrupts. Then
    when `jailhosue cell linux xx` to start 2nd linux again, the issue
    trigger.

    In order to disable irqs before requesting them, both UCR1 and UCR2 irqs
    should be disabled, so here fix that, disable the Ageing Timer interrupt
    in UCR2 as UCR1 does.

    Fixes: 8a61f0c70ae6 ("serial: imx: Disable irqs before requesting them")
    Suggested-by: Sherry Sun <sherry.sun@nxp.com>
    Reviewed-by: Sherry Sun <sherry.sun@nxp.com>
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Acked-by: Jason Liu <jason.hui.liu@nxp.com>
    Link: https://lore.kernel.org/r/20230206013016.29352-1-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b0ed9d76ed9a64ff6587213e17f1d18053e52a7
Author: Shenwei Wang <shenwei.wang@nxp.com>
Date:   Tue Feb 7 10:24:20 2023 -0600

    serial: fsl_lpuart: fix RS485 RTS polariy inverse issue

    [ Upstream commit 3957b9501a5a8fa709ae4a47483714491471f6db ]

    The previous 'commit 846651eca073 ("serial: fsl_lpuart: RS485 RTS
```

```
     polariy is inverse")' only fixed the inverse issue on lpuart 8bit
     platforms.

     This is a follow-up patch to fix the RS485 polarity inverse
     issue on lpuart 32bit platforms.

     Fixes: 03895cf41d18 ("tty: serial: fsl_lpuart: Add support for RS-485")
     Reported-by: Sherry Sun <sherry.sun@nxp.com>
     Signed-off-by: Shenwei Wang <shenwei.wang@nxp.com>
     Link: https://lore.kernel.org/r/20230207162420.3647904-1-shenwei.wang@nxp.com
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 209e4aa9a7b636d8aaa1297e1d089ee2ed91d73f
Author: Mustafa Ismail <mustafa.ismail@intel.com>
Date:   Tue Feb 7 14:19:38 2023 -0600

    RDMA/irdma: Cap MSIX used to online CPUs + 1

    [ Upstream commit 9cd9842c46996ef62173c36619c746f57416bcb0 ]

    The irdma driver can use a maximum number of msix vectors equal
    to num_online_cpus() + 1 and the kernel warning stack below is shown
    if that number is exceeded.

    The kernel throws a warning as the driver tries to update the affinity
    hint with a CPU mask greater than the max CPU IDs. Fix this by capping
    the MSIX vectors to num_online_cpus() + 1.

     WARNING: CPU: 7 PID: 23655 at include/linux/cpumask.h:106
irdma_cfg_ceq_vector+0x34c/0x3f0 [irdma]
     RIP: 0010:irdma_cfg_ceq_vector+0x34c/0x3f0 [irdma]
     Call Trace:
     irdma_rt_init_hw+0xa62/0x1290 [irdma]
     ? irdma_alloc_local_mac_entry+0x1a0/0x1a0 [irdma]
     ? __is_kernel_percpu_address+0x63/0x310
     ? rcu_read_lock_held_common+0xe/0xb0
     ? irdma_lan_unregister_qset+0x280/0x280 [irdma]
     ? irdma_request_reset+0x80/0x80 [irdma]
     ? ice_get_qos_params+0x84/0x390 [ice]
     irdma_probe+0xa40/0xfc0 [irdma]
     ? rcu_read_lock_bh_held+0xd0/0xd0
     ? irdma_remove+0x140/0x140 [irdma]
     ? rcu_read_lock_sched_held+0x62/0xe0
     ? down_write+0x187/0x3d0
     ? auxiliary_match_id+0xf0/0x1a0
     ? irdma_remove+0x140/0x140 [irdma]
     auxiliary_bus_probe+0xa6/0x100
     __driver_probe_device+0x4a4/0xd50
     ? __device_attach_driver+0x2c0/0x2c0
     driver_probe_device+0x4a/0x110
     __driver_attach+0x1aa/0x350
     bus_for_each_dev+0x11d/0x1b0
     ? subsys_dev_iter_init+0xe0/0xe0
     bus_add_driver+0x3b1/0x610
     driver_register+0x18e/0x410
     ? 0xffffffffc0b88000
     irdma_init_module+0x50/0xaa [irdma]
     do_one_initcall+0x103/0x5f0
     ? perf_trace_initcall_level+0x420/0x420
     ? do_init_module+0x4e/0x700
     ? __kasan_kmalloc+0x7d/0xa0
     ? kmem_cache_alloc_trace+0x188/0x2b0
     ? kasan_unpoison+0x21/0x50
     do_init_module+0x1d1/0x700
```

```
       load_module+0x3867/0x5260
       ? layout_and_allocate+0x3990/0x3990
       ? rcu_read_lock_held_common+0xe/0xb0
       ? rcu_read_lock_sched_held+0x62/0xe0
       ? rcu_read_lock_bh_held+0xd0/0xd0
       ? __vmalloc_node_range+0x46b/0x890
       ? lock_release+0x5c8/0xba0
       ? alloc_vm_area+0x120/0x120
       ? selinux_kernel_module_from_file+0x2a5/0x300
       ? __inode_security_revalidate+0xf0/0xf0
       ? __do_sys_init_module+0x1db/0x260
      __do_sys_init_module+0x1db/0x260
       ? load_module+0x5260/0x5260
       ? do_syscall_64+0x22/0x450
      do_syscall_64+0xa5/0x450
       entry_SYSCALL_64_after_hwframe+0x66/0xdb

    Fixes: 44d9e52977a1 ("RDMA/irdma: Implement device initialization definitions")
    Signed-off-by: Mustafa Ismail <mustafa.ismail@intel.com>
    Signed-off-by: Shiraz Saleem <shiraz.saleem@intel.com>
    Signed-off-by: Sindhu Devale <sindhu.devale@intel.com>
    Link: https://lore.kernel.org/r/20230207201938.1329-1-sindhu.devale@intel.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 99a5bdd62e111518e7e91027c078e6c3391004e6
Author: Mark Tomlinson <mark.tomlinson@alliedtelesis.co.nz>
Date:   Tue Feb 7 16:33:37 2023 +1300

    usb: max-3421: Fix setting of I/O pins

    [ Upstream commit a7efe3fc7cbe27c6eb2c2a3ab612194f8f800f4c ]

    To update the I/O pins, the registers are read/modified/written. The
    read operation incorrectly always read the first register. Although
    wrong, there wasn't any impact as all the output pins are always
    written, and the inputs are read only anyway.

    Fixes: 2d53139f3162 ("Add support for using a MAX3421E chip as a host driver.")
    Signed-off-by: Mark Tomlinson <mark.tomlinson@alliedtelesis.co.nz>
    Link: https://lore.kernel.org/r/20230207033337.18112-1-
mark.tomlinson@alliedtelesis.co.nz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ddc77eefb2a567b705c3c86ab2ddabe43cadf1b
Author: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
Date:   Thu Feb 2 10:48:50 2023 -0800

    RDMA/cxgb4: Fix potential null-ptr-deref in pass_establish()

    [ Upstream commit 283861a4c52c1ea4df3dd1b6fc75a50796ce3524 ]

    If get_ep_from_tid() fails to lookup non-NULL value for ep, ep is
    dereferenced later regardless of whether it is empty.
    This patch adds a simple sanity check to fix the issue.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 944661dd97f4 ("RDMA/iw_cxgb4: atomically lookup ep and get a reference")
    Signed-off-by: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
    Link: https://lore.kernel.org/r/20230202184850.29882-1-n.zhandarovich@fintech.ru
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 140ec56617b5d08a6727f30c1857bbe58d718a40
Author: Bernard Metzler <bmt@zurich.ibm.com>
Date:   Thu Feb 2 11:10:00 2023 +0100

    RDMA/siw: Fix user page pinning accounting

    [ Upstream commit 65a8fc30fb6722fc25adec6d7dd5b53b0bb85820 ]

    To avoid racing with other user memory reservations, immediately
    account full amount of pages to be pinned.

    Fixes: 2251334dcac9 ("rdma/siw: application buffer management")
    Reported-by: Jason Gunthorpe <jgg@nvidia.com>
    Suggested-by: Alistair Popple <apopple@nvidia.com>
    Reviewed-by: Alistair Popple <apopple@nvidia.com>
    Signed-off-by: Bernard Metzler <bmt@zurich.ibm.com>
    Link: https://lore.kernel.org/r/20230202101000.402990-1-bmt@zurich.ibm.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a2041a86bdbe00edd4b7f45d4fa0d1fd4b4f435
Author: Andreas Kemnade <andreas@kemnade.info>
Date:   Sat Jan 21 12:16:21 2023 +0100

    power: supply: remove faulty cooling logic

    [ Upstream commit c85c191694cb1cf290b11059b3d2de8a2732ffd0 ]

    The rn5t618 power driver fails to register
    a cooling device because POWER_SUPPLY_PROP_CHARGE_CONTROL_LIMIT_MAX
    is missing but availability is not checked before registering
    cooling device. After improved error checking in the thermal
    code, the registration of the power supply fails entirely.

    Checking for availability of _MAX before registering cooling device
    fixes the rn5t618 problem. But the whole logic feels questionable.

    First, the logic is inverted here:
    the code tells: max_current = max_cooling but
    0 = max_cooling, so there needs to be some inversion
    in the code which cannot be found. Comparing with other
    cooling devices, it can be found that value for fan speed is not
    inverted, value for cpufreq cooling is inverted (similar situation
    as here lowest frequency = max cooling)

    Second, analyzing usage of _MAX: it is seems that maximum capabilities
    of charging controller are specified and not of the battery. Probably
    there is not too much mismatch in the drivers actually implementing
    that. So nothing has exploded yet.  So there is no easy and safe way
    to specifify a max cooling value now.

    Conclusion for now (as a regression fix) just remove the cooling device
    registration and do it properly later on.

    Fixes: e49a1e1ee078 ("thermal/core: fix error code in
__thermal_cooling_device_register()")
    Fixes: 952aeeb3ee28 ("power_supply: Register power supply for thermal cooling device")
    Signed-off-by: Andreas Kemnade <andreas@kemnade.info>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit adcc71bda55094329fcfe1368cc7a415e777b596
Author: Lu Baolu <baolu.lu@linux.intel.com>
Date:   Tue Jan 31 15:37:33 2023 +0800
```

```
    iommu/vt-d: Set No Execute Enable bit in PASID table entry

    [ Upstream commit e06d24435596c8afcaa81c0c498f5b0ec4ee2b7c ]

    Setup No Execute Enable bit (Bit 133) of a scalable mode PASID entry.
    This is to allow the use of XD bit of the first level page table.

    Fixes: ddf09b6d43ec ("iommu/vt-d: Setup pasid entries for iova over first level")
    Signed-off-by: Ashok Raj <ashok.raj@intel.com>
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Link: https://lore.kernel.org/r/20230126095438.354205-1-baolu.lu@linux.intel.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6ad90408113aa1bd6b4dab93f6f8f7cbb6d8612
Author: Sergio Paracuellos <sergio.paracuellos@gmail.com>
Date:   Sat Dec 31 08:40:41 2022 +0100

    PCI: mt7621: Delay phy ports initialization

    [ Upstream commit 0cb2a8f3456ff1cc51d571e287a48e8fddc98ec2 ]

    Some devices like ZBT WE1326 and ZBT WF3526-P and some Netgear models need
    to delay phy port initialization after calling the mt7621_pcie_init_port()
    driver function to get into reliable boots for both warm and hard resets.

    The delay required to detect the ports seems to be in the range [75-100]
    milliseconds.

    If the ports are not detected the controller is not functional.

    There is no datasheet or something similar to really understand why this
    extra delay is needed only for these devices and it is not for most of
    the boards that are built on mt7621 SoC.

    This issue has been reported by openWRT community and the complete
    discussion is in [0]. The 100 milliseconds delay has been tested in all
    devices to validate it.

    Add the extra 100 milliseconds delay to fix the issue.

    [0]: https://github.com/openwrt/openwrt/pull/11220

    Link: https://lore.kernel.org/r/20221231074041.264738-1-sergio.paracuellos@gmail.com
    Fixes: 2bdd5238e756 ("PCI: mt7621: Add MediaTek MT7621 PCIe host controller driver")
    Signed-off-by: Sergio Paracuellos <sergio.paracuellos@gmail.com>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01b685ca4d710aa6c81e503782b53bf14d370a38
Author: Chunfeng Yun <chunfeng.yun@mediatek.com>
Date:   Wed Jan 18 16:43:41 2023 +0800

    phy: mediatek: remove temporary variable @mask_

    [ Upstream commit 49094d928618309877c50e589f23e639a3b0c453 ]

    Remove the temporary @mask_, this may cause build warning when use clang
    compiler for powerpc, but can't reproduce it when compile for arm64.
    the build warning is caused by:

    "warning: result of comparison of constant 18446744073709551615 with
    expression of type (aka 'unsigned long') is always false
    [-Wtautological-constant-out-of-range-compare]"
```

```
    More information provided in below lore link.

    After removing @mask_, there is a "CHECK:MACRO_ARG_REUSE" when run
    checkpatch.pl, but due to @mask is constant, no reuse problem will happen.

    Link: https://lore.kernel.org/lkml/202212160357.jJuesD8n-lkp@intel.com/t/
    Fixes: 84513eccd678 ("phy: mediatek: fix build warning of FIELD_PREP()")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Chunfeng Yun <chunfeng.yun@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20230118084343.26913-1-chunfeng.yun@mediatek.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 25b32360c9b816e4f653462ff3ffc3eb8335222f
Author: Udipto Goswami <quic_ugoswami@quicinc.com>
Date:   Wed Feb 1 18:53:08 2023 +0530

    usb: gadget: configfs: Restrict symlink creation is UDC already binded

    [ Upstream commit 89e7252d6c7e7eeb31971cd7df987316ecc64ff5 ]

    During enumeration or composition switch,a userspace process
    agnostic of the conventions of configs can try to create function
    symlinks even after the UDC is bound to current config which is
    not correct. Potentially it can create duplicates within the
    current config.

    Prevent this by adding a check if udc_name already exists, then bail
    out of cfg_link.

    Following is an example:

    Step1:
    ln -s X1 ffs.a
    -->cfg_link
    --> usb_get_function(ffs.a)
            ->ffs_alloc

            CFG->FUNC_LIST: <ffs.a>
            C->FUNCTION: <empty>

    Step2:
    echo udc.name > /config/usb_gadget/g1/UDC
    --> UDC_store
            ->composite_bind
            ->usb_add_function

            CFG->FUNC_LIST: <empty>
            C->FUNCTION: <ffs.a>

    Step3:
    ln -s Y1 ffs.a
    -->cfg_link
    -->usb_get_function(ffs.a)
            ->ffs_alloc

            CFG->FUNC_LIST: <ffs.a>
            C->FUNCTION: <ffs.a>

    both the lists corresponds to the same function instance ffs.a
    but the usb_function* pointer is different because in step 3
    ffs_alloc has created a new reference to usb_function* for
    ffs.a and added it to cfg_list.
```

```
    Step4:
    Now a composition switch involving <ffs.b,ffs.a> is executed.

    the composition switch will involve 3 things:
            1. unlinking the previous functions existing
            2. creating new symlinks
            3. writing UDC

    However, the composition switch is generally taken care by
    userspace process which creates the symlinks in its own
    nomenclature(X*) and removes only those.
    So it won't be able to remove Y1 which user had created
    by own.

    Due to this the new symlinks cannot be created for ffs.a
    since the entry already exists in CFG->FUNC_LIST.

    The state of the CFG->FUNC_LIST is as follows:
            CFG->FUNC_LIST: <ffs.a>

    Fixes: 88af8bbe4ef7 ("usb: gadget: the start of the configfs interface")
    Signed-off-by: Krishna Kurapati PSSNV <quic_kriskura@quicinc.com>
    Signed-off-by: Udipto Goswami <quic_ugoswami@quicinc.com>
    Link: https://lore.kernel.org/r/20230201132308.31523-1-quic_ugoswami@quicinc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4aa0500b9b21fe7ac5cd1351fc69eca0e749e135
Author: Dan Carpenter <error27@gmail.com>
Date:   Tue Jan 24 18:20:46 2023 +0300

    usb: musb: mediatek: don't unregister something that wasn't registered

    [ Upstream commit ba883de971d1ad018f3083d9195b8abe54d87407 ]

    This function only calls mtk_otg_switch_init() when the ->port_mode
    is MUSB_OTG so the clean up code should only call mtk_otg_switch_exit()
    for that mode.

    Fixes: 0990366bab3c ("usb: musb: Add support for MediaTek musb controller")
    Signed-off-by: Dan Carpenter <error27@gmail.com>
    Link: https://lore.kernel.org/r/Y8/3TqpqiSr0RxFH@kili
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef71422a6ab4db0966b4c2e28fe1cd5ff530e292
Author: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
Date:   Wed Feb 1 09:21:03 2023 -0800

    RDMA/cxgb4: add null-ptr-check after ip_dev_find()

    [ Upstream commit ef42520240aacfc0d46c8d780c051d135a8dc9b7 ]

    ip_dev_find() may return NULL and assign it to pdev which is
    dereferenced later.
    Fix this by checking the return value of ip_dev_find() for NULL
    similar to the way it is done with other instances of said function.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 1cab775c3e75 ("RDMA/cxgb4: Fix LE hash collision bug for passive open
connection")
    Signed-off-by: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
    Link: https://lore.kernel.org/r/20230201172103.17261-1-n.zhandarovich@fintech.ru
```

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

commit 9331593d14360e62108b17bdcc510ff7e5d75b96
Author: Sherry Sun <sherry.sun@nxp.com>
Date:   Mon Jan 30 14:44:46 2023 +0800

    tty: serial: fsl_lpuart: Fix the wrong RXWATER setting for rx dma case

    [ Upstream commit 9ad9df8447547febe9dd09b040f4528a09e495f0 ]

    The RXWATER value must be greater than 0 according to the LPUART
    reference manual. And when the number of datawords in the receive
    FIFO is greater than RXWATER, an interrupt or a DMA request is
    generated, so no need to set the different value for lpuart interrupt
    case and dma case. Here delete the wrong RXWATER setting for dma case
    directly.

    Fixes: 42b68768e51b ("serial: fsl_lpuart: DMA support for 32-bit variant")
    Signed-off-by: Sherry Sun <sherry.sun@nxp.com>
    Link: https://lore.kernel.org/r/20230130064449.9564-4-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df7c8aba7309f4dc55df94e06b67f576c0f52406
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sun Jan 29 19:23:08 2023 +0100

    usb: early: xhci-dbc: Fix a potential out-of-bound memory access

    [ Upstream commit a4a97ab3db5c081eb6e7dba91306adefb461e0bd ]

    If xdbc_bulk_write() fails, the values in 'buf' can be anything. So the
    string is not guaranteed to be NULL terminated when xdbc_trace() is called.

    Reserve an extra byte, which will be zeroed automatically because 'buf' is
    a static variable, in order to avoid troubles, should it happen.

    Fixes: aeb9dd1de98c ("usb/early: Add driver for xhci debug capability")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Link:
https://lore.kernel.org/r/d6a7562c5e839a195cee85db6dc81817f9372cb1.1675016180.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f44b50c7700068db4d544b7f47bab33e324ba6c3
Author: Ivan Bornyakov <i.bornyakov@metrotek.ru>
Date:   Fri Dec 30 12:29:21 2022 +0300

    fpga: microchip-spi: rewrite status polling in a time measurable way

    [ Upstream commit 88e705697e801299a13ecaf2ba54599964fe711c ]

    Original busy loop with retries count in mpf_poll_status() is not too
    reliable, as it takes different times on different systems. Replace it
    with read_poll_timeout() macro.

    While at it, fix polling stop condition to met function's original
    intention declared in the comment. The issue with original polling stop
    condition is that it stops if any of mask bits is set, while intention
    was to stop if all mask bits is set. This was not noticible because only
    MPF_STATUS_READY is passed as mask argument and it is BIT(1).

    Fixes: 5f8d4a900830 ("fpga: microchip-spi: add Microchip MPF FPGA manager")
```

```
    Signed-off-by: Ivan Bornyakov <i.bornyakov@metrotek.ru>
    Reviewed-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Acked-by: Conor Dooley <conor.dooley@microchip.com>
    Acked-by: Xu Yilun <yilun.xu@intel.com>
    Link: https://lore.kernel.org/r/20221230092922.18822-3-i.bornyakov@metrotek.ru
    Signed-off-by: Xu Yilun <yilun.xu@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71f18e099d37ea1c89af8cced6bea18a5888ceba
Author: Ivan Bornyakov <i.bornyakov@metrotek.ru>
Date:   Fri Dec 30 12:29:20 2022 +0300

    fpga: microchip-spi: move SPI I/O buffers out of stack

    [ Upstream commit 1da53d23a41c5f77963984d8da5623ed56918ada ]

    As spi-summary doc says:
     > I/O buffers use the usual Linux rules, and must be DMA-safe.
     > You'd normally allocate them from the heap or free page pool.
     > Don't use the stack, or anything that's declared "static".

    Replace spi_write() with spi_write_then_read(), which is dma-safe for
    on-stack buffers. Use cacheline aligned buffers for transfers used in
    spi_sync_transfer().

    Although everything works OK with stack-located I/O buffers, better
    follow the doc to be safe.

    Fixes: 5f8d4a900830 ("fpga: microchip-spi: add Microchip MPF FPGA manager")
    Signed-off-by: Ivan Bornyakov <i.bornyakov@metrotek.ru>
    Acked-by: Conor Dooley <conor.dooley@microchip.com>
    Acked-by: Xu Yilun <yilun.xu@intel.com>
    Link: https://lore.kernel.org/r/20221230092922.18822-2-i.bornyakov@metrotek.ru
    Signed-off-by: Xu Yilun <yilun.xu@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 156269103acdf2a8de4122d30d54405869cbff8b
Author: Serge Semin <Sergey.Semin@baikalelectronics.ru>
Date:   Fri Jan 13 20:13:46 2023 +0300

    dmaengine: dw-edma: Fix missing src/dst address of interleaved xfers

    [ Upstream commit 13b6299cf66165a442089fa895a7f70250703584 ]

    Interleaved DMA transfer support was added by 85e7518f42c8 ("dmaengine:
    dw-edma: Add device_prep_interleave_dma() support"), but depending on the
    selected channel, either source or destination address are left
    uninitialized which was obviously wrong.

    Initialize the destination address of the eDMA burst descriptors for
    DEV_TO_MEM interleaved operations and the source address for MEM_TO_DEV
    operations.

    Link: https://lore.kernel.org/r/20230113171409.30470-5-Sergey.Semin@baikalelectronics.ru
    Fixes: 85e7518f42c8 ("dmaengine: dw-edma: Add device_prep_interleave_dma() support")
    Tested-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Serge Semin <Sergey.Semin@baikalelectronics.ru>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Acked-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5217979df0ffd5617e5c57f1f405f74b62977b4f
Author: Fabian Vogt <fabian@ritter-vogt.de>
```

```
Date:   Mon Jan 23 08:35:06 2023 +0100

    fotg210-udc: Add missing completion handler

    [ Upstream commit e55f67391fa986f7357edba0ca59e668d99c3a5f ]

    This is used when responding to GET_STATUS requests. Without this, it
    crashes on completion.

    Fixes: b84a8dee23fd ("usb: gadget: add Faraday fotg210_udc driver")
    Signed-off-by: Fabian Vogt <fabian@ritter-vogt.de>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20230123073508.2350402-2-linus.walleij@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b59cecc74fd7f2893f6e71c493748b35489110ef
Author: Yi Liu <yi.l.liu@intel.com>
Date:   Fri Jan 20 04:20:40 2023 -0800

    iommufd: Add three missing structures in ucmd_buffer

    [ Upstream commit 84798f2849942bb5e8817417adfdfa6241df2835 ]

    struct iommu_ioas_copy, struct iommu_option and struct iommu_vfio_ioas are
    missed in ucmd_buffer. Although they are smaller than the size of
    ucmd_buffer, it is safer to list them in ucmd_buffer explicitly.

    Fixes: aad37e71d5c4 ("iommufd: IOCTLs for the io_pagetable")
    Fixes: d624d6652a65 ("iommufd: vfio container FD ioctl compatibility")
    Link: https://lore.kernel.org/r/20230120122040.280219-1-yi.l.liu@intel.com
    Signed-off-by: Yi Liu <yi.l.liu@intel.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 642f7cc52f87a9e3b97c93c40f9f0579747d2ac9
Author: Nicolin Chen <nicolinc@nvidia.com>
Date:   Thu Jan 19 23:42:04 2023 -0800

    selftests: iommu: Fix test_cmd_destroy_access() call in user_copy

    [ Upstream commit 9fabbdf338b701f2d763d9edbc3e82ce1e7fa1b4 ]

    The test_cmd_destroy_access() should end with a semicolon, so add one.
    There is a test_ioctl_destroy(ioas_id) following already, so drop one.

    Fixes: 57f0988706fe ("iommufd: Add a selftest")
    Link: https://lore.kernel.org/r/20230120074204.1368-1-nicolinc@nvidia.com
    Signed-off-by: Nicolin Chen <nicolinc@nvidia.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5d0492d1d934642bdfd2057acc1b56f4b57be465
Author: Chen Zhongjin <chenzhongjin@huawei.com>
Date:   Fri Nov 11 09:53:26 2022 +0800

    firmware: dmi-sysfs: Fix null-ptr-deref in dmi_sysfs_register_handle

    [ Upstream commit 18e126e97c961f7a93823795c879d7c085fe5098 ]

    KASAN reported a null-ptr-deref error:

    KASAN: null-ptr-deref in range [0x0000000000000008-0x000000000000000f]
    CPU: 0 PID: 1373 Comm: modprobe
```

```
    Hardware name: QEMU Standard PC (i440FX + PIIX, 1996)
    RIP: 0010:dmi_sysfs_entry_release
    ...
    Call Trace:
     <TASK>
     kobject_put
     dmi_sysfs_register_handle (drivers/firmware/dmi-sysfs.c:540) dmi_sysfs
     dmi_decode_table (drivers/firmware/dmi_scan.c:133)
     dmi_walk (drivers/firmware/dmi_scan.c:1115)
     dmi_sysfs_init (drivers/firmware/dmi-sysfs.c:149) dmi_sysfs
     do_one_initcall (init/main.c:1296)
     ...
    Kernel panic - not syncing: Fatal exception
    Kernel Offset: 0x4000000 from 0xffffffff81000000
    ---[ end Kernel panic - not syncing: Fatal exception ]---

    It is because previous patch added kobject_put() to release the memory
    which will call  dmi_sysfs_entry_release() and list_del().

    However, list_add_tail(entry->list) is called after the error block,
    so the list_head is uninitialized and cannot be deleted.

    Move error handling to after list_add_tail to fix this.

    Fixes: 660ba678f999 ("firmware: dmi-sysfs: Fix memory leak in
dmi_sysfs_register_handle")
    Signed-off-by: Chen Zhongjin <chenzhongjin@huawei.com>
    Link: https://lore.kernel.org/r/20221111015326.251650-2-chenzhongjin@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ee3112f35e6ad573b5d6caa9bd70502685ae49e
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Nov 15 11:16:38 2022 +0800

    drivers: base: transport_class: fix resource leak when transport_add_device() fails

    [ Upstream commit e5da06b27ff5a148e42265c8e306670a9d913969 ]

    The normal call sequence of using transport class is:

    Add path:
    transport_setup_device()
      transport_setup_classdev()  // call sas_host_setup() here
    transport_add_device()        // if fails, need call transport_destroy_device()
    transport_configure_device()

    Remove path:
    transport_remove_device()
      transport_remove_classdev  // call sas_host_remove() here
    transport_destroy_device()

    If transport_add_device() fails, need call transport_destroy_device()
    to free memory, but in this case, ->remove() is not called, and the
    resources allocated in ->setup() are leaked. So fix these leaks by
    calling ->remove() in transport_add_class_device() if it returns error.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20221115031638.3816551-1-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84545227b2e738e59fb1db5aa153d710d14ca7bf
Author: Yang Yingliang <yangyingliang@huawei.com>
```

```
Date:   Thu Nov 10 18:23:07 2022 +0800

    drivers: base: transport_class: fix possible memory leak

    [ Upstream commit a86367803838b369fe5486ac18771d14723c258c ]

    Current some drivers(like iscsi) call transport_register_device()
    failed, they don't call transport_destroy_device() to release the
    memory allocated in transport_setup_device(), because they don't
    know what was done, it should be internal thing to release the
    resource in register function. So fix this leak by calling destroy
    function inside register function.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20221110102307.3492557-1-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a9de90951bbeaed775e4b8d1b16b4d359e82bf5
Author: Hanjun Guo <guohanjun@huawei.com>
Date:   Tue Nov 22 15:37:28 2022 +0800

    driver core: location: Free struct acpi_pld_info *pld before return false

    [ Upstream commit 0d150f967e8410e1e6712484543eec709356a65d ]

    struct acpi_pld_info *pld should be freed before the return of allocation
    failure, to prevent memory leak, add the ACPI_FREE() to fix it.

    Fixes: bc443c31def5 ("driver core: location: Check for allocations failure")
    Signed-off-by: Hanjun Guo <guohanjun@huawei.com>
    Link: https://lore.kernel.org/r/1669102648-11517-1-git-send-email-guohanjun@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f39d21154db87545d8f0b25d13c326f37cc32239
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Wed Nov 23 09:20:42 2022 +0800

    driver core: fix resource leak in device_add()

    [ Upstream commit 6977b1a5d67097eaa4d02b0c126c04cc6e8917c0 ]

    When calling kobject_add() failed in device_add(), it will call
    cleanup_glue_dir() to free resource. But in kobject_add(),
    dev->kobj.parent has been set to NULL. This will cause resource leak.

    The process is as follows:
    device_add()
            get_device_parent()
                    class_dir_create_and_add()
                            kobject_add()           //kobject_get()
            ...
            dev->kobj.parent = kobj;
            ...
            kobject_add()           //failed, but set dev->kobj.parent = NULL
            ...
            glue_dir = get_glue_dir(dev)    //glue_dir = NULL, and goto
                                            //"Error" label
            ...
            cleanup_glue_dir()      //becaues glue_dir is NULL, not call
                                    //kobject_put()

    The preceding problem may cause insmod mac80211_hwsim.ko to failed.
```

```
      sysfs: cannot create duplicate filename '/devices/virtual/mac80211_hwsim'
      Call Trace:
      <TASK>
      dump_stack_lvl+0x8e/0xd1
      sysfs_warn_dup.cold+0x1c/0x29
      sysfs_create_dir_ns+0x224/0x280
      kobject_add_internal+0x2aa/0x880
      kobject_add+0x135/0x1a0
      get_device_parent+0x3d7/0x590
      device_add+0x2aa/0x1cb0
      device_create_groups_vargs+0x1eb/0x260
      device_create+0xdc/0x110
      mac80211_hwsim_new_radio+0x31e/0x4790 [mac80211_hwsim]
      init_mac80211_hwsim+0x48d/0x1000 [mac80211_hwsim]
      do_one_initcall+0x10f/0x630
      do_init_module+0x19f/0x5e0
      load_module+0x64b7/0x6eb0
      __do_sys_finit_module+0x140/0x200
      do_syscall_64+0x35/0x80
      entry_SYSCALL_64_after_hwframe+0x46/0xb0
      </TASK>
      kobject_add_internal failed for mac80211_hwsim with -EEXIST, don't try to
      register things with the same name in the same directory.

      Fixes: cebf8fd16900 ("driver core: fix race between creating/querying glue dir and its
  cleanup")
      Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
      Link: https://lore.kernel.org/r/20221123012042.335252-1-shaozhengchao@huawei.com
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 67173e889c1a6601538a6cf9e8c96c88cb0de028
  Author: Yang Yingliang <yangyingliang@huawei.com>
  Date:   Wed Jan 4 17:57:02 2023 +0800

      iommu/exynos: Fix error handling in exynos_iommu_init()

      [ Upstream commit 53719876e064643a6e58b5e6067a149a0fd191ec ]

      If platform_driver_register() fails, it don't need unregister and call
      kmem_cache_free() to free the memory allocated before calling register.

      Fixes: bbc4d205d93f ("iommu/exynos: Fix driver initialization sequence")
      Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
      Reviewed-by: Sam Protsenko <semen.protsenko@linaro.org>
      Link: https://lore.kernel.org/r/20230104095702.2591122-1-yangyingliang@huawei.com
      Signed-off-by: Joerg Roedel <jroedel@suse.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 464765a6016764ebac7f2354cf96a9ba7be00c06
  Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
  Date:   Fri Nov 25 18:45:02 2022 +0100

      misc: fastrpc: Fix an error handling path in fastrpc_rpmsg_probe()

      [ Upstream commit 11819ed2b70da94acc41fec34178a011c4d3d25d ]

      If of_platform_populate() fails, some resources need to be freed as already
      done in the other error handling paths.

      Fixes: 278d56f970ae ("misc: fastrpc: Reference count channel context")
      Fixes: 3abe3ab3cdab ("misc: fastrpc: add secure domain support")
      Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
      Link:
  https://lore.kernel.org/r/b909d2f3273b794ea0f1f78d14bc24affb08ea5f.1669398274.git.christophe
```

```
.jaillet@wanadoo.fr
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 23fbe2eda72c9e1aa85298d56898ce1b41dfd624
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Wed Dec 28 18:05:00 2022 +0200

    misc/mei/hdcp: Use correct macros to initialize uuid_le

    [ Upstream commit 512ba04d8211dd1a54dd36adc3ecc527a28069c5 ]

    GUID_INIT() is for internal guid_t type and shouldn't be used
    for the uuid_le. I.o.w. relying on the implementation details
    is layering violation. Use correct macros to initialize uuid_le.

    Fixes: 64e9bbdd9588 ("misc/mei/hdcp: Client driver for HDCP application")
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Acked-by: Tomas Winkler <tomas.winkler@intel.com>
    Link: https://lore.kernel.org/r/20221228160500.21220-1-andriy.shevchenko@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1df27cde906c9d17c7ed546cdc483c94221aad80
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Wed Dec 28 18:05:58 2022 +0200

    mei: pxp: Use correct macros to initialize uuid_le

    [ Upstream commit 503b676dde2780330c90a10b37a844686601a784 ]

    GUID_INIT() is for internal guid_t type and shouldn't be used
    for the uuid_le. I.o.w. relying on the implementation details
    is layering violation. Use correct macros to initialize uuid_le.

    Fixes: c2004ce99ed7 ("mei: pxp: export pavp client to me client bus")
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Acked-by: Tomas Winkler <tomas.winkler@intel.com>
    Link: https://lore.kernel.org/r/20221228160558.21311-1-andriy.shevchenko@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 91b8e4f61f8f4594ee65368c8d89e6fdc29d3fb1
Author: George Kennedy <george.kennedy@oracle.com>
Date:   Mon Nov 28 15:18:25 2022 -0500

    VMCI: check context->notify_page after call to get_user_pages_fast() to avoid GPF

    [ Upstream commit 1a726cb47fd204109c767409fa9ca15a96328f14 ]

    The call to get_user_pages_fast() in vmci_host_setup_notify() can return
    NULL context->notify_page causing a GPF. To avoid GPF check if
    context->notify_page == NULL and return error if so.

    general protection fault, probably for non-canonical address
        0xe0009d1000000060: 0000 [#1] PREEMPT SMP KASAN NOPTI
    KASAN: maybe wild-memory-access in range [0x0005088000000300-
        0x0005088000000307]
    CPU: 2 PID: 26180 Comm: repro_34802241 Not tainted 6.1.0-rc4 #1
    Hardware name: Red Hat KVM, BIOS 1.15.0-2.module+el8.6.0 04/01/2014
    RIP: 0010:vmci_ctx_check_signal_notify+0x91/0xe0
    Call Trace:
     <TASK>
     vmci_host_unlocked_ioctl+0x362/0x1f40
     __x64_sys_ioctl+0x1a1/0x230
```

```
      do_syscall_64+0x3a/0x90
      entry_SYSCALL_64_after_hwframe+0x63/0xcd

    Fixes: a1d88436d53a ("VMCI: Fix two UVA mapping bugs")
    Reported-by: syzkaller <syzkaller@googlegroups.com>
    Signed-off-by: George Kennedy <george.kennedy@oracle.com>
    Reviewed-by: Vishnu Dasa <vdasa@vmware.com>
    Link: https://lore.kernel.org/r/1669666705-24012-1-git-send-email-
george.kennedy@oracle.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f2cdcce25cbd3fb3082689188b82a3ebc62afb23
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Nov 29 10:36:02 2022 -0600

    firmware: stratix10-svc: fix error handle while alloc/add device failed

    [ Upstream commit d66a4c20ae55ac88136b4a3befd944c093ffa677 ]

    If add device "stratix10-rsu" failed in stratix10_svc_drv_probe(),
    the 'svc_fifo' and 'genpool' need be freed in the error path.

    If allocate or add device "intel-fcs" failed in stratix10_svc_drv_probe(),
    the device "stratix10-rsu" need be unregistered in the error path.

    Fixes: e6281c26674e ("firmware: stratix10-svc: Add support for FCS")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Dinh Nguyen <dinguyen@kernel.org>
    Link: https://lore.kernel.org/r/20221129163602.462369-2-dinguyen@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06358cb52d59ed3b5b52bc9b9fcd0e71acdb2441
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Nov 29 10:36:01 2022 -0600

    firmware: stratix10-svc: add missing gen_pool_destroy() in stratix10_svc_drv_probe()

    [ Upstream commit 9175ee1a99d57ec07d66ff572e1d5a724477ab37 ]

    In error path in stratix10_svc_drv_probe(), gen_pool_destroy() should be called
    to destroy the memory pool that created by svc_create_memory_pool().

    Fixes: 7ca5ce896524 ("firmware: add Intel Stratix10 service layer driver")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Dinh Nguyen <dinguyen@kernel.org>
    Link: https://lore.kernel.org/r/20221129163602.462369-1-dinguyen@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8a313090dfd166347e94d41a7371d39bd04036de
Author: Xiongfeng Wang <wangxiongfeng2@huawei.com>
Date:   Tue Nov 22 19:40:35 2022 +0800

    applicom: Fix PCI device refcount leak in applicom_init()

    [ Upstream commit ce4273d89c52167d6fe20572136c58117eae0657 ]

    As comment of pci_get_class() says, it returns a pci_device with its
    refcount increased and decreased the refcount for the input parameter
    @from if it is not NULL.

    If we break the loop in applicom_init() with 'dev' not NULL, we need to
    call pci_dev_put() to decrease the refcount. Add the missing
```

```
    pci_dev_put() to avoid refcount leak.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Xiongfeng Wang <wangxiongfeng2@huawei.com>
    Link: https://lore.kernel.org/r/20221122114035.24194-1-wangxiongfeng2@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2e4047f9064cdd3004aa910277543ad7f9bc6017
Author: Yuan Can <yuancan@huawei.com>
Date:   Thu Nov 10 02:00:30 2022 +0000

    eeprom: idt_89hpesx: Fix error handling in idt_init()

    [ Upstream commit d717a3ab282f51ec45142f911f7ef8a55c057de5 ]

    A problem about idt_89hpesx create debugfs failed is triggered with the
    following log given:

     [ 4973.269647] debugfs: Directory 'idt_csr' with parent '/' already present!

    The reason is that idt_init() returns i2c_add_driver() directly without
    checking its return value, if i2c_add_driver() failed, it returns without
    destroy the newly created debugfs, resulting the debugfs of idt_csr can
    never be created later.

     idt_init()
       debugfs_create_dir() # create debugfs directory
       i2c_add_driver()
         driver_register()
           bus_add_driver()
             priv = kzalloc(...) # OOM happened
       # return without destroy debugfs directory

    Fix by removing debugfs when i2c_add_driver() returns error.

    Fixes: cfad6425382e ("eeprom: Add IDT 89HPESx EEPROM/CSR driver")
    Signed-off-by: Yuan Can <yuancan@huawei.com>
    Acked-by: Serge Semin <fancer.lancer@gmail.com>
    Link: https://lore.kernel.org/r/20221110020030.47711-1-yuancan@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4c55fb0d0e5c5c4c3e0cab9d963fa87d16608d71
Author: Duoming Zhou <duoming@zju.edu.cn>
Date:   Wed Jan 18 22:10:00 2023 +0800

    Revert "char: pcmcia: cm4000_cs: Replace mdelay with usleep_range in set_protocol"

    [ Upstream commit 70fae37a09268455b8ab4f64647086b61da6f39c ]

    This reverts commit be826ada52f1fcabed5b5217c94609ebf5967211.

    The function monitor_card() is a timer handler that runs in an
    atomic context, but it calls usleep_range() that can sleep.
    As a result, the sleep-in-atomic-context bugs will happen.
    The process is shown below:

        (atomic context)
    monitor_card()
      set_protocol()
        usleep_range() //sleep

    The origin commit c1986ee9bea3 ("[PATCH] New Omnikey Cardman
    4000 driver") works fine.
```

```
    Fixes: be826ada52f1 ("char: pcmcia: cm4000_cs: Replace mdelay with usleep_range in
set_protocol")
    Signed-off-by: Duoming Zhou <duoming@zju.edu.cn>
    Link: https://lore.kernel.org/r/20230118141000.5580-1-duoming@zju.edu.cn
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f15b81c0464f3a44ac44215eb273e4adec6fd898
Author: Yi Yang <yiyang13@huawei.com>
Date:   Sat Nov 26 10:08:52 2022 +0800

    serial: tegra: Add missing clk_disable_unprepare() in tegra_uart_hw_init()

    [ Upstream commit 38f28cfe9d08e3a47ef008798b275fef8118fc20 ]

    Add the missing clk_disable_unprepare() before return from
    tegra_uart_hw_init() in the error handling path.
    When request_irq() fails in tegra_uart_startup(), 'tup->uart_clk'
    has been enabled, fix it by adding clk_disable_unprepare().

    Fixes: cc9ca4d95846 ("serial: tegra: Only print FIFO error message when an error
occurs")
    Fixes: d781ec21bae6 ("serial: tegra: report clk rate errors")
    Signed-off-by: Yi Yang <yiyang13@huawei.com>
    Link: https://lore.kernel.org/r/20221126020852.113378-1-yiyang13@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a293ae8cc381b4a13b258beb00d9188b5d3f3507
Author: Sherry Sun <sherry.sun@nxp.com>
Date:   Fri Nov 25 18:19:53 2022 +0800

    tty: serial: fsl_lpuart: clear LPUART Status Register in lpuart32_shutdown()

    [ Upstream commit 4029dfc034febb54f6dd8ea83568accc943bc088 ]

    The LPUART Status Register needs to be cleared when closing the uart
    port to get a clean environment when reopening the uart.

    Fixes: 380c966c093e ("tty: serial: fsl_lpuart: add 32-bit register interface support")
    Signed-off-by: Sherry Sun <sherry.sun@nxp.com>
    Link: https://lore.kernel.org/r/20221125101953.18753-4-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14a9a6b6dc8105451a579cabf178f305e81a6baf
Author: Sherry Sun <sherry.sun@nxp.com>
Date:   Fri Nov 25 18:19:52 2022 +0800

    tty: serial: fsl_lpuart: disable Rx/Tx DMA in lpuart32_shutdown()

    [ Upstream commit 1d4bd0e4ae4ba95892bef919a8d4d3f08f122d7e ]

    UARTBAUD_RDMAE and UARTBAUD_TDMAE are enabled in lpuart32_startup(), but
    lpuart32_shutdown() not disable them, only free the dma ring buffer and
    release the dma channels, so here disable the Rx/Tx DMA first in
    lpuart32_shutdown().

    Fixes: 42b68768e51b ("serial: fsl_lpuart: DMA support for 32-bit variant")
    Signed-off-by: Sherry Sun <sherry.sun@nxp.com>
    Link: https://lore.kernel.org/r/20221125101953.18753-3-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8a9232f56ea72bc210611b1fec90eba869b3a084
Author: Yicong Yang <yangyicong@hisilicon.com>
Date:   Thu Jan 12 19:22:00 2023 +0800

    hwtracing: hisi_ptt: Only add the supported devices to the filters list

    [ Upstream commit b8d976c7d41a28c0fccf22c7113be9a29dc07e5c ]

    The PTT device can only support the devices on the same PCIe core,
    within BDF range [lower_bdf, upper_bdf]. It's not correct to assume
    the devices on the root bus are from the same PCIe core, there are
    cases that root ports from different PCIe core are sharing the same
    bus. So check when initializing the filters list.

    Fixes: ff0de066b463 ("hwtracing: hisi_ptt: Add trace function support for HiSilicon PCIe
Tune and Trace device")
    Signed-off-by: Yicong Yang <yangyicong@hisilicon.com>
    Reviewed-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230112112201.16283-1-yangyicong@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cdf0f19cb953d99c2d3b9b6c131cd6052d018aac
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Jan 3 10:49:07 2023 +0800

    PCI: endpoint: pci-epf-vntb: Add epf_ntb_mw_bar_clear() num_mws kernel-doc

    [ Upstream commit fd858402c6d0a80e0b543886b9f7865c6d76d5d6 ]

    8e4bfbe644a6 ("PCI: endpoint: pci-epf-vntb: fix error handle in
    epf_ntb_mw_bar_init()") added a "num_mws" parameter to
    epf_ntb_mw_bar_clear() but failed to add kernel-doc for num_mws.

    Add kernel-doc for num_mws on epf_ntb_mw_bar_clear().

    Fixes: 8e4bfbe644a6 ("PCI: endpoint: pci-epf-vntb: fix error handle in
epf_ntb_mw_bar_init()")
    Link: https://lore.kernel.org/r/20230103024907.293853-1-yangyingliang@huawei.com
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec4b5955729b1d2cfbc33eb3d17fc9f9c2e2aac8
Author: Bjorn Helgaas <bhelgaas@google.com>
Date:   Fri Dec 16 10:21:26 2022 -0600

    PCI: switchtec: Return -EFAULT for copy_to_user() errors

    [ Upstream commit ddc10938e08cd7aac63d8385f7305f7889df5179 ]

    switchtec_dev_read() didn't handle copy_to_user() errors correctly: it
    assigned "rc = -EFAULT", but actually returned either "size", -ENXIO, or
    -EBADMSG instead.

    Update the failure cases to unlock mrpc_mutex and return -EFAULT directly.

    Link: https://lore.kernel.org/r/20221216162126.207863-3-helgaas@kernel.org
    Fixes: 080b47def5e5 ("MicroSemi Switchtec management interface driver")
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3c277b7fdf76a7902afecc0cf522042306dade5
Author: Alexey V. Vissarionov <gremlin@altlinux.org>
```

```
Date:   Sun Dec 18 06:33:47 2022 +0300

    PCI/IOV: Enlarge virtfn sysfs name buffer

    [ Upstream commit ea0b5aa5f184cf8293c93163f0fb00505190d431 ]

    The sysfs link name "virtfn%u" constructed by pci_iov_sysfs_link() requires
    17 bytes to contain the longest possible string.  Increase VIRTFN_ID_LEN to
    accommodate that.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    [bhelgaas: commit log, comment at #define]
    Fixes: dd7cc44d0bce ("PCI: add SR-IOV API for Physical Function driver")
    Link: https://lore.kernel.org/r/20221218033347.23743-1-gremlin@altlinux.org
    Signed-off-by: Alexey V. Vissarionov <gremlin@altlinux.org>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 226e6fbbf54991740e1e694f163ad28a605daa0c
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Mon Jan 2 22:29:32 2023 +0200

    usb: typec: intel_pmc_mux: Don't leak the ACPI device reference count

    [ Upstream commit c3194949ae8fcbe2b7e38670e7c6a5cfd2605edc ]

    When acpi_dev_get_memory_resources() fails, the reference count is
    left bumped. Drop it as it's done in the other error paths.

    Fixes: 43d596e32276 ("usb: typec: intel_pmc_mux: Check the port status before connect")
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Link: https://lore.kernel.org/r/20230102202933.15968-1-andriy.shevchenko@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2551b9d4dfcb3c7aa16b528c766dd20f766be9fe
Author: Mao Jinlong <quic_jinlmao@quicinc.com>
Date:   Tue Jan 10 11:07:35 2023 +0000

    coresight: cti: Add PM runtime call in enable_store

    [ Upstream commit eff674a9b86a6ffdd10c3af3863545acf7f1ce4f ]

    In commit 6746eae4bbad ("coresight: cti: Fix hang in cti_disable_hw()")
    PM runtime calls are removed from cti_enable_hw/cti_disable_hw. When
    enabling CTI by writing enable sysfs node, clock for accessing CTI
    register won't be enabled. Device will crash due to register access
    issue. Add PM runtime call in enable_store to fix this issue.

    Fixes: 6746eae4bbad ("coresight: cti: Fix hang in cti_disable_hw()")
    Signed-off-by: Mao Jinlong <quic_jinlmao@quicinc.com>
    [Change to only call pm_runtime_put if a disable happened]
    Tested-by: Jinlong Mao <quic_jinlmao@quicinc.com>
    Signed-off-by: James Clark <james.clark@arm.com>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230110110736.2709917-3-james.clark@arm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14d3414b9d7bf22cd76f93c76f5cec845e23193b
Author: James Clark <james.clark@arm.com>
Date:   Tue Jan 10 11:07:34 2023 +0000

    coresight: cti: Prevent negative values of enable count
```

```
    [ Upstream commit 3244fb6dbbf1ffc114cdf382cc167bdd8c18088a ]

    Writing 0 to the enable control repeatedly results in a negative value
    for enable_req_count. After this, writing 1 to the enable control
    appears to not work until the count returns to positive.

    Change it so that it's impossible for enable_req_count to be < 0.
    Return an error to indicate that the disable request was invalid.

    Fixes: 835d722ba10a ("coresight: cti: Initial CoreSight CTI Driver")
    Tested-by: Jinlong Mao <quic_jinlmao@quicinc.com>
    Signed-off-by: James Clark <james.clark@arm.com>
    Reviewed-by: Mike Leach <mike.leach@linaro.org>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230110110736.2709917-2-james.clark@arm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 065a789b413ab69b6a67c6af4970d997ac50c96c
Author: Junhao He <hejunhao3@huawei.com>
Date:   Sat Jan 14 17:16:32 2023 +0800

    coresight: etm4x: Fix accesses to TRCSEQRSTEVR and TRCSEQSTR

    [ Upstream commit 589d928248b72f8377d45904a14bcf686aa8bbeb ]

    The TRCSEQRSTEVR and TRCSEQSTR registers are not implemented if the
    TRCIDR5.NUMSEQSTATE == 0. Skip accessing the registers in such cases.

    Fixes: 2e1cdfe184b5 ("coresight-etm4x: Adding CoreSight ETM4x driver")
    Signed-off-by: Junhao He <hejunhao3@huawei.com>
    Reviewed-by: Mike Leach <mike.leach@linaro.org>
    Reviewed-by: Anshuman Khandual <anshuman.khandual@arm.com>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230114091632.60095-1-hejunhao3@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 43a44b62e5c239ddcfe30575abaa817b10f97b3a
Author: Ricardo Ribalda <ribalda@chromium.org>
Date:   Thu Jan 5 14:52:56 2023 +0100

    media: uvcvideo: Refactor power_line_frequency_controls_limited

    [ Upstream commit 3aa8628eb78a63d0bf93e159846e9c92bccc8f69 ]

    Move the control mapping to uvc_ctrl.c. This way we do not have
    references to UVC controls or V4L2 controls in uvc_driver.c.

    This also fixes a bug introduced in commit 382075604a68 ("media:
    uvcvideo: Limit power line control for Quanta UVC Webcam"). The
    offending commit caused the power line control menu entries to have
    incorrect indices compared to the V4L2_CID_POWER_LINE_FREQUENCY_*
    enumeration. Now that the limited mapping reuses the correct menu_info
    array, the indices correctly map to the V4L2 control specification.

    Fixes: 382075604a68 ("media: uvcvideo: Limit power line control for Quanta UVC Webcam")
    Signed-off-by: Ricardo Ribalda <ribalda@chromium.org>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 93f6ffdad02c3d0f0bca2e5d6e16a2e1277f6760
Author: Ricardo Ribalda <ribalda@chromium.org>
Date:   Thu Jan 5 14:52:55 2023 +0100
```

```
    media: uvcvideo: Refactor uvc_ctrl_mappings_uvcXX

    [ Upstream commit 96a160b068e09b4ed5a32e2179e5219fc3545eca ]

    Convert the array of structs into an array of pointers, that way the
    mappings can be reused.

    Signed-off-by: Ricardo Ribalda <ribalda@chromium.org>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Stable-dep-of: 3aa8628eb78a ("media: uvcvideo: Refactor
power_line_frequency_controls_limited")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 42cbbc6b2c39b02e07cbd24dc2155d4edb99dd04
Author: Ricardo Ribalda <ribalda@chromium.org>
Date:   Thu Jan 5 14:52:54 2023 +0100

    media: uvcvideo: Implement mask for V4L2_CTRL_TYPE_MENU

    [ Upstream commit 40140eda661ea4be219ef194a4f43b7b5e3bbd27 ]

    Replace the count with a mask field that lets us choose not only the max
    value, but also the minimum value and what values are valid in between.

    Suggested-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Ricardo Ribalda <ribalda@chromium.org>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Stable-dep-of: 3aa8628eb78a ("media: uvcvideo: Refactor
power_line_frequency_controls_limited")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90f016587b58c0a1fcc41213eecf6db51ec153eb
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Tue Jan 3 15:36:19 2023 +0100

    media: uvcvideo: Check for INACTIVE in uvc_ctrl_is_accessible()

    [ Upstream commit 9f582f0418ed1c18f92c9e4628075d6ec9a7d9fb ]

    Check for inactive controls in uvc_ctrl_is_accessible().

    Use the new value for the master_id controls if present, otherwise
    use the existing value to determine if it is OK to set the control.
    Doing this here avoids attempting to set an inactive control, which
    will return an error from the USB device, which returns an invalid
    errorcode.

    This fixes:
      warn: v4l2-test-controls.cpp(483): s_ctrl returned EIO
      warn: v4l2-test-controls.cpp(483): s_ctrl returned EIO
    test VIDIOC_G/S_CTRL: OK
      warn: v4l2-test-controls.cpp(739): s_ext_ctrls returned EIO
      warn: v4l2-test-controls.cpp(739): s_ext_ctrls returned EIO
      warn: v4l2-test-controls.cpp(816): s_ext_ctrls returned EIO
    test VIDIOC_G/S/TRY_EXT_CTRLS: OK

    Tested with:
    v4l2-ctl -c auto_exposure=1
    OK
    v4l2-ctl -c exposure_time_absolute=251
    OK
    v4l2-ctl -c auto_exposure=3
    OK
```

```
      v4l2-ctl -c exposure_time_absolute=251
      VIDIOC_S_EXT_CTRLS: failed: Input/output error
      exposure_time_absolute: Input/output error
      ERROR
      v4l2-ctl -c auto_exposure=3,exposure_time_absolute=251,auto_exposure=1
      v4l2-ctl -C auto_exposure,exposure_time_absolute
      auto_exposure: 1
      exposure_time_absolute: 251

      Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
      Reviewed-by: Ricardo Ribalda <ribalda@chromium.org>
      Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
      Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 98395e597198a0f5f911873eec6931e34940be44
  Author: Al Viro <viro@zeniv.linux.org.uk>
  Date:   Mon Jan 9 00:14:02 2023 -0500

      alpha/boot/tools/objstrip: fix the check for ELF header

      [ Upstream commit 1878787797cbb019eeefe6f905514dcd557302b6 ]

      Just memcmp() with ELFMAG - that's the normal way to do it in userland
      code, which that thing is.  Besides, that has the benefit of actually
      building - str_has_prefix() is *NOT* present in <string.h>.

      Fixes: 5f14596e55de "alpha: Replace strncmp with str_has_prefix"
      Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit e6ca1549afb2b495e61a8fb53ae942193ef823e0
  Author: Wang Hai <wanghai38@huawei.com>
  Date:   Tue Dec 20 09:21:43 2022 +0800

      kobject: Fix slab-out-of-bounds in fill_kobj_path()

      [ Upstream commit 3bb2a01caa813d3a1845d378bbe4169ef280d394 ]

      In kobject_get_path(), if kobj->name is changed between calls
      get_kobj_path_length() and fill_kobj_path() and the length becomes
      longer, then fill_kobj_path() will have an out-of-bounds bug.

      The actual current problem occurs when the ixgbe probe.

      In ixgbe_mii_bus_init(), if the length of netdev->dev.kobj.name
      length becomes longer, out-of-bounds will occur.

      cpu0                                            cpu1
      ixgbe_probe
       register_netdev(netdev)
        netdev_register_kobject
         device_add
          kobject_uevent // Sending ADD events
                                                      systemd-udevd // rename netdev
                                                       dev_change_name
                                                        device_rename
                                                         kobject_rename
       ixgbe_mii_bus_init                                 |
        mdiobus_register                                  |
         __mdiobus_register                               |
          device_register                                 |
           device_add                                     |
            kobject_uevent                                |
             kobject_get_path                             |
```

```
             len = get_kobj_path_length // old name  |
             path = kzalloc(len, gfp_mask);          |
                                                     kobj->name = name;
                                                     /* name length becomes
                                                      * longer
                                                      */
             fill_kobj_path /* kobj path length is
                             * longer than path,
                             * resulting in out of
                             * bounds when filling path
                             */
```

This is the kasan report:

```
==================================================================
BUG: KASAN: slab-out-of-bounds in fill_kobj_path+0x50/0xc0
Write of size 7 at addr ff1100090573d1fd by task kworker/28:1/673

 Workqueue: events work_for_cpu_fn
 Call Trace:
 <TASK>
 dump_stack_lvl+0x34/0x48
 print_address_description.constprop.0+0x86/0x1e7
 print_report+0x36/0x4f
 kasan_report+0xad/0x130
 kasan_check_range+0x35/0x1c0
 memcpy+0x39/0x60
 fill_kobj_path+0x50/0xc0
 kobject_get_path+0x5a/0xc0
 kobject_uevent_env+0x140/0x460
 device_add+0x5c7/0x910
 __mdiobus_register+0x14e/0x490
 ixgbe_probe.cold+0x441/0x574 [ixgbe]
 local_pci_probe+0x78/0xc0
 work_for_cpu_fn+0x26/0x40
 process_one_work+0x3b6/0x6a0
 worker_thread+0x368/0x520
 kthread+0x165/0x1a0
 ret_from_fork+0x1f/0x30
```

This reproducer triggers that bug:

```
while:
do
    rmmod ixgbe
    sleep 0.5
    modprobe ixgbe
    sleep 0.5
```

When calling fill_kobj_path() to fill path, if the name length of
kobj becomes longer, return failure and retry. This fixes the problem.

Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
Signed-off-by: Wang Hai <wanghai38@huawei.com>
Link: https://lore.kernel.org/r/20221220012143.52141-1-wanghai38@huawei.com
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17982304806c5c10924e73f7ca5556e0d7378452
Author: Yang Yingliang <yangyingliang@huawei.com>
Date: Mon Dec 5 11:49:04 2022 +0800

    driver core: fix potential null-ptr-deref in device_add()

    [Upstream commit f6837f34a34973ef6600c08195ed300e24e97317]

```
    I got the following null-ptr-deref report while doing fault injection test:

    BUG: kernel NULL pointer dereference, address: 0000000000000058
    CPU: 2 PID: 278 Comm: 37-i2c-ds2482 Tainted: G    B   W        N 6.1.0-rc3+
    RIP: 0010:klist_put+0x2d/0xd0
    Call Trace:
     <TASK>
     klist_remove+0xf1/0x1c0
     device_release_driver_internal+0x196/0x210
     bus_remove_device+0x1bd/0x240
     device_add+0xd3d/0x1100
     w1_add_master_device+0x476/0x490 [wire]
     ds2482_probe+0x303/0x3e0 [ds2482]

    This is how it happened:

    w1_alloc_dev()
      // The dev->driver is set to w1_master_driver.
      memcpy(&dev->dev, device, sizeof(struct device));
      device_add()
        bus_add_device()
        dpm_sysfs_add() // It fails, calls bus_remove_device.

        // error path
        bus_remove_device()
          // The dev->driver is not null, but driver is not bound.
          __device_release_driver()
            klist_remove(&dev->p->knode_driver) <-- It causes null-ptr-deref.

        // normal path
        bus_probe_device() // It's not called yet.
          device_bind_driver()

    If dev->driver is set, in the error path after calling bus_add_device()
    in device_add(), bus_remove_device() is called, then the device will be
    detached from driver. But device_bind_driver() is not called yet, so it
    causes null-ptr-deref while access the 'knode_driver'. To fix this, set
    dev->driver to null in the error path before calling bus_remove_device().

    Fixes: 57eee3d23e88 ("Driver core: Call device_pm_add() after bus_add_device() in
device_add()")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20221205034904.2077765-1-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aec6fe239b49d4c20f98e3f1dba161e4d5d24dd8
Author: Richard Fitzgerald <rf@opensource.cirrus.com>
Date:   Fri Dec 2 16:18:10 2022 +0000

    soundwire: cadence: Don't overflow the command FIFOs

    [ Upstream commit 7cbfee2e2e40d2be54196362a845a3ea0a3f877d ]

    The command FIFOs in the Cadence IP can be configured during design
    up to 32 entries, and the code in cadence_master.c was assuming the
    full 32-entry FIFO. But all current Intel implementations use an 8-entry
    FIFO.

    Up to now the longest message used was 6 entries so this wasn't
    causing any problem. But future Cirrus Logic codecs have downloadable
    firmware or tuning blobs. It is more efficient for the codec driver to
    issue long transfers that can take advantage of any queuing in the
    Soundwire controller and avoid the overhead of repeatedly writing the
```

```
    page registers.

    Signed-off-by: Richard Fitzgerald <rf@opensource.cirrus.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Fixes: 2f52a5177caa ("soundwire: cdns: Add cadence library")
    Link: https://lore.kernel.org/r/20221202161812.4186897-2-rf@opensource.cirrus.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a7b051b0a71ce90df3f862ce02a9af2020541e02
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 28 17:34:02 2022 +0800

    i2c: qcom-geni: change i2c_master_hub to static

    [ Upstream commit 597688792c4d9939e1900f5840ca18804e9d4290 ]

    i2c_master_hub is only used in i2c-qcom-geni.c now,
    change it to static.

    Fixes: cacd9643eca7 ("i2c: qcom-geni: add support for I2C Master Hub variant")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Reviewed-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8637b7105b2ffc7f70599b2d9d8cd449458085b
Author: Hanna Hawa <hhhawa@amazon.com>
Date:   Mon Dec 19 17:23:45 2022 +0000

    i2c: designware: fix i2c_dw_clk_rate() return size to be u32

    [ Upstream commit f2e1fa99550dd7a882229e2c2cd9ecab4ce773d0 ]

    Make i2c_dw_clk_rate() to return u32 instead of unsigned long, as the
    function return the value of get_clk_rate_khz() which returns u32.

    Fixes: b33af11de236 ("i2c: designware: Do not require clock when SSCN and FFCN are
provided")
    Signed-off-by: Hanna Hawa <hhhawa@amazon.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Acked-by: Jarkko Nikula <jarkko.nikula@linux.intel.com>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2a45906e654026850c0a2600080123660290332f
Author: Gaosheng Cui <cuigaosheng1@huawei.com>
Date:   Wed Nov 23 09:41:21 2022 +0800

    usb: gadget: fusb300_udc: free irq on the error path in fusb300_probe()

    [ Upstream commit a8d3392e0e5cfeb03f0cea1f2bc3f5f183c1deb4 ]

    When request_irq(ires1->start) failed in w5300_hw_probe(), irq
    ires->start has not been freed, and on the clean_up3 error path,
    we also need to free ires1->start irq, fix it.

    In addition, We should add free_irq in fusb300_remove(), and give
    the lables a proper name so that they can be understood easily,
    so add free_irq in fusb300_remove(), and update clean_up3 to
    err_alloc_request.

    Fixes: 0fe6f1d1f612 ("usb: udc: add Faraday fusb300 driver")
    Signed-off-by: Gaosheng Cui <cuigaosheng1@huawei.com>
    Link: https://lore.kernel.org/r/20221123014121.1989721-1-cuigaosheng1@huawei.com
```

```
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b621272c19c7cf151d6949627a70d0c4eb9e9bbe
Author: Ferry Toth <ftoth@exalondelft.nl>
Date:   Wed Dec 7 21:03:38 2022 +0200

    iio: light: tsl2563: Do not hardcode interrupt trigger type

    [ Upstream commit 027641b52fe37b64af61025298ce160c8b9b7a73 ]

    Instead of hardcoding IRQ trigger type to IRQF_TRIGGER_RAISING,
    let's respect the settings specified in the firmware description.
    To be compatible with the older firmware descriptions, if trigger
    type is not set up there, we'll set it to default (raising edge).

    Fixes: 388be4883952 ("staging:iio: tsl2563 abi fixes and interrupt handling")
    Fixes: bdab1001738f ("staging:iio:light:tsl2563 remove old style event registration.")
    Signed-off-by: Ferry Toth <ftoth@exalondelft.nl>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20221207190348.9347-1-andriy.shevchenko@linux.intel.com
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8abd2ff2256a2a99c11c7ecdcb5512429933620f
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Dec 23 11:29:00 2022 +0400

    RDMA/hns: Fix refcount leak in hns_roce_mmap

    [ Upstream commit cf6a05c8494a8ae7fec8e5f1229b45ca5b4bcd30 ]

    rdma_user_mmap_entry_get_pgoff() takes the reference.
    Add missing rdma_user_mmap_entry_put() to release the reference.

    Fixes: 0045e0d3f42e ("RDMA/hns: Support direct wqe of userspace")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Acked-by Haoyue Xu <xuhaoyue1@hisilicon.com>
    Link: https://lore.kernel.org/r/20221223072900.802728-1-linmq006@gmail.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e66a55e7a7028557f09ffb57783b804fb89192f
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Nov 21 16:23:53 2022 +0100

    dmaengine: HISI_DMA should depend on ARCH_HISI

    [ Upstream commit dcca9d045c0852584ad092123c7f6e6526a633b1 ]

    The HiSilicon DMA Engine is only present on HiSilicon SoCs.  Hence add a
    dependency on ARCH_HISI, to prevent asking the user about this driver
    when configuring a kernel without HiSilicon SoC support.

    Fixes: e9f08b65250d73ab ("dmaengine: hisilicon: Add Kunpeng DMA engine support")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/r/363a1816d36cd3cf604d88ec90f97c75f604de64.1669044190.git.geert+rene
sas@glider.be
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 410f0f46ffca4d0102470c1e0c747ecfece4204c
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Tue Dec 20 16:11:39 2022 +0400
```

```
    RDMA/erdma: Fix refcount leak in erdma_mmap

    [ Upstream commit ee84146c05ad2316b9a7222d0ec4413e0bf30eeb ]

    rdma_user_mmap_entry_get() take reference, we should release it when not
    need anymore, add the missing rdma_user_mmap_entry_put() in the error
    path to fix it.

    Fixes: 155055771704 ("RDMA/erdma: Add verbs implementation")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Link: https://lore.kernel.org/r/20221220121139.1540564-1-linmq006@gmail.com
    Acked-by: Cheng Xu <chengyou@linux.alibaba.com>
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59079d129d461890750f6b2a0e14c9f650159a69
Author: Fenghua Yu <fenghua.yu@intel.com>
Date:   Fri Dec 9 09:21:41 2022 -0800

    dmaengine: idxd: Set traffic class values in GRPCFG on DSA 2.0

    [ Upstream commit 9735bde36487da43d3c3fc910df49639f72decbf ]

    On DSA/IAX 1.0, TC-A and TC-B in GRPCFG are set as 1 to have best
    performance and cannot be changed through sysfs knobs unless override
    option is given.

    The same values should be set on DSA 2.0 as well.

    Fixes: ea7c8f598c32 ("dmaengine: idxd: restore traffic class defaults after wq reset")
    Fixes: ade8a86b512c ("dmaengine: idxd: Set defaults for GRPCFG traffic class")
    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20221209172141.562648-1-fenghua.yu@intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a62a5e79202967176a9c1a04e477860779accd6c
Author: Qiheng Lin <linqiheng@huawei.com>
Date:   Thu Dec 8 14:15:55 2022 +0800

    mfd: pcf50633-adc: Fix potential memleak in pcf50633_adc_async_read()

    [ Upstream commit 8b450dcff23aa254844492831a8e2b508a9d522d ]

    `req` is allocated in pcf50633_adc_async_read(), but
    adc_enqueue_request() could fail to insert the `req` into queue.
    We need to check the return value and free it in the case of failure.

    Fixes: 08c3e06a5eb2 ("mfd: PCF50633 adc driver")
    Signed-off-by: Qiheng Lin <linqiheng@huawei.com>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20221208061555.8776-1-linqiheng@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65899de920275ec3116dd0703ea763db4a261e01
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Wed Nov 30 17:25:41 2022 -0800

    mfd: cs5535: Don't build on UML

    [ Upstream commit 5ec32a3e4053c1a726b45381d56aa9e39eaf3911 ]

    The cs5535-mfd driver uses CPU-specific data that is not available
```

```
    for ARCH=um builds, so don't allow it to be built for UML.

    Prevents these build errors:

    In file included from ../arch/x86/include/asm/olpc.h:7,
                     from ../drivers/mfd/cs5535-mfd.c:17:
    ../arch/x86/include/asm/geode.h: In function ‘is_geode_gx’:
    ../arch/x86/include/asm/geode.h:16:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86_vendor’
       16 |         return ((boot_cpu_data.x86_vendor == X86_VENDOR_NSC) &&
    ../arch/x86/include/asm/geode.h:16:46: error: ‘X86_VENDOR_NSC’ undeclared (first use in
 this function); did you mean ‘X86_VENDOR_ANY’?
       16 |         return ((boot_cpu_data.x86_vendor == X86_VENDOR_NSC) &&
    ../arch/x86/include/asm/geode.h:17:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86’
       17 |                 (boot_cpu_data.x86 == 5) &&
    ../arch/x86/include/asm/geode.h:18:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86_model’
       18 |                 (boot_cpu_data.x86_model == 5));
    ../arch/x86/include/asm/geode.h: In function ‘is_geode_lx’:
    ../arch/x86/include/asm/geode.h:23:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86_vendor’
       23 |         return ((boot_cpu_data.x86_vendor == X86_VENDOR_AMD) &&
    ../arch/x86/include/asm/geode.h:23:46: error: ‘X86_VENDOR_AMD’ undeclared (first use in
 this function); did you mean ‘X86_VENDOR_ANY’?
       23 |         return ((boot_cpu_data.x86_vendor == X86_VENDOR_AMD) &&
    ../arch/x86/include/asm/geode.h:24:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86’
       24 |                 (boot_cpu_data.x86 == 5) &&
    ../arch/x86/include/asm/geode.h:25:31: error: ‘struct cpuinfo_um’ has no member named
 ‘x86_model’
       25 |                 (boot_cpu_data.x86_model == 10));

    Fixes: 68f5d3f3b654 ("um: add PCI over virtio emulation driver")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20221201012541.11809-1-rdunlap@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit c475cf7218fb1d05b680f3981a494ff5da13fde8
 Author: Tom Fitzhenry <tom@tom-fitzhenry.me.uk>
 Date:   Mon Jan 2 22:11:47 2023 +1100

    mfd: rk808: Re-add rk808-clkout to RK818

    [ Upstream commit 5d69b181cd0db10dc8327d28ce837b3623cd531a ]

    Fixes RK818 (e.g. on Pinephone Pro) to register its clock, without which
    dependent devices (e.g. wifi/BT, via sdio-wifi-pwrseq) fail to probe.

    This line was removed in commit 3633daacea2e
    ("mfd: rk808: Permit having multiple PMIC instances"), but only from RK818.

    Fixes: 3633daacea2e ("mfd: rk808: Permit having multiple PMIC instances")
    Signed-off-by: Tom Fitzhenry <tom@tom-fitzhenry.me.uk>
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Acked-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20230102111147.2580861-1-tom@tom-fitzhenry.me.uk
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit ee0c54cfe39bdddeb26796d9227bcac49f256cbe
 Author: Ondrej Mosnacek <omosnace@redhat.com>
 Date:   Fri Feb 10 15:58:23 2023 +0100
```

```
    sysctl: fix proc_dobool() usability

    [ Upstream commit f1aa2eb5ea05ccd1fd92d235346e60e90a1ed949 ]

    Currently proc_dobool expects a (bool *) in table->data, but sizeof(int)
    in table->maxsize, because it uses do_proc_dointvec() directly.

    This is unsafe for at least two reasons:
    1. A sysctl table definition may use { .data = &variable, .maxsize =
       sizeof(variable) }, not realizing that this makes the sysctl unusable
       (see the Fixes: tag) and that they need to use the completely
       counterintuitive sizeof(int) instead.
    2. proc_dobool() will currently try to parse an array of values if given
       .maxsize >= 2*sizeof(int), but will try to write values of type bool
       by offsets of sizeof(int), so it will not work correctly with neither
       an (int *) nor a (bool *). There is no .maxsize validation to prevent
       this.

    Fix this by:
    1. Constraining proc_dobool() to allow only one value and .maxsize ==
       sizeof(bool).
    2. Wrapping the original struct ctl_table in a temporary one with .data
       pointing to a local int variable and .maxsize set to sizeof(int) and
       passing this one to proc_dointvec(), converting the value to/from
       bool as needed (using proc_dou8vec_minmax() as an example).
    3. Extending sysctl_check_table() to enforce proc_dobool() expectations.
    4. Fixing the proc_dobool() docstring (it was just copy-pasted from
       proc_douintvec, apparently...).
    5. Converting all existing proc_dobool() users to set .maxsize to
       sizeof(bool) instead of sizeof(int).

    Fixes: 83efeeeb3d04 ("tty: Allow TIOCSTI to be disabled")
    Fixes: a2071573d634 ("sysctl: introduce new proc handler proc_dobool")
    Signed-off-by: Ondrej Mosnacek <omosnace@redhat.com>
    Acked-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Luis Chamberlain <mcgrof@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68a3c4e04664ac5a02e6950965462329b602bcd2
Author: Masami Hiramatsu (Google) <mhiramat@kernel.org>
Date:   Tue Feb 21 08:49:16 2023 +0900

    selftests/ftrace: Fix probepoint testcase to ignore __pfx_* symbols

    [ Upstream commit 96cd93af794cf3ef83ae1ad7291160029d7b525e ]

    Fix kprobe probepoint testcase to ignore __pfx_* prefix symbols. Those are
    introduced by commit b341b20d648b ("x86: Add prefix symbols for function
    padding") for identifying PADDING_BYTES of NOPs. Since kprobe events can
    not probe these prefix symbols, this testcase has to skip those symbols.

    Link:
https://lore.kernel.org/all/167309835609.640500.9664678940260305746.stgit@devnote3/

    Fixes: b341b20d648b ("x86: Add prefix symbols for function padding")
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Reviewed-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Acked-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2db826d7093833e2068ada98605b8947bfba4a3
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Wed Feb 15 14:00:58 2023 +0100

    objtool: add UACCESS exceptions for __tsan_volatile_read/write
```

```
    [ Upstream commit d5d469247264e56960705dc5ae7e1d014861fe40 ]

    A lot of the tsan helpers are already excempt from the UACCESS warnings,
    but some more functions were added that need the same thing:

    kernel/kcsan/core.o: warning: objtool: __tsan_volatile_read16+0x0: call to
__tsan_unaligned_read16() with UACCESS enabled
    kernel/kcsan/core.o: warning: objtool: __tsan_volatile_write16+0x0: call to
__tsan_unaligned_write16() with UACCESS enabled
    vmlinux.o: warning: objtool: __tsan_unaligned_volatile_read16+0x4: call to
__tsan_unaligned_read16() with UACCESS enabled
    vmlinux.o: warning: objtool: __tsan_unaligned_volatile_write16+0x4: call to
__tsan_unaligned_write16() with UACCESS enabled

    As Marco points out, these functions don't even call each other
    explicitly but instead gcc (but not clang) notices the functions
    being identical and turns one symbol into a direct branch to the
    other.

    Link: https://lkml.kernel.org/r/20230215130058.3836177-4-arnd@kernel.org
    Fixes: 75d75b7a4d54 ("kcsan: Support distinguishing volatile accesses")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Acked-by: Marco Elver <elver@google.com>
    Cc: Alexander Potapenko <glider@google.com>
    Cc: Andrey Konovalov <andreyknvl@gmail.com>
    Cc: Andrey Ryabinin <ryabinin.a.a@gmail.com>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: Josh Poimboeuf <jpoimboe@kernel.org>
    Cc: Kuan-Ying Lee <Kuan-Ying.Lee@mediatek.com>
    Cc: Peter Zijlstra (Intel) <peterz@infradead.org>
    Cc: Vincenzo Frascino <vincenzo.frascino@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e4974fc546ffae464333135c0eeb7a428a2402f
Author: Kajol Jain <kjain@linux.ibm.com>
Date:   Wed Feb 15 15:08:27 2023 +0530

    perf tests stat_all_metrics: Change true workload to sleep workload for system wide
check

    [ Upstream commit f9fa0778ee7349a9aa3d2ea10e9f2ab843a0b44e ]

    Testcase stat_all_metrics.sh fails in powerpc:

    98: perf all metrics test : FAILED!

    Logs with verbose:

      [command]# ./perf test 98 -vv
       98: perf all metrics test                                          :
       --- start ---
      test child forked, pid 13262
      Testing BRU_STALL_CPI
      Testing COMPLETION_STALL_CPI
       ----
      Testing TOTAL_LOCAL_NODE_PUMPS_P23
      Metric 'TOTAL_LOCAL_NODE_PUMPS_P23' not printed in:
      Error:
      Invalid event (hv_24x7/PM_PB_LNS_PUMP23,chip=3/) in per-thread mode, enable system
wide with '-a'.
      Testing TOTAL_LOCAL_NODE_PUMPS_RETRIES_P01
      Metric 'TOTAL_LOCAL_NODE_PUMPS_RETRIES_P01' not printed in:
      Error:
```

```
       Invalid event (hv_24x7/PM_PB_RTY_LNS_PUMP01,chip=3/) in per-thread mode, enable system
  wide with '-a'.
       ----
```

    Based on above logs, we could see some of the hv-24x7 metric events
    fails, and logs suggest to run the metric event with -a option. This
    change happened after the commit a4b8cfcabb1d90ec ("perf stat: Delay
    metric parsing"), which delayed the metric parsing phase and now before
    metric parsing phase perf tool identifies, whether target is system-wide
    or not. With this change, perf_event_open will fails with workload
    monitoring for uncore events as expected.

    The perf all metric test case fails as some of the hv-24x7 metric events
    may need bigger workload with system wide monitoring to get the data.
    Fix this issue by changing current system wide check from true workload
    to sleep 0.01 workload.

    Result with the patch changes in powerpc:

      98: perf all metrics test : Ok

    Fixes: a4b8cfcabb1d90ec ("perf stat: Delay metric parsing")
    Suggested-by: Ian Rogers <irogers@google.com>
    Reviewed-by: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
    Signed-off-by: Kajol Jain <kjain@linux.ibm.com>
    Tested-by: Disha Goel <disgoel@linux.ibm.com>
    Tested-by: Ian Rogers <irogers@google.com>
    Cc: Madhavan Srinivasan <maddy@linux.ibm.com>
    Cc: Nageswara R Sastry <rnsastry@linux.ibm.com>
    Cc: linuxppc-dev@lists.ozlabs.org
    Link: https://lore.kernel.org/r/20230215093827.124921-1-kjain@linux.ibm.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c8ffa6de2c44db77cf48790a512d371ec075ddc
Author: Arnd Bergmann <arnd@arndb.de>
Date: Mon Feb 6 20:40:57 2023 +0100

    printf: fix errname.c list

    [Upstream commit 0c2baf6509af1d11310ae4c1c839481a6e9a4bc4]

    On most architectures, gcc -Wextra warns about the list of error
    numbers containing both EDEADLK and EDEADLOCK:

```
    lib/errname.c:15:67: warning: initialized field overwritten [-Woverride-init]
       15 | #define E(err) [err + BUILD_BUG_ON_ZERO(err <= 0 || err > 300)] = "-" #err
          |                                                                   ^~~
    lib/errname.c:172:2: note: in expansion of macro 'E'
      172 |  E(EDEADLK), /* EDEADLOCK */
          |  ^
```

    On parisc, a similar error happens with -ECANCELLED, which is an
    alias for ECANCELED.

    Make the EDEADLK printing conditional on the number being distinct
    from EDEADLOCK, and remove the -ECANCELLED bit completely as it
    can never be hit.

    To ensure these are correct, add static_assert lines that verify
    all the remaining aliases are in fact identical to the canonical
    name.

    Fixes: 57f5677e535b ("printf: add support for printing symbolic error names")
    Cc: Petr Mladek <pmladek@suse.com>

```
    Suggested-by: Rasmus Villemoes <linux@rasmusvillemoes.dk>
    Acked-by: Uwe Kleine-König <uwe@kleine-koenig.org>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Link: https://lore.kernel.org/all/20210514213456.745039-1-arnd@kernel.org/
    Link: https://lore.kernel.org/all/20210927123409.1109737-1-arnd@kernel.org/
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Sergey Senozhatsky <senozhatsky@chromium.org>
    Acked-by: Rasmus Villemoes <linux@rasmusvillemoes.dk>
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Signed-off-by: Petr Mladek <pmladek@suse.com>
    Link: https://lore.kernel.org/r/20230206194126.380350-1-arnd@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 476ffae1d766bafe117afcc461f29d5d7b51cae1
Author: Yang Jihong <yangjihong1@huawei.com>
Date:   Wed Feb 15 12:23:24 2023 +0000

    perf record: Fix segfault with --overwrite and --max-size

    [ Upstream commit 91621be65d6812cd74b2ea09573ff9ee0cbf5666 ]

    When --overwrite and --max-size options of perf record are used
    together, a segmentation fault occurs. The following is an example:

      # perf record -e sched:sched* --overwrite --max-size 1K -a -- sleep 1
      [ perf record: Woken up 1 times to write data ]
      perf: Segmentation fault
      Obtained 12 stack frames.
      ./perf/perf(+0x197673) [0x55f99710b673]
      /lib/x86_64-linux-gnu/libc.so.6(+0x3ef0f) [0x7fa45f3cff0f]
      ./perf/perf(+0x8eb40) [0x55f997002b40]
      ./perf/perf(+0x1f6882) [0x55f99716a882]
      ./perf/perf(+0x794c2) [0x55f996fed4c2]
      ./perf/perf(+0x7b7c7) [0x55f996fef7c7]
      ./perf/perf(+0x9074b) [0x55f99700474b]
      ./perf/perf(+0x12e23c) [0x55f9970a223c]
      ./perf/perf(+0x12e54a) [0x55f9970a254a]
      ./perf/perf(+0x7db60) [0x55f996ff1b60]
      /lib/x86_64-linux-gnu/libc.so.6(__libc_start_main+0xe6) [0x7fa45f3b2c86]
      ./perf/perf(+0x7dfe9) [0x55f996ff1fe9]
      Segmentation fault (core dumped)

    backtrace of the core file is as follows:

      (gdb) bt
      #0  record__bytes_written (rec=0x55f99755a200 <record>) at builtin-record.c:234
      #1  record__output_max_size_exceeded (rec=0x55f99755a200 <record>) at builtin-
record.c:242
      #2  record__write (map=0x0, size=12816, bf=0x55f9978da2e0, rec=0x55f99755a200
<record>) at builtin-record.c:263
      #3  process_synthesized_event (tool=tool@entry=0x55f99755a200 <record>,
event=event@entry=0x55f9978da2e0, sample=sample@entry=0x0,
machine=machine@entry=0x55f997893658) at builtin-record.c:618
      #4  0x000055f99716a883 in __perf_event__synthesize_id_index
(tool=tool@entry=0x55f99755a200 <record>, process=process@entry=0x55f997002aa0
<process_synthesized_event>, evlist=0x55f9978928b0, machine=machine@entry=0x55f997893658,
          from=from@entry=0) at util/synthetic-events.c:1895
      #5  0x000055f99716a91f in perf_event__synthesize_id_index
(tool=tool@entry=0x55f99755a200 <record>, process=process@entry=0x55f997002aa0
<process_synthesized_event>, evlist=<optimized out>, machine=machine@entry=0x55f997893658)
          at util/synthetic-events.c:1905
      #6  0x000055f996fed4c3 in record__synthesize (tail=tail@entry=true, rec=0x55f99755a200
<record>) at builtin-record.c:1997
      #7  0x000055f996fef7c8 in __cmd_record (argc=argc@entry=2,
argv=argv@entry=0x7ffc67551260, rec=0x55f99755a200 <record>) at builtin-record.c:2802
```

```
      #8  0x000055f99700474c in cmd_record (argc=<optimized out>, argv=0x7ffc67551260) at
builtin-record.c:4258
      #9  0x000055f9970a223d in run_builtin (p=0x55f997564d88 <commands+264>, argc=10,
argv=0x7ffc67551260) at perf.c:330
      #10 0x000055f9970a254b in handle_internal_command (argc=10, argv=0x7ffc67551260) at
perf.c:384
      #11 0x000055f996ff1b61 in run_argv (argcp=<synthetic pointer>, argv=<synthetic
pointer>) at perf.c:428
      #12 main (argc=<optimized out>, argv=0x7ffc67551260) at perf.c:562

    The reason is that record__bytes_written accesses the freed memory rec->thread_data,
    The process is as follows:
      __cmd_record
        -> record__free_thread_data
          -> zfree(&rec->thread_data)         // free rec->thread_data
        -> record__synthesize
          -> perf_event__synthesize_id_index
            -> process_synthesized_event
              -> record__write
                -> record__bytes_written      // access rec->thread_data

    We add a member variable "thread_bytes_written" in the struct "record"
    to save the data size written by the threads.

    Fixes: 6d57581659f72299 ("perf record: Add support for limit perf output file size")
    Signed-off-by: Yang Jihong <yangjihong1@huawei.com>
    Acked-by: Namhyung Kim <namhyung@kernel.org>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Jiwei Sun <jiwei.sun@windriver.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Link: https://lore.kernel.org/r/CAM9d7ci_TRrqBQVQNW8=GwakUr7SsZpYxaaty-
S4bxF8zJWyqw@mail.gmail.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd6bb993741b655736f61d109871a57a1af3966c
Author: Guillaume Tucker <guillaume.tucker@collabora.com>
Date:   Thu Feb 9 09:55:36 2023 +0100

    selftests: use printf instead of echo -ne

    [ Upstream commit 9e34fad00fc889abbb99d751a4c22cf2bded10df ]

    Rather than trying to guess which implementation of "echo" to run with
    support for "-ne" options, use "printf" instead of "echo -ne".  It
    handles escape characters as a standard feature and it is widespread
    among modern shells.

    Reported-by: "kernelci.org bot" <bot@kernelci.org>
    Suggested-by: David Laight <David.Laight@ACULAB.COM>
    Fixes: 3297a4df805d ("kselftests: Enable the echo command to print newlines in
Makefile")
    Fixes: 79c16b1120fe ("selftests: find echo binary to use -ne options")
    Signed-off-by: Guillaume Tucker <guillaume.tucker@collabora.com>
    Reviewed-by: Guenter Roeck <groeck@chromium.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d58872b3f20622a7a3b253fe2d182c481f627b55
Author: Masami Hiramatsu (Google) <mhiramat@kernel.org>
Date:   Sun Jan 22 08:32:50 2023 +0900
```

```
    selftests/ftrace: Fix bash specific "==" operator

    [ Upstream commit 1e6b485c922fbedf41d5a9f4e6449c5aeb923a32 ]

    Since commit a1d6cd88c897 ("selftests/ftrace: event_triggers: wait
    longer for test_event_enable") introduced bash specific "=="
    comparation operator, that test will fail when we run it on a
    posix-shell. `checkbashisms` warned it as below.

    possible bashism in ftrace/func_event_triggers.tc line 45 (should be 'b = a'):
            if [ "$e" == $val ]; then

    This replaces it with "=".

    Fixes: a1d6cd88c897 ("selftests/ftrace: event_triggers: wait longer for
test_event_enable")
    Signed-off-by: Masami Hiramatsu (Google) <mhiramat@kernel.org>
    Reviewed-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32e490889304a0a0b94313be7058ff954db2ff81
Author: Guillaume Tucker <guillaume.tucker@collabora.com>
Date:   Fri Feb 3 16:26:03 2023 +0100

    selftests: find echo binary to use -ne options

    [ Upstream commit 4ebe33398c40c1118b4d8546978036c0e0032d1b ]

    Find the actual echo binary using $(which echo) and use it for
    formatted output with -ne.  On some systems, the default echo command
    doesn't handle the -e option and the output looks like this (arm64
    build):

    -ne Emit Tests for alsa

    -ne Emit Tests for amd-pstate

    -ne Emit Tests for arm64

    This is for example the case with the KernelCI Docker images
    e.g. kernelci/gcc-10:x86-kselftest-kernelci.  With the actual echo
    binary (e.g. in /bin/echo), the output is formatted as expected (x86
    build this time):

    Emit Tests for alsa
    Emit Tests for amd-pstate
    Skipping non-existent dir: arm64

    Only the install target is using "echo -ne" so keep the $ECHO variable
    local to it.

    Reported-by: "kernelci.org bot" <bot@kernelci.org>
    Fixes: 3297a4df805d ("kselftests: Enable the echo command to print newlines in
Makefile")
    Signed-off-by: Guillaume Tucker <guillaume.tucker@collabora.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2aa10cb545b6b4ade5795958eed1591efe01d1c
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sat Feb 4 16:43:57 2023 -0800

    sparc: allow PM configs for sparc32 COMPILE_TEST
```

```
    [ Upstream commit 7be6a87c2473957090995b7eb541e31d57a2c801 ]

    When doing randconfig builds for sparc32 with COMPILE_TEST, some
    (non-Sparc) drivers cause kconfig warnings with the Kconfig symbols PM,
    PM_GENERIC_DOMAINS, or PM_GENERIC_DOMAINS_OF.

    This is due to arch/sparc/Kconfig not using the PM Kconfig for
    Sparc32:

      if SPARC64
      source "kernel/power/Kconfig"
      endif

    Arnd suggested adding "|| COMPILE_TEST" to the conditional,
    instead of trying to track down every driver that selects
    any of these PM symbols.

    Fixes the following kconfig warnings:

    WARNING: unmet direct dependencies detected for PM
      Depends on [n]: SPARC64 [=n]
      Selected by [y]:
      - SUN20I_PPU [=y] && (ARCH_SUNXI || COMPILE_TEST [=y])

    WARNING: unmet direct dependencies detected for PM
      Depends on [n]: SPARC64 [=n]
      Selected by [y]:
      - SUN20I_PPU [=y] && (ARCH_SUNXI || COMPILE_TEST [=y])

    WARNING: unmet direct dependencies detected for PM_GENERIC_DOMAINS
      Depends on [n]: SPARC64 [=n] && PM [=y]
      Selected by [y]:
      - QCOM_GDSC [=y] && COMMON_CLK [=y] && PM [=y]
      - SUN20I_PPU [=y] && (ARCH_SUNXI || COMPILE_TEST [=y])
      - MESON_GX_PM_DOMAINS [=y] && (ARCH_MESON || COMPILE_TEST [=y]) && PM [=y] && OF [=y]
      - BCM2835_POWER [=y] && (ARCH_BCM2835 || COMPILE_TEST [=y] && OF [=y]) && PM [=y]
      - BCM_PMB [=y] && (ARCH_BCMBCA || COMPILE_TEST [=y] && OF [=y]) && PM [=y]
      - ROCKCHIP_PM_DOMAINS [=y] && (ARCH_ROCKCHIP || COMPILE_TEST [=y]) && PM [=y]
      Selected by [m]:
      - ARM_SCPI_POWER_DOMAIN [=m] && (ARM_SCPI_PROTOCOL [=m] || COMPILE_TEST [=y] && OF
[=y]) && PM [=y]
      - MESON_EE_PM_DOMAINS [=m] && (ARCH_MESON || COMPILE_TEST [=y]) && PM [=y] && OF [=y]
      - QCOM_AOSS_QMP [=m] && (ARCH_QCOM || COMPILE_TEST [=y]) && MAILBOX [=y] && COMMON_CLK
[=y] && PM [=y]

    WARNING: unmet direct dependencies detected for PM_GENERIC_DOMAINS_OF
      Depends on [n]: SPARC64 [=n] && PM_GENERIC_DOMAINS [=y] && OF [=y]
      Selected by [y]:
      - MESON_GX_PM_DOMAINS [=y] && (ARCH_MESON || COMPILE_TEST [=y]) && PM [=y] && OF [=y]
      Selected by [m]:
      - MESON_EE_PM_DOMAINS [=m] && (ARCH_MESON || COMPILE_TEST [=y]) && PM [=y] && OF [=y]

    Link: https://lkml.kernel.org/r/20230205004357.29459-1-rdunlap@infradead.org
    Fixes: bdde6b3c8ba4 ("sparc64: Hibernation support")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Acked-by: Sam Ravnborg <sam@ravnborg.org>
    Cc: "David S. Miller" <davem@davemloft.net>
    Cc: Kirill Tkhai <tkhai@yandex.ru>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b16aa691ef807a12b2e57bf5912ae13867d45eb7
  Author: Ian Rogers <irogers@google.com>
  Date:   Wed Feb 8 22:44:47 2023 -0800
```

```
    perf stat: Avoid merging/aggregating metric counts twice

    [ Upstream commit 37f322cd58d81a9d46456531281c908de9ef6e42 ]

    The added perf_stat_merge_counters combines uncore counters. When
    metrics are enabled, the counts are merged into a metric_leader via the
    stat-shadow saved_value logic. As the leader now is passed an aggregated
    count, it leads to all counters being added together twice and counts
    appearing approximately doubled in metrics.

    This change disables the saved_value merging of counts for evsels that
    are merged. It is recommended that later changes remove the saved_value
    entirely as the two layers of aggregation in the code is confusing.

    Fixes: 942c5593393d9418 ("perf stat: Add perf_stat_merge_counters()")
    Reported-by: Perry Taylor <perry.taylor@intel.com>
    Signed-off-by: Ian Rogers <irogers@google.com>
    Acked-by: Namhyung Kim <namhyung@kernel.org>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Eduard Zingerman <eddyz87@gmail.com>
    Cc: Florian Fischer <florian.fischer@muhq.space>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: James Clark <james.clark@arm.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Stephane Eranian <eranian@google.com>
    Cc: Xing Zhengjun <zhengjun.xing@linux.intel.com>
    Link: https://lore.kernel.org/r/20230209064447.83733-1-irogers@google.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c65e272d1623abd1e1d129270b306f61ce379df3
Author: Yicong Yang <yangyicong@hisilicon.com>
Date:   Tue Feb 7 11:50:57 2023 +0800

    perf tools: Fix auto-complete on aarch64

    [ Upstream commit ffd1240e8f0814262ceb957dbe961f6e0aef1e7a ]

    On aarch64 CPU related events are not under event_source/devices/cpu/events,
    they're under event_source/devices/armv8_pmuv3_0/events on my machine.
    Using current auto-complete script will generate below error:

      [root@localhost bin]# perf stat -e
      ls: cannot access '/sys/bus/event_source/devices/cpu/events': No such file or
directory

    Fix this by not testing /sys/bus/event_source/devices/cpu/events on
    aarch64 machine.

    Fixes: 74cd5815d9af6e6c ("perf tool: Improve bash command line auto-complete for
multiple events with comma")
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Yicong Yang <yangyicong@hisilicon.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Jin Yao <yao.jin@linux.intel.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: linux-arm-kernel@lists.infradead.org
```

```
      Cc: linuxarm@huawei.com
      Cc: prime.zeng@hisilicon.com
      Link: https://lore.kernel.org/r/20230207035057.43394-1-yangyicong@huawei.com
      Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 0294495a4530517fd0803631a89b7af93e2fbb47
  Author: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
  Date:   Thu Jan 5 17:47:42 2023 +0530

      perf test bpf: Skip test if kernel-debuginfo is not present

      [ Upstream commit 34266f904abd45731bdade2e92d0536c092ee9bc ]

      Perf BPF filter test fails in environment where "kernel-debuginfo"
      is not installed.

      Test failure logs:

        <<>>
        42: BPF filter                            :
        42.1: Basic BPF filtering                 : Ok
        42.2: BPF pinning                         : Ok
        42.3: BPF prologue generation             : FAILED!
        <<>>

      Enabling verbose option provided debug logs, which says debuginfo
      needs to be installed. Snippet of verbose logs:

        <<>>
        42.3: BPF prologue generation                                      :
        --- start ---
        test child forked, pid 28218
        <<>>
        Rebuild with CONFIG_DEBUG_INFO=y, or install an appropriate debuginfo
        package.
        bpf_probe: failed to convert perf probe events
        Failed to add events selected by BPF
        test child finished with -1
        ---- end ----
        BPF filter subtest 3: FAILED!
        <<>>

      Here the subtest "BPF prologue generation" failed and logs shows
      debuginfo is needed. After installing kernel-debuginfo package, testcase
      passes.

      The "BPF prologue generation" subtest failed because, the do_test()
      returns TEST_FAIL without checking the error type returned by
      parse_events_load_bpf_obj().

      parse_events_load_bpf_obj() can also return error of type -ENODATA
      incase kernel-debuginfo package is not installed. Fix this by adding
      check for -ENODATA error.

      Test result after the patch changes:

      Test failure logs:

        <<>>
        42: BPF filter                :
        42.1: Basic BPF filtering     : Ok
        42.2: BPF pinning             : Ok
        42.3: BPF prologue generation  : Skip (clang/debuginfo isn't installed or environment
  missing BPF support)
```

```
        <<>>

    Fixes: ba1fae431e74bb42 ("perf test: Add 'perf test BPF'")
    Signed-off-by: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
    Cc: Andi Kleen <ak@linux.intel.com>
    Cc: Disha Goel <disgoel@linux.ibm.com>
    Cc: Ian Rogers <irogers@google.com>
    Cc: James Clark <james.clark@arm.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Kajol Jain <kjain@linux.ibm.com>
    Cc: Madhavan Srinivasan <maddy@linux.vnet.ibm.com>
    Cc: Michael Ellerman <mpe@ellerman.id.au>
    Cc: Nageswara R Sastry <rnsastry@linux.ibm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Wang Nan <wangnan0@huawei.com>
    Cc: linuxppc-dev@lists.ozlabs.org
    Link: http://lore.kernel.org/linux-perf-users/Y7bIk77mdE4j8Jyi@kernel.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a40e866214e8604bc7a3a61815712f4525386f47
Author: Ian Rogers <irogers@google.com>
Date:   Thu Jan 26 15:36:43 2023 -0800

    perf jevents: Correct bad character encoding

    [ Upstream commit d2e3dc829e389d686194d06f0a64adda4158faae ]

    A character encoding issue added a "3D" character that breaks the
    metrics test.

    Fixes: 40769665b63d8c84 ("perf jevents: Parse metrics during conversion")
    Reviewed-by: Kajol Jain <kjain@linux.ibm.com>
    Signed-off-by: Ian Rogers <irogers@google.com>
    Cc: Adrian Hunter <adrian.hunter@intel.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Caleb Biggers <caleb.biggers@intel.com>
    Cc: Florian Fischer <florian.fischer@muhq.space>
    Cc: Ian Rogers <irogers@google.com>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: James Clark <james.clark@arm.com>
    Cc: Jing Zhang <renyu.zj@linux.alibaba.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: John Garry <john.g.garry@oracle.com>
    Cc: Kan Liang <kan.liang@linux.intel.com>
    Cc: Kang Minchul <tegongkang@gmail.com>
    Cc: Kim Phillips <kim.phillips@amd.com>
    Cc: Leo Yan <leo.yan@linaro.org>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Mike Leach <mike.leach@linaro.org>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Perry Taylor <perry.taylor@intel.com>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Ravi Bangoria <ravi.bangoria@amd.com>
    Cc: Rob Herring <robh@kernel.org>
    Cc: Sandipan Das <sandipan.das@amd.com>
    Cc: Stephane Eranian <eranian@google.com>
    Cc: Will Deacon <will@kernel.org>
    Cc: Xing Zhengjun <zhengjun.xing@linux.intel.com>
    Cc: linux-arm-kernel@lists.infradead.org
    Cc: linuxppc-dev@lists.ozlabs.org
    Link: https://lore.kernel.org/r/20230126233645.200509-14-irogers@google.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 31d8fc3756b140ba23517a0cced368130bf37b96
Author: Namhyung Kim <namhyung@kernel.org>
Date:   Wed Jan 25 11:24:31 2023 -0800

    perf stat: Hide invalid uncore event output for aggr mode

    [ Upstream commit dd15480a3d67b9cf04a1f6f5d60f1c0dc018e22f ]

    The current display code for perf stat iterates given cpus and build the
    aggr map to collect the event data for the aggregation mode.

    But uncore events have their own cpu maps and it won't guarantee that
    it'd match to the aggr map.  For example, per-package uncore events
    would generate a single value for each socket.  When user asks per-core
    aggregation mode, the output would contain 0 values for other cores.

    Thus it needs to check the uncore PMU's cpumask and if it matches to the
    current aggregation id.

    Before:
      $ sudo ./perf stat -a --per-core -e power/energy-pkg/ sleep 1

       Performance counter stats for 'system wide':

      S0-D0-C0              1              3.73 Joules power/energy-pkg/
      S0-D0-C1              0      <not counted> Joules power/energy-pkg/
      S0-D0-C2              0      <not counted> Joules power/energy-pkg/
      S0-D0-C3              0      <not counted> Joules power/energy-pkg/

            1.001404046 seconds time elapsed

      Some events weren't counted. Try disabling the NMI watchdog:
            echo 0 > /proc/sys/kernel/nmi_watchdog
            perf stat ...
            echo 1 > /proc/sys/kernel/nmi_watchdog

    The core 1, 2 and 3 should not be printed because the event is handled
    in a cpu in the core 0 only.  With this change, the output becomes like
    below.

    After:
      $ sudo ./perf stat -a --per-core -e power/energy-pkg/ sleep 1

       Performance counter stats for 'system wide':

      S0-D0-C0              1              2.09 Joules power/energy-pkg/

    Fixes: b897613510890d6e ("perf stat: Update event skip condition for system-wide per-
thread mode and merged uncore and hybrid events")
    Signed-off-by: Namhyung Kim <namhyung@kernel.org>
    Tested-by: Ian Rogers <irogers@google.com>
    Acked-by: Ian Rogers <irogers@google.com>
    Cc: Adrian Hunter <adrian.hunter@intel.com>
    Cc: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Michael Petlan <mpetlan@redhat.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Link: https://lore.kernel.org/r/20230125192431.2929677-1-namhyung@kernel.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ea05aea5f5e7ddea8ef6e64266eca6e457a4dae9
Author: Namhyung Kim <namhyung@kernel.org>
```

```
Date:   Mon Jan 30 18:33:48 2023 -0800

    perf intel-pt: Do not try to queue auxtrace data on pipe

    [ Upstream commit aeb802f872a7c42e4381f36041e77d1745908255 ]

    When it processes AUXTRACE_INFO, it calls to auxtrace_queue_data() to
    collect AUXTRACE data first.  That won't work with pipe since it needs
    lseek() to read the scattered aux data.

      $ perf record -o- -e intel_pt// true | perf report -i- --itrace=i100
      # To display the perf.data header info, please use --header/--header-only options.
      #
      0x4118 [0xa0]: failed to process type: 70
      Error:
      failed to process sample

    For the pipe mode, it can handle the aux data as it gets.  But there's
    no guarantee it can get the aux data in time.  So the following warning
    will be shown at the beginning:

      WARNING: Intel PT with pipe mode is not recommended.
               The output cannot relied upon.  In particular,
               time stamps and the order of events may be incorrect.

    Fixes: dbd134322e74f19d ("perf intel-pt: Add support for decoding AUX area samples")
    Reviewed-by: Adrian Hunter <adrian.hunter@intel.com>
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Namhyung Kim <namhyung@kernel.org>
    Cc: Adrian Hunter <adrian.hunter@intel.com>
    Cc: Ian Rogers <irogers@google.com>
    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Leo Yan <leo.yan@linaro.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Stephane Eranian <eranian@google.com>
    Link: https://lore.kernel.org/r/20230131023350.1903992-3-namhyung@kernel.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4523045c8f3d9514eb79d97883d20f59b941866d
Author: Namhyung Kim <namhyung@kernel.org>
Date:   Mon Jan 30 18:33:47 2023 -0800

    perf inject: Use perf_data__read() for auxtrace

    [ Upstream commit 1746212daeba95e9ae1639227dc0c3591d41deeb ]

    In copy_bytes(), it reads the data from the (input) fd and writes it to
    the output file.  But it does with the read(2) unconditionally which
    caused a problem of mixing buffered vs unbuffered I/O together.

    You can see the problem when using pipes.

      $ perf record -e intel_pt// -o- true | perf inject -b > /dev/null
      [ perf record: Woken up 1 times to write data ]
      [ perf record: Captured and wrote 0.000 MB - ]
      0x45c0 [0x30]: failed to process type: 71

    It should use perf_data__read() to honor the 'use_stdio' setting.

    Fixes: 601366678c93618f ("perf data: Allow to use stdio functions for pipe mode")
    Reviewed-by: Adrian Hunter <adrian.hunter@intel.com>
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Namhyung Kim <namhyung@kernel.org>
```

```
    Cc: Ian Rogers <irogers@google.com>
    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: Leo Yan <leo.yan@linaro.org>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Stephane Eranian <eranian@google.com>
    Link: https://lore.kernel.org/r/20230131023350.1903992-2-namhyung@kernel.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 745e536ed3c64020a24a084dbffaa79a68698b68
Author: Andreas Ziegler <br015@umbiko.net>
Date:   Tue Jan 3 11:33:59 2023 +0100

    tools/tracing/rtla: osnoise_hist: use total duration for average calculation

    [ Upstream commit fe137a4fe0e77eb95396cfc5c3dd7df404421aa4 ]

    Sampled durations must be weighted by observed quantity, to arrive at a correct
    average duration value.

    Perform calculation of total duration by summing (duration * count).

    Link: https://lkml.kernel.org/r/20230103103400.275566-2-br015@umbiko.net

    Fixes: 829a6c0b5698 ("rtla/osnoise: Add the hist mode")

    Signed-off-by: Andreas Ziegler <br015@umbiko.net>
    Acked-by: Daniel Bristot de Oliveira <bristot@kernel.org>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3012112688f99279930bb06487281c6325a7ff5e
Author: Henning Schild <henning.schild@siemens.com>
Date:   Fri Oct 7 17:33:23 2022 +0200

    leds: simatic-ipc-leds-gpio: Make sure we have the GPIO providing driver

    [ Upstream commit c64964ebee2a415384385205950ee7a05f78451e ]

    If we register a "leds-gpio" platform device for GPIO pins that do not
    exist we get a -EPROBE_DEFER and the probe will be tried again later.
    If there is no driver to provide that pin we will poll forever and also
    create a lot of log messages.

    So check if that GPIO driver is configured, if so it will come up
    eventually. If not, we exit our probe function early and do not even
    bother registering the "leds-gpio". This method was chosen over "Kconfig
    depends" since this way we can add support for more devices and GPIO
    backends more easily without "depends":ing on all GPIO backends.

    Fixes: a6c80bec3c93 ("leds: simatic-ipc-leds-gpio: Add GPIO version of Siemens driver")
    Signed-off-by: Henning Schild <henning.schild@siemens.com>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20221007153323.1326-1-henning.schild@siemens.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f99e2d065db1e5c9b225bcf2bb57c4b6c9e8d581
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Wed Dec 28 11:32:38 2022 +0200

    leds: is31fl319x: Wrap mutex_destroy() for devm_add_action_or_rest()
```

```
    [ Upstream commit a82c7cf803d98751cd3ddb35828faad925d71982 ]

    Clang complains that devm_add_action() takes a parameter with a wrong type:

    warning: cast from 'void (*)(struct mutex *)' to 'void (*)(void *)' converts to
 incompatible function type [-Wcast-function-type-strict]
        err = devm_add_action(dev, (void (*)(void *))mutex_destroy, &is31->lock);
                              ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
        1 warning generated.

    It appears that the commit e1af5c815586 ("leds: is31fl319x: Fix devm vs.
    non-devm ordering") missed two things:

    - whilst the commit mentions devm_add_action_or_reset() the actual change
      utilised devm_add_action() call by mistake
    - strictly speaking the parameter is not compatible by type

    Fix both issues by switching to devm_add_action_or_reset() and adding a
    wrapper for mutex_destroy() call.

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: e1af5c815586 ("leds: is31fl319x: Fix devm vs. non-devm ordering")
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Vincent Knecht <vincent.knecht@mailoo.org>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20221228093238.82713-1-andriy.shevchenko@linux.intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit ddf3e82164afd9381b1d52c9f00b3878f7b6d308
 Author: Miaoqian Lin <linmq006@gmail.com>
 Date:   Tue Dec 20 16:18:07 2022 +0400

    leds: led-core: Fix refcount leak in of_led_get()

    [ Upstream commit da1afe8e6099980fe1e2fd7436dca284af9d3f29 ]

    class_find_device_by_of_node() calls class_find_device(), it will take
    the reference, use the put_device() to drop the reference when not need
    anymore.

    Fixes: 699a8c7c4bd3 ("leds: Add of_led_get() and led_put()")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20221220121807.1543790-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 3992d60e46bab15a288078bc4dd7d3967c12c86d
 Author: Ian Rogers <irogers@google.com>
 Date:   Thu Jan 5 00:26:09 2023 -0800

    perf llvm: Fix inadvertent file creation

    [ Upstream commit 9f19aab47ced012eddef1e2bc96007efc7713b61 ]

    The LLVM template is first echo-ed into command_out and then
    command_out executed. The echo surrounds the template with double
    quotes, however, the template itself may contain quotes. This is
    generally innocuous but in tools/perf/tests/bpf-script-test-prologue.c
    we see:
    ...
    SEC("func=null_lseek file->f_mode offset orig")
    ...
    where the first double quote ends the double quote of the echo, then
    the > redirects output into a file called f_mode.
```

```
    To avoid this inadvertent behavior substitute redirects and similar
    characters to be ASCII control codes, then substitute the output in
    the echo back again.

    Fixes: 5eab5a7ee032acaa ("perf llvm: Display eBPF compiling command in debug output")
    Signed-off-by: Ian Rogers <irogers@google.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Andrii Nakryiko <andrii@kernel.org>
    Cc: bpf@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Cc: llvm@lists.linux.dev
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Nathan Chancellor <nathan@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Tom Rix <trix@redhat.com>
    Link: https://lore.kernel.org/r/20230105082609.344538-1-irogers@google.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3f1e860761152c991bc4f50ea18daa5268a8608
Author: Andreas Gruenbacher <agruenba@redhat.com>
Date:   Wed Feb 1 15:08:50 2023 +0100

    gfs2: jdata writepage fix

    [ Upstream commit cbb60951ce18c9b6e91d2eb97deb41d8ff616622 ]

    The ->writepage() and ->writepages() operations are supposed to write
    entire pages.  However, on filesystems with a block size smaller than
    PAGE_SIZE, __gfs2_jdata_writepage() only adds the first block to the
    current transaction instead of adding the entire page.  Fix that.

    Fixes: 18ec7d5c3f43 ("[GFS2] Make journaled data files identical to normal files on
disk")
    Signed-off-by: Andreas Gruenbacher <agruenba@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 85c32ee09e334282e41c38f37baae84eced13ba3
Author: Shyam Prasad N <sprasad@microsoft.com>
Date:   Fri Feb 10 16:34:31 2023 +0000

    cifs: use tcon allocation functions even for dummy tcon

    [ Upstream commit df57109bd50b9ed6911f3c2aa914189fe4c1fe2c ]

    In smb2_reconnect_server, we allocate a dummy tcon for
    calling reconnect for just the session. This should be
    allocated using tconInfoAlloc, and not kmalloc.

    Fixes: 3663c9045f51 ("cifs: check reconnects for channels of active tcons too")
    Signed-off-by: Shyam Prasad N <sprasad@microsoft.com>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d0c4f5f618f58eba03385363717703bee873c64
Author: Zhang Xiaoxu <zhangxiaoxu5@huawei.com>
Date:   Fri Nov 18 16:42:08 2022 +0800

    cifs: Fix warning and UAF when destroy the MR list

    [ Upstream commit 3e161c2791f8e661eed24a2c624087084d910215 ]
```

```
If the MR allocate failed, the MR recovery work not initialized
and list not cleared. Then will be warning and UAF when release
the MR:

  WARNING: CPU: 4 PID: 824 at kernel/workqueue.c:3066 __flush_work.isra.0+0xf7/0x110
  CPU: 4 PID: 824 Comm: mount.cifs Not tainted 6.1.0-rc5+ #82
  RIP: 0010:__flush_work.isra.0+0xf7/0x110
  Call Trace:
   <TASK>
   __cancel_work_timer+0x2ba/0x2e0
   smbd_destroy+0x4e1/0x990
   _smbd_get_connection+0x1cbd/0x2110
   smbd_get_connection+0x21/0x40
   cifs_get_tcp_session+0x8ef/0xda0
   mount_get_conns+0x60/0x750
   cifs_mount+0x103/0xd00
   cifs_smb3_do_mount+0x1dd/0xcb0
   smb3_get_tree+0x1d5/0x300
   vfs_get_tree+0x41/0xf0
   path_mount+0x9b3/0xdd0
   __x64_sys_mount+0x190/0x1d0
   do_syscall_64+0x35/0x80
   entry_SYSCALL_64_after_hwframe+0x46/0xb0

  BUG: KASAN: use-after-free in smbd_destroy+0x4fc/0x990
  Read of size 8 at addr ffff88810b156a08 by task mount.cifs/824
  CPU: 4 PID: 824 Comm: mount.cifs Tainted: G        W          6.1.0-rc5+ #82
  Call Trace:
   dump_stack_lvl+0x34/0x44
   print_report+0x171/0x472
   kasan_report+0xad/0x130
   smbd_destroy+0x4fc/0x990
   _smbd_get_connection+0x1cbd/0x2110
   smbd_get_connection+0x21/0x40
   cifs_get_tcp_session+0x8ef/0xda0
   mount_get_conns+0x60/0x750
   cifs_mount+0x103/0xd00
   cifs_smb3_do_mount+0x1dd/0xcb0
   smb3_get_tree+0x1d5/0x300
   vfs_get_tree+0x41/0xf0
   path_mount+0x9b3/0xdd0
   __x64_sys_mount+0x190/0x1d0
   do_syscall_64+0x35/0x80
   entry_SYSCALL_64_after_hwframe+0x46/0xb0

  Allocated by task 824:
   kasan_save_stack+0x1e/0x40
   kasan_set_track+0x21/0x30
   __kasan_kmalloc+0x7a/0x90
   _smbd_get_connection+0x1b6f/0x2110
   smbd_get_connection+0x21/0x40
   cifs_get_tcp_session+0x8ef/0xda0
   mount_get_conns+0x60/0x750
   cifs_mount+0x103/0xd00
   cifs_smb3_do_mount+0x1dd/0xcb0
   smb3_get_tree+0x1d5/0x300
   vfs_get_tree+0x41/0xf0
   path_mount+0x9b3/0xdd0
   __x64_sys_mount+0x190/0x1d0
   do_syscall_64+0x35/0x80
   entry_SYSCALL_64_after_hwframe+0x46/0xb0

  Freed by task 824:
   kasan_save_stack+0x1e/0x40
```

```
        kasan_set_track+0x21/0x30
        kasan_save_free_info+0x2a/0x40
        ____kasan_slab_free+0x143/0x1b0
        __kmem_cache_free+0xc8/0x330
        _smbd_get_connection+0x1c6a/0x2110
        smbd_get_connection+0x21/0x40
        cifs_get_tcp_session+0x8ef/0xda0
        mount_get_conns+0x60/0x750
        cifs_mount+0x103/0xd00
        cifs_smb3_do_mount+0x1dd/0xcb0
        smb3_get_tree+0x1d5/0x300
        vfs_get_tree+0x41/0xf0
        path_mount+0x9b3/0xdd0
        __x64_sys_mount+0x190/0x1d0
        do_syscall_64+0x35/0x80
        entry_SYSCALL_64_after_hwframe+0x46/0xb0

    Let's initialize the MR recovery work before MR allocate to prevent
    the warning, remove the MRs from the list to prevent the UAF.

    Fixes: c7398583340a ("CIFS: SMBD: Implement RDMA memory registration")
    Acked-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Reviewed-by: Tom Talpey <tom@talpey.com>
    Signed-off-by: Zhang Xiaoxu <zhangxiaoxu5@huawei.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04b7e13b8a13264282f874db5378fc3d3253cfac
Author: Zhang Xiaoxu <zhangxiaoxu5@huawei.com>
Date:   Fri Nov 18 16:42:07 2022 +0800

    cifs: Fix lost destroy smbd connection when MR allocate failed

    [ Upstream commit e9d3401d95d62a9531082cd2453ed42f2740e3fd ]

    If the MR allocate failed, the smb direct connection info is NULL,
    then smbd_destroy() will directly return, then the connection info
    will be leaked.

    Let's set the smb direct connection info to the server before call
    smbd_destroy().

    Fixes: c7398583340a ("CIFS: SMBD: Implement RDMA memory registration")
    Signed-off-by: Zhang Xiaoxu <zhangxiaoxu5@huawei.com>
    Acked-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Reviewed-by: David Howells <dhowells@redhat.com>
    Reviewed-by: Tom Talpey <tom@talpey.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1dcc9fad55477b942cc76992246107f7e5f1904
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Tue Feb 14 10:07:59 2023 -0500

    NFSD: copy the whole verifier in nfsd_copy_write_verifier

    [ Upstream commit 90d2175572470ba7f55da8447c72ddd4942923c4 ]

    Currently, we're only memcpy'ing the first __be32. Ensure we copy into
    both words.

    Fixes: 91d2e9b56cf5 ("NFSD: Clean up the nfsd_net::nfssvc_boot field")
    Reported-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit e01e63be5c3b6c333ce8ec2935d4e13f155cb1c5
Author: Jeff Layton <jlayton@kernel.org>
Date:   Tue Feb 7 12:02:46 2023 -0500

    nfsd: don't fsync nfsd_files on last close

    [ Upstream commit 4c475eee02375ade6e864f1db16976ba0d96a0a2 ]

    Most of the time, NFSv4 clients issue a COMMIT before the final CLOSE of
    an open stateid, so with NFSv4, the fsync in the nfsd_file_free path is
    usually a no-op and doesn't block.

    We have a customer running knfsd over very slow storage (XFS over Ceph
    RBD). They were using the "async" export option because performance was
    more important than data integrity for this application. That export
    option turns NFSv4 COMMIT calls into no-ops. Due to the fsync in this
    codepath however, their final CLOSE calls would still stall (since a
    CLOSE effectively became a COMMIT).

    I think this fsync is not strictly necessary. We only use that result to
    reset the write verifier. Instead of fsync'ing all of the data when we
    free an nfsd_file, we can just check for writeback errors when one is
    acquired and when it is freed.

    If the client never comes back, then it'll never see the error anyway
    and there is no point in resetting it. If an error occurs after the
    nfsd_file is removed from the cache but before the inode is evicted,
    then it will reset the write verifier on the next nfsd_file_acquire,
    (since there will be an unseen error).

    The only exception here is if something else opens and fsyncs the file
    during that window. Given that local applications work with this
    limitation today, I don't see that as an issue.

    Link: https://bugzilla.redhat.com/show_bug.cgi?id=2166658
    Fixes: ac3a2585f018 ("nfsd: rework refcounting in filecache")
    Reported-and-tested-by: Pierguido Lambri <plambri@redhat.com>
    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 6c63d608bce8e8c0091d83ee0a52a9d41cddba74
Author: Jeff Layton <jlayton@kernel.org>
Date:   Fri Feb 3 13:18:34 2023 -0500

    nfsd: fix courtesy client with deny mode handling in nfs4_upgrade_open

    [ Upstream commit dcd779dc46540e174a6ac8d52fbed23593407317 ]

    The nested if statements here make no sense, as you can never reach
    "else" branch in the nested statement. Fix the error handling for
    when there is a courtesy client that holds a conflicting deny mode.

    Fixes: 3d6942715180 ("NFSD: add support for share reservation conflict to courteous
server")
    Reported-by: 張智諺 <cc85nod@gmail.com>
    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Reviewed-by: Dai Ngo <dai.ngo@oracle.com>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c6678eed8d1437e0ada01ed4fe1e9e095390632e
Author: Dai Ngo <dai.ngo@oracle.com>
Date:   Tue Jan 31 11:12:29 2023 -0800
```

```
    NFSD: fix problems with cleanup on errors in nfsd4_copy

    [ Upstream commit 81e722978ad21072470b73d8f6a50ad62c7d5b7d ]

    When nfsd4_copy fails to allocate memory for async_copy->cp_src, or
    nfs4_init_copy_state fails, it calls cleanup_async_copy to do the
    cleanup for the async_copy which causes page fault since async_copy
    is not yet initialized.

    This patche rearranges the order of initializing the fields in
    async_copy and adds checks in cleanup_async_copy to skip un-initialized
    fields.

    Fixes: ce0887ac96d3 ("NFSD add nfs4 inter ssc to nfsd4_copy")
    Fixes: 87689df69491 ("NFSD: Shrink size of struct nfsd4_copy")
    Signed-off-by: Dai Ngo <dai.ngo@oracle.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 8f565846fbe8182961498d4cbe618b15076a683b
Author: Jeff Layton <jlayton@kernel.org>
Date:   Tue Jan 17 14:38:31 2023 -0500

    nfsd: clean up potential nfsd_file refcount leaks in COPY codepath

    [ Upstream commit 6ba434cb1a8d403ea9aad1b667c3ea3ad8b3191f ]

    There are two different flavors of the nfsd4_copy struct. One is
    embedded in the compound and is used directly in synchronous copies. The
    other is dynamically allocated, refcounted and tracked in the client
    struture. For the embedded one, the cleanup just involves releasing any
    nfsd_files held on its behalf. For the async one, the cleanup is a bit
    more involved, and we need to dequeue it from lists, unhash it, etc.

    There is at least one potential refcount leak in this code now. If the
    kthread_create call fails, then both the src and dst nfsd_files in the
    original nfsd4_copy object are leaked.

    The cleanup in this codepath is also sort of weird. In the async copy
    case, we'll have up to four nfsd_file references (src and dst for both
    flavors of copy structure). They are both put at the end of
    nfsd4_do_async_copy, even though the ones held on behalf of the embedded
    one outlive that structure.

    Change it so that we always clean up the nfsd_file refs held by the
    embedded copy structure before nfsd4_copy returns. Rework
    cleanup_async_copy to handle both inter and intra copies. Eliminate
    nfsd4_cleanup_intra_ssc since it now becomes a no-op.

    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Stable-dep-of: 81e722978ad2 ("NFSD: fix problems with cleanup on errors in nfsd4_copy")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 593468682e9b6eeb2cd48fb15060a31c0a3dd08c
Author: Benjamin Coddington <bcodding@redhat.com>
Date:   Fri Jan 27 11:18:56 2023 -0500

    nfsd: fix race to check ls_layouts

    [ Upstream commit fb610c4dbc996415d57d7090957ecddd4fd64fb6 ]

    Its possible for __break_lease to find the layout's lease before we've
```

```
    added the layout to the owner's ls_layouts list.  In that case, setting
    ls_recalled = true without actually recalling the layout will cause the
    server to never send a recall callback.

    Move the check for ls_layouts before setting ls_recalled.

    Fixes: c5c707f96fc9 ("nfsd: implement pNFS layout recalls")
    Signed-off-by: Benjamin Coddington <bcodding@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 6c3c05402547aaca3edb23327b50f01a881831b9
Author: Dai Ngo <dai.ngo@oracle.com>
Date:   Mon Jan 23 21:34:13 2023 -0800

    NFSD: fix leaked reference count of nfsd4_ssc_umount_item

    [ Upstream commit 34e8f9ec4c9ac235f917747b23a200a5e0ec857b ]

    The reference count of nfsd4_ssc_umount_item is not decremented
    on error conditions. This prevents the laundromat from unmounting
    the vfsmount of the source file.

    This patch decrements the reference count of nfsd4_ssc_umount_item
    on error.

    Fixes: f4e44b393389 ("NFSD: delay unmount source's export after inter-server copy
completed.")
    Signed-off-by: Dai Ngo <dai.ngo@oracle.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 0517ebb80044efd80c5b200272312349cf3e199d
Author: Dai Ngo <dai.ngo@oracle.com>
Date:   Sun Dec 18 16:55:53 2022 -0800

    NFSD: enhance inter-server copy cleanup

    [ Upstream commit df24ac7a2e3a9d0bc68f1756a880e50bfe4b4522 ]

    Currently nfsd4_setup_inter_ssc returns the vfsmount of the source
    server's export when the mount completes. After the copy is done
    nfsd4_cleanup_inter_ssc is called with the vfsmount of the source
    server and it searches nfsd_ssc_mount_list for a matching entry
    to do the clean up.

    The problems with this approach are (1) the need to search the
    nfsd_ssc_mount_list and (2) the code has to handle the case where
    the matching entry is not found which looks ugly.

    The enhancement is instead of nfsd4_setup_inter_ssc returning the
    vfsmount, it returns the nfsd4_ssc_umount_item which has the
    vfsmount embedded in it. When nfsd4_cleanup_inter_ssc is called
    it's passed with the nfsd4_ssc_umount_item directly to do the
    clean up so no searching is needed and there is no need to handle
    the 'not found' case.

    Signed-off-by: Dai Ngo <dai.ngo@oracle.com>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    [ cel: adjusted whitespace and variable/function names ]
    Reviewed-by: Olga Kornievskaia <kolga@netapp.com>
    Stable-dep-of: 34e8f9ec4c9a ("NFSD: fix leaked reference count of
nfsd4_ssc_umount_item")
```

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

commit 86defba5a34c47e3090fa679478c962f7650d114
Author: Asahi Lina <lina@asahilina.net>
Date:   Sun Feb 5 21:51:24 2023 +0900

    drm/shmem-helper: Fix locking for drm_gem_shmem_get_pages_sgt()

    [ Upstream commit ddddedaa0db99481c5e5abe628ad54f65e8765bc ]

    Other functions touching shmem->sgt take the pages lock, so do that here
    too. drm_gem_shmem_get_pages() & co take the same lock, so move to the
    _locked() variants to avoid recursive locking.

    Discovered while auditing locking to write the Rust abstractions.

    Fixes: 2194a63a818d ("drm: Add library for shmem backed GEM objects")
    Fixes: 4fa3d66f132b ("drm/shmem: Do dma_unmap_sg before purging pages")
    Signed-off-by: Asahi Lina <lina@asahilina.net>
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230205125124.2260-1-
lina@asahilina.net
    (cherry picked from commit aa8c85affe3facd3842c8912186623415931cc72)
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94f48aed2c9f244d7626416b6b028ed956763f0e
Author: Orlando Chamberlain <orlandoch.dev@gmail.com>
Date:   Thu Feb 16 21:34:51 2023 +1100

    ALSA: hda/hdmi: Register with vga_switcheroo on Dual GPU Macbooks

    [ Upstream commit 5beb5627a2481aade9aa630b7ebb7f99442321b6 ]

    Commit 586bc4aab878 ("ALSA: hda/hdmi - fix vgaswitcheroo detection for
    AMD") caused only AMD gpu's with PX to have their audio component register
    with vga_switcheroo. This meant that Apple Macbooks with apple-gmux as the
    gpu switcher no longer had the audio client registering, so when the gpu is
    powered off by vga_switcheroo snd_hda_intel is unaware that it should have
    suspended the device:

    amdgpu: switched off
    snd_hda_intel 0000:03:00.1:
        Unable to change power state from D3hot to D0, device inaccessible
    snd_hda_intel 0000:03:00.1: CORB reset timeout#2, CORBRP = 65535

    To resolve this, we use apple_gmux_detect() and register a
    vga_switcheroo audio client when apple-gmux is detected.

    Fixes: 586bc4aab878 ("ALSA: hda/hdmi - fix vgaswitcheroo detection for AMD")
    Link: https://lore.kernel.org/all/20230210044826.9834-9-orlandoch.dev@gmail.com/
    Signed-off-by: Orlando Chamberlain <orlandoch.dev@gmail.com>
    Link: https://lore.kernel.org/r/20230216103450.12925-1-orlandoch.dev@gmail.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc2d28d6ce1399f452be9d266a547a01da58cb14
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 00:01:44 2023 +0000

    hid: bigben_probe(): validate report count

    [ Upstream commit b94335f899542a0da5fafc38af8edcaf90195843 ]
```

```
    bigben_probe() does not validate that the output report has the
    needed report values in the first field.
    A malicious device registering a report with one field and a single
    value causes an head OOB write in bigben_worker() when
    accessing report_field->value[1] to report_field->value[7].
    Use hid_validate_values() which takes care of all the needed checks.

    Fixes: 256a90ed9e46 ("HID: hid-bigbenff: driver for BigBen Interactive PS3OFMINIPAD
gamepad")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230211-bigben-oob-v1-1-d2849688594c@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90289e71514e9533a9c44d694e2b492be9ed2b77
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 19:00:01 2023 +0000

    HID: bigben: use spinlock to safely schedule workers

    [ Upstream commit 76ca8da989c7d97a7f76c75d475fe95a584439d7 ]

    Use spinlocks to deal with workers introducing a wrapper
    bigben_schedule_work(), and several spinlock checks.
    Otherwise, bigben_set_led() may schedule bigben->worker after the
    structure has been freed, causing a use-after-free.

    Fixes: 4eb1b01de5b9 ("HID: hid-bigbenff: fix race condition for scheduled work during
removal")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230125-hid-unregister-leds-v4-3-
7860c5763c38@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f409a2265765005e8b4f777aaa4a41b7afba64f2
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 19:00:00 2023 +0000

    HID: bigben_worker() remove unneeded check on report_field

    [ Upstream commit 27d2a2fd844ec7da70d19fabb482304fd1e0595b ]

    bigben_worker() checks report_field to be non-NULL.
    The check has been added in commit
    918aa1ef104d ("HID: bigbenff: prevent null pointer dereference")
    to prevent a NULL pointer crash.
    However, the true root cause was a missing check for output
    reports, patched in commit
    c7bf714f8755 ("HID: check empty report_list in bigben_probe()"),
    where the type-confused report list_entry was overlapping with
    a NULL pointer, which was then causing the crash.

    Fixes: 918aa1ef104d ("HID: bigbenff: prevent null pointer dereference")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230125-hid-unregister-leds-v4-2-
7860c5763c38@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bcb202bdec3cc1eac81d485c521f5156edc5f61
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 18:59:59 2023 +0000

    HID: bigben: use spinlock to protect concurrent accesses
```

```
    [ Upstream commit 9fefb6201c4f8dd9f58c581b2a66e5cde2895ea2 ]

    bigben driver has a worker that may access data concurrently.
    Proct the accesses using a spinlock.

    Fixes: 256a90ed9e46 ("HID: hid-bigbenff: driver for BigBen Interactive PS3OFMINIPAD
gamepad")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230125-hid-unregister-leds-v4-1-
7860c5763c38@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5d822c5f5cd47d513086452e80944e3ec1de00ef
Author: Lucas Tanure <lucas.tanure@collabora.com>
Date:   Wed Feb 15 13:28:51 2023 +0000

    ASoC: soc-dapm.h: fixup warning struct snd_pcm_substream not declared

    [ Upstream commit fdff966bfde7cf0c85562d2bfb1ff1ba83da5f7b ]

    Add struct snd_pcm_substream forward declaration

    Fixes: 078a85f2806f ("ASoC: dapm: Only power up active channels from a DAI")
    Signed-off-by: Lucas Tanure <lucas.tanure@collabora.com>
    Reviewed-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20230215132851.1626881-1-lucas.tanure@collabora.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7f0e04663c4222f34398e8d9aab33ea87e3ee51d
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Wed Feb 15 14:01:28 2023 +0100

    spi: synquacer: Fix timeout handling in synquacer_spi_transfer_one()

    [ Upstream commit e6a0b671880207566e1ece983bf989dde60bc1d7 ]

    wait_for_completion_timeout() never returns a <0 value. It returns either
    on timeout or a positive value (at least 1, or number of jiffies left
    till timeout)

    So, fix the error handling path and return -ETIMEDOUT should a timeout
    occur.

    Fixes: b0823ee35cf9 ("spi: Add spi driver for Socionext SynQuacer platform")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Acked-by: Jassi Brar <jaswinder.singh@linaro.org>
    Link:
https://lore.kernel.org/r/c2040bf3cfa201fd8890cfab14fa5a701ffeca14.1676466072.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0873f089822920b603fabf366fcf9aa6df432974
Author: Lucas De Marchi <lucas.demarchi@intel.com>
Date:   Mon Feb 6 08:54:09 2023 -0800

    drm/i915: Fix GEN8_MISCCPCTL

    [ Upstream commit 6a8b2e4984f73f8d00c8c16b87a8b115d34088e4 ]

    Register 0x9424 is not replicated on any platform, so it shouldn't be
```

```
    declared with REG_MCR(). Declaring it with _MMIO() is basically
    duplicate of the GEN7 version, so just remove the GEN8 and change all
    the callers to use the right functions.

    Old versions of the gen8 bspec page used to contain a table with MCR
    registers, apparently implying 0x9400 - 0x94ff registers were
    replicated. However that table went away and there is no information
    related to the ranges for gen8 anymore. Moreover the current behavior of
    the driver wouldn't do anything special for 0x9424 since there is no
    equivalent table in intel_gt_mcr.c: the driver would just fallback to
    intel_uncore_{read,write}(). Therefore, do not care about the possible
    special case for gen8 and just use the register as non-MCR for all the
    platforms.

    One place doing read + write is also converted to intel_uncore_rmw().

    v2: Reword commit message adding the justification wrt gen8

    Fixes: a9e69428b1b4 ("drm/i915: Define MCR registers explicitly")
    Cc: Balasubramani Vivekanandan <balasubramani.vivekanandan@intel.com>
    Cc: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Cc: Gustavo Sousa <gustavo.sousa@intel.com>
    Cc: Matt Atwood <matthew.s.atwood@intel.com>
    Cc: Ashutosh Dixit <ashutosh.dixit@intel.com>
    Signed-off-by: Lucas De Marchi <lucas.demarchi@intel.com>
    Reviewed-by: Matt Roper <matthew.d.roper@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230206165410.3056073-1-
lucas.demarchi@intel.com
    (cherry picked from commit 869bace73ae2b4227e57ee3fd994bfa7d4808938)
    Signed-off-by: Jani Nikula <jani.nikula@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 34802bbf1fd433d910769551c7a395918dafd19e
Author: Matt Roper <matthew.d.roper@intel.com>
Date:   Wed Feb 1 14:28:28 2023 -0800

    drm/i915/pvc: Annotate two more workaround/tuning registers as MCR

    [ Upstream commit effc0905d741b4138806747407baf8de98390c72 ]

    XEHPC_LNCFMISCCFGREG0 and XEHPC_L3SCRUB are both in MCR register ranges
    on PVC (with HALFBSLICE and L3BANK replication respectively), so they
    should be explicitly declared as MCR registers and use MCR-aware
    workaround handlers.

    The workarounds/tuning settings should still be applied properly on PVC
    even without the MCR annotation, but readback verification on
    CONFIG_DRM_I915_DEBUG_GEM builds could potentitally give false positive
    "workaround lost on load" warnings on parts fused such that a unicast
    read targets a terminated register instance.

    Fixes: a9e69428b1b4 ("drm/i915: Define MCR registers explicitly")
    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Reviewed-by: Gustavo Sousa <gustavo.sousa@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230201222831.608281-1-
matthew.d.roper@intel.com
    (cherry picked from commit 4039e44237e8ebb06f0e4af549fbedf7c41df9db)
    Signed-off-by: Jani Nikula <jani.nikula@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d483c9a2a4394967412a713a8674ea725fbce1a3
Author: Wayne Boyer <wayne.boyer@intel.com>
Date:   Wed Nov 30 09:07:23 2022 -0800

    drm/i915/pvc: Implement recommended caching policy
```

```
    [ Upstream commit e3995e08a39a41691742b380023a0d480247afb0 ]

    As per the performance tuning guide, set the HOSTCACHEEN bit to
    implement the recommended caching policy on PVC.

    Signed-off-by: Wayne Boyer <wayne.boyer@intel.com>
    Reviewed-by: Matt Roper <matthew.d.roper@intel.com>
    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221130170723.2460014-1-
wayne.boyer@intel.com
    Stable-dep-of: effc0905d741 ("drm/i915/pvc: Annotate two more workaround/tuning
registers as MCR")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d5f35b9d0d651efdeae3495613358e3e4392a16
Author: NeilBrown <neilb@suse.de>
Date:   Wed Feb 8 15:45:38 2023 +1100

    NFS: fix disabling of swap

    [ Upstream commit 5bab56fff53ce161ed859d9559a10361d4f79578 ]

    When swap is activated to a file on an NFSv4 mount we arrange that the
    state manager thread is always present as starting a new thread requires
    memory allocations that might block waiting for swap.

    Unfortunately the code for allowing the state manager thread to exit when
    swap is disabled was not tested properly and does not work.
    This can be seen by examining /proc/fs/nfsfs/servers after disabling swap
    and unmounting the filesystem.  The servers file will still list one
    entry.  Also a "ps" listing will show the state manager thread is still
    present.

    There are two problems.
     1/ rpc_clnt_swap_deactivate() doesn't walk up the ->cl_parent list to
        find the primary client on which the state manager runs.

     2/ The thread is not woken up properly and it immediately goes back to
        sleep without checking whether it is really needed.  Using
        nfs4_schedule_state_manager() ensures a proper wake-up.

    Reported-by: Olga Kornievskaia <aglo@umich.edu>
    Fixes: 4dc73c679114 ("NFSv4: keep state manager thread active if swap is enabled")
    Signed-off-by: NeilBrown <neilb@suse.de>
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1edd69a3236b5460bf83bcd32d9732baec91aed4
Author: Benjamin Coddington <bcodding@redhat.com>
Date:   Tue Feb 14 08:18:23 2023 -0500

    nfs4trace: fix state manager flag printing

    [ Upstream commit b46d80bd2d6e7e063c625a20de54248afe8d4889 ]

    __print_flags wants a mask, not the enum value.  Add two more flags.

    Fixes: 511ba52e4c01 ("NFS4: Trace state recovery operation")
    Signed-off-by: Benjamin Coddington <bcodding@redhat.com>
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69562613d32b3a54a75b8b3c5517bdc8f1edb8a8
Author: Mike Snitzer <snitzer@kernel.org>
```

```
Date:   Tue Feb 14 13:06:05 2023 -0500

    dm: remove flush_scheduled_work() during local_exit()

    [ Upstream commit 0b22ff5360f5c4e11050b89206370fdf7dc0a226 ]

    Commit acfe0ad74d2e1 ("dm: allocate a special workqueue for deferred
    device removal") switched from using system workqueue to a single
    workqueue local to DM.  But it didn't eliminate the call to
    flush_scheduled_work() that was introduced purely for the benefit of
    deferred device removal with commit 2c140a246dc ("dm: allow remove to
    be deferred").

    Since DM core uses its own workqueue (and queue_work) there is no need
    to call flush_scheduled_work() from local_exit().  local_exit()'s
    destroy_workqueue(deferred_remove_workqueue) handles flushing work
    started with queue_work().

    Fixes: acfe0ad74d2e1 ("dm: allocate a special workqueue for deferred device removal")
    Signed-off-by: Mike Snitzer <snitzer@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc169f8298fd4720cba0916118ff568890697299
Author: Steffen Aschbacher <steffen.aschbacher@stihl.de>
Date:   Mon Feb 13 09:38:05 2023 +0200

    ASoC: tlv320adcx140: fix 'ti,gpio-config' DT property init

    [ Upstream commit 771725efe5e2e5396dd9d1220437e5f9d6b9ca9d ]

    When the 'ti,gpio-config' property is not defined, the
    device_property_count_u32() will return an error, rather than zero.

    The current check, only handles a return value of zero, which assumes that
    the property is defined and has nothing defined.

    This change extends the check to also check for an error case (most likely
    to be hit by the case that the 'ti,gpio-config' is not defined).

    In case that the 'ti,gpio-config' and the returned 'gpio_count' is not
    correct, there is a 'if (gpio_count != ADCX140_NUM_GPIO_CFGS)' check, a few
    lines lower that will return -EINVAL.
    This means that someone tried to define 'ti,gpio-config', but with the
    wrong number of GPIOs.

    Fixes: d5214321498a ("ASoC: tlv320adcx140: Add support for configuring GPIO pin")
    Signed-off-by: Steffen Aschbacher <steffen.aschbacher@stihl.de>
    Signed-off-by: Alexandru Ardelean <alex@shruggie.ro>
    Link: https://lore.kernel.org/r/20230213073805.14640-1-alex@shruggie.ro
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c51d9f7a5d43ade7539f5308b2ed21e7849fbdab
Author: Vadim Pasternak <vadimp@nvidia.com>
Date:   Sun Feb 12 16:57:30 2023 +0200

    hwmon: (mlxreg-fan) Return zero speed for broken fan

    [ Upstream commit a1ffd3c46267ee5c807acd780e15df9bb692223f ]

    Currently for broken fan driver returns value calculated based on error
    code (0xFF) in related fan speed register.
    Thus, for such fan user gets fan{n}_fault to 1 and fan{n}_input with
    misleading value.
```

```
    Add check for fan fault prior return speed value and return zero if
    fault is detected.

    Fixes: 65afb4c8e7e4 ("hwmon: (mlxreg-fan) Add support for Mellanox FAN driver")
    Signed-off-by: Vadim Pasternak <vadimp@nvidia.com>
    Link: https://lore.kernel.org/r/20230212145730.24247-1-vadimp@nvidia.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b516cdf32e0f00aaed0db148a3330c2ae2163f8
Author: William Zhang <william.zhang@broadcom.com>
Date:   Thu Feb 9 12:02:41 2023 -0800

    spi: bcm63xx-hsspi: Fix multi-bit mode setting

    [ Upstream commit 811ff802aaf878ebbbaeac0307a0164fa21e7d40 ]

    Currently the driver always sets the controller to dual data bit mode
    for both tx and rx data in the profile mode control register even for
    single data bit transfer. Luckily the opcode is set correctly according
    to SPI transfer data bit width so it does not actually cause issues.

    This change fixes the problem by setting tx and rx data bit mode field
    correctly according to the actual SPI transfer tx and rx data bit width.

    Fixes: 142168eba9dc ("spi: bcm63xx-hsspi: add bcm63xx HSSPI driver")
    Signed-off-by: William Zhang <william.zhang@broadcom.com>
    Link: https://lore.kernel.org/r/20230209200246.141520-11-william.zhang@broadcom.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3042f6af979564aa48d917d59cdb16e573f92d45
Author: Bastien Nocera <hadess@hadess.net>
Date:   Mon Jan 16 14:09:37 2023 +0100

    HID: logitech-hidpp: Hard-code HID++ 1.0 fast scroll support

    [ Upstream commit 719acb4d3b7accc9cfbaf21c1c2d51dc7384aee2 ]

    HID++ 1.0 devices only export whether Fast Scrolling is enabled, not
    whether they are capable of it. Reinstate the original quirks for the 3
    supported mice so fast scrolling works again on those devices.

    Fixes: 908d325e1665 ("HID: logitech-hidpp: Detect hi-res scrolling support")
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216903
    Signed-off-by: Bastien Nocera <hadess@hadess.net>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Link: https://lore.kernel.org/r/20230116130937.391441-1-hadess@hadess.net
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 232a9748000e122ebae999722e8ce204e4447c5b
Author: Hamza Mahfooz <hamza.mahfooz@amd.com>
Date:   Wed Feb 8 13:25:49 2023 -0500

    drm/amd/display: don't call dc_interrupt_set() for disabled crtcs

    [ Upstream commit 4936458bf989d168f5a89015dd81067c4c2bdc64 ]

    As made mention of in commit 4ea7fc09539b ("drm/amd/display: Do not
    program interrupt status on disabled crtc"), we shouldn't program
    disabled crtcs. So, filter out disabled crtcs in dm_set_vupdate_irq()
    and dm_set_vblank().

    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Fixes: 589d2739332d ("drm/amd/display: Use crtc enable/disable_vblank hooks")
```

```
    Fixes: d2574c33bb71 ("drm/amd/display: In VRR mode, do DRM core vblank handling at end
of vblank. (v2)")
    Signed-off-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 335ed52e265bdb7f0b66ee09885632f3407b7ccb
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Feb 9 12:28:04 2023 +0000

    ASoC: codecs: lpass: fix incorrect mclk rate

    [ Upstream commit e7621434378c40b62ef858c14ae6415fb6469a8e ]

    For some reason we ended up with incorrect mclk rate which should be
    1920000 instead of 96000, So far we were getting lucky as the same clk
    is set to 192000 by wsa and va macro. This issue is discovered when there
    is no wsa macro active and only rx or tx path is tested.
    Fix this by setting correct rate.

    Fixes: c39667ddcfc5 ("ASoC: codecs: lpass-tx-macro: add support for lpass tx macro")
    Fixes: af3d54b99764 ("ASoC: codecs: lpass-rx-macro: add support for lpass rx macro")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230209122806.18923-7-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4b322e30dfbb1b0ae012c9593d8307520295ff81
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Feb 9 12:28:03 2023 +0000

    ASoC: codecs: lpass: register mclk after runtime pm

    [ Upstream commit 1dc3459009c33e335f0d62b84dd39a6bbd7fd5d2 ]

    move mclk out registration after runtime pm is enabled so that the
    clk framework can resume the codec if it requires to enable the mclk out.

    Fixes: c96baa2949b2 ("ASoC: codecs: wsa-macro: add runtime pm support")
    Fixes: 72ad25eabda0 ("ASoC: codecs: va-macro: add runtime pm support")
    Fixes: 366ff79ed539 ("ASoC: codecs: rx-macro: add runtime pm support")
    Fixes: 1fb83bc5cf64 ("ASoC: codecs: tx-macro: add runtime pm support")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230209122806.18923-6-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0ce16b073dd3287580435ccda555d6733fc957af
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Feb 9 12:28:01 2023 +0000

    ASoC: qcom: q6apm-dai: Add SNDRV_PCM_INFO_BATCH flag

    [ Upstream commit aa759f3f9f4394a3af65ad1772fca6cb9dd9e4cc ]

    At the moment, playing audio with PulseAudio with the qdsp6 driver
    results in distorted sound. It seems like its timer-based scheduling
    does not work properly with qdsp6 since setting tsched=0 in
    the PulseAudio configuration avoids the issue.

    Apparently this happens when the pointer() callback is not accurate
    enough. There is a SNDRV_PCM_INFO_BATCH flag that can be used to stop
    PulseAudio from using timer-based scheduling by default.

    According to https://www.alsa-project.org/pipermail/alsa-devel/2014-March/073816.html:
```

```
    The flag is being used in the sense explained in the previous audio
    meeting -- the data transfer granularity isn't fine enough but aligned
    to the period size (or less).

    q6apm-dai reports the position as multiple of

    prtd->pcm_count = snd_pcm_lib_period_bytes(substream)

    so it indeed just a multiple of the period size.

    Therefore adding the flag here seems appropriate and makes audio
    work out of the box.

    Comment log inspired by Stephan Gerhold sent for q6asm-dai.c few years back.

    Fixes: 9b4fe0f1cd79 ("ASoC: qdsp6: audioreach: add q6apm-dai support")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230209122806.18923-4-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 74f751f86226717fe0b8b43817191ea183ba0cf5
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Feb 9 12:28:00 2023 +0000

    ASoC: qcom: q6apm-dai: fix race condition while updating the position pointer

    [ Upstream commit 84222ef54bfd8f043c23c8603fd5257a64b00780 ]

    It is noticed that the position pointer value seems to get a get corrupted
    due to missing locking between updating and reading.

    Fix this by adding a spinlock around the position pointer.

    Fixes: 9b4fe0f1cd79 ("ASoC: qdsp6: audioreach: add q6apm-dai support")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230209122806.18923-3-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit deaab3c632a1e5bd789b5eb3949ad0615801eab2
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Feb 9 12:27:59 2023 +0000

    ASoC: qcom: q6apm-lpass-dai: unprepare stream if its already prepared

    [ Upstream commit c2ac3aec474da0455df79c4a182f19687bc98d1d ]

    prepare callback can be called multiple times, so unprepare the stream
    if its already prepared.

    Without this DSP is not happy to setting the params on a already
    prepared graph.

    Fixes: 9b4fe0f1cd79 ("ASoC: qdsp6: audioreach: add q6apm-dai support")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230209122806.18923-2-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c55922aae2cb53758bd5fecbf59097f5f702116f
Author: Dmitry Torokhov <dmitry.torokhov@gmail.com>
Date:   Tue Feb 7 15:03:30 2023 -0800
```

```
    HID: retain initial quirks set up when creating HID devices

    [ Upstream commit 03a86105556e23650e4470c09f91cf7c360d5e28 ]

    In certain circumstances, such as when creating I2C-connected HID
    devices, we want to pass and retain some quirks (axis inversion, etc).
    The source of such quirks may be device tree, or DMI data, or something
    else not readily available to the HID core itself and therefore cannot
    be reconstructed easily. To allow this, introduce "initial_quirks" field
    in hid_device structure and use it when determining the final set of
    quirks.

    This fixes the problem with i2c-hid setting up device-tree sourced
    quirks too late and losing them on device rebind, and also allows to
    sever the tie between hid-code and i2c-hid when applying DMI-based
    quirks.

    Fixes: b60d3c803d76 ("HID: i2c-hid-of: Expose the touchscreen-inverted properties")
    Fixes: a2f416bf062a ("HID: multitouch: Add quirks for flipped axes")
    Reviewed-by: Guenter Roeck <groeck@chromium.org>
    Tested-by: Allen Ballway <ballway@chromium.org>
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Reviewed-by: Alistair Francis <alistair@alistair23.me>
    Link: https://lore.kernel.org/r/Y+LYwu3Zs13hdVDy@google.com
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1b0a173b98910899a1039f597a04e223b412369
Author: Allen Ballway <ballway@chromium.org>
Date:   Tue Jan 10 20:25:50 2023 +0000

    HID: multitouch: Add quirks for flipped axes

    [ Upstream commit a2f416bf062a38bb76cccd526d2d286b8e4db4d9 ]

    Certain touchscreen devices, such as the ELAN9034, are oriented
    incorrectly and report touches on opposite points on the X and Y axes.
    For example, a 100x200 screen touched at (10,20) would report (90, 180)
    and vice versa.

    This is fixed by adding device quirks to transform the touch points
    into the correct spaces, from X -> MAX(X) - X, and Y -> MAX(Y) - Y.

    Signed-off-by: Allen Ballway <ballway@chromium.org>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Stable-dep-of: 03a86105556e ("HID: retain initial quirks set up when creating HID
devices")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 047bad4f9d767828a2dffefc5737fb4f9d5e6255
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Sat Jan 28 19:08:32 2023 +0800

    scsi: aic94xx: Add missing check for dma_map_single()

    [ Upstream commit 32fe45274edb5926abc0fac7263d9f889d02d9cf ]

    Add check for dma_map_single() and return error if it fails in order to
    avoid invalid DMA address.

    Fixes: 2908d778ab3e ("[SCSI] aic94xx: new driver")
    Link: https://lore.kernel.org/r/20230128110832.6792-1-jiasheng@iscas.ac.cn
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Jason Yan <yanaijie@huawei.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
```

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

commit 28137ea3eb05a87329a7154a8ff410d9e8bcc0a5
Author: Tomas Henzl <thenzl@redhat.com>
Date:   Tue Feb 7 16:21:59 2023 +0100

    scsi: mpt3sas: Fix a memory leak

    [ Upstream commit 54dd96015e8d7a2a07359e2dfebf05b529d1780c ]

    Add a forgotten kfree().

    Fixes: dbec4c9040ed ("scsi: mpt3sas: lockless command submission")
    Link: https://lore.kernel.org/r/20230207152159.18627-1-thenzl@redhat.com
    Signed-off-by: Tomas Henzl <thenzl@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5bbd770b79311c6a9cd2efd3e2e9404995df119f
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Feb 6 20:36:02 2023 +0100

    drm/amdgpu: fix enum odm_combine_mode mismatch

    [ Upstream commit 087bad7eb1f6945f8232f132953ecc2bda8bd38d ]

    A conversion from 'bool' to 'enum odm_combine_mode' was incomplete,
    and gcc warns about this with many instances of

    display/dc/dml/dcn20/display_mode_vba_20.c:3899:44: warning: implicit conversion from
'enum <anonymous>' to 'enum
    odm_combine_mode' [-Wenum-conversion]
     3899 |     locals->ODMCombineEnablePerState[i][k] = false;

    Change the ones that we get a warning for, using the same numerical
    values to leave the behavior unchanged.

    Fixes: 5fc11598166d ("drm/amd/display: expand dml structs")
    Link: https://lore.kernel.org/all/20201026210039.3884312-3-arnd@kernel.org/
    Link: https://lore.kernel.org/all/20210927100659.1431744-1-arnd@kernel.org/
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 87ca47acc793ec061df95373f1023a23ffe8e886
Author: Jaroslav Kysela <perex@perex.cz>
Date:   Thu Feb 2 10:20:13 2023 +0100

    ALSA: hda: Fix the control element identification for multiple codecs

    [ Upstream commit d045bceff5a904bd79d71dede9f927c00ce4906f ]

    Some motherboards have multiple HDA codecs connected to the serial bus.
    The current code may create multiple mixer controls with the almost
    identical identification.

    The current code use id.device field from the control element structure
    to store the codec address to avoid such clashes for multiple codecs.
    Unfortunately, the user space do not handle this correctly. For mixer
    controls, only name and index are used for the identifiers.

    This patch fixes this problem to compose the index using the codec
    address as an offset in case, when the control already exists. It is
    really unlikely that one codec will create 10 similar controls.
```

```
    This patch adds new kernel module parameter 'ctl_dev_id' to allow
    select the old behaviour, too. The CONFIG_SND_HDA_CTL_DEV_ID Kconfig
    option sets the default value.

    BugLink: https://github.com/alsa-project/alsa-lib/issues/294
    BugLink: https://github.com/alsa-project/alsa-lib/issues/205
    Fixes: 54d174031576 ("[ALSA] hda-codec - Fix connection list parsing")
    Fixes: 1afe206ab699 ("ALSA: hda - Try to find an empty control index when it's
occupied")
    Signed-off-by: Jaroslav Kysela <perex@perex.cz>
    Link: https://lore.kernel.org/r/20230202092013.4066998-1-perex@perex.cz
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bba4779cf9ee9aaeb9877460655660d1bef2072c
Author: Jonathan Cormier <jcormier@criticallink.com>
Date:   Thu Jan 26 17:32:25 2023 -0500

    hwmon: (ltc2945) Handle error case in ltc2945_value_store

    [ Upstream commit 178b01eccfb0b8149682f61388400bd3d903dddc ]

    ltc2945_val_to_reg errors were not being handled
    which would have resulted in register being set to
    0 (clamped) instead of being left alone.

    Fixes: 6700ce035f83 ("hwmon: Driver for Linear Technologies LTC2945")

    Signed-off-by: Jonathan Cormier <jcormier@criticallink.com>
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 81722f9a804c5ba8fbc34b8369c048c78872966a
Author: Eugene Shalygin <eugene.shalygin@gmail.com>
Date:   Sat Jan 21 12:17:28 2023 +0100

    hwmon: (asus-ec-sensors) add missing mutex path

    [ Upstream commit e2de0e6abd91b05411cb1f0953115dbbbe9b11ce ]

    Add missing mutex path for ProArt X570-CREATOR WIFI.

    Fixes: de8fbac5e59e (hwmon: (asus-ec-sensors) implement locking via the ACPI global
lock)
    Signed-off-by: Eugene Shalygin <eugene.shalygin@gmail.com>
    Link: https://lore.kernel.org/r/20230121111728.168514-2-eugene.shalygin@gmail.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 706722858c6102ccb7bca9c92a6e3e00b9d34d38
Author: Jerome Neanne <jneanne@baylibre.com>
Date:   Fri Feb 3 15:01:19 2023 +0100

    regulator: tps65219: use generic set_bypass()

    [ Upstream commit 0365df81145a4cfaae5f4da896160de512256e6d ]

    Due to wrong interpretation of the specification,
    custom implementation was used instead of standard regmap helper.
    LINK: https://lore.kernel.org/all/c2014039-f1e8-6976-33d6-52e2dd4e7b66@baylibre.com/

    Fixes: c12ac5fc3e0a ("regulator: drivers: Add TI TPS65219 PMIC regulators support")

    Regulator does NOT require to be off to be switched to bypass mode.
```

```
    Signed-off-by: Jerome Neanne <jneanne@baylibre.com>
    Link: https://lore.kernel.org/r/20230203140119.13029-1-jneanne@baylibre.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3004ac2b0ccb58cb092c29ee5b6a29b387cbd99
Author: Jerome Brunet <jbrunet@baylibre.com>
Date:   Thu Feb 2 19:36:46 2023 +0100

    ASoC: dt-bindings: meson: fix gx-card codec node regex

    [ Upstream commit 480b26226873c88e482575ceb0d0a38d76e1be57 ]

    'codec' is a valid node name when there is a single codec
    in the link. Fix the node regular expression to apply this.

    Fixes: fd00366b8e41 ("ASoC: meson: gx: add sound card dt-binding documentation")
    Signed-off-by: Jerome Brunet <jbrunet@baylibre.com>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230202183653.486216-3-jbrunet@baylibre.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 43755deadfe4285e770614b7a6a62f5d478839c3
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Thu Feb 2 09:34:19 2023 -0700

    ASoC: mchp-spdifrx: Fix uninitialized use of mr in mchp_spdifrx_hw_params()

    [ Upstream commit 218674a45930c700486d27b765bf2f1b43f8cbf7 ]

    Clang warns:

      ../sound/soc/atmel/mchp-spdifrx.c:455:3: error: variable 'mr' is uninitialized when
used here [-Werror,-Wuninitialized]
                      mr |= SPDIFRX_MR_ENDIAN_BIG;
                      ^~
      ../sound/soc/atmel/mchp-spdifrx.c:432:8: note: initialize the variable 'mr' to silence
this warning
              u32 mr;
                  ^
                   = 0
      1 error generated.

    Zero initialize mr so that these bitwise OR and assignment operation
    works unconditionally.

    Fixes: fa09fa60385a ("ASoC: mchp-spdifrx: fix controls which rely on rsr register")
    Link: https://github.com/ClangBuiltLinux/linux/issues/1797
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20230202-mchp-spdifrx-fix-uninit-mr-v1-1-
629a045d7a2f@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e83d62dd1c3cc4ec4523c87122253fca661156c0
Author: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
Date:   Wed Feb 1 01:59:52 2023 +0000

    ASoC: rsnd: fixup #endif position

    [ Upstream commit 49123b51cd896e00b256a27c2ce9e6bfe1bbc22f ]

    commit 1f9c82b5ab83ff2 ("ASoC: rsnd: add debugfs support") added
```

```
    CONFIG_DEBUG_FS related definitions on rsnd.h, but it should be
    added inside of RSND_H. This patch fixup it.

    Fixes: 1f9c82b5ab83 ("ASoC: rsnd: add debugfs support")
    Signed-off-by: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
    Link: https://lore.kernel.org/r/877cx26t7r.wl-kuninori.morimoto.gx@renesas.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03f7ac1aaf585afacaeb678eb77e689feeaedbb6
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Fri Jan 27 23:14:55 2023 +0100

    accel: fix CONFIG_DRM dependencies

    [ Upstream commit 9f20c9f4b1e17e83e9ccc247cfdc0b61041bff3d ]

    At the moment, accel drivers can be built-in even with CONFIG_DRM=m,
    but this causes a link failure:

    x86_64-linux-ld: drivers/accel/ivpu/ivpu_drv.o: in function `ivpu_dev_init':
    ivpu_drv.c:(.text+0x1535): undefined reference to `drmm_kmalloc'
    x86_64-linux-ld: ivpu_drv.c:(.text+0x1562): undefined reference to `drmm_kmalloc'
    x86_64-linux-ld: drivers/accel/ivpu/ivpu_drv.o: in function `ivpu_remove':
    ivpu_drv.c:(.text+0x1faa): undefined reference to `drm_dev_unregister'
    x86_64-linux-ld: drivers/accel/ivpu/ivpu_drv.o: in function `ivpu_probe':
    ivpu_drv.c:(.text+0x1fef): undefined reference to `__devm_drm_dev_alloc'

    The problem is that DRM_ACCEL is a 'bool' symbol, so driver that
    only depend on DRM_ACCEL but not also on DRM do not see the restriction
    to =m configs.

    To ensure that each accel driver has an implied dependency on CONFIG_DRM,
    enclose the entire Kconfig file in an if/endif check.

    Fixes: 8bf4889762a8 ("drivers/accel: define kconfig and register a new major")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Stanislaw Gruszka <stanislaw.gruszka@linux.intel.com>
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Jacek Lawrynowicz <jacek.lawrynowicz@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230127221504.2522909-1-
arnd@kernel.org
    (cherry picked from commit 3524c96a121952f214271622bb372661ced86101)
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 73745e903ab27547cca5847b4666bcdf36dbcbe7
Author: Daniel Golle <daniel@makrotopia.org>
Date:   Mon Jan 30 02:04:57 2023 +0000

    regmap: apply reg_base and reg_downshift for single register ops

    [ Upstream commit 697c3892d825fb78f42ec8e53bed065dd728db3e ]

    reg_base and reg_downshift currently don't have any effect if used with
    a regmap_bus or regmap_config which only offers single register
    operations (ie. reg_read, reg_write and optionally reg_update_bits).

    Fix that and take them into account also for regmap_bus with only
    reg_read and read_write operations by applying reg_base and
    reg_downshift in _regmap_bus_reg_write, _regmap_bus_reg_read.

    Also apply reg_base and reg_downshift in _regmap_update_bits, but only
    in case the operation is carried out with a reg_update_bits call
```

```
      defined in either regmap_bus or regmap_config.

      Fixes: 0074f3f2b1e43d ("regmap: allow a defined reg_base to be added to every address")
      Fixes: 86fc59ef818beb ("regmap: add configurable downshift for addresses")
      Signed-off-by: Daniel Golle <daniel@makrotopia.org>
      Tested-by: Colin Foster <colin.foster@in-advantage.com>
      Link: https://lore.kernel.org/r/Y9clyVS3tQEHlUhA@makrotopia.org
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit eb9501a0f132ca834185545e7d0f10c48ec53b53
  Author: Mike Snitzer <snitzer@kernel.org>
  Date:   Wed Jan 25 13:01:46 2023 -0500

      dm: improve shrinker debug names

      [ Upstream commit c87791bcc455a91e51ca9800faaacc21c8d67785 ]

      Commit e33c267ab70d ("mm: shrinkers: provide shrinkers with names")
      chose some fairly bad names for DM's shrinkers.

      Fixes: e33c267ab70d ("mm: shrinkers: provide shrinkers with names")
      Signed-off-by : Mike Snitzer <snitzer@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 10e49b1fb2e6cad14b0f97145373cf8f215cba44
  Author: Claudiu Beznea <claudiu.beznea@microchip.com>
  Date:   Mon Jan 30 14:06:43 2023 +0200

      ASoC: mchp-spdifrx: disable all interrupts in mchp_spdifrx_dai_remove()

      [ Upstream commit aaecdc32b7e35b4f9b457fb3509414aa9a932589 ]

      CSC interrupts which might be used in controls are on bits 8 and 9 of
      SPDIFRX_IDR register. Thus disable all the interrupts that are exported
      by driver.

      Fixes: ef265c55c1ac ("ASoC: mchp-spdifrx: add driver for SPDIF RX")
      Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
      Link: https://lore.kernel.org/r/20230130120647.638049-5-claudiu.beznea@microchip.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit ce1cb602c2ab3e55440e1cd5e307705996d9faf0
  Author: Claudiu Beznea <claudiu.beznea@microchip.com>
  Date:   Mon Jan 30 14:06:42 2023 +0200

      ASoC: mchp-spdifrx: fix controls that works with completion mechanism

      [ Upstream commit d3681df44e856aab523a6eb7ba15b5e41efcbb1c ]

      Channel status get and channel subcode get controls relies on data
      returned by controls when certain IRQs are raised. To achieve that
      completions are used b/w controls and interrupt service routine. The
      concurrent accesses to these controls are protected by
      struct snd_card::controls_rwsem.

      Issues identified:
      - reinit_completion() may be called while waiting for completion
        which should be avoided
      - in case of multiple threads waiting, the complete() call in interrupt
        will signal only one waiting thread per interrupt which may lead to
        timeout for the others
      - in case of channel status get as the CSC interrupt is not refcounted
        ISR may disable interrupt for threads that were just enabled it.
```

```
    To solve these the access to controls were protected by a mutex. Along
    with this there is no need for spinlock to protect the software cache
    reads/updates b/w controls and ISR as the update is happening only when
    requested from control, and only one reader can reach the control.

    Fixes: ef265c55c1ac ("ASoC: mchp-spdifrx: add driver for SPDIF RX")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20230130120647.638049-4-claudiu.beznea@microchip.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c5cb63e83998495e6c886a8261257358b22628f3
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Mon Jan 30 14:06:41 2023 +0200

    ASoC: mchp-spdifrx: fix return value in case completion times out

    [ Upstream commit a4c4161d6eae3ef5f486d1638ef452d9bc1376b0 ]

    wait_for_completion_interruptible_timeout() returns 0 in case of
    timeout. Check this into account when returning from function.

    Fixes: ef265c55c1ac ("ASoC: mchp-spdifrx: add driver for SPDIF RX")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20230130120647.638049-3-claudiu.beznea@microchip.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbf4ced8885f1434cfee96f353de6af8c9c6e7db
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Mon Jan 30 14:06:40 2023 +0200

    ASoC: mchp-spdifrx: fix controls which rely on rsr register

    [ Upstream commit fa09fa60385abbf99342494b280da8b4aebbc0e9 ]

    The SPDIFRX block is clocked by 2 clocks: peripheral and generic clocks.
    Peripheral clock feeds user interface (registers) and generic clock feeds
    the receiver.

    To enable the receiver the generic clock needs to be enabled and also the
    ENABLE bit of MCHP_SPDIFRX_MR register need to be set.

    The signal control exported by mchp-spdifrx driver reports wrong status
    when the receiver is disabled. This can happen when requesting the signal
    and the capture was not previously started. To solve this the receiver
    needs to be enabled (by enabling generic clock and setting ENABLE bit of
    MR register) before reading the signal status.

    As with this fix there are 2 paths now that need to control the generic
    clock and ENABLE bit of SPDIFRX_MR register (one path though controls, one
    path though configuration) a mutex has been introduced. We can't rely on
    subsystem locking as the controls are protected by
    struct snd_card::controls_rwsem semaphore and configuration is protected
    by a different lock (embedded in snd_pcm_stream_lock_irq()).

    The introduction of mutex is also extended to other controls which rely on
    SPDIFRX_RSR.ULOCK bit as it has been discovered experimentally that having
    both clocks enabled but not the receiver (through ENABLE bit of SPDIFRX.MR)
    leads to inconsistent values of SPDIFRX_RSR.ULOCK. Thus on some controls we
    rely on software state (dev->trigger_enabled protected by mutex) to
    retrieve proper values.

    Fixes: ef265c55c1ac ("ASoC: mchp-spdifrx: add driver for SPDIF RX")
```

```
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Link: https://lore.kernel.org/r/20230130120647.638049-2-claudiu.beznea@microchip.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4551a7c0b5fb3f5b0642b628edc43167f78c5ce4
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Jan 30 15:01:40 2023 +0100

    spi: dw_bt1: fix MUX_MMIO dependencies

    [ Upstream commit d4bde04318c0d33705e9a77d4c7df72f262011e0 ]

    Selecting a symbol with additional dependencies requires
    adding the same dependency here:

    WARNING: unmet direct dependencies detected for MUX_MMIO
      Depends on [n]: MULTIPLEXER [=y] && OF [=n]
      Selected by [y]:
      - SPI_DW_BT1 [=y] && SPI [=y] && SPI_MASTER [=y] && SPI_DESIGNWARE [=y] &&
(MIPS_BAIKAL_T1 || COMPILE_TEST [=y])

    Drop the 'select' here to avoid the problem. Anyone using
    the dw-bt1 SPI driver should make sure they include the
    mux driver as well now.

    Fixes: 7218838109fe ("spi: dw-bt1: Fix undefined devm_mux_control_get symbol")
    Fixes: abf00907538e ("spi: dw: Add Baikal-T1 SPI Controller glue driver")
    Link: https://lore.kernel.org/all/20221218192523.c6vnfo26ua6xqf26@mobilestation/
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Serge Semin <fancer.lancer@gmail.com>
    Link: https://lore.kernel.org/r/20230130140156.3620863-1-arnd@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 95e40af87921f7b7ba8085ac73c614bb412669ad
Author: Amadeusz Sławiński <amadeuszx.slawinski@linux.intel.com>
Date:   Sat Jan 28 00:11:01 2023 +0100

    ASoC: topology: Properly access value coming from topology file

    [ Upstream commit c5d184c92df2b631fb81fe2ce6e96bfc5ba720e5 ]

    When accessing values coming from topology, le32_to_cpu should be used.
    One of recent commits missed that.

    Fixes: 86e2d14b6d1a ("ASoC: topology: Add header payload_size verification")
    Reviewed-by: Cezary Rojewski <cezary.rojewski@intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Amadeusz Sławiński <amadeuszx.slawinski@linux.intel.com>
    Link: https://lore.kernel.org/r/20230127231111.937721-2-
amadeuszx.slawinski@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffc045aa002b3b63b1776a7e7cb2edc20db1b59b
Author: Haibo Chen <haibo.chen@nxp.com>
Date:   Tue Dec 20 17:02:47 2022 +0800

    gpio: vf610: connect GPIO label to dev name

    [ Upstream commit 6f8ecb7f85f441eb7d78ba2a4df45ee8a821934e ]

    Current GPIO label is fixed, so can't distinguish different GPIO
    controllers through labels. Use dev name instead.
```

```
    Fixes: 7f2691a19627 ("gpio: vf610: add gpiolib/IRQ chip driver for Vybrid")
    Signed-off-by: Clark Wang <xiaoning.wang@nxp.com>
    Signed-off-by: Haibo Chen <haibo.chen@nxp.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Bartosz Golaszewski <bartosz.golaszewski@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a23b5130d3785ac736de0f6ef166ce711cbd3fc
Author: Bartosz Golaszewski <bartosz.golaszewski@linaro.org>
Date:   Fri Dec 30 11:00:44 2022 +0100

    gpio: pca9570: rename platform_data to chip_data

    [ Upstream commit a3f7c1d6ddcbd487964c58ff246506a781e5be8f ]

    By convention platform_data refers to structures passed to drivers by
    code that registers devices. When talking about model-specific data
    structures associated with OF compatibles, we usually call them chip_data.

    In order to avoid confusion rename all mentions of platform_data to
    chip_data.

    Fixes: fbb19fe17eae ("gpio: pca9570: add slg7xl45106 support")
    Signed-off-by: Bartosz Golaszewski <bartosz.golaszewski@linaro.org>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b9398f17186a0a9ea1c3227393d688bb4efe147c
Author: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
Date:   Wed Jan 18 17:18:28 2023 +0800

    dt-bindings: display: mediatek: Fix the fallback for mediatek,mt8186-disp-ccorr

    [ Upstream commit 137272ef1b0f17a1815fec00d583515a0163f85e ]

    The mt8186-disp-ccorr is not fully compatible with the mt8183-disp-ccorr
    implementation. It causes a crash when system resumes if it binds to the
    device.

    We should use mt8192-disp-ccorr as fallback of mt8186-disp-ccorr.

    Fixes: 8a26ea19d4dc ("dt-bindings: display: mediatek: add MT8186 SoC binding")
    Signed-off-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Reviewed-by: Rob Herring <robh@kernel.org>
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/20230118091829.755-10-
allen-kh.cheng@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71ad9a78e109fcef6d49915e300f329bab8465fb
Author: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
Date:   Mon Jan 23 23:17:20 2023 +0000

    ASoC: soc-compress.c: fixup private_data on snd_soc_new_compress()

    [ Upstream commit ffe4c0f0bfaa571a676a0e946d4a6a0607f94294 ]

    commit d3268a40d4b19f ("ASoC: soc-compress.c: fix NULL dereference")
    enables DPCM capture, but it should independent from playback.
    This patch fixup it.

    Fixes: d3268a40d4b1 ("ASoC: soc-compress.c: fix NULL dereference")
```

```
    Link: https://lore.kernel.org/r/87tu0i6j7j.wl-kuninori.morimoto.gx@renesas.com
    Acked-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Acked-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Signed-off-by: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
    Link: https://lore.kernel.org/r/871qnkvo1s.wl-kuninori.morimoto.gx@renesas.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3887c771576c5d740c5c5b8bf654a8ab8020b7d
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Tue Nov 22 09:39:49 2022 -0500

    drm/mediatek: Clean dangling pointer on bind error path

    [ Upstream commit 36aa8c61af55675ed967900fbe5deb32d776f051 ]

    mtk_drm_bind() can fail, in which case drm_dev_put() is called,
    destroying the drm_device object. However a pointer to it was still
    being held in the private object, and that pointer would be passed along
    to DRM in mtk_drm_sys_prepare() if a suspend were triggered at that
    point, resulting in a panic. Clean the pointer when destroying the
    object in the error path to prevent this from happening.

    Fixes: 119f5173628a ("drm/mediatek: Add DRM Driver for Mediatek SoC MT8173.")
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/20221122143949.3493104-
1-nfraprado@collabora.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d569ae98ee5490585929be69fea68047679b7b2
Author: ruanjinjie <ruanjinjie@huawei.com>
Date:   Mon Dec 5 17:51:15 2022 +0800

    drm/mediatek: mtk_drm_crtc: Add checks for devm_kcalloc

    [ Upstream commit 5bf1e3bd7da625ccf9a22c8cb7d65271e6e47f4c ]

    As the devm_kcalloc may return NULL, the return value needs to be checked
    to avoid NULL poineter dereference.

    Fixes: 31c5558dae05 ("drm/mediatek: Refactor plane init")
    Signed-off-by: ruanjinjie <ruanjinjie@huawei.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/20221205095115.2905090-
1-ruanjinjie@huawei.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7df662fb47b9c076ca94fc8ddb7ccf3b9a628b7d
Author: Rob Clark <robdclark@chromium.org>
Date:   Thu Jan 19 15:12:55 2023 -0800

    drm/mediatek: Drop unbalanced obj unref

    [ Upstream commit 4deef811828e87e26a978d5d6433b261d4713849 ]

    In the error path, mtk_drm_gem_object_mmap() is dropping an obj
    reference that it doesn't own.

    Fixes: 119f5173628a ("drm/mediatek: Add DRM Driver for Mediatek SoC MT8173.")
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/20230119231255.2883365-
1-robdclark@gmail.com/
```

```
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9434894b62ca0ad3b583f457178f2378fa98d860
Author: Miles Chen <miles.chen@mediatek.com>
Date:   Wed Jan 11 10:44:41 2023 +0800

    drm/mediatek: Use NULL instead of 0 for NULL pointer

    [ Upstream commit 4744cde06f57dd6fbaac468663b1fe2f653eaa16 ]

    Use NULL for NULL pointer to fix the following sparse warning:
    drivers/gpu/drm/mediatek/mtk_drm_gem.c:265:27: sparse: warning: Using plain integer as
NULL pointer

    Fixes: 3df64d7b0a4f ("drm/mediatek: Implement gem prime vmap/vunmap function")
    Signed-off-by: Miles Chen <miles.chen@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/20230111024443.24559-1-
miles.chen@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f88d4566a02cdc1a3c061d348282449af96af5f3
Author: Xinlei Lee <xinlei.lee@mediatek.com>
Date:   Tue Jan 10 13:54:51 2023 +0800

    drm/mediatek: dsi: Reduce the time of dsi from LP11 to sending cmd

    [ Upstream commit 91aeaed2c1147e3b1157dc084d23f190856a6c23 ]

    According to Figure 16 Turnaround Procedure on page 36 in [1], you
    can see the status of LP-00 -> LP10 -> LP11. This state can correspond
    to the state of DSI from LP00 -> LP11 in mtk_dsi_lane_ready function
    in mtk_dsi.c.

    LP-00 -> LP10 -> LP11 takes about 2*TLPX time (refer to [1] page 51
    to see that TLPX is 50ns)

    The delay at the end of the mtk_dsi_lane_ready function should be
    greater than the 2*TLPX specified by the DSI spec, and less than
    the time specified by the DSI_RX (generally 6ms to 40ms), to avoid
    problems caused by the RX specification

    [1]:mipi_D-PHY_specification_v1-1

    Fixes: 39e8d062b03c ("drm/mediatek: Keep dsi as LP00 before dcs cmds transfer")
    Signed-off-by: Xinlei Lee <xinlei.lee@mediatek.com>
    Acked-by: Sam Ravnborg <sam@ravnborg.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://patchwork.kernel.org/project/linux-mediatek/patch/1673330093-6771-2-git-
send-email-xinlei.lee@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 187e821c8d19ed81c14cc9d22ff3ec13c0857c54
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Thu Dec 29 21:18:30 2022 +0200

    drm/msm/dpu: set pdpu->is_rt_pipe early in dpu_plane_sspp_atomic_update()

    [ Upstream commit 1d233b1cb149ec78c20fac58331b27bb460f9558 ]

    The function dpu_plane_sspp_atomic_update() updates pdpu->is_rt_pipe
    flag, but after the commit 854f6f1c653b ("drm/msm/dpu: update the qos
```

```
    remap only if the client type changes") it sets the flag late, after all
    the qos functions have updated QoS programming. Move the flag update
    back to the place where it happened before the mentioned commit to let
    the pipe be programmed according to its current RT/non-RT state.

    Fixes: 854f6f1c653b ("drm/msm/dpu: update the qos remap only if the client type
changes")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/516239/
    Link: https://lore.kernel.org/r/20221229191856.3508092-2-dmitry.baryshkov@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7db51a15a99d954eed093960a39be4255a65866d
Author: Matt Roper <matthew.d.roper@intel.com>
Date:   Wed Jan 25 15:41:59 2023 -0800

    drm/i915/xehp: Annotate a couple more workaround registers as MCR

    [ Upstream commit 7649a5d1f263b8cc5c2827ef0443ead9bee9ae0e ]

    GAMSTLB_CTRL and GAMCNTRL_CTRL became multicast/replicated registers on
    Xe_HP.  They should be defined accordingly and use MCR-aware operations.

    These registers have only been used for some dg2/xehpsdv workarounds, so
    this fix is mostly just for consistency/future-proofing; even lacking
    the MCR annotation, workarounds will always be properly applied in a
    multicast manner on these platforms.

    Cc: Gustavo Sousa <gustavo.sousa@intel.com>
    Fixes: 58bc2453ab8a ("drm/i915: Define multicast registers as a new type")
    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Reviewed-by: Gustavo Sousa <gustavo.sousa@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230125234159.3015385-3-
matthew.d.roper@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c543b2d9e622b2e265796f14ab52673ea90b1f5
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
Date:   Mon Jan 2 22:18:11 2023 +0000

    pinctrl: renesas: rzg2l: Fix configuring the GPIO pins as interrupts

    [ Upstream commit 00dfe29887761405ccd23cc0aa07cb86623bb2b7 ]

    On the RZ/G2UL SoC we have less number of pins compared to RZ/G2L and also
    the pin configs are completely different. This patch makes sure we use the
    appropriate pin configs for each SoC (which is passed as part of the OF
    data) while configuring the GPIO pin as interrupts instead of using
    rzg2l_gpio_configs[] for all the SoCs.

    Fixes: bfc69bdbaad1 ("pinctrl: renesas: rzg2l: Add RZ/G2UL support")
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Acked-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20230102221815.273719-3-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76d433431d858522303725a463ae05b689f3fe93
Author: Matt Roper <matthew.d.roper@intel.com>
Date:   Wed Jan 25 15:41:57 2023 -0800
```

```
    drm/i915/xehp: GAM registers don't need to be re-applied on engine resets

    [ Upstream commit 49cbda6386efc5f023f396dca087aaba5d4f885e ]

    Register reset characteristics (i.e., whether the register maintains or
    loses its value on engine reset) is an important factor that determines
    which wa_list we want to add workarounds to.  We recently found out that
    the bspec documentation for the Xe_HP's "GAM" registers in the 0xC800 -
    0xCFFF range was misleading; these registers do not actually lose their
    value on engine resets as the documentation implied.  This means there's
    no need to re-apply workarounds touching these registers after a reset,
    and the corresponding workarounds should be moved from the 'engine'
    lists back to the 'gt' list.

    v2:
     - Don't add Wa_18018781329 to xehpsdv; the original condition didn't
       include that platform.  (Gustavo)
     - Move the MTL code to the GT function as-is for now; we'll take care
       of the additional fixes needed in a follow-up patch.

    Cc: Gustavo Sousa <gustavo.sousa@intel.com>
    Fixes: edf176f48d87 ("drm/i915/dg2: Move misplaced 'ctx' & 'gt' wa's to engine wa list")
    Fixes: b2006061ae28 ("drm/i915/xehpsdv: Move render/compute engine reset domains related
workarounds")
    Fixes: 41bb543f5598 ("drm/i915/mtl: Add initial gt workarounds")
    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Reviewed-by: Gustavo Sousa <gustavo.sousa@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230125234159.3015385-1-
matthew.d.roper@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5008c43a7187d2c83949e8d0159aea18f3eccb4f
Author: Matt Roper <matthew.d.roper@intel.com>
Date:   Thu Jan 5 15:44:08 2023 -0800

    drm/i915/mtl: Add initial gt workarounds

    [ Upstream commit 41bb543f5598fb44e0e8dbd723f5821be83b466b ]

    This patch introduces initial gt workarounds for the MTL platform.

    v2: drop redundant/stale comments specifying wa platforms affected
    (Lucas).
    v3: drop additional redundant stale comments (MattR)

    Bspec: 66622

    Signed-off-by: Matt Roper <matthew.d.roper@intel.com>
    Signed-off-by: Matt Atwood <matthew.s.atwood@intel.com>
    Reviewed-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Radhakrishna Sripada <radhakrishna.sripada@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230105234408.277750-1-
matthew.s.atwood@intel.com
    Stable-dep-of: 49cbda6386ef ("drm/i915/xehp: GAM registers don't need to be re-applied
on engine resets")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c4fa13e40f945d6e27babd3c48b05b2968eb4a7b
Author: Mikko Perttunen <mperttunen@nvidia.com>
Date:   Thu Jan 19 15:39:01 2023 +0200

    drm/tegra: firewall: Check for is_addr_reg existence in IMM check

    [ Upstream commit 1b5c09de25e8c08655c270a70e5e74e93b6bad1f ]
```

```
    In the IMM opcode check, don't call is_addr_reg if it's not set.

    Fixes: 8cc95f3fd35e ("drm/tegra: Add job firewall")
    Signed-off-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9ef73d7d43d1115018d2c63632b2db1fc9c8e47
Author: Mikko Perttunen <mperttunen@nvidia.com>
Date:   Thu Jan 19 15:39:00 2023 +0200

    gpu: host1x: Don't skip assigning syncpoints to channels

    [ Upstream commit eb258cc1fd458e584082be987dbc6ec42668c05e ]

    The code to write the syncpoint channel assignment register
    incorrectly skips the write if hypervisor registers are not available.

    The register, however, is within the guest aperture so remove the
    check and assign syncpoints properly even on virtualized systems.

    Fixes: c3f52220f276 ("gpu: host1x: Enable Tegra186 syncpoint protection")
    Signed-off-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 40fa34580c87d78c7d8f3acc80b62df24612c62f
Author: Mikko Perttunen <mperttunen@nvidia.com>
Date:   Thu Jan 19 15:38:59 2023 +0200

    gpu: host1x: Fix mask for syncpoint increment register

    [ Upstream commit 79aad29c7d2d2cd64790115d3a6ebac28c00a8ec ]

    On Tegra186+, the syncpoint ID has 10 bits of space. To allow
    using more than 256 syncpoints, fix the mask.

    Fixes: 9abdd497cd0a ("gpu: host1x: Tegra234 device data and headers")
    Signed-off-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dcc5206bb4c93c1981765c18e7a070c4dfcff97b
Author: Guodong Liu <Guodong.Liu@mediatek.com>
Date:   Wed Jan 18 14:20:36 2023 +0800

    pinctrl: mediatek: Initialize variable *buf to zero

    [ Upstream commit 2e34f82ba214134ecf590fbe0cdbd87401645a8a ]

    Coverity spotted that *buf is not initialized to zero in
    mtk_pctrl_dbg_show. Using uninitialized variable *buf as argument to %s
    when calling seq_printf. Fix this coverity by initializing *buf as zero.

    Fixes: 184d8e13f9b1 ("pinctrl: mediatek: Add support for pin configuration dump via
debugfs.")
    Signed-off-by: Guodong Liu <Guodong.Liu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20230118062036.26258-3-Guodong.Liu@mediatek.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e24f7b04307ba42a68e4d3012ac9b61fe429fe7d
Author: Guodong Liu <Guodong.Liu@mediatek.com>
Date:   Wed Jan 18 14:20:35 2023 +0800
```

```
    pinctrl: mediatek: Initialize variable pullen and pullup to zero

    [ Upstream commit a298c70a10c604a6b3df5a0aa56597b705ba0f6b ]

    Coverity spotted that pullen and pullup is not initialized to zero in
    mtk_pctrl_show_one_pin. The uninitialized variable pullen is used in
    assignment statement "rsel = pullen;" in mtk_pctrl_show_one_pin, and
    Uninitialized variable pullup is used when calling scnprintf. Fix this
    coverity by initializing pullen and pullup as zero.

    Fixes: 184d8e13f9b1 ("pinctrl: mediatek: Add support for pin configuration dump via
debugfs.")
    Signed-off-by: Guodong Liu <Guodong.Liu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20230118062036.26258-2-Guodong.Liu@mediatek.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff01e46c5a6327dc63d5eacba99722fa17df17e1
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Fri Jan 13 23:53:50 2023 +0200

    pinctrl: bcm2835: Remove of_node_put() in bcm2835_of_gpio_ranges_fallback()

    [ Upstream commit 2d578dd27871372f7159dd3206149ec616700d87 ]

    Remove wrong of_node_put() in bcm2835_of_gpio_ranges_fallback(),
    there is no counterpart of_node_get() for it.

    Fixes: d2b67744fd99 ("pinctrl: bcm2835: implement hook for missing gpio-ranges")
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Reviewed-by: Stefan Wahren <stefan.wahren@i2se.com>
    Tested-by: Stefan Wahren <stefan.wahren@i2se.com>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Reviewed-by: Florian Fainelli <f.fainelli@gmail.com>
    Acked-by: Bartosz Golaszewski <bartosz.golaszewski@linaro.org>
    Link: https://lore.kernel.org/r/20230113215352.44272-3-andriy.shevchenko@linux.intel.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f607dcdb990cf0afe3916f643bf68b1727264836
Author: farah kassabri <fkassabri@habana.ai>
Date:   Sun Jan 8 17:33:44 2023 +0200

    habanalabs: bugs fixes in timestamps buff alloc

    [ Upstream commit 1693fef9e95dbe8ab767d208a02328fff13fbb94 ]

    use argument instead of fixed GFP value for allocation
    in Timestamps buffers alloc function.
    change data type of size to size_t.

    Fixes: 9158bf69e74f ("habanalabs: Timestamps buffers registration")
    Signed-off-by: farah kassabri <fkassabri@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37ff771ed008b9cbffd0eab77985968364694ce3
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Dec 6 15:48:19 2022 +0800

    drm/msm/mdp5: Add check for kzalloc
```

```
    [ Upstream commit 13fcfcb2a9a4787fe4e49841d728f6f2e9fa6911 ]

    As kzalloc may fail and return NULL pointer,
    it should be better to check the return value
    in order to avoid the NULL pointer dereference.

    Fixes: 1cff7440a86e ("drm/msm: Convert to using __drm_atomic_helper_crtc_reset() for
reset.")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/514154/
    Link: https://lore.kernel.org/r/20221206074819.18134-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c745fd463ccf21943cfc35cb8842a33dda2ae259
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Dec 6 16:02:36 2022 +0800

    drm/msm/dpu: Add check for pstates

    [ Upstream commit 93340e10b9c5fc86730d149636e0aa8b47bb5a34 ]

    As kzalloc may fail and return NULL pointer,
    it should be better to check pstates
    in order to avoid the NULL pointer dereference.

    Fixes: 25fdd5933e4c ("drm/msm: Add SDM845 DPU support")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/514160/
    Link: https://lore.kernel.org/r/20221206080236.43687-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a52e5a002d18bffabff66f6f59a74f8e9aac5afe
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Dec 6 16:05:17 2022 +0800

    drm/msm/dpu: Add check for cstate

    [ Upstream commit c96988b7d99327bb08bd9efd29a203b22cd88ace ]

    As kzalloc may fail and return NULL pointer,
    it should be better to check cstate
    in order to avoid the NULL pointer dereference
    in __drm_atomic_helper_crtc_reset.

    Fixes: 1cff7440a86e ("drm/msm: Convert to using __drm_atomic_helper_crtc_reset() for
reset.")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/514163/
    Link: https://lore.kernel.org/r/20221206080517.43786-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf9432d473e260b4dad6c70cb055320f36c0b3ea
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 04:01:52 2023 +0200

    drm/msm: use strscpy instead of strncpy

    [ Upstream commit d7fd8634f48d76aa799ed57beb7d87dab91bde80 ]
```

```
    Using strncpy can result in non-NULL-terminated destination string. Use
    strscpy instead. This fixes following warning:

    drivers/gpu/drm/msm/msm_fence.c: In function ‘msm_fence_context_alloc’:
    drivers/gpu/drm/msm/msm_fence.c:25:9: warning: ‘strncpy’ specified bound 32 equals
 destination size [-Wstringop-truncation]
       25 |         strncpy(fctx->name, name, sizeof(fctx->name));
          |         ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    Fixes: f97decac5f4c ("drm/msm: Support multiple ringbuffers")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/518787/
    Link: https://lore.kernel.org/r/20230118020152.1689213-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9c46159f7a4c8203c01ef0fabe9baa7272ab542
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Mon Jan 16 12:30:55 2023 +0200

    drm/msm/dpu: sc7180: add missing WB2 clock control

    [ Upstream commit 255f056181ac69278dcd8778f0be32a21b2d2a6a ]

    Add missing DPU_CLK_CTRL_WB2 to sc7180_mdp clocks array.

    Fixes: 51e4d60e6ba5 ("drm/msm/dpu: add writeback support for sc7180")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Tested-by: Jessica Zhang <quic_jesszhan@quicinc.com> # Trogdor (sc7180)
    Patchwork: https://patchwork.freedesktop.org/patch/518504/
    Link: https://lore.kernel.org/r/20230116103055.780767-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aeee3dcc835b121befb350a6dcdee8f35d96b1d4
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Thu Jan 12 15:42:13 2023 -0800

    scsi: ufs: exynos: Fix DMA alignment for PAGE_SIZE != 4096

    [ Upstream commit 86bd0c4a2a5dc4265884648cb92c681646509692 ]

    The Exynos UFS controller only supports scatter/gather list elements that
    are aligned on a 4 KiB boundary. Fix DMA alignment in case PAGE_SIZE !=
    4096. Rename UFSHCD_QUIRK_ALIGN_SG_WITH_PAGE_SIZE into
    UFSHCD_QUIRK_4KB_DMA_ALIGNMENT.

    Cc: Kiwoong Kim <kwmad.kim@samsung.com>
    Fixes: 2b2bfc8aa519 ("scsi: ufs: Introduce a quirk to allow only page-aligned sg
 entries")
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Reviewed-by: Alim Akhtar <alim.akhtar@samsung.com>
    Tested-by: Alim Akhtar <alim.akhtar@samsung.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 697d0cdb9b48a87d46cbce12681b4d13e3ca8456
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Jan 20 22:00:53 2023 +0100

    drm/msm/dsi: Allow 2 CTRLs on v2.5.0
```

```
    [ Upstream commit 1ae654ded7c5a19dc13f57a4fe4434fef879b6f9 ]

    v2.5.0 support was originally added for SC7280, but this hw is also
    present on SM8350, which has one more DSI host. Bump up the dsi count
    and fill in the register of the secondary host to allow it to probe.

    This should not have any adverse effects on SC7280, as the secondary
    CTRL will only be touched if it's defined, anyway.

    Fixes: 65c391b31994 ("drm/msm/dsi: Add DSI support for SC7280")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/519513/
    Link: https://lore.kernel.org/r/20230120210101.2146852-1-konrad.dybcio@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8799435762cef75499f869825b743f16373e5771
Author: Jagan Teki <jagan@amarulasolutions.com>
Date:   Mon Dec 12 20:27:44 2022 +0530

    drm: exynos: dsi: Fix MIPI_DSI*_NO_* mode flags

    [ Upstream commit 996e1defca34485dd2bd70b173f069aab5f21a65 ]

    HFP/HBP/HSA/EOT_PACKET modes in Exynos DSI host specifies
    0 = Enable and 1 = Disable.

    The logic for checking these mode flags was correct before
    the MIPI_DSI*_NO_* mode flag conversion.

    This patch is trying to fix this MIPI_DSI*_NO_* mode flags handling
    Exynos DSI host and update the mode_flags in relevant panel drivers.

    Fixes: 0f3b68b66a6d ("drm/dsi: Add _NO_ to MIPI_DSI_* flags disabling features")
    Reviewed-by: Marek Vasut <marex@denx.de>
    Reviewed-by: Nicolas Boichat <drinkcat@chromium.org>
    Reported-by: Sébastien Szymanski <sebastien.szymanski@armadeus.com>
    Signed-off-by: Jagan Teki <jagan@amarulasolutions.com>
    Reviewed-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Signed-off-by: Marek Vasut <marex@denx.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221212145745.15387-1-
jagan@amarulasolutions.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e819faf4a8446c030c503fecfe47f713cd59f19c
Author: Daniel Mentz <danielmentz@google.com>
Date:   Mon Jan 16 17:49:07 2023 -0500

    drm/mipi-dsi: Fix byte order of 16-bit DCS set/get brightness

    [ Upstream commit c9d27c6be518b4ef2966d9564654ef99292ea1b3 ]

    The MIPI DCS specification demands that brightness values are sent in
    big endian byte order. It also states that one parameter (i.e. one byte)
    shall be sent/received for 8 bit wide values, and two parameters shall
    be used for values that are between 9 and 16 bits wide.

    Add new functions to properly handle 16-bit brightness in big endian,
    since the two 8- and 16-bit cases are distinct from each other.

    [richard: use separate functions instead of switch/case]
    [richard: split into 16-bit component]

    Fixes: 1a9d759331b8 ("drm/dsi: Implement DCS set/get display brightness")
```

```
    Signed-off-by: Daniel Mentz <danielmentz@google.com>
    Link:
 https://android.googlesource.com/kernel/msm/+/754affd62d0ee268c686c53169b1dbb7deac8550
    [richard: fix 16-bit brightness_get]
    Signed-off-by: Richard Acayan <mailingradian@gmail.com>
    Tested-by: Caleb Connolly <caleb@connolly.tech>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230116224909.23884-2-
 mailingradian@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit afcdc3d5d7ea20c91067bc8b14fb46b1348a2243
 Author: Randy Dunlap <rdunlap@infradead.org>
 Date:   Sat Jan 14 10:57:36 2023 -0800

    regulator: tps65219: use IS_ERR() to detect an error pointer

    [ Upstream commit 2bbba115c3c9a647bcb3201b014fcc3728fe75c8 ]

    Fix pointer comparison to integer warning from gcc & sparse:

    GCC:
    ../drivers/regulator/tps65219-regulator.c:370:26: warning: ordered comparison of pointer
 with integer zero [-Wextra]
      370 |                 if (rdev < 0) {
          |                          ^

    sparse warning:
    drivers/regulator/tps65219-regulator.c:370:26: sparse: error: incompatible types for
 operation (<):
    drivers/regulator/tps65219-regulator.c:370:26: sparse:    struct regulator_dev *
 [assigned] rdev
    drivers/regulator/tps65219-regulator.c:370:26: sparse:    int

    Fixes: c12ac5fc3e0a ("regulator: drivers: Add TI TPS65219 PMIC regulators support")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Jerome Neanne <jneanne@baylibre.com>
    Cc: Tony Lindgren <tony@atomide.com>
    Cc: linux-omap@vger.kernel.org
    Cc: Liam Girdwood <lgirdwood@gmail.com>
    Cc: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20230114185736.2076-1-rdunlap@infradead.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit ef0c9dfb7051513578854ad6e77b8fe0716e0872
 Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
 Date:   Wed Jan 18 10:16:51 2023 +0200

    drm/bridge: lt9611: pass a pointer to the of node

    [ Upstream commit b0a7f8736789935f62d6df32d441cdf05a5c05d2 ]

    Pass a pointer to the OF node while registering lt9611 MIPI device.

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-7-
 dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit ba6c79b56d8a98b91cd147bf412b713130523c7b
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 10:16:50 2023 +0200

    drm/bridge: lt9611: fix clock calculation

    [ Upstream commit 2576eb26494eb0509dd9ceb0cd27771a7a5e3674 ]

    Instead of having several fixed values for the pcr register, calculate
    it before programming. This allows the bridge to support most of the
    display modes.

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-6-
dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6b4324a84693e77cd9d5679b4fb85484233d791
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 10:16:49 2023 +0200

    drm/bridge: lt9611: fix programming of video modes

    [ Upstream commit ad188aa47edaa033a270e1a3efae43836ff47569 ]

    Program the upper part of the hfront_porch into the proper register.

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-5-
dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a660eb89db983080d63a882ae273d4feceb5ee0a
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 10:16:48 2023 +0200

    drm/bridge: lt9611: fix polarity programming

    [ Upstream commit 0b157efa384ea417304b1da284ee2f603c607fc3 ]

    Fix programming of hsync and vsync polarities

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-4-
dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9701d566238f0dc6308b9f5b7bf0981760fb2c83
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 10:16:47 2023 +0200

    drm/bridge: lt9611: fix HPD reenablement

    [ Upstream commit a7790f6bd38f3642b60ae3504a2c749135b89451 ]

    The driver will reset the bridge in the atomic_pre_enable(). However
```

```
    this will also drop the HPD interrupt state. Instead of resetting the
    bridge, properly wake it up. This fixes the HPD interrupt delivery after
    the disable/enable cycle.

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-3-
dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 56e7d7dcc206eacfd7035364260fcfdcade35728
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Jan 18 10:16:46 2023 +0200

    drm/bridge: lt9611: fix sleep mode setup

    [ Upstream commit ae2d329f104b75a0a78dcaded29fe6283289cdf9 ]

    On atomic_post_disable the bridge goes to the low power state. However
    the code disables too much of the chip, so the HPD event is not being
    detected and delivered to the host. Reduce the power saving in order to
    get the HPD event.

    Fixes: 23278bf54afe ("drm/bridge: Introduce LT9611 DSI to HDMI bridge")
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230118081658.2198520-2-
dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9fe3644c720ac87d150f0bba5a4ae86cae55afaf
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Tue Jan 10 00:15:55 2023 +0100

    drm/msm/dpu: Disallow unallocated resources to be returned

    [ Upstream commit abc40122d9a69f56c04efb5a7485795f5ac799d1 ]

    In the event that the topology requests resources that have not been
    created by the system (because they are typically not represented in
    dpu_mdss_cfg ^1), the resource(s) in global_state (in this case DSC
    blocks, until their allocation/assignment is being sanity-checked in
    "drm/msm/dpu: Reject topologies for which no DSC blocks are available")
    remain NULL but will still be returned out of
    dpu_rm_get_assigned_resources, where the caller expects to get an array
    containing num_blks valid pointers (but instead gets these NULLs).

    To prevent this from happening, where null-pointer dereferences
    typically result in a hard-to-debug platform lockup, num_blks shouldn't
    increase past NULL blocks and will print an error and break instead.
    After all, max_blks represents the static size of the maximum number of
    blocks whereas the actual amount varies per platform.

    ^1: which can happen after a git rebase ended up moving additions to
    _dpu_cfg to a different struct which has the same patch context.

    Fixes: bb00a452d6f7 ("drm/msm/dpu: Refactor resource manager")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/517636/
    Link: https://lore.kernel.org/r/20230109231556.344977-1-marijn.suijten@somainline.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
```

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

commit abadeef9a82f86e6668759259adbcf8be6f44a9b
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Mon Dec 12 17:11:17 2022 +0800

    drm/msm/gem: Add check for kmalloc

    [ Upstream commit d839f0811a31322c087a859c2b181e2383daa7be ]

    Add the check for the return value of kmalloc in order to avoid
    NULL pointer dereference in copy_from_user.

    Fixes: 20224d715a88 ("drm/msm/submit: Move copy_from_user ahead of locking bos")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/514678/
    Link: https://lore.kernel.org/r/20221212091117.43511-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad7da026e8a4b412bfc37e1341c417b2d4b9d9f6
Author: Leo Liu <leo.liu@amd.com>
Date:   Mon Jan 16 09:55:57 2023 -0500

    drm/amdgpu: Use the sched from entity for amdgpu_cs trace

    [ Upstream commit cf22ef78f22ce4df4757472c5dbd33c430c5b659 ]

    The problem is that base sched hasn't been assigned yet at this moment,
    causing something like "ring=0" all the time from trace.

    mpv:cs0-3473    [002] ..... 129.047431: amdgpu_cs: ring=0, dw=48, fences=0
    mpv:cs0-3473    [002] ..... 129.089125: amdgpu_cs: ring=0, dw=48, fences=0
    mpv:cs0-3473    [002] ..... 129.130987: amdgpu_cs: ring=0, dw=48, fences=0
    mpv:cs0-3473    [002] ..... 129.172478: amdgpu_cs: ring=0, dw=48, fences=0

    Fixes: 4624459c84d7 ("drm/amdgpu: add gang submit frontend v6")
    Signed-off-by: Leo Liu <leo.liu@amd.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59473f6b70c470cb6db0fbfe9bb174314810e338
Author: Alexey V. Vissarionov <gremlin@altlinux.org>
Date:   Tue Jan 17 14:15:23 2023 +0300

    ALSA: hda/ca0132: minor fix for allocation size

    [ Upstream commit 3ee0fe7fa39b14d1cea455b7041f2df933bd97d2 ]

    Although the "dma_chan" pointer occupies more or equal space compared
    to "*dma_chan", the allocation size should use the size of variable
    itself.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 01ef7dbffb41 ("ALSA: hda - Update CA0132 codec to load DSP firmware binary")
    Signed-off-by: Alexey V. Vissarionov <gremlin@altlinux.org>
    Link: https://lore.kernel.org/r/20230117111522.GA15213@altlinux.org
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5fef23c1c0edceb44d16e64e7818f27d48b5bc38
Author: Akhil P Oommen <quic_akhilpo@quicinc.com>
```

```
Date:   Wed Dec 21 20:39:56 2022 +0530

    drm/msm/adreno: Fix null ptr access in adreno_gpu_cleanup()

    [ Upstream commit dbeedbcb268d055d8895aceca427f897e12c2b50 ]

    Fix the below kernel panic due to null pointer access:
    [   18.504431] Unable to handle kernel NULL pointer dereference at virtual address
0000000000000048
    [   18.513464] Mem abort info:
    [   18.516346]   ESR = 0x0000000096000005
    [   18.520204]   EC = 0x25: DABT (current EL), IL = 32 bits
    [   18.525706]   SET = 0, FnV = 0
    [   18.528878]   EA = 0, S1PTW = 0
    [   18.532117]   FSC = 0x05: level 1 translation fault
    [   18.537138] Data abort info:
    [   18.540110]   ISV = 0, ISS = 0x00000005
    [   18.544060]   CM = 0, WnR = 0
    [   18.547109] user pgtable: 4k pages, 39-bit VAs, pgdp=0000000112826000
    [   18.553738] [0000000000000048] pgd=0000000000000000, p4d=0000000000000000,
pud=0000000000000000
    [   18.562690] Internal error: Oops: 0000000096000005 [#1] PREEMPT SMP
    **Snip**
    [   18.696758] Call trace:
    [   18.699278]  adreno_gpu_cleanup+0x30/0x88
    [   18.703396]  a6xx_destroy+0xc0/0x130
    [   18.707066]  a6xx_gpu_init+0x308/0x424
    [   18.710921]  adreno_bind+0x178/0x288
    [   18.714590]  component_bind_all+0xe0/0x214
    [   18.718797]  msm_drm_bind+0x1d4/0x614
    [   18.722566]  try_to_bring_up_aggregate_device+0x16c/0x1b8
    [   18.728105]  __component_add+0xa0/0x158
    [   18.732048]  component_add+0x20/0x2c
    [   18.735719]  adreno_probe+0x40/0xc0
    [   18.739300]  platform_probe+0xb4/0xd4
    [   18.743068]  really_probe+0xfc/0x284
    [   18.746738]  __driver_probe_device+0xc0/0xec
    [   18.751129]  driver_probe_device+0x48/0x110
    [   18.755421]  __device_attach_driver+0xa8/0xd0
    [   18.759900]  bus_for_each_drv+0x90/0xdc
    [   18.763843]  __device_attach+0xfc/0x174
    [   18.767786]  device_initial_probe+0x20/0x2c
    [   18.772090]  bus_probe_device+0x40/0xa0
    [   18.776032]  deferred_probe_work_func+0x94/0xd0
    [   18.780686]  process_one_work+0x190/0x3d0
    [   18.784805]  worker_thread+0x280/0x3d4
    [   18.788659]  kthread+0x104/0x1c0
    [   18.791981]  ret_from_fork+0x10/0x20
    [   18.795654] Code: f9400408 aa0003f3 aa1f03f4 91142015 (f9402516)
    [   18.801913] ---[ end trace 0000000000000000 ]---
    [   18.809039] Kernel panic - not syncing: Oops: Fatal exception

    Fixes: 17e822f7591f ("drm/msm: fix unbalanced pm_runtime_enable in adreno_gpu_{init,
cleanup}")
    Signed-off-by: Akhil P Oommen <quic_akhilpo@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/515605/
    Link:
https://lore.kernel.org/r/20221221203925.v2.1.Ib978de92c4bd000b515486aad72e96c2481f84d0@chan
geid
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8efc9f6379be6d21c8ed25eaf1c81882200c3e31
Author: Marek Vasut <marex@denx.de>
Date:   Sun Oct 16 02:35:56 2022 +0200
```

```
    drm/bridge: tc358767: Set default CLRSIPO count

    [ Upstream commit 01338bb82fed40a6a234c2b36a92367c8671adf0 ]

    The current CLRSIPO count is still marginal and does not work with high
    DSI clock rates in burst mode. Increase it further to allow the DSI link
    to work at up to 1Gbps lane speed. This returns the counts to defaults
    as provided by datasheet.

    Fixes: ea6490b02240b ("drm/bridge: tc358767: increase CLRSIPO count")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Acked-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221016003556.406441-1-
marex@denx.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 28dcf69c8d2a149b5a9482f8a37524a47ca9cfa1
Author: Shengjiu Wang <shengjiu.wang@nxp.com>
Date:   Mon Jan 16 15:07:54 2023 +0800

    ASoC: fsl_sai: initialize is_dsp_mode flag

    [ Upstream commit a23924b7dd7b748fff8e305e1daf590fed2af21b ]

    Initialize is_dsp_mode flag in the beginning of function
    fsl_sai_set_dai_fmt_tr().

    When the DAIFMT is DAIFMT_DSP_B the first time, is_dsp_mode is
    true, then the second time DAIFMT is DAIFMT_I2S, is_dsp_mode
    still true, which is a wrong state. So need to initialize
    is_dsp_mode flag every time.

    Fixes: a3f7dcc9cc03 ("ASoC: fsl-sai: Add SND_SOC_DAIFMT_DSP_A/B support.")
    Signed-off-by: Shengjiu Wang <shengjiu.wang@nxp.com>
    Reviewed-by: Iuliana Prodan <iuliana.prodan@nxp.com>
    Link: https://lore.kernel.org/r/1673852874-32200-1-git-send-email-shengjiu.wang@nxp.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0bd119bd900b947eaee4d257c05793983384af06
Author: Quinn Tran <qutran@marvell.com>
Date:   Wed Dec 21 20:39:31 2022 -0800

    scsi: qla2xxx: edif: Fix clang warning

    [ Upstream commit 2f5fab1b6c3a8efc93ba52c28539c45a8d0142ad ]

    clang warning:

      drivers/scsi/qla2xxx/qla_edif_bsg.h:93:12: warning: field remote_pid
      within 'struct app_pinfo_req' is less aligned than 'port_id_t' and is
      usually due to 'struct app_pinfo_req' being packed, which can lead to
      unaligned accesses [-Wunaligned-access]
      port_id_t remote_pid;
                ^
      2 warnings generated.

    Remove u32 field in remote_pid to silence warning.

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: 7ebb336e45ef ("scsi: qla2xxx: edif: Add start + stop bsgs")
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
```

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

commit f2dde125ae9849b84f46a98abd98f655148821ab
Author: Quinn Tran <qutran@marvell.com>
Date:   Mon Dec 19 03:07:42 2022 -0800

    scsi: qla2xxx: Fix exchange oversubscription for management commands

    [ Upstream commit 5f63a163ed2f12c34dd4ae9b2757962ec7bb86e5 ]

    Add resource checking for management (non-I/O) commands.

    Fixes: 89c72f4245a8 ("scsi: qla2xxx: Add IOCB resource tracking")
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3f3a950e30da30950d498cbb8f6be444b9ec935
Author: Quinn Tran <qutran@marvell.com>
Date:   Mon Dec 19 03:07:41 2022 -0800

    scsi: qla2xxx: Fix exchange oversubscription

    [ Upstream commit 41e5afe51f75f2858f5563145348f6c26d307b8f ]

    In large environment, it is possible to experience command timeout and
    escalation of path recovery. Currently the driver does not track the number
    of exchanges/commands sent to FW. If there is a delay for commands at the
    head of the queue, then this will create back pressure for commands at the
    back of the queue.

    Check for exchange availability before command submission.

    Fixes: 89c72f4245a8 ("scsi: qla2xxx: Add IOCB resource tracking")
    Signed-off-by: Quinn Tran <qutran@marvell.com>
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf71f10f89930a2b8232889432a50c3cefad99a2
Author: Abel Vesa <abel.vesa@linaro.org>
Date:   Sat Dec 31 16:27:20 2022 +0200

    drm/panel-edp: fix name for IVO product id 854b

    [ Upstream commit 9cce08cadc6ce8670280d0a042cf0b6d2987d9f9 ]

    The actual name is R133NW4K-R0.

    Fixes: 0f9fa5f58c78 ("drm/panel-edp: add IVO M133NW4J-R3 panel entry")
    Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
    Reviewed-by: Johan Hovold <johan+linaro@kernel.org>
    Tested-by: Steev Klimaszewski <steev@kali.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221231142721.338643-1-
abel.vesa@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfd8ce7f1dbc743ddd28922413370711a22cd4dc
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sat Jun 18 02:33:25 2022 +0300
```

```
    drm/msm: clean event_thread->worker in case of an error

    [ Upstream commit c79bb6b92defdcb834ceeeed9c1cf591beb1b71a ]

    If worker creation fails, nullify the event_thread->worker, so that
    msm_drm_uninit() doesn't try accessing invalid memory location. While we
    are at it, remove duplicate assignment to the ret variable.

    Fixes: 1041dee2178f ("drm/msm: use kthread_create_worker instead of kthread_run")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/490106/
    Link: https://lore.kernel.org/r/20220617233328.1143665-2-dmitry.baryshkov@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3df634fc05f62eae3708192313effc42785a797f
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:25 2022 +0100

    drm/vc4: hdmi: Correct interlaced timings again

    [ Upstream commit 771d6539f27bd55f43d8a95d53a7eeaaffa2681c ]

    The back porch timings were correct, only the sync offset was wrong.
    Correct timing is now reported for 1080i and 576i, but the h offset is
    incorrect for 480i for non-obvious reasons.

    Fixes: fb10dc451c0f ("drm/vc4: hdmi: Correct HDMI timing registers for interlaced
modes")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-14-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 633c84bee0e5d799dd03fd5c3fb13b88a0361d0e
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:17 2022 +0100

    drm/vc4: hvs: Fix colour order for xRGB1555 on HVS5

    [ Upstream commit 902973dc1a049c0d7bf0c222b8f2b3876f01b4a2 ]

    Same as the xRGB8888 formats, HVS5 has managed to swap the colour
    channels for the xRGB1555 formats as well. Add the relevant
    config for pixel_order_hvs5.

    Fixes: c54619b0bfb3 ("drm/vc4: Add support for the BCM2711 HVS5")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-6-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7aa0c9e8f833ffa260af35e2e01b3caae231908b
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:15 2022 +0100

    drm/vc4: hvs: Correct interrupt masking bit assignment for HVS5

    [ Upstream commit 87551ec650bb87d35f1b29bba6a2430896e08da0 ]

    HVS5 has moved the interrupt enable bits around within the
    DISPCTRL register, therefore the configuration has to be updated
    to account for this.
```

```
    Fixes: c54619b0bfb3 ("drm/vc4: Add support for the BCM2711 HVS5")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-4-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6e80c98aa3d25bf47293bede8e73466f71a24b2
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:14 2022 +0100

    drm/vc4: hvs: SCALER_DISPBKGND_AUTOHS is only valid on HVS4

    [ Upstream commit 982ee94486863a41c6af9f2ab3f6681f72bc5c48 ]

    The bit used for SCALER_DISPBKGND_AUTOHS in SCALER_DISPBKGNDX
    has been repurposed on HVS5 to configure whether a display can
    win back-to-back arbitration wins for the COB.

    This is not desirable, therefore only select this bit on HVS4,
    and explicitly clear it on HVS5.

    Fixes: c54619b0bfb3 ("drm/vc4: Add support for the BCM2711 HVS5")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-3-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89016ccd1fd88280ab2783f3026b0e6a0a4bafbc
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:13 2022 +0100

    drm/vc4: hvs: Set AXI panic modes

    [ Upstream commit df993fced230daa8452892406f3180c93ebf7e7b ]

    The HVS can change AXI request mode based on how full the COB
    FIFOs are.
    Until now the vc4 driver has been relying on the firmware to
    have set these to sensible values.

    With HVS channel 2 now being used for live video, change the
    panic mode for all channels to be explicitly set by the driver,
    and the same for all channels.

    Fixes: c54619b0bfb3 ("drm/vc4: Add support for the BCM2711 HVS5")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-2-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 190526c50f7fdb90fb0ba8f64ce523596912ff10
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Wed Dec 7 12:53:12 2022 +0100

    drm/vc4: hvs: Configure the HVS COB allocations

    [ Upstream commit 1284365e3e8cc33849f90cae20bca36fd7544e24 ]

    The HVS Composite Output Buffer (COB) is the memory used to
    generate the output pixel data.
    Until now the vc4 driver has been relying on the firmware to
    have set these to sensible values.

    In testing triple screen support it has been noted that only
    1 line was being assigned to HVS channel 2. Whilst that is fine
```

```
    for the transposer (TXP), and indeed needed as only some pixels
    have an alpha channel, it is insufficient to run a live display.

    Split the COB more evenly between the 3 HVS channels.

    Fixes: c54619b0bfb3 ("drm/vc4: Add support for the BCM2711 HVS5")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20221207-rpi-hvs-crtc-misc-v1-1-1f8e0770798b@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c40b34e3462aab12af3dba77d2e1602afc72e80
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 15:28:45 2023 +0400

    pinctrl: rockchip: Fix refcount leak in rockchip_pinctrl_parse_groups

    [ Upstream commit c818ae563bf99457f02e8170aabd6b174f629f65 ]

    of_find_node_by_phandle() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: d3e5116119bd ("pinctrl: add pinctrl driver for Rockchip SoCs")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Link: https://lore.kernel.org/r/20230102112845.3982407-1-linmq006@gmail.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ae053d1eb87875d56f95b6a123a69827225a70e
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 12:24:56 2023 +0400

    pinctrl: stm32: Fix refcount leak in stm32_pctrl_get_irq_domain

    [ Upstream commit dcef18c8ac40aa85bb339f64c1dd31dd458b06fb ]

    of_irq_find_parent() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: d86f4d71e42a ("pinctrl: stm32: check irq controller availability at probe")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Link: https://lore.kernel.org/r/20230102082503.3944927-1-linmq006@gmail.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8272f663257103b8628a601413593c28909f759
Author: Adam Skladowski <a39.skl@gmail.com>
Date:   Sat Dec 31 17:42:50 2022 +0100

    pinctrl: qcom: pinctrl-msm8976: Correct function names for wcss pins

    [ Upstream commit a7cc0e2685082a0d79baec02df184dfa83cbfac3 ]

    Adjust names of function for wcss pins, also fix third gpio in bt group.

    Fixes: bcd11493f0ab ("pinctrl: qcom: Add a pinctrl driver for MSM8976 and 8956")
    Signed-off-by: Adam Skladowski <a39.skl@gmail.com>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Link: https://lore.kernel.org/r/20221231164250.74550-1-a39.skl@gmail.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae5ca116a0c0ba9fc4123b1f1ec3c4f4d0d01b3f
```

```
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Fri Jan 6 10:30:11 2023 +0800

    drm/msm/hdmi: Add missing check for alloc_ordered_workqueue

    [ Upstream commit afe4cb96153a0d8003e4e4ebd91b5c543e10df84 ]

    Add check for the return value of alloc_ordered_workqueue as it may return
    NULL pointer and cause NULL pointer dereference in `hdmi_hdcp.c` and
    `hdmi_hpd.c`.

    Fixes: c6a57a50ad56 ("drm/msm/hdmi: add hdmi hdcp support (V3)")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/517211/
    Link: https://lore.kernel.org/r/20230106023011.3985-1-jiasheng@iscas.ac.cn
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5ee51b19855c5dd72aca57b8014f3b70d7798733
Author: Hui Tang <tanghui20@huawei.com>
Date:   Sat Nov 19 13:55:18 2022 +0800

    drm/msm/dpu: check for null return of devm_kzalloc() in dpu_writeback_init()

    [ Upstream commit 21e9a838f505178e109ccb3bf19d7808eb0326f4 ]

    Because of the possilble failure of devm_kzalloc(), dpu_wb_conn might
    be NULL and will cause null pointer dereference later.

    Therefore, it might be better to check it and directly return -ENOMEM.

    Fixes: 77b001acdcfe ("drm/msm/dpu: add the writeback connector layer")
    Signed-off-by: Hui Tang <tanghui20@huawei.com>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/512277/
    Link: https://lore.kernel.org/r/20221119055518.179937-1-tanghui20@huawei.com
    [DB: fixed typo in commit message]
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41dbd5fa93d7d7496ef08598fe7ff07adb04a210
Author: Armin Wolf <W_Armin@gmx.de>
Date:   Sat Dec 24 05:18:53 2022 +0100

    hwmon: (ftsteutates) Fix scaling of measurements

    [ Upstream commit ca8fd8c16a8b77dfcf7f6ce52d2c863220693a78 ]

    A user complained that the ftsteutates driver was displaying
    bogus values since its introduction. This happens because the
    sensor measurements need to be scaled in order to produce
    meaningful results:
    - the fan speed needs to be multiplied by 60 since its in RPS
    - the temperature is in degrees celsius and needs an offset of 64
    - the voltage is in 1/256 of 3.3V

    The offical datasheet says the voltage needs to be divided by 256,
    but this is likely an off-by-one-error, since even the BIOS
    devides by 255 (otherwise 3.3V could not be measured).

    The voltage channels additionally need a board-specific multiplier,
    however this can be done by the driver since its board-specific.

    The reason the missing scaling of measurements is the way Fujitsu
```

```
    used this driver when it was still out-of-tree. Back then, all
    scaling was done in userspace by libsensors, even the generic one.

    Tested on a Fujitsu DS3401-B1.

    Fixes: 08426eda58e0 ("hwmon: Add driver for FTS BMC chip "Teutates"")
    Signed-off-by: Armin Wolf <W_Armin@gmx.de>
    Link: https://lore.kernel.org/r/20221224041855.83981-2-W_Armin@gmx.de
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f483eda89a2e717daa7b15c8ca228a209f65617b
Author: Maíra Canal <mcanal@igalia.com>
Date:   Thu Dec 29 16:46:38 2022 -0300

    drm/vc4: drop all currently held locks if deadlock happens

    [ Upstream commit 479d4f0be4237ba33bc9432787aeb62c90e30f95 ]

    If vc4_hdmi_reset_link() returns -EDEADLK, it means that a deadlock
    happened in the locking context. This situation should be addressed by
    dropping all currently held locks and block until the contended lock
    becomes available. Currently, vc4 is not dealing with the deadlock
    properly, producing the following output when PROVE_LOCKING is enabled:

    [  825.612809] ------------[ cut here ]------------
    [  825.612852] WARNING: CPU: 1 PID: 116 at drivers/gpu/drm/drm_modeset_lock.c:276
drm_modeset_drop_locks+0x60/0x68 [drm]
    [  825.613458] Modules linked in: 8021q mrp garp stp llc
    raspberrypi_cpufreq brcmfmac brcmutil crct10dif_ce hci_uart cfg80211
    btqca btbcm bluetooth vc4 raspberrypi_hwmon snd_soc_hdmi_codec cec
    clk_raspberrypi ecdh_generic drm_display_helper ecc rfkill
    drm_dma_helper drm_kms_helper pwm_bcm2835 bcm2835_thermal bcm2835_rng
    rng_core i2c_bcm2835 drm fuse ip_tables x_tables ipv6
    [  825.613735] CPU: 1 PID: 116 Comm: kworker/1:2 Tainted: G        W 6.1.0-rc6-01399-
g941aae326315 #3
    [  825.613759] Hardware name: Raspberry Pi 3 Model B Rev 1.2 (DT)
    [  825.613777] Workqueue: events output_poll_execute [drm_kms_helper]
    [  825.614038] pstate: 60000005 (nZCv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
    [  825.614063] pc : drm_modeset_drop_locks+0x60/0x68 [drm]
    [  825.614603] lr : drm_helper_probe_detect+0x120/0x1b4 [drm_kms_helper]
    [  825.614829] sp : ffff800008313bf0
    [  825.614844] x29: ffff800008313bf0 x28: ffffcd7778b8b000 x27: 0000000000000000
    [  825.614883] x26: 0000000000000001 x25: 0000000000000001 x24: ffff677cc35c2758
    [  825.614920] x23: ffffcd7707d01430 x22: ffffcd7707c3edc7 x21: 0000000000000001
    [  825.614958] x20: 0000000000000000 x19: ffff800008313c10 x18: 000000000000b6d3
    [  825.614995] x17: ffffcd777835e214 x16: ffffcd7777cef870 x15: fffff81000000000
    [  825.615033] x14: 0000000000000000 x13: 0000000000000099 x12: 0000000000000002
    [  825.615070] x11: 72917988020af800 x10: 72917988020af800 x9 : 72917988020af800
    [  825.615108] x8 : ffff677cc665e0a8 x7 : d00a8c180000110c x6 : ffffcd77774c0054
    [  825.615145] x5 : 0000000000000000 x4 : 0000000000000001 x3 : 0000000000000000
    [  825.615181] x2 : ffff677cc55e1880 x1 : ffffcd7777cef8ec x0 : ffff800008313c10
    [  825.615219] Call trace:
    [  825.615232]  drm_modeset_drop_locks+0x60/0x68 [drm]
    [  825.615773]  drm_helper_probe_detect+0x120/0x1b4 [drm_kms_helper]
    [  825.616003]  output_poll_execute+0xe4/0x224 [drm_kms_helper]
    [  825.616233]  process_one_work+0x2b4/0x618
    [  825.616264]  worker_thread+0x24c/0x464
    [  825.616288]  kthread+0xec/0x110
    [  825.616310]  ret_from_fork+0x10/0x20
    [  825.616335] irq event stamp: 7634
    [  825.616349] hardirqs last  enabled at (7633): [<ffffcd777831ee90>]
_raw_spin_unlock_irq+0x3c/0x78
    [  825.616384] hardirqs last disabled at (7634): [<ffffcd7778315a78>]
__schedule+0x134/0x9f0
```

```
    [  825.616411] softirqs last  enabled at (7630): [<ffffcd7707aacea0>]
local_bh_enable+0x4/0x30 [ipv6]
    [  825.617019] softirqs last disabled at (7618): [<ffffcd7707aace70>]
local_bh_disable+0x4/0x30 [ipv6]
    [  825.617586] ---[ end trace 0000000000000000 ]---

    Therefore, deal with the deadlock as suggested by [1], using the
    function drm_modeset_backoff().

    [1] https://docs.kernel.org/gpu/drm-kms.html?highlight=kms#kms-locking

    Fixes: 6bed2ea3cb38 ("drm/vc4: hdmi: Reset link on hotplug")
    Reported-by: Stefan Wahren <stefan.wahren@i2se.com>
    Signed-off-by: Maíra Canal <mcanal@igalia.com>
    Tested-by: Stefan Wahren <stefan.wahren@i2se.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221229194638.178712-1-
mcanal@igalia.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 86db720693decf89f00e9457dcc1824ec70102dd
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Fri Dec 16 20:30:05 2022 +0100

    drm/ast: Init iosys_map pointer as I/O memory for damage handling

    [ Upstream commit b1def7fadfa544bd2467581ce40b659583eb7e79 ]

    Ast hardware scans out the primary plane from video memory, which
    is in I/O-memory space. Hence init the damage handler's iosys_map
    pointer as I/O memory.

    Not all platforms support accessing I/O memory as system memory,
    although it's usually not a problem in ast's x86-based systems.

    The error report is at [1].

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Reviewed-by: Jocelyn Falempe <jfalempe@redhat.com>
    Fixes: f2fa5a99ca81 ("drm/ast: Convert ast to SHMEM")
    Cc: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: Jocelyn Falempe <jfalempe@redhat.com>
    Cc: Dave Airlie <airlied@redhat.com>
    Cc: dri-devel@lists.freedesktop.org
    Link: https://lore.kernel.org/lkml/202212170111.eInM0unS-lkp@intel.com/T/#u # 1
    Link: https://patchwork.freedesktop.org/patch/msgid/20221216193005.30280-1-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 38fad2dcd97aabbaccd76e3e1fcb9f05105f706e
Author: Liang He <windhl@126.com>
Date:   Wed Jul 20 23:22:27 2022 +0800

    gpu: ipu-v3: common: Add of_node_put() for reference returned by
of_graph_get_port_by_id()

    [ Upstream commit 9afdf98cfdfa2ba8ec068cf08c5fcdc1ed8daf3f ]

    In ipu_add_client_devices(), we need to call of_node_put() for
    reference returned by of_graph_get_port_by_id() in fail path.

    Fixes: 17e052175039 ("gpu: ipu-v3: Do not bail out on missing optional port nodes")
    Signed-off-by: Liang He <windhl@126.com>
    Reviewed-by: Philipp Zabel <p.zabel@pengutronix.de>
```

```
    Link: https://lore.kernel.org/r/20220720152227.1288413-1-windhl@126.com
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220720152227.1288413-1-
windhl@126.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7427f147a315e619f6449887bcb101041e300557
Author: Randolph Sapp <rs@ti.com>
Date:   Thu Dec 1 18:18:03 2022 -0600

    drm: tidss: Fix pixel format definition

    [ Upstream commit 2df0433b18f2735a49d2c3a968b40fa2881137c0 ]

    There was a long-standing bug from a typo that created 2 ARGB1555 and
    ABGR1555 pixel format entries. Weston 10 has a sanity check that alerted
    me to this issue.

    According to the Supported Pixel Data formats table we have the later
    entries should have been for Alpha-X instead.

    Signed-off-by: Randolph Sapp <rs@ti.com>
    Fixes: 32a1795f57eecc ("drm/tidss: New driver for TI Keystone platform Display
SubSystem")
    Reviewed-by: Aradhya Bhatia <a-bhatia1@ti.com>
    Acked-by: Andrew Davis <afd@ti.com>
    Signed-off-by: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221202001803.1765805-1-rs@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d2d17ba47686ac14dd7d5b98a3ce5aa46e9bc081
Author: Pin-yen Lin <treapking@chromium.org>
Date:   Wed Nov 9 17:52:27 2022 +0800

    drm/bridge: it6505: Guard bridge power in IRQ handler

    [ Upstream commit 5eb9a4314053bda7642643f70f49a2b415920812 ]

    Add a pair of pm_runtime_get_if_in_use and pm_runtime_put_sync in the
    interrupt handler to make sure the bridge won't be powered off during
    the interrupt handlings. Also remove the irq_lock mutex because it's not
    guarding anything now.

    Fixes: ab28896f1a83 ("drm/bridge: it6505: Improve synchronization between extcon
subsystem")
    Signed-off-by: Pin-yen Lin <treapking@chromium.org>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221109095227.3320919-1-
treapking@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 72044b9881736f2030bd9bfa86af92958b929ee8
Author: Dave Stevenson <dave.stevenson@raspberrypi.com>
Date:   Thu Dec 1 09:42:52 2022 +0100

    drm/vc4: dpi: Fix format mapping for RGB565

    [ Upstream commit 0870d86eac8a9abd89a0be1b719d5dc5bac936f0 ]

    The mapping is incorrect for RGB565_1X16 as it should be
    DPI_FORMAT_18BIT_666_RGB_1 instead of DPI_FORMAT_18BIT_666_RGB_3.

    Fixes: 08302c35b59d ("drm/vc4: Add DPI driver")
    Signed-off-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
```

```
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Link: https://lore.kernel.org/r/20221013-rpi-dpi-improvements-v3-7-
eb76e26a772d@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a31ef669bc020b2af690eb7e564b83373c594be
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Nov 28 09:19:37 2022 +0100

    drm/modes: Use strscpy() to copy command-line mode name

    [ Upstream commit 0f9aa074c92dd9274b811c1c3fa93736814a4b0d ]

    The mode name in struct drm_cmdline_mode can hold 32 characters at most,
    which can easily get overrun. Switch to strscpy() to prevent such a
    thing.

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527354 ("Security best practices violations")
    Fixes: a7ab155397dd ("drm/modes: Switch to named mode descriptors")
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://lore.kernel.org/r/20221128081938.742410-2-maxime@cerno.tech
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57031c474c3a920ea73afeb5dc352e537f5793ee
Author: Yuan Can <yuancan@huawei.com>
Date:   Tue Nov 1 06:51:56 2022 +0000

    drm/vkms: Fix null-ptr-deref in vkms_release()

    [ Upstream commit 2fe2a8f40c21161ffe7653cc234e7934db5b7cc5 ]

    A null-ptr-deref is triggered when it tries to destroy the workqueue in
    vkms->output.composer_workq in vkms_release().

     KASAN: null-ptr-deref in range [0x0000000000000118-0x000000000000011f]
     CPU: 5 PID: 17193 Comm: modprobe Not tainted 6.0.0-11331-gd465bff130bf #24
     RIP: 0010:destroy_workqueue+0x2f/0x710
     ...
     Call Trace:
      <TASK>
      ? vkms_config_debugfs_init+0x50/0x50 [vkms]
      __devm_drm_dev_alloc+0x15a/0x1c0 [drm]
      vkms_init+0x245/0x1000 [vkms]
      do_one_initcall+0xd0/0x4f0
      do_init_module+0x1a4/0x680
      load_module+0x6249/0x7110
      __do_sys_finit_module+0x140/0x200
      do_syscall_64+0x35/0x80
      entry_SYSCALL_64_after_hwframe+0x46/0xb0

    The reason is that an OOM happened which triggers the destroy of the
    workqueue, however, the workqueue is alloced in the later process,
    thus a null-ptr-deref happened. A simple call graph is shown as below:

     vkms_init()
      vkms_create()
        devm_drm_dev_alloc()
          __devm_drm_dev_alloc()
            devm_drm_dev_init()
              devm_add_action_or_reset()
                devm_add_action() # an error happened
                devm_drm_dev_init_release()
```

```
                  drm_dev_put()
                    kref_put()
                      drm_dev_release()
                        vkms_release()
                          destroy_workqueue() # null-ptr-deref happened
          vkms_modeset_init()
            vkms_output_init()
              vkms_crtc_init() # where the workqueue get allocated
```

    Fix this by checking if composer_workq is NULL before passing it to
    the destroy_workqueue() in vkms_release().

    Fixes: 6c234fe37c57 ("drm/vkms: Implement CRC debugfs API")
    Signed-off-by: Yuan Can <yuancan@huawei.com>
    Reviewed-by: Melissa Wen <mwen@igalia.com>
    Signed-off-by: Melissa Wen <melissa.srw@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221101065156.41584-3-
yuancan@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07ab77154d6fd2d67e465ab5ce30083709950f02
Author: Yuan Can <yuancan@huawei.com>
Date: Tue Nov 1 06:51:55 2022 +0000

    drm/vkms: Fix memory leak in vkms_init()

    [Upstream commit 0d0b368b9d104b437e1f4850ae94bdb9a3601e89]

    A memory leak was reported after the vkms module install failed.

    unreferenced object 0xffff88810bc28520 (size 16):
      comm "modprobe", pid 9662, jiffies 4298009455 (age 42.590s)
      hex dump (first 16 bytes):
        01 01 00 64 81 88 ff ff 00 00 dc 0a 81 88 ff ff ...d............
      backtrace:
        [<00000000e7561ff8>] kmalloc_trace+0x27/0x60
        [<000000000b1954a0>] 0xffffffffc45200a9
        [<00000000abbf1da0>] do_one_initcall+0xd0/0x4f0
        [<000000001505ee87>] do_init_module+0x1a4/0x680
        [<00000000958079ad>] load_module+0x6249/0x7110
        [<00000000117e4696>] __do_sys_finit_module+0x140/0x200
        [<00000000f74b12d2>] do_syscall_64+0x35/0x80
        [<000000008fc6fcde>] entry_SYSCALL_64_after_hwframe+0x46/0xb0

    The reason is that the vkms_init() returns without checking the return
    value of vkms_create(), and if the vkms_create() failed, the config
    allocated at the beginning of vkms_init() is leaked.

     vkms_init()
       config = kmalloc(...) # config allocated
       ...
       return vkms_create() # vkms_create failed and config is leaked

    Fix this problem by checking return value of vkms_create() and free the
    config if error happened.

    Fixes: 2df7af93fdad ("drm/vkms: Add vkms_config type")
    Signed-off-by: Yuan Can <yuancan@huawei.com>
    Reviewed-by: Melissa Wen <mwen@igalia.com>
    Signed-off-by: Melissa Wen <melissa.srw@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221101065156.41584-2-
yuancan@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d658773c24652751e6987bd34cb710bc026466e

```
Author: Yuan Can <yuancan@huawei.com>
Date:   Tue Nov 8 09:12:26 2022 +0000

    drm/bridge: megachips: Fix error handling in i2c_register_driver()

    [ Upstream commit 4ecff954c370b82bce45bdca2846c5c5563e8a8a ]

    A problem about insmod megachips-stdpxxxx-ge-b850v3-fw.ko failed is
    triggered with the following log given:

    [ 4497.981497] Error: Driver 'stdp4028-ge-b850v3-fw' is already registered, aborting...
    insmod: ERROR: could not insert module megachips-stdpxxxx-ge-b850v3-fw.ko: Device or
resource busy

    The reason is that stdpxxxx_ge_b850v3_init() returns i2c_add_driver()
    directly without checking its return value, if i2c_add_driver() failed,
    it returns without calling i2c_del_driver() on the previous i2c driver,
    resulting the megachips-stdpxxxx-ge-b850v3-fw can never be installed
    later.
    A simple call graph is shown as below:

     stdpxxxx_ge_b850v3_init()
       i2c_add_driver(&stdp4028_ge_b850v3_fw_driver)
       i2c_add_driver(&stdp2690_ge_b850v3_fw_driver)
         i2c_register_driver()
           driver_register()
             bus_add_driver()
               priv = kzalloc(...) # OOM happened
       # return without delete stdp4028_ge_b850v3_fw_driver

    Fix by calling i2c_del_driver() on stdp4028_ge_b850v3_fw_driver when
    i2c_add_driver() returns error.

    Fixes: fcfa0ddc18ed ("drm/bridge: Drivers for megachips-stdpxxxx-ge-b850v3-fw (LVDS-
DP++)")
    Signed-off-by: Yuan Can <yuancan@huawei.com>
    Reviewed-by: Andrzej Hajda <andrzej.hajda@intel.com>
    Tested-by: Ian Ray <ian.ray@ge.com>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221108091226.114524-1-
yuancan@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ebc0aad294551e9e1c7be1f163fabd939e3ee334
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Nov 21 16:59:55 2022 +0100

    drm: mxsfb: DRM_MXSFB should depend on ARCH_MXS || ARCH_MXC

    [ Upstream commit 7783cc67862f9166c901bfa0f80b717aa8d354dd ]

    Freescale/NXP i.MX LCDIF and eLCDIF LCD controllers are only present on
    Freescale/NXP i.MX SoCs.  Hence add a dependency on ARCH_MXS ||
    ARCH_MXC, to prevent asking the user about this driver when configuring
    a kernel without Freescale/NXP i.MX support.

    Fixes: 45d59d704080cc0c ("drm: Add new driver for MXSFB controller")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Marek Vasut <marex@denx.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/98e74779ca2bc575d91afff03369e86b080c01ac.16690
46358.git.geert+renesas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 61a2d66dd413203cfe14e4d4f73171c06c4e69e6
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Nov 21 17:00:48 2022 +0100

    drm: mxsfb: DRM_IMX_LCDIF should depend on ARCH_MXC

    [ Upstream commit 10ef5f2992006720318b9886961820155b3750fd ]

    The Freescale/NXP i.MX LCDIFv3 LCD controller is only present on
    Freescale/NXP i.MX SoCs.  Hence add a dependency on ARCH_MXC, to prevent
    asking the user about this driver when configuring a kernel without
    Freescale/NXP i.MX support.

    Fixes: 9db35bb349a0ef32 ("drm: lcdif: Add support for i.MX8MP LCDIF variant")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Marek Vasut <marex@denx.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/6103c1aa65a7888c12d351ae63f29850f29f42b9.16690
46403.git.geert+renesas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3649b7f581e511fbfd6b0e8f06745132cfbc3c66
Author: Frieder Schrempf <frieder.schrempf@kontron.de>
Date:   Tue Nov 22 09:12:18 2022 +0100

    drm/bridge: ti-sn65dsi83: Fix delay after reset deassert to match spec

    [ Upstream commit 4b03d5e0d3e86ee492d54254927d020dc0fe8acf ]

    The datasheet specifies a delay of 10 milliseconds, but the current
    driver only waits for 1 ms. Fix this to make sure the initialization
    sequence meets the spec.

    Fixes: ceb515ba29ba ("drm/bridge: ti-sn65dsi83: Add TI SN65DSI83 and SN65DSI84 driver")
    Signed-off-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Reviewed-by: Alexander Stein <alexander.stein@ew.tq-group.com>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221122081219.20143-1-
frieder@fris.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4af49c71ffdcc27032c620949c9e5077ddfc9953
Author: Geert Uytterhoeven <geert@linux-m68k.org>
Date:   Wed Nov 23 17:43:10 2022 +0100

    drm/fourcc: Add missing big-endian XRGB1555 and RGB565 formats

    [ Upstream commit 6fb6c979ca628583d4d0c59a0f8ff977e581ecc0 ]

    As of commit eae06120f1974e1a ("drm: refuse ADDFB2 ioctl for broken
    bigendian drivers"), drivers must set the
    quirk_addfb_prefer_host_byte_order quirk to make the drm_mode_addfb()
    compat code work correctly on big-endian machines.

    While that works fine for big-endian XRGB8888 and ARGB8888, which are
    mapped to the existing little-endian BGRX8888 and BGRA8888 formats, it
    does not work for big-endian XRGB1555 and RGB565, as the latter are not
    listed in the format database.

    Fix this by adding the missing formats.  Limit this to big-endian
    platforms, as there is currently no need to support these formats on
    little-endian platforms.

    Fixes: 6960e6da9cec3f66 ("drm: fix drm_mode_addfb() on big endian machines.")
```

```
    Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Signed-off-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link:
https://patchwork.freedesktop.org/patch/msgid/3ee1f8144feb96c28742b22384189f1f83bcfc1a.16692
21671.git.geert@linux-m68k.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 961620ad67611a7320a49f4b6f3c5e2906833a03
Author: Shang XiaoJing <shangxiaojing@huawei.com>
Date:   Fri Nov 18 10:16:51 2022 +0800

    drm: Fix potential null-ptr-deref due to drmm_mode_config_init()

    [ Upstream commit 834c23e4f798dcdc8af251b3c428ceef94741991 ]

    drmm_mode_config_init() will call drm_mode_create_standard_properties()
    and won't check the ret value. When drm_mode_create_standard_properties()
    failed due to alloc, property will be a NULL pointer and may causes the
    null-ptr-deref. Fix the null-ptr-deref by adding the ret value check.

    Found null-ptr-deref while testing insert module bochs:
    general protection fault, probably for non-canonical address
        0xdffffc000000000c: 0000 [#1] SMP KASAN PTI
    KASAN: null-ptr-deref in range [0x0000000000000060-0x0000000000000067]
    CPU: 3 PID: 249 Comm: modprobe Not tainted 6.1.0-rc1+ #364
    Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS
    rel-1.15.0-0-g2dd4b9b3f840-prebuilt.qemu.org 04/01/2014
    RIP: 0010:drm_object_attach_property+0x73/0x3c0 [drm]
    Call Trace:
     <TASK>
     __drm_connector_init+0xb6c/0x1100 [drm]
     bochs_pci_probe.cold.11+0x4cb/0x7fe [bochs]
     pci_device_probe+0x17d/0x340
     really_probe+0x1db/0x5d0
     __driver_probe_device+0x1e7/0x250
     driver_probe_device+0x4a/0x120
     __driver_attach+0xcd/0x2c0
     bus_for_each_dev+0x11a/0x1b0
     bus_add_driver+0x3d7/0x500
     driver_register+0x18e/0x320
     do_one_initcall+0xc4/0x3e0
     do_init_module+0x1b4/0x630
     load_module+0x5dca/0x7230
     __do_sys_finit_module+0x100/0x170
     do_syscall_64+0x3f/0x90
     entry_SYSCALL_64_after_hwframe+0x63/0xcd
    RIP: 0033:0x7ff65af9f839

    Fixes: 6b4959f43a04 ("drm/atomic: atomic plane properties")
    Signed-off-by: Shang XiaoJing <shangxiaojing@huawei.com>
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20221118021651.2460-1-
shangxiaojing@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 52901f5f0b3ea6abe222bbde10712d9dedc739fb
Author: Jiri Pirko <jiri@nvidia.com>
Date:   Mon Feb 20 14:23:36 2023 +0100

    sefltests: netdevsim: wait for devlink instance after netns removal

    [ Upstream commit f922c7b1c1c45740d329bf248936fdb78c0cff6e ]

    When devlink instance is put into network namespace and that network
    namespace gets deleted, devlink instance is moved back into init_ns.
```

```
    This is done as a part of cleanup_net() routine. Since cleanup_net()
    is called asynchronously from workqueue, there is no guarantee that
    the devlink instance move is done after "ip netns del" returns.

    So fix this race by making sure that the devlink instance is present
    before any other operation.

    Reported-by: Amir Tzin <amirtz@nvidia.com>
    Fixes: b74c37fd35a2 ("selftests: netdevsim: add tests for devlink reload with
resources")
    Signed-off-by: Jiri Pirko <jiri@nvidia.com>
    Reviewed-by: Pavan Chebbi <pavan.chebbi@broadcom.com>
    Link: https://lore.kernel.org/r/20230220132336.198597-1-jiri@resnulli.us
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6b0e84ade750ccfad862aee0cf7aaf52e3300880
Author: Roxana Nicolescu <roxana.nicolescu@canonical.com>
Date:   Mon Feb 20 12:04:00 2023 +0100

    selftest: fib_tests: Always cleanup before exit

    [ Upstream commit b60417a9f2b890a8094477b2204d4f73c535725e ]

    Usage of `set -e` before executing a command causes immediate exit
    on failure, without cleanup up the resources allocated at setup.
    This can affect the next tests that use the same resources,
    leading to a chain of failures.

    A simple fix is to always call cleanup function when the script exists.
    This approach is already used by other existing tests.

    Fixes: 1056691b2680 ("selftests: fib_tests: Make test results more verbose")
    Signed-off-by: Roxana Nicolescu <roxana.nicolescu@canonical.com>
    Link: https://lore.kernel.org/r/20230220110400.26737-2-roxana.nicolescu@canonical.com
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd753c108db26ab83e53b8867a0b980cfab193ce
Author: Leon Romanovsky <leon@kernel.org>
Date:   Sun Feb 19 11:09:10 2023 +0200

    net/mlx5e: Align IPsec ASO result memory to be as required by hardware

    [ Upstream commit f2b6cfda76d2119871e10fa01ecdc7178401ef22 ]

    Hardware requires an alignment to 64 bytes to return ASO data. Missing
    this alignment caused to unpredictable results while ASO events were
    generated.

    Fixes: 8518d05b8f9a ("net/mlx5e: Create Advanced Steering Operation object for IPsec")
    Reported-by: Emeel Hakim <ehakim@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Link:
https://lore.kernel.org/r/de0302c572b90c9224a72868d4e0d657b6313c4b.1676797613.git.leon@kerne
l.org
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b0cb2744dc73b9bb2dec975ff5f9b178bd8c8f9e
Author: Kees Cook <keescook@chromium.org>
Date:   Sat Feb 18 10:38:50 2023 -0800

    net/mlx4_en: Introduce flexible array to silence overflow warning
```

```
    [ Upstream commit f8f185e39b4de91bc5235e5be0d829bea69d9b06 ]

    The call "skb_copy_from_linear_data(skb, inl + 1, spc)" triggers a FORTIFY
    memcpy() warning on ppc64 platform:

    In function ‘fortify_memcpy_chk’,
        inlined from ‘skb_copy_from_linear_data’ at ./include/linux/skbuff.h:4029:2,
        inlined from ‘build_inline_wqe’ at drivers/net/ethernet/mellanox/mlx4/en_tx.c:722:4,
        inlined from ‘mlx4_en_xmit’ at drivers/net/ethernet/mellanox/mlx4/en_tx.c:1066:3:
    ./include/linux/fortify-string.h:513:25: error: call to ‘__write_overflow_field’
 declared with
    attribute warning: detected write beyond size of field (1st parameter); maybe use
 struct_group()?
    [-Werror=attribute-warning]
      513 |                         __write_overflow_field(p_size_field, size);
          |                         ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    Same behaviour on x86 you can get if you use "__always_inline" instead of
    "inline" for skb_copy_from_linear_data() in skbuff.h

    The call here copies data into inlined tx destricptor, which has 104
    bytes (MAX_INLINE) space for data payload. In this case "spc" is known
    in compile-time but the destination is used with hidden knowledge
    (real structure of destination is different from that the compiler
    can see). That cause the fortify warning because compiler can check
    bounds, but the real bounds are different.  "spc" can't be bigger than
    64 bytes (MLX4_INLINE_ALIGN), so the data can always fit into inlined
    tx descriptor. The fact that "inl" points into inlined tx descriptor is
    determined earlier in mlx4_en_xmit().

    Avoid confusing the compiler with "inl + 1" constructions to get to past
    the inl header by introducing a flexible array "data" to the struct so
    that the compiler can see that we are not dealing with an array of inl
    structs, but rather, arbitrary data following the structure. There are
    no changes to the structure layout reported by pahole, and the resulting
    machine code is actually smaller.

    Reported-by: Josef Oskera <joskera@redhat.com>
    Link: https://lore.kernel.org/lkml/20230217094541.2362873-1-joskera@redhat.com
    Fixes: f68f2ff91512 ("fortify: Detect struct member overflows in memcpy() at compile-
 time")
    Cc: Yishai Hadas <yishaih@nvidia.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Tariq Toukan <tariqt@nvidia.com>
    Link: https://lore.kernel.org/r/20230218183842.never.954-kees@kernel.org
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 3304fb3994e3ab6a272c3f249ba8a6b1e7b062ee
 Author: Horatiu Vultur <horatiu.vultur@microchip.com>
 Date:   Fri Feb 17 22:09:17 2023 +0100

    net: lan966x: Fix possible deadlock inside PTP

    [ Upstream commit 3a70e0d4c9d74cb00f7c0ec022f5599f9f7ba07d ]

    When doing timestamping in lan966x and having PROVE_LOCKING
    enabled the following warning is shown.

    ========================================================
    WARNING: possible irq lock inversion dependency detected
    6.2.0-rc7-01749-gc54e1f7f7e36 #2786 Tainted: G                N
    --------------------------------------------------------
    swapper/0/0 just changed the state of lock:
    c2609f50 (_xmit_ETHER#2){+.-.}-{2:2}, at: sch_direct_xmit+0x16c/0x2e8
```

```
     but this lock took another, SOFTIRQ-unsafe lock in the past:
      (&lan966x->ptp_ts_id_lock){+.+.}-{2:2}

     and interrupts could create inverse lock ordering between them.

     other info that might help us debug this:
      Possible interrupt unsafe locking scenario:

            CPU0                    CPU1
            ----                    ----
       lock(&lan966x->ptp_ts_id_lock);
                                    local_irq_disable();
                                    lock(_xmit_ETHER#2);
                                    lock(&lan966x->ptp_ts_id_lock);
       <Interrupt>
         lock(_xmit_ETHER#2);

      *** DEADLOCK ***

     5 locks held by swapper/0/0:
      #0: c1001e18 ((&ndev->rs_timer)){+.-.}-{0:0}, at: call_timer_fn+0x0/0x33c
      #1: c105e7c4 (rcu_read_lock){....}-{1:2}, at: ndisc_send_skb+0x134/0x81c
      #2: c105e7d8 (rcu_read_lock_bh){....}-{1:2}, at: ip6_finish_output2+0x17c/0xc64
      #3: c105e7d8 (rcu_read_lock_bh){....}-{1:2}, at: __dev_queue_xmit+0x4c/0x1224
      #4: c3056174 (dev->qdisc_tx_busylock ?: &qdisc_tx_busylock){+...}-{2:2}, at:
  __dev_queue_xmit+0x354/0x1224

     the shortest dependencies between 2nd lock and 1st lock:
      -> (&lan966x->ptp_ts_id_lock){+.+.}-{2:2} {
         HARDIRQ-ON-W at:
                             lock_acquire.part.0+0xb0/0x248
                             _raw_spin_lock+0x38/0x48
                             lan966x_ptp_irq_handler+0x164/0x2a8
                             irq_thread_fn+0x1c/0x78
                             irq_thread+0x130/0x278
                             kthread+0xec/0x110
                             ret_from_fork+0x14/0x28
         SOFTIRQ-ON-W at:
                             lock_acquire.part.0+0xb0/0x248
                             _raw_spin_lock+0x38/0x48
                             lan966x_ptp_irq_handler+0x164/0x2a8
                             irq_thread_fn+0x1c/0x78
                             irq_thread+0x130/0x278
                             kthread+0xec/0x110
                             ret_from_fork+0x14/0x28
         INITIAL USE at:
                            lock_acquire.part.0+0xb0/0x248
                            _raw_spin_lock_irqsave+0x4c/0x68
                            lan966x_ptp_txtstamp_request+0x128/0x1cc
                            lan966x_port_xmit+0x224/0x43c
                            dev_hard_start_xmit+0xa8/0x2f0
                            sch_direct_xmit+0x108/0x2e8
                            __dev_queue_xmit+0x41c/0x1224
                            packet_sendmsg+0xdb4/0x134c
                            __sys_sendto+0xd0/0x154
                            sys_send+0x18/0x20
                            ret_fast_syscall+0x0/0x1c
       }
       ... key      at: [<c174ba0c>] __key.2+0x0/0x8
       ... acquired at:
        _raw_spin_lock_irqsave+0x4c/0x68
        lan966x_ptp_txtstamp_request+0x128/0x1cc
        lan966x_port_xmit+0x224/0x43c
        dev_hard_start_xmit+0xa8/0x2f0
        sch_direct_xmit+0x108/0x2e8
```

```
          __dev_queue_xmit+0x41c/0x1224
          packet_sendmsg+0xdb4/0x134c
          __sys_sendto+0xd0/0x154
          sys_send+0x18/0x20
          ret_fast_syscall+0x0/0x1c

       -> (_xmit_ETHER#2){+.-.}-{2:2} {
          HARDIRQ-ON-W at:
                            lock_acquire.part.0+0xb0/0x248
                            _raw_spin_lock+0x38/0x48
                            netif_freeze_queues+0x38/0x68
                            dev_deactivate_many+0xac/0x388
                            dev_deactivate+0x38/0x6c
                            linkwatch_do_dev+0x70/0x8c
                            __linkwatch_run_queue+0xd4/0x1e8
                            linkwatch_event+0x24/0x34
                            process_one_work+0x284/0x744
                            worker_thread+0x28/0x4bc
                            kthread+0xec/0x110
                            ret_from_fork+0x14/0x28
          IN-SOFTIRQ-W at:
                            lock_acquire.part.0+0xb0/0x248
                            _raw_spin_lock+0x38/0x48
                            sch_direct_xmit+0x16c/0x2e8
                            __dev_queue_xmit+0x41c/0x1224
                            ip6_finish_output2+0x5f4/0xc64
                            ndisc_send_skb+0x4cc/0x81c
                            addrconf_rs_timer+0xb0/0x2f8
                            call_timer_fn+0xb4/0x33c
                            expire_timers+0xb4/0x10c
                            run_timer_softirq+0xf8/0x2a8
                            __do_softirq+0xd4/0x5fc
                            __irq_exit_rcu+0x138/0x17c
                            irq_exit+0x8/0x28
                            __irq_svc+0x90/0xbc
                            arch_cpu_idle+0x30/0x3c
                            default_idle_call+0x44/0xac
                            do_idle+0xc8/0x138
                            cpu_startup_entry+0x18/0x1c
                            rest_init+0xcc/0x168
                            arch_post_acpi_subsys_init+0x0/0x8
          INITIAL USE at:
                           lock_acquire.part.0+0xb0/0x248
                           _raw_spin_lock+0x38/0x48
                           netif_freeze_queues+0x38/0x68
                           dev_deactivate_many+0xac/0x388
                           dev_deactivate+0x38/0x6c
                           linkwatch_do_dev+0x70/0x8c
                           __linkwatch_run_queue+0xd4/0x1e8
                           linkwatch_event+0x24/0x34
                           process_one_work+0x284/0x744
                           worker_thread+0x28/0x4bc
                           kthread+0xec/0x110
                           ret_from_fork+0x14/0x28
        }
        ... key      at: [<c175974c>] netdev_xmit_lock_key+0x8/0x1c8
        ... acquired at:
          __lock_acquire+0x978/0x2978
          lock_acquire.part.0+0xb0/0x248
          _raw_spin_lock+0x38/0x48
          sch_direct_xmit+0x16c/0x2e8
          __dev_queue_xmit+0x41c/0x1224
          ip6_finish_output2+0x5f4/0xc64
          ndisc_send_skb+0x4cc/0x81c
          addrconf_rs_timer+0xb0/0x2f8
```

```
        call_timer_fn+0xb4/0x33c
        expire_timers+0xb4/0x10c
        run_timer_softirq+0xf8/0x2a8
        __do_softirq+0xd4/0x5fc
        __irq_exit_rcu+0x138/0x17c
        irq_exit+0x8/0x28
        __irq_svc+0x90/0xbc
        arch_cpu_idle+0x30/0x3c
        default_idle_call+0x44/0xac
        do_idle+0xc8/0x138
        cpu_startup_entry+0x18/0x1c
        rest_init+0xcc/0x168
        arch_post_acpi_subsys_init+0x0/0x8

    stack backtrace:
    CPU: 0 PID: 0 Comm: swapper/0 Tainted: G                N 6.2.0-rc7-01749-gc54e1f7f7e36
#2786
    Hardware name: Generic DT based system
     unwind_backtrace from show_stack+0x10/0x14
     show_stack from dump_stack_lvl+0x58/0x70
     dump_stack_lvl from mark_lock.part.0+0x59c/0x93c
     mark_lock.part.0 from __lock_acquire+0x978/0x2978
     __lock_acquire from lock_acquire.part.0+0xb0/0x248
     lock_acquire.part.0 from _raw_spin_lock+0x38/0x48
     _raw_spin_lock from sch_direct_xmit+0x16c/0x2e8
     sch_direct_xmit from __dev_queue_xmit+0x41c/0x1224
     __dev_queue_xmit from ip6_finish_output2+0x5f4/0xc64
     ip6_finish_output2 from ndisc_send_skb+0x4cc/0x81c
     ndisc_send_skb from addrconf_rs_timer+0xb0/0x2f8
     addrconf_rs_timer from call_timer_fn+0xb4/0x33c
     call_timer_fn from expire_timers+0xb4/0x10c
     expire_timers from run_timer_softirq+0xf8/0x2a8
     run_timer_softirq from __do_softirq+0xd4/0x5fc
     __do_softirq from __irq_exit_rcu+0x138/0x17c
     __irq_exit_rcu from irq_exit+0x8/0x28
     irq_exit from __irq_svc+0x90/0xbc
    Exception stack(0xc1001f20 to 0xc1001f68)
    1f20: ffffffff ffffffff 00000001 c011f840 c100e000 c100e000 c1009314 c1009370
    1f40: c10f0c1a c0d5e564 c0f5da8c 00000000 00000000 c1001f70 c010f0bc c010f0c0
    1f60: 600f0013 ffffffff
     __irq_svc from arch_cpu_idle+0x30/0x3c
     arch_cpu_idle from default_idle_call+0x44/0xac
     default_idle_call from do_idle+0xc8/0x138
     do_idle from cpu_startup_entry+0x18/0x1c
     cpu_startup_entry from rest_init+0xcc/0x168
     rest_init from arch_post_acpi_subsys_init+0x0/0x8

    Fix this by using spin_lock_irqsave/spin_lock_irqrestore also
    inside lan966x_ptp_irq_handler.

    Fixes: e85a96e48e33 ("net: lan966x: Add support for ptp interrupts")
    Signed-off-by: Horatiu Vultur <horatiu.vultur@microchip.com>
    Link: https://lore.kernel.org/r/20230217210917.2649365-1-horatiu.vultur@microchip.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f310445f9873bafdfdef69d6bb2b89d46bf88cd0
Author: Doug Berger <opendmb@gmail.com>
Date:   Thu Feb 16 11:41:28 2023 -0800

    net: bcmgenet: fix MoCA LED control

    [ Upstream commit a7515af9fb8f0890fe540b108def4a86b9e8330a ]

    When the bcmgenet_mii_config() code was refactored it was missed
```

```
    that the LED control for the MoCA interface got overwritten by
    the port_ctrl value. Its previous programming is restored here.

    Fixes: 4f8d81b77e66 ("net: bcmgenet: Refactor register access in bcmgenet_mii_config")
    Signed-off-by: Doug Berger <opendmb@gmail.com>
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5370647dd745bb3d8f37057006be207ddd8e9314
Author: Shigeru Yoshida <syoshida@redhat.com>
Date:   Fri Feb 17 01:37:10 2023 +0900

    l2tp: Avoid possible recursive deadlock in l2tp_tunnel_register()

    [ Upstream commit 9ca5e7ecab064f1f47da07f7c1ddf40e4bc0e5ac ]

    When a file descriptor of pppol2tp socket is passed as file descriptor
    of UDP socket, a recursive deadlock occurs in l2tp_tunnel_register().
    This situation is reproduced by the following program:

    int main(void)
    {
            int sock;
            struct sockaddr_pppol2tp addr;

            sock = socket(AF_PPPOX, SOCK_DGRAM, PX_PROTO_OL2TP);
            if (sock < 0) {
                    perror("socket");
                    return 1;
            }

            addr.sa_family = AF_PPPOX;
            addr.sa_protocol = PX_PROTO_OL2TP;
            addr.pppol2tp.pid = 0;
            addr.pppol2tp.fd = sock;
            addr.pppol2tp.addr.sin_family = PF_INET;
            addr.pppol2tp.addr.sin_port = htons(0);
            addr.pppol2tp.addr.sin_addr.s_addr = inet_addr("192.168.0.1");
            addr.pppol2tp.s_tunnel = 1;
            addr.pppol2tp.s_session = 0;
            addr.pppol2tp.d_tunnel = 0;
            addr.pppol2tp.d_session = 0;

            if (connect(sock, (const struct sockaddr *)&addr, sizeof(addr)) < 0) {
                    perror("connect");
                    return 1;
            }

            return 0;
    }

    This program causes the following lockdep warning:

     ============================================
     WARNING: possible recursive locking detected
     6.2.0-rc5-00205-gc96618275234 #56 Not tainted
     --------------------------------------------
     repro/8607 is trying to acquire lock:
     ffff8880213c8130 (sk_lock-AF_PPPOX){+.+.}-{0:0}, at: l2tp_tunnel_register+0x2b7/0x11c0

     but task is already holding lock:
     ffff8880213c8130 (sk_lock-AF_PPPOX){+.+.}-{0:0}, at: pppol2tp_connect+0xa82/0x1a30

     other info that might help us debug this:
```

```
       Possible unsafe locking scenario:

             CPU0
             ----
        lock(sk_lock-AF_PPPOX);
        lock(sk_lock-AF_PPPOX);

       *** DEADLOCK ***

       May be due to missing lock nesting notation

      1 lock held by repro/8607:
       #0: ffff8880213c8130 (sk_lock-AF_PPPOX){+.+.}-{0:0}, at: pppol2tp_connect+0xa82/0x1a30

      stack backtrace:
      CPU: 0 PID: 8607 Comm: repro Not tainted 6.2.0-rc5-00205-gc96618275234 #56
      Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.16.1-2.fc37 04/01/2014
      Call Trace:
       <TASK>
       dump_stack_lvl+0x100/0x178
       __lock_acquire.cold+0x119/0x3b9
       ? lockdep_hardirqs_on_prepare+0x410/0x410
       lock_acquire+0x1e0/0x610
       ? l2tp_tunnel_register+0x2b7/0x11c0
       ? lock_downgrade+0x710/0x710
       ? __fget_files+0x283/0x3e0
       lock_sock_nested+0x3a/0xf0
       ? l2tp_tunnel_register+0x2b7/0x11c0
       l2tp_tunnel_register+0x2b7/0x11c0
       ? sprintf+0xc4/0x100
       ? l2tp_tunnel_del_work+0x6b0/0x6b0
       ? debug_object_deactivate+0x320/0x320
       ? lockdep_init_map_type+0x16d/0x7a0
       ? lockdep_init_map_type+0x16d/0x7a0
       ? l2tp_tunnel_create+0x2bf/0x4b0
       ? l2tp_tunnel_create+0x3c6/0x4b0
       pppol2tp_connect+0x14e1/0x1a30
       ? pppol2tp_put_sk+0xd0/0xd0
       ? aa_sk_perm+0x2b7/0xa80
       ? aa_af_perm+0x260/0x260
       ? bpf_lsm_socket_connect+0x9/0x10
       ? pppol2tp_put_sk+0xd0/0xd0
       __sys_connect_file+0x14f/0x190
       __sys_connect+0x133/0x160
       ? __sys_connect_file+0x190/0x190
       ? lockdep_hardirqs_on+0x7d/0x100
       ? ktime_get_coarse_real_ts64+0x1b7/0x200
       ? ktime_get_coarse_real_ts64+0x147/0x200
       ? __audit_syscall_entry+0x396/0x500
       __x64_sys_connect+0x72/0xb0
       do_syscall_64+0x38/0xb0
       entry_SYSCALL_64_after_hwframe+0x63/0xcd

     This patch fixes the issue by getting/creating the tunnel before
     locking the pppol2tp socket.

     Fixes: 0b2c59720e65 ("l2tp: close all race conditions in l2tp_tunnel_register()")
     Cc: Cong Wang <cong.wang@bytedance.com>
     Signed-off-by: Shigeru Yoshida <syoshida@redhat.com>
     Reviewed-by: Guillaume Nault <gnault@redhat.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit d06d871894324cfc316cbeaaf4104f86b9c99e3e
  Author: Jakub Sitnicki <jakub@cloudflare.com>
```

```
Date:   Thu Feb 16 13:43:40 2023 +0100

    selftests/net: Interpret UDP_GRO cmsg data as an int value

    [ Upstream commit 436864095a95fcc611c20c44a111985fa9848730 ]

    Data passed to user-space with a (SOL_UDP, UDP_GRO) cmsg carries an
    int (see udp_cmsg_recv), not a u16 value, as strace confirms:

      recvmsg(8, {msg_name=...,
                  msg_iov=[{iov_base="\0\0..."..., iov_len=96000}],
                  msg_iovlen=1,
                  msg_control=[{cmsg_len=20,         <-- sizeof(cmsghdr) + 4
                                cmsg_level=SOL_UDP,
                                cmsg_type=0x68}],    <-- UDP_GRO
                                msg_controllen=24,
                                msg_flags=0}, 0) = 11200

    Interpreting the data as an u16 value won't work on big-endian platforms.
    Since it is too late to back out of this API decision [1], fix the test.

    [1]: https://lore.kernel.org/netdev/20230131174601.203127-1-jakub@cloudflare.com/

    Fixes: 3327a9c46352 ("selftests: add functionals test for UDP GRO")
    Suggested-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: Jakub Sitnicki <jakub@cloudflare.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c1001caf8e4910eb47071d173e4582fb771c7c3
Author: D. Wythe <alibuda@linux.alibaba.com>
Date:   Thu Feb 16 14:39:05 2023 +0800

    net/smc: fix application data exception

    [ Upstream commit 475f9ff63ee8c296aa46c6e9e9ad9bdd301c6bdf ]

    There is a certain probability that following
    exceptions will occur in the wrk benchmark test:

    Running 10s test @ http://11.213.45.6:80
      8 threads and 64 connections
      Thread Stats   Avg      Stdev     Max   +/- Stdev
        Latency     3.72ms   13.94ms 245.33ms   94.17%
        Req/Sec     1.96k   713.67     5.41k    75.16%
      155262 requests in 10.10s, 23.10MB read
    Non-2xx or 3xx responses: 3

    We will find that the error is HTTP 400 error, which is a serious
    exception in our test, which means the application data was
    corrupted.

    Consider the following scenarios:

    CPU0                            CPU1

    buf_desc->used = 0;
                                    cmpxchg(buf_desc->used, 0, 1)
                                    deal_with(buf_desc)

    memset(buf_desc->cpu_addr,0);

    This will cause the data received by a victim connection to be cleared,
    thus triggering an HTTP 400 error in the server.
```

```
    This patch exchange the order between clear used and memset, add
    barrier to ensure memory consistency.

    Fixes: 1c5526968e27 ("net/smc: Clear memory when release and reuse buffer")
    Signed-off-by: D. Wythe <alibuda@linux.alibaba.com>
    Reviewed-by: Wenjia Zhang <wenjia@linux.ibm.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c764cc271d3aa6528ae1b3394babf34ac01f775
Author: D. Wythe <alibuda@linux.alibaba.com>
Date:   Thu Feb 16 14:37:36 2023 +0800

    net/smc: fix potential panic dues to unprotected smc_llc_srv_add_link()

    [ Upstream commit e40b801b3603a8f90b46acbacdea3505c27f01c0 ]

    There is a certain chance to trigger the following panic:

    PID: 5900   TASK: ffff88c1c8af4100  CPU: 1   COMMAND: "kworker/1:48"
     #0 [ffff9456c1cc79a0] machine_kexec at ffffffff870665b7
     #1 [ffff9456c1cc79f0] __crash_kexec at ffffffff871b4c7a
     #2 [ffff9456c1cc7ab0] crash_kexec at ffffffff871b5b60
     #3 [ffff9456c1cc7ac0] oops_end at ffffffff87026ce7
     #4 [ffff9456c1cc7ae0] page_fault_oops at ffffffff87075715
     #5 [ffff9456c1cc7b58] exc_page_fault at ffffffff87ad0654
     #6 [ffff9456c1cc7b80] asm_exc_page_fault at ffffffff87c00b62
        [exception RIP: ib_alloc_mr+19]
        RIP: ffffffffc0c9cce3  RSP: ffff9456c1cc7c38  RFLAGS: 00010202
        RAX: 0000000000000000  RBX: 0000000000000002  RCX: 0000000000000004
        RDX: 0000000000000010  RSI: 0000000000000000  RDI: 0000000000000000
        RBP: ffff88c1ea281d00   R8: 000000020a34ffff   R9: ffff88c1350bbb20
        R10: 0000000000000000  R11: 0000000000000001  R12: 0000000000000000
        R13: 0000000000000010  R14: ffff88c1ab040a50  R15: ffff88c1ea281d00
        ORIG_RAX: ffffffffffffffff  CS: 0010  SS: 0018
     #7 [ffff9456c1cc7c60] smc_ib_get_memory_region at ffffffffc0aff6df [smc]
     #8 [ffff9456c1cc7c88] smcr_buf_map_link at ffffffffc0b0278c [smc]
     #9 [ffff9456c1cc7ce0] __smc_buf_create at ffffffffc0b03586 [smc]

    The reason here is that when the server tries to create a second link,
    smc_llc_srv_add_link() has no protection and may add a new link to
    link group. This breaks the security environment protected by
    llc_conf_mutex.

    Fixes: 2d2209f20189 ("net/smc: first part of add link processing as SMC server")
    Signed-off-by: D. Wythe <alibuda@linux.alibaba.com>
    Reviewed-by: Larysa Zaremba <larysa.zaremba@intel.com>
    Reviewed-by: Wenjia Zhang <wenjia@linux.ibm.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3c8b56b525cddadf70c91ac2b68b48a34427f45
Author: Florian Fainelli <f.fainelli@gmail.com>
Date:   Fri Dec 16 15:09:34 2022 -0800

    irqchip/irq-bcm7120-l2: Set IRQ_LEVEL for level triggered interrupts

    [ Upstream commit 13a157b38ca5b4f9eed81442b8821db293755961 ]

    When support for the interrupt controller was added with a5042de2688d,
    we forgot to update the flags to be set to contain IRQ_LEVEL. While the
    flow handler is correct, the output from /proc/interrupts does not show
    such interrupts as being level triggered when they are, correct that.
```

```
    Fixes: a5042de2688d ("irqchip: bcm7120-l2: Add Broadcom BCM7120-style Level 2 interrupt
controller")
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Reviewed-by: Philippe Mathieu-Daudé <philmd@linaro.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20221216230934.2478345-3-f.fainelli@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76659c67e60922116f79d132ddc5802dc2dfcc76
Author: Florian Fainelli <f.fainelli@gmail.com>
Date:   Fri Dec 16 15:09:33 2022 -0800

    irqchip/irq-brcmstb-l2: Set IRQ_LEVEL for level triggered interrupts

    [ Upstream commit 94debe03e8afa1267f95a9001786a6aa506b9ff3 ]

    When support for the level triggered interrupt controller flavor was
    added with c0ca7262088e, we forgot to update the flags to be set to
    contain IRQ_LEVEL. While the flow handler is correct, the output from
    /proc/interrupts does not show such interrupts as being level triggered
    when they are, correct that.

    Fixes: c0ca7262088e ("irqchip/brcmstb-l2: Add support for the BCM7271 L2 controller")
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Reviewed-by: Philippe Mathieu-Daudé <philmd@linaro.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20221216230934.2478345-2-f.fainelli@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7c0cacd2fee5d6d63e26b6bd22031354df8f180
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Feb 15 20:59:52 2023 -0800

    bpf: Fix global subprog context argument resolution logic

    [ Upstream commit d384dce281ed1b504fae2e279507827638d56fa3 ]

    KPROBE program's user-facing context type is defined as typedef
    bpf_user_pt_regs_t. This leads to a problem when trying to passing
    kprobe/uprobe/usdt context argument into global subprog, as kernel
    always strip away mods and typedefs of user-supplied type, but takes
    expected type from bpf_ctx_convert as is, which causes mismatch.

    Current way to work around this is to define a fake struct with the same
    name as expected typedef:

      struct bpf_user_pt_regs_t {};

      __noinline my_global_subprog(struct bpf_user_pt_regs_t *ctx) { ... }

    This patch fixes the issue by resolving expected type, if it's not
    a struct. It still leaves the above work-around working for backwards
    compatibility.

    Fixes: 91cc1a99740e ("bpf: Annotate context types")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Stanislav Fomichev <sdf@google.com>
    Link: https://lore.kernel.org/bpf/20230216045954.3002473-2-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f296248dafb4eefd8ea3621eec45c5ec522f7e0d
Author: Hengqi Chen <hengqi.chen@gmail.com>
Date:   Tue Feb 14 15:26:33 2023 +0000
```

```
    LoongArch, bpf: Use 4 instructions for function address in JIT

    [ Upstream commit 64f50f6575721ef03d001e907455cbe3baa2a5b1 ]

    This patch fixes the following issue of function calls in JIT, like:

      [   29.346981] multi-func JIT bug 105 != 103

    The issus can be reproduced by running the "inline simple bpf_loop call"
    verifier test.

    This is because we are emiting 2-4 instructions for 64-bit immediate moves.
    During the first pass of JIT, the placeholder address is zero, emiting two
    instructions for it. In the extra pass, the function address is in XKVRANGE,
    emiting four instructions for it. This change the instruction index in
    JIT context. Let's always use 4 instructions for function address in JIT.
    So that the instruction sequences don't change between the first pass and
    the extra pass for function calls.

    Fixes: 5dc615520c4d ("LoongArch: Add BPF JIT support")
    Signed-off-by: Hengqi Chen <hengqi.chen@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Tested-by: Tiezhu Yang <yangtiezhu@loongson.cn>
    Link: https://lore.kernel.org/bpf/20230214152633.2265699-1-hengqi.chen@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffe19750e68d0bb21e8110b398346eef20b156a7
Author: Maciej Fijalkowski <maciej.fijalkowski@intel.com>
Date:   Wed Feb 15 15:33:09 2023 +0100

    xsk: check IFF_UP earlier in Tx path

    [ Upstream commit 1596dae2f17ec5c6e8c8f0e3fec78c5ae55c1e0b ]

    Xsk Tx can be triggered via either sendmsg() or poll() syscalls. These
    two paths share a call to common function xsk_xmit() which has two
    sanity checks within. A pseudo code example to show the two paths:

    __xsk_sendmsg() :                       xsk_poll():
    if (unlikely(!xsk_is_bound(xs)))        if (unlikely(!xsk_is_bound(xs)))
        return -ENXIO;                          return mask;
    if (unlikely(need_wait))                (...)
        return -EOPNOTSUPP;                 xsk_xmit()
    mark napi id
    (...)
    xsk_xmit()

    xsk_xmit():
    if (unlikely(!(xs->dev->flags & IFF_UP)))
            return -ENETDOWN;
    if (unlikely(!xs->tx))
            return -ENOBUFS;

    As it can be observed above, in sendmsg() napi id can be marked on
    interface that was not brought up and this causes a NULL ptr
    dereference:

    [31757.505631] BUG: kernel NULL pointer dereference, address: 0000000000000018
    [31757.512710] #PF: supervisor read access in kernel mode
    [31757.517936] #PF: error_code(0x0000) - not-present page
    [31757.523149] PGD 0 P4D 0
    [31757.525726] Oops: 0000 [#1] PREEMPT SMP NOPTI
    [31757.530154] CPU: 26 PID: 95641 Comm: xdpsock Not tainted 6.2.0-rc5+ #40
    [31757.536871] Hardware name: Intel Corporation S2600WFT/S2600WFT, BIOS
SE5C620.86B.02.01.0008.031920191559 03/19/2019
```

```
    [31757.547457] RIP: 0010:xsk_sendmsg+0xde/0x180
    [31757.551799] Code: 00 75 a2 48 8b 00 a8 04 75 9b 84 d2 74 69 8b 85 14 01 00 00 85 c0
 75 1b 48 8b 85 28 03 00 00 48 8b 80 98 00 00 00 48 8b 40 20 <8b> 40 18 89 85 14 01 00 00 8b
 bd 14 01 00 00 81 ff 00 01 00 00 0f
    [31757.570840] RSP: 0018:ffffc90034f27dc0 EFLAGS: 00010246
    [31757.576143] RAX: 0000000000000000 RBX: ffffc90034f27e18 RCX: 0000000000000000
    [31757.583389] RDX: 0000000000000001 RSI: ffffc90034f27e18 RDI: ffff88984cf3c100
    [31757.590631] RBP: ffff88984714a800 R08: ffff88984714a800 R09: 0000000000000000
    [31757.597877] R10: 0000000000000001 R11: 0000000000000000 R12: 00000000fffffffa
    [31757.605123] R13: 0000000000000000 R14: 0000000000000003 R15: 0000000000000000
    [31757.612364] FS:  00007fb4c5931180(0000) GS:ffff88afdfa00000(0000)
 knlGS:0000000000000000
    [31757.620571] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [31757.626406] CR2: 0000000000000018 CR3: 000000184b41c003 CR4: 0000000007706e0
    [31757.633648] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [31757.640894] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [31757.648139] PKRU: 55555554
    [31757.650894] Call Trace:
    [31757.653385]  <TASK>
    [31757.655524]  sock_sendmsg+0x8f/0xa0
    [31757.659077]  ? sockfd_lookup_light+0x12/0x70
    [31757.663416]  __sys_sendto+0xfc/0x170
    [31757.667051]  ? do_sched_setscheduler+0xdb/0x1b0
    [31757.671658]  __x64_sys_sendto+0x20/0x30
    [31757.675557]  do_syscall_64+0x38/0x90
    [31757.679197]  entry_SYSCALL_64_after_hwframe+0x72/0xdc
    [31757.687969] Code: 8e f6 ff 44 8b 4c 24 2c 4c 8b 44 24 20 41 89 c4 44 8b 54 24 28 48
 8b 54 24 18 b8 2c 00 00 00 48 8b 74 24 10 8b 7c 24 08 0f 05 <48> 3d 00 f0 ff ff 77 3a 44 89
 e7 48 89 44 24 08 e8 b5 8e f6 ff 48
    [31757.707007] RSP: 002b:00007ffd49c73c70 EFLAGS: 00000293 ORIG_RAX: 000000000000002c
    [31757.714694] RAX: ffffffffffffffda RBX: 000055a996565380 RCX: 00007fb4c5727c16
    [31757.721939] RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000003
    [31757.729184] RBP: 0000000000000040 R08: 0000000000000000 R09: 0000000000000000
    [31757.736429] R10: 0000000000000040 R11: 0000000000000293 R12: 0000000000000000
    [31757.743673] R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000000
    [31757.754940]  </TASK>

    To fix this, let's make xsk_xmit a function that will be responsible for
    generic Tx, where RCU is handled accordingly and pull out sanity checks
    and xs->zc handling. Populate sanity checks to __xsk_sendmsg() and
    xsk_poll().

    Fixes: ca2e1a627035 ("xsk: Mark napi_id on sendmsg()")
    Fixes: 18b1ab7aa76b ("xsk: Fix race at socket teardown")
    Signed-off-by: Maciej Fijalkowski <maciej.fijalkowski@intel.com>
    Reviewed-by: Alexander Lobakin <aleksander.lobakin@intel.com>
    Link: https://lore.kernel.org/r/20230215143309.13145-1-maciej.fijalkowski@intel.com
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 6d3c6d2304bedf082cba5c9bf79c8d0f8d41e88e
 Author: Frank Jungclaus <frank.jungclaus@esd.eu>
 Date:   Thu Feb 16 20:04:49 2023 +0100

    can: esd_usb: Make use of can_change_state() and relocate checking skb for NULL

    [ Upstream commit 9684b000a86299b5968fef8ffbf1484def37452a ]

    Start a rework initiated by Vincents remarks "You should not report
    the greatest of txerr and rxerr but the one which actually increased."
    [1] and "As far as I understand, those flags should be set only when
    the threshold is reached" [2] .

    Therefore make use of can_change_state() to (among others) set the
```

```
    flags CAN_ERR_CRTL_[RT]X_WARNING and CAN_ERR_CRTL_[RT]X_PASSIVE,
    maintain CAN statistic counters for error_warning, error_passive and
    bus_off.

    Relocate testing alloc_can_err_skb() for NULL to the end of
    esd_usb_rx_event(), to have things like can_bus_off(),
    can_change_state() working even in out of memory conditions.

    Fixes: 96d8e90382dc ("can: Add driver for esd CAN-USB/2 device")
    Signed-off-by: Frank Jungclaus <frank.jungclaus@esd.eu>
    Link: [1] https://lore.kernel.org/all/CAMZ6RqKGBWe15aMkf8-QLf-
cOQg99GQBebSm+1wEzTqHgvmNuw@mail.gmail.com/
    Link: [2] https://lore.kernel.org/all/CAMZ6Rq+QBO1yTX_o6GV0yhdBj-
RzZSRGWDZBS0fs7zbSTy4hmA@mail.gmail.com/
    Link: https://lore.kernel.org/all/20230216190450.3901254-3-frank.jungclaus@esd.eu
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e18193fb92e0f132abd146d21c1e71cc6c8311a
Author: Frank Jungclaus <frank.jungclaus@esd.eu>
Date:   Thu Feb 16 20:04:48 2023 +0100

    can: esd_usb: Move mislocated storage of SJA1000_ECC_SEG bits in case of a bus error

    [ Upstream commit 118469f88180438ef43dee93d71f77c00e7b425d ]

    Move the supply for cf->data[3] (bit stream position of CAN error), in
    case of a bus- or protocol-error, outside of the "switch (ecc &
    SJA1000_ECC_MASK){}"-statement, because this bit stream position is
    independent of the error type.

    Fixes: 96d8e90382dc ("can: Add driver for esd CAN-USB/2 device")
    Signed-off-by: Frank Jungclaus <frank.jungclaus@esd.eu>
    Link: https://lore.kernel.org/all/20230216190450.3901254-2-frank.jungclaus@esd.eu
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc6a048611355f37e966b23263eab1a0d5e95dd7
Author: Ilya Leoshkevich <iii@linux.ibm.com>
Date:   Sat Jan 28 01:06:29 2023 +0100

    selftests/bpf: Fix xdp_do_redirect on s390x

    [ Upstream commit 06c1865b0b0c7820ea53af2394dd7aff31100295 ]

    s390x cache line size is 256 bytes, so skb_shared_info must be aligned
    on a much larger boundary than for x86. This makes the maximum packet
    size smaller.

    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230128000650.1516334-11-iii@linux.ibm.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Stable-dep-of: 6c20822fada1 ("bpf, test_run: fix &xdp_frame misplacement for
LIVE_FRAMES")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5d447e04290e78bdc1a3a6c321320d384e09c2f1
Author: Hou Tao <houtao1@huawei.com>
Date:   Wed Feb 15 16:21:31 2023 +0800

    bpf: Zeroing allocated object from slab in bpf memory allocator

    [ Upstream commit 997849c4b969034e225153f41026657def66d286 ]

    Currently the freed element in bpf memory allocator may be immediately
```

```
    reused, for htab map the reuse will reinitialize special fields in map
    value (e.g., bpf_spin_lock), but lookup procedure may still access
    these special fields, and it may lead to hard-lockup as shown below:

     NMI backtrace for cpu 16
     CPU: 16 PID: 2574 Comm: htab.bin Tainted: G             L     6.1.0+ #1
     Hardware name: QEMU Standard PC (i440FX + PIIX, 1996),
     RIP: 0010:queued_spin_lock_slowpath+0x283/0x2c0
     ......
     Call Trace:
      <TASK>
      copy_map_value_locked+0xb7/0x170
      bpf_map_copy_value+0x113/0x3c0
      __sys_bpf+0x1c67/0x2780
      __x64_sys_bpf+0x1c/0x20
      do_syscall_64+0x30/0x60
      entry_SYSCALL_64_after_hwframe+0x46/0xb0
     ......
      </TASK>

    For htab map, just like the preallocated case, these is no need to
    initialize these special fields in map value again once these fields
    have been initialized. For preallocated htab map, these fields are
    initialized through __GFP_ZERO in bpf_map_area_alloc(), so do the
    similar thing for non-preallocated htab in bpf memory allocator. And
    there is no need to use __GFP_ZERO for per-cpu bpf memory allocator,
    because __alloc_percpu_gfp() does it implicitly.

    Fixes: 0fd7c5d43339 ("bpf: Optimize call_rcu in non-preallocated hash map.")
    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Link: https://lore.kernel.org/r/20230215082132.3856544-2-houtao@huaweicloud.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 229f98bd46dbb26e5aa60e92522f0ff6ffbe26f5
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Feb 15 10:07:05 2023 +0100

    wifi: mac80211: pass 'sta' to ieee80211_rx_data_set_sta()

    [ Upstream commit 0d846bdc11101ac0ba4d89c2be359af08cb9379b ]

    There's at least one case in ieee80211_rx_for_interface()
    where we might pass &((struct sta_info *)NULL)->sta to it
    only to then do container_of(), and then checking the
    result for NULL, but checking the result of container_of()
    for NULL looks really odd.

    Fix this by just passing the struct sta_info * instead.

    Fixes: e66b7920aa5a ("wifi: mac80211: fix initialization of rx->link and rx->link_sta")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 55bb5f3d472c819d719db993ade82c447dc2df91
Author: Alexei Starovoitov <ast@kernel.org>
Date:   Tue Feb 14 15:50:51 2023 -0800

    selftests/bpf: Fix map_kptr test.

    [ Upstream commit 62d101d5f422cde39b269f7eb4cbbe2f1e26f9d4 ]

    The compiler is optimizing out majority of unref_ptr read/writes, so the test
    wasn't testing much. For example, one could delete '__kptr' tag from
    'struct prog_test_ref_kfunc __kptr *unref_ptr;' and the test would still "pass".
```

```
    Convert it to volatile stores. Confirmed by comparing bpf asm before/after.

    Fixes: 2cbc469a6fc3 ("selftests/bpf: Add C tests for kptr")
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Stanislav Fomichev <sdf@google.com>
    Acked-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230214235051.22938-1-alexei.starovoitov@gmail.com
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68e675a9b69cfc34dd915d91a4650e3ee53421f4
Author: Yongqin Liu <yongqin.liu@linaro.org>
Date:   Fri Feb 10 22:15:07 2023 +0800

    thermal/drivers/hisi: Drop second sensor hi3660

    [ Upstream commit 15cc25829a97c3957e520e971868aacc84341317 ]

    The commit 74c8e6bffbe1 ("driver core: Add __alloc_size hint to devm
    allocators") exposes a panic "BRK handler: Fatal exception" on the
    hi3660_thermal_probe funciton.
    This is because the function allocates memory for only one
    sensors array entry, but tries to fill up a second one.

    Fix this by removing the unneeded second access.

    Fixes: 7d3a2a2bbadb ("thermal/drivers/hisi: Fix number of sensors on hi3660")
    Signed-off-by: Yongqin Liu <yongqin.liu@linaro.org>
    Link: https://lore.kernel.org/linux-mm/20221101223321.1326815-5-keescook@chromium.org/
    Link: https://lore.kernel.org/r/20230210141507.71014-1-yongqin.liu@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7f3e65f809c4db4420f12d6200ffa152a0d554f4
Author: Vincent Guittot <vincent.guittot@linaro.org>
Date:   Thu Feb 2 11:28:12 2023 +0100

    tools/lib/thermal: Fix thermal_sampling_exit()

    [ Upstream commit a29cbd76aaf63f5493e962aa2fbaadcdc4615143 ]

    thermal_sampling_init() suscribes to THERMAL_GENL_SAMPLING_GROUP_NAME group
    so thermal_sampling_exit() should unsubscribe from the same group.

    Fixes: 47c4b0de080a ("tools/lib/thermal: Add a thermal library")
    Signed-off-by: Vincent Guittot <vincent.guittot@linaro.org>
    Link: https://lore.kernel.org/r/20230202102812.453357-1-vincent.guittot@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 53050740afb9af4d6c3f668bcedf53d4acd239f4
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Tue Feb 14 20:08:15 2023 +0100

    wifi: mac80211: fix off-by-one link setting

    [ Upstream commit cf08e29db760b144bde51e2444f3430c75763e26 ]

    The convention for find_first_bit() is 0-based, while ffs()
    is 1-based, so this is now off-by-one. I cannot reproduce the
    gcc-9 problem, but since the -1 is now removed, I'm hoping it
    will still avoid the original issue.
```

```
    Reported-by: Alexander Lobakin <alexandr.lobakin@intel.com>
    Fixes: 1d8d4af43474 ("wifi: mac80211: avoid u32_encode_bits() warning")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97cb35a38de66bd126fce58089cb06f9ea2423ff
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Tue Feb 14 14:20:21 2023 +0100

    wifi: mac80211: avoid u32_encode_bits() warning

    [ Upstream commit 1d8d4af4347420d657be448f8be4c39c558f3b5d ]

    gcc-9 triggers a false-postive warning in ieee80211_mlo_multicast_tx()
    for u32_encode_bits(ffs(links) - 1, ...), since ffs() can return zero
    on an empty bitmask, and the negative argument to u32_encode_bits()
    is then out of range:

    In file included from include/linux/ieee80211.h:21,
                     from include/net/cfg80211.h:23,
                     from net/mac80211/tx.c:23:
    In function 'u32_encode_bits',
        inlined from 'ieee80211_mlo_multicast_tx' at net/mac80211/tx.c:4437:17,
        inlined from 'ieee80211_subif_start_xmit' at net/mac80211/tx.c:4485:3:
    include/linux/bitfield.h:177:3: error: call to '__field_overflow' declared with
attribute error: value doesn't fit into mask
      177 |   __field_overflow();     \
          |   ^~~~~~~~~~~~~~~~~~
    include/linux/bitfield.h:197:2: note: in expansion of macro '____MAKE_OP'
      197 |  ____MAKE_OP(u##size,u##size,,)
          |  ^~~~~~~~~~~
    include/linux/bitfield.h:200:1: note: in expansion of macro '__MAKE_OP'
      200 | __MAKE_OP(32)
          | ^~~~~~~~

    Newer compiler versions do not cause problems with the zero argument
    because they do not consider this a __builtin_constant_p().
    It's also harmless since the hweight16() check already guarantees
    that this cannot be 0.

    Replace the ffs() with an equivalent find_first_bit() check that
    matches the later for_each_set_bit() style and avoids the warning.

    Fixes: 963d0e8d08d9 ("wifi: mac80211: optionally implement MLO multicast TX")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20230214132025.1532147-1-arnd@kernel.org
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 975433e43d997811616e8a167f2c18086809abda
Author: Andrei Otcheretianski <andrei.otcheretianski@intel.com>
Date:   Tue Feb 14 12:10:48 2023 +0200

    wifi: mac80211: Don't translate MLD addresses for multicast

    [ Upstream commit daf8fb4295dccc032515cdc1bd3873370063542b ]

    MLD address translation should be done only for individually addressed
    frames. Otherwise, AAD calculation would be wrong and the decryption
    would fail.

    Fixes: e66b7920aa5ac ("wifi: mac80211: fix initialization of rx->link and rx->link_sta")
    Signed-off-by: Andrei Otcheretianski <andrei.otcheretianski@intel.com>
    Link: https://lore.kernel.org/r/20230214101048.792414-1-andrei.otcheretianski@intel.com
```

```
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f92f8ae69cb7f3a267dc49cc7da62422c517f079
Author: Karthikeyan Periyasamy <quic_periyasa@quicinc.com>
Date:   Mon Feb 6 21:33:30 2023 +0530

    wifi: mac80211: fix non-MLO station association

    [ Upstream commit aaacf1740f2f95e0c5449ff3bbcff252d69cf952 ]

    Non-MLO station frames are dropped in Rx path due to the condition
    check in ieee80211_rx_is_valid_sta_link_id(). In multi-link AP scenario,
    non-MLO stations try to connect in any of the valid links in the ML AP,
    where the station valid_links and link_id params are valid in the
    ieee80211_sta object. But ieee80211_rx_is_valid_sta_link_id() always
    return false for the non-MLO stations by the assumption taken is
    valid_links and link_id are not valid in non-MLO stations object
    (ieee80211_sta), this assumption is wrong. Due to this assumption,
    non-MLO station frames are dropped which leads to failure in association.

    Fix it by removing the condition check and allow the link validation
    check for the non-MLO stations.

    Fixes: e66b7920aa5a ("wifi: mac80211: fix initialization of rx->link and rx->link_sta")
    Signed-off-by: Karthikeyan Periyasamy <quic_periyasa@quicinc.com>
    Link: https://lore.kernel.org/r/20230206160330.1613-1-quic_periyasa@quicinc.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1380e660c8b8e3dd3047d9e7b3b57b85578a62bf
Author: Shayne Chen <shayne.chen@mediatek.com>
Date:   Thu Feb 9 19:06:59 2023 +0800

    wifi: mac80211: make rate u32 in sta_set_rate_info_rx()

    [ Upstream commit 59336e07b287d91dc4ec265e07724e8f7e3d0209 ]

    The value of last_rate in ieee80211_sta_rx_stats is degraded from u32 to
    u16 after being assigned to rate variable, which causes information loss
    in STA_STATS_FIELD_TYPE and later bitfields.

    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Link: https://lore.kernel.org/r/20230209110659.25447-1-shayne.chen@mediatek.com
    Fixes: 41cbb0f5a295 ("mac80211: add support for HE")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 51ac31696da7c01e560dee41df3da6677c910085
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Wed Jan 25 12:01:02 2023 +0100

    wifi: mac80211: move color collision detection report in a delayed work

    [ Upstream commit 9288188438d85e22c23cfd6657ee8a801babc83c ]

    Move color collision report in a dedicated delayed work and do not run
    it in interrupt context in order to rate-limit the number of events
    reported to userspace. Moreover grab wdev mutex in
    ieee80211_color_collision_detection_work routine since it is required
    by cfg80211_obss_color_collision_notify().

    Tested-by: Nicolas Cavallari <nicolas.cavallari@green-communications.fr>
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Fixes: 5f9404abdf2a ("mac80211: add support for BSS color change")
```

```
    Link:
https://lore.kernel.org/r/3f6cf60c892ad40c1cca4a55d62b1224ef1c6ce9.1674644379.git.lorenzo@ke
rnel.org
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d33c0c56371bd84a1bdb6947da66698ac6812261
Author: Eric Farman <farman@linux.ibm.com>
Date:   Fri Feb 10 18:42:27 2023 +0100

    vfio/ccw: remove WARN_ON during shutdown

    [ Upstream commit 1c06bb87afb2d95b8e9f4f2e3d0d6772c68f3e76 ]

    The logic in vfio_ccw_sch_shutdown() always assumed that the input
    subchannel would point to a vfio_ccw_private struct, without checking
    that one exists. The blamed commit put in a check for this scenario,
    to prevent the possibility of a missing private.

    The trouble is that check was put alongside a WARN_ON(), presuming
    that such a scenario would be a cause for concern. But this can be
    triggered by binding a subchannel to vfio-ccw, and rebooting the
    system before starting the mdev (via "mdevctl start" or similar)
    or after stopping it. In those cases, shutdown doesn't need to
    worry because either the private was never allocated, or it was
    cleaned up by vfio_ccw_mdev_remove().

    Remove the WARN_ON() piece of this check, since there are plausible
    scenarios where private would be NULL in this path.

    Fixes: 9e6f07cd1eaa ("vfio/ccw: create a parent struct")
    Signed-off-by: Eric Farman <farman@linux.ibm.com>
    Reviewed-by: Matthew Rosato <mjrosato@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230210174227.2256424-1-farman@linux.ibm.com
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c456a638336076f14a8f374e510507ed4f64020e
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Mon Feb 6 14:01:53 2023 +0800

    crypto: crypto4xx - Call dma_unmap_page when done

    [ Upstream commit bcdda4301bdc4955d45f7e1ffefb6207967b067e ]

    In crypto4xx_cipher_done, we should be unmapping the dst page, not
    mapping it.

    This was flagged by a sparse warning about the unused addr variable.
    While we're at it, also fix a sparse warning regarding the unused
    ctx variable in crypto4xx_ahash_done (by actually using it).

    Fixes: 049359d65527 ("crypto: amcc - Add crypt4xx driver")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Tested-by: Christian Lamparter <chunkeey@gmail.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e13d3af2d09d46c95e3ab1db968bd059e97722f8
Author: Alexander Lobakin <alobakin@pm.me>
Date:   Mon Feb 6 13:18:50 2023 +0800

    crypto: octeontx2 - Fix objects shared between several modules

    [ Upstream commit 72bc4e71dbeedee0a446bcbc37c9bb25449072b7 ]
```

```
    cn10k_cpt.o, otx2_cptlf.o and otx2_cpt_mbox_common.o are linked
    into both rvu_cptpf and rvu_cptvf modules:

    > scripts/Makefile.build:252: ./drivers/crypto/marvell/octeontx2/Makefile:
    > cn10k_cpt.o is added to multiple modules: rvu_cptpf rvu_cptvf
    > scripts/Makefile.build:252: ./drivers/crypto/marvell/octeontx2/Makefile:
    > otx2_cptlf.o is added to multiple modules: rvu_cptpf rvu_cptvf
    > scripts/Makefile.build:252: ./drivers/crypto/marvell/octeontx2/Makefile:
    > otx2_cpt_mbox_common.o is added to multiple modules: rvu_cptpf rvu_cptvf

    Despite they're build under the same Kconfig option
    (CONFIG_CRYPTO_DEV_OCTEONTX2_CPT), it's better do link the common
    code into a standalone module and export the shared functions. Under
    certain circumstances, this can lead to the same situation as fixed
    by commit 637a642f5ca5 ("zstd: Fixing mixed module-builtin objects").
    Plus, those three common object files are relatively big to duplicate
    them several times.

    Introduce the new module, rvu_cptcommon, to provide the common
    functions to both modules.

    Fixes: 19d8e8c7be15 ("crypto: octeontx2 - add virtual function driver support")
    Suggested-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Alexander Lobakin <alobakin@pm.me>
    Reviewed-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9333e96dfc094fe30eb409e50118a1837d8542a7
Author: Werner Sembach <wse@tuxedocomputers.com>
Date:   Mon Feb 13 19:16:53 2023 +0100

    ACPI: resource: Do IRQ override on all TongFang GMxRGxx

    [ Upstream commit 17bb7046e7ce038a73ee97eaa804e0300c5199e2 ]

    Apply commit 7592b79ba4a9 ("ACPI: resource: do IRQ override on XMG Core 15")
    override for all vendors using this mainboard.

    Signed-off-by: Werner Sembach <wse@tuxedocomputers.com>
    Fixes: 9946e39fe8d0 ("ACPI: resource: skip IRQ override on AMD Zen platforms")
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8816409ad451a8e0d3079847c2f4052e0b6b7ad1
Author: Adam Niederer <adam.niederer@gmail.com>
Date:   Sat Feb 11 15:13:33 2023 -0500

    ACPI: resource: Add IRQ overrides for MAINGEAR Vector Pro 2 models

    [ Upstream commit cb18703c179713056bd7e3bdfc2260ab4e8658f0 ]

    Fix a regression introduced by commit 9946e39fe8d0 ("ACPI: resource: skip
    IRQ override on AMD Zen platforms") on MAINGEAR Vector Pro 2 systems, which
    causes the built-in keyboard to not work. This restores the functionality
    by adding an IRQ override.

    No other IRQs were being overridden before, so this should be all that is
    needed for these systems. I have personally tested this on the 15" model
    (MG-VCP2-15A3070T), and I have confirmation that the issue is present on
    the 17" model (MG-VCP2-17A3070T).

    Fixes: 9946e39fe8d0 ("ACPI: resource: skip IRQ override on AMD Zen platforms")
    Signed-off-by: Adam Niederer <adam.niederer@gmail.com>
```

```
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ce9dc8801f9ef1560baa2770bab9ba6d1624724
Author: Ilya Leoshkevich <iii@linux.ibm.com>
Date:   Thu Feb 9 00:12:11 2023 +0100

    selftests/bpf: Fix out-of-srctree build

    [ Upstream commit 0b0757244754ea1d0721195c824770f5576e119e ]

    Building BPF selftests out of srctree fails with:

      make: *** No rule to make target '/linux-build//ima_setup.sh', needed by
'ima_setup.sh'.  Stop.

    The culprit is the rule that defines convenient shorthands like
    "make test_progs", which builds $(OUTPUT)/test_progs. These shorthands
    make sense only for binaries that are built though; scripts that live
    in the source tree do not end up in $(OUTPUT).

    Therefore drop $(TEST_PROGS) and $(TEST_PROGS_EXTENDED) from the rule.

    The issue exists for a while, but it became a problem only after commit
    d68ae4982cb7 ("selftests/bpf: Install all required files to run selftests"),
    which added dependencies on these scripts.

    Fixes: 03dcb78460c2 ("selftests/bpf: Add simple per-test targets to Makefile")
    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230208231211.283606-1-iii@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b93bf957a66566fe23868a7162c6136d7055cb2
Author: Zong-Zhe Yang <kevin_yang@realtek.com>
Date:   Thu Jan 19 14:43:41 2023 +0800

    wifi: rtw89: fix parsing offset for MCC C2H

    [ Upstream commit 24d72944d79e6795ba4330c114de0387386bf3aa ]

    A 8-byte offset is missed during parsing C2Hs (chip to host packets)
    of MCC (multi-channel concurrent) series.
    So, we fix it.

    Fixes: ef9dff4cb491 ("wifi: rtw89: mac: process MCC related C2H")
    Signed-off-by: Zong-Zhe Yang <kevin_yang@realtek.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230119064342.65391-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47858de82451781401d71030fe31d6d31d2021cd
Author: Dan Carpenter <error27@gmail.com>
Date:   Mon Feb 6 17:41:33 2023 +0300

    wifi: mwifiex: fix loop iterator in mwifiex_update_ampdu_txwinsize()

    [ Upstream commit 3cfb7df24cee0f5fdc4cc5d3176cab9aadfcb430 ]

    This code re-uses "i" to be the iterator for both the inside and outside
    loops.  It means the outside loop will exit earlier than intended.

    Fixes: d219b7eb3792 ("mwifiex: handle BT coex event to adjust Rx BA window size")
    Signed-off-by: Dan Carpenter <error27@gmail.com>
```

```
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/Y+ERnaDaZD7RtLvX@kili
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c517df92622ea39f6d9977df870ff5cbd419bd58
Author: Hector Martin <marcan@marcan.st>
Date:   Sun Feb 12 15:38:13 2023 +0900

    wifi: brcmfmac: pcie: Perform correct BCM4364 firmware selection

    [ Upstream commit 6a142f70774fd10350a52a10ba1297d52da46780 ]

    This chip exists in two revisions (B2=r3 and B3=r4) on different
    platforms, and was added without regard to doing proper firmware
    selection or differentiating between them. Fix this to have proper
    per-revision firmwares and support Apple NVRAM selection.

    Revision B2 is present on at least these Apple T2 Macs:

    kauai:    MacBook Pro 15" (Touch/2018-2019)
    maui:     MacBook Pro 13" (Touch/2018-2019)
    lanai:    Mac mini (Late 2018)
    ekans:    iMac Pro 27" (5K, Late 2017)

    And these non-T2 Macs:

    nihau:    iMac 27" (5K, 2019)

    Revision B3 is present on at least these Apple T2 Macs:

    bali:     MacBook Pro 16" (2019)
    trinidad: MacBook Pro 13" (2020, 4 TB3)
    borneo:   MacBook Pro 16" (2019, 5600M)
    kahana:   Mac Pro (2019)
    kahana:   Mac Pro (2019, Rack)
    hanauma:  iMac 27" (5K, 2020)
    kure:     iMac 27" (5K, 2020, 5700/XT)

    Also fix the firmware interface for 4364, from BCA to WCC.

    Fixes: 24f0bd136264 ("brcmfmac: add the BRCM 4364 found in MacBook Pro 15,2")
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Reviewed-by: Arend van Spriel <arend.vanspriel@broadcom.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230212063813.27622-5-marcan@marcan.st
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 45e863cd00d02e80c20d211ddaf5440bae2a66e2
Author: Hector Martin <marcan@marcan.st>
Date:   Sun Feb 12 15:38:12 2023 +0900

    wifi: brcmfmac: pcie: Add IDs/properties for BCM4377

    [ Upstream commit bf8bbd903f07aabfef8d4cb3b513163a3e974d2b ]

    This chip is present on at least these Apple T2 Macs:

    * tahiti:  MacBook Pro 13" (2020, 2 TB3)
    * formosa: MacBook Pro 13" (Touch/2019)
    * fiji:    MacBook Air 13" (Scissor, 2020)

    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Reviewed-by: Arend van Spriel <arend.vanspriel@broadcom.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
```

```
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230212063813.27622-4-marcan@marcan.st
    Stable-dep-of: 6a142f70774f ("wifi: brcmfmac: pcie: Perform correct BCM4364 firmware
selection")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ff7d7162d040bd821181f2da465e1ade9167e06
Author: Hector Martin <marcan@marcan.st>
Date:   Sun Feb 12 15:38:11 2023 +0900

    wifi: brcmfmac: pcie: Add IDs/properties for BCM4355

    [ Upstream commit 69005e67ce54cb837b8218b002c1bb868c83b7a9 ]

    This chip is present on at least these Apple T2 Macs:

    * hawaii: MacBook Air 13" (Late 2018)
    * hawaii: MacBook Air 13" (True Tone, 2019)

    Users report seeing PCI revision ID 12 for this chip, which Arend
    reports should be revision C2, but Apple has the firmware tagged as
    revision C1. Assume the right cutoff point for firmware versions is
    revision ID 11 then, and leave older revisions using the non-versioned
    firmware filename (Apple only uses C1 firmware builds).

    Reviewed-by: Arend van Spriel <arend.vanspriel@broadcom.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230212063813.27622-3-marcan@marcan.st
    Stable-dep-of: 6a142f70774f ("wifi: brcmfmac: pcie: Perform correct BCM4364 firmware
selection")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc9c1b13968de12a4f405658235e738609c5fca2
Author: Hector Martin <marcan@marcan.st>
Date:   Sun Feb 12 15:38:10 2023 +0900

    wifi: brcmfmac: Rename Cypress 89459 to BCM4355

    [ Upstream commit 54f01f56cf63ebb92ac37450d65c7e4da379d4ca ]

    The commit that introduced support for this chip incorrectly claimed it
    is a Cypress-specific part, while in actuality it is just a variant of
    BCM4355 silicon (as evidenced by the chip ID).

    The relationship between Cypress products and Broadcom products isn't
    entirely clear but given what little information is available and prior
    art in the driver, it seems the convention should be that originally
    Broadcom parts should retain the Broadcom name.

    Thus, rename the relevant constants and firmware file. Also rename the
    specific 89459 PCIe ID to BCM43596, which seems to be the original
    subvariant name for this PCI ID (as defined in the out-of-tree bcmdhd
    driver).

    Since Cypress added this part and will presumably be providing its
    supported firmware, we keep the CYW designation for this device.

    We also drop the RAW device ID in this commit. We don't do this for the
    other chips since apparently some devices with them exist in the wild,
    but there is already a 4355 entry with the Broadcom subvendor and WCC
    firmware vendor, so adding a generic fallback to Cypress seems
    redundant (no reason why a device would have the raw device ID *and* an
    explicitly programmed subvendor).
```

```
    Fixes: dce45ded7619 ("brcmfmac: Support 89459 pcie")
    Reviewed-by: Arend van Spriel <arend.vanspriel@broadcom.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230212063813.27622-2-marcan@marcan.st
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 878a7c8357764e08bc778bcb26127fc12a4b36b7
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Thu Feb 9 09:07:48 2023 +0800

    wifi: iwl4965: Add missing check for create_singlethread_workqueue()

    [ Upstream commit 26e6775f75517ad6844fe5b79bc5f3fa8c22ee61 ]

    Add the check for the return value of the create_singlethread_workqueue()
    in order to avoid NULL pointer dereference.

    Fixes: b481de9ca074 ("[IWLWIFI]: add iwlwifi wireless drivers")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Acked-by: Stanislaw Gruszka <stf_xl@wp.pl>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230209010748.45454-1-jiasheng@iscas.ac.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17e07d6587c55015956862ef3b101fd45fa49fbc
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Feb 8 14:30:32 2023 +0800

    wifi: iwl3945: Add missing check for create_singlethread_workqueue

    [ Upstream commit 1fdeb8b9f29dfd64805bb49475ac7566a3cb06cb ]

    Add the check for the return value of the create_singlethread_workqueue
    in order to avoid NULL pointer dereference.

    Fixes: b481de9ca074 ("[IWLWIFI]: add iwlwifi wireless drivers")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Acked-by: Stanislaw Gruszka <stf_xl@wp.pl>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230208063032.42763-2-jiasheng@iscas.ac.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a74b3da825bf3984afa8b0e532ab3b0afe6c8a02
Author: Matt Evans <mev@rivosinc.com>
Date:   Wed Feb 1 19:49:42 2023 +0000

    clocksource/drivers/riscv: Patch riscv_clock_next_event() jump before first use

    [ Upstream commit 225b9596cb0227c1c1b1e4a836dad43595c3e61a ]

    A static key is used to select between SBI and Sstc timer usage in
    riscv_clock_next_event(), but currently the direction is resolved
    after cpuhp_setup_state() is called (which sets the next event).  The
    first event will therefore fall through the sbi_set_timer() path; this
    breaks Sstc-only systems.  So, apply the jump patching before first
    use.

    Fixes: 9f7a8ff6391f ("RISC-V: Prefer sstc extension if available")
    Signed-off-by: Matt Evans <mev@rivosinc.com>
    Reviewed-by: Palmer Dabbelt <palmer@rivosinc.com>
    Acked-by: Palmer Dabbelt <palmer@rivosinc.com>
    Reviewed-by: Anup Patel <anup@brainfault.org>
    Link: https://lore.kernel.org/r/CDDAB2D0-264E-42F3-8E31-BA210BEB8EC1@rivosinc.com
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
```

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

commit 9fef20a044140c864ba696b7be2850adb173752c
Author: Conor Dooley <conor.dooley@microchip.com>
Date:   Tue Jan 3 19:41:00 2023 +0530

    RISC-V: time: initialize hrtimer based broadcast clock event device

    [ Upstream commit 8b3b8fbb4896984b5564789a42240e4b3caddb61 ]

    Similarly to commit 022eb8ae8b5e ("ARM: 8938/1: kernel: initialize
    broadcast hrtimer based clock event device"), RISC-V needs to initiate
    hrtimer based broadcast clock event device before C3STOP can be used.
    Otherwise, the introduction of C3STOP for the RISC-V arch timer in
    commit 232ccac1bd9b ("clocksource/drivers/riscv: Events are stopped
    during CPU suspend") leaves us without any broadcast timer registered.
    This prevents the kernel from entering oneshot mode, which breaks timer
    behaviour, for example clock_nanosleep().

    A test app that sleeps each cpu for 6, 5, 4, 3 ms respectively, HZ=250
    & C3STOP enabled, the sleep times are rounded up to the next jiffy:
    == CPU: 1 ==      == CPU: 2 ==      == CPU: 3 ==      == CPU: 4 ==
    Mean: 7.974992    Mean: 7.976534    Mean: 7.962591    Mean: 3.952179
    Std Dev: 0.154374 Std Dev: 0.156082 Std Dev: 0.171018 Std Dev: 0.076193
    Hi: 9.472000      Hi: 10.495000     Hi: 8.864000      Hi: 4.736000
    Lo: 6.087000      Lo: 6.380000      Lo: 4.872000      Lo: 3.403000
    Samples: 521      Samples: 521      Samples: 521      Samples: 521

    Link: https://lore.kernel.org/linux-riscv/YzYTNQRxLr7Q9JR0@spud/
    Fixes: 232ccac1bd9b ("clocksource/drivers/riscv: Events are stopped during CPU suspend")
    Suggested-by: Samuel Holland <samuel@sholland.org>
    Signed-off-by: Conor Dooley <conor.dooley@microchip.com>
    Signed-off-by: Anup Patel <apatel@ventanamicro.com>
    Reviewed-by: Samuel Holland <samuel@sholland.org>
    Acked-by: Palmer Dabbelt <palmer@rivosinc.com>
    Link: https://lore.kernel.org/r/20230103141102.772228-2-apatel@ventanamicro.com
    Signed-off-by: Daniel Lezcano <daniel.lezcano@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4db81f2eecc1bedba01f7d3d079bd0e9c930d8fc
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Wed Feb 8 17:08:25 2023 -0800

    m68k: /proc/hardware should depend on PROC_FS

    [ Upstream commit 1e5b5df65af99013b4d31607ddb3ca5731dbe44d ]

    When CONFIG_PROC_FS is not set, there is a build error for an unused
    function. Make PROC_HARDWARE depend on PROC_FS to prevent this error.

    In file included from ../arch/m68k/kernel/setup.c:3:
    ../arch/m68k/kernel/setup_mm.c:477:12: error: 'hardware_proc_show' defined but not used
[-Werror=unused-function]
      477 | static int hardware_proc_show(struct seq_file *m, void *v)
          |            ^~~~~~~~~~~~~~~~~~

    Fixes: 66d857b08b8c ("m68k: merge m68k and m68knommu arch directories") # v3.0
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Reviewed-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Link: https://lore.kernel.org/r/20230209010825.24136-1-rdunlap@infradead.org
    Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6805246ee32dea17ed3d25f20a21c8b7d7b6a669
Author: Herbert Xu <herbert@gondor.apana.org.au>
```

```
Date:   Tue Jan 31 16:02:04 2023 +0800

    crypto: rsa-pkcs1pad - Use akcipher_request_complete

    [ Upstream commit 564cabc0ca0bdfa8f0fc1ae74b24d0a7554522c5 ]

    Use the akcipher_request_complete helper instead of calling the
    completion function directly.  In fact the previous code was buggy
    in that EINPROGRESS was never passed back to the original caller.

    Fixes: 3d5b1ecdea6f ("crypto: rsa - RSA padding algorithm")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 02d1ad9418bf3726c030109e5cee096cf4eed1f5
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Thu Feb 9 12:26:23 2023 +0000

    rds: rds_rm_zerocopy_callback() correct order for list_add_tail()

    [ Upstream commit 68762148d1b011d47bc2ceed7321739b5aea1e63 ]

    rds_rm_zerocopy_callback() uses list_add_tail() with swapped
    arguments. This links the list head with the new entry, losing
    the references to the remaining part of the list.

    Fixes: 9426bbc6de99 ("rds: use list structure to track information for zerocopy
completion notification")
    Suggested-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8f0d9e6c53c997107d928bed811876f71dd13866
Author: Oleksandr Tyshchenko <oleksandr_tyshchenko@epam.com>
Date:   Wed Feb 8 17:36:49 2023 +0200

    xen/grant-dma-iommu: Implement a dummy probe_device() callback

    [ Upstream commit 2062f9fb6445451b189595e295765c69f43bc12e ]

    Update stub IOMMU driver (which main purpose is to reuse generic
    IOMMU device-tree bindings by Xen grant DMA-mapping layer on Arm)
    according to the recent changes done in the following
    commit 57365a04c921 ("iommu: Move bus setup to IOMMU device registration").

    With probe_device() callback being called during IOMMU device registration,
    the uninitialized callback just leads to the "kernel NULL pointer
    dereference" issue during boot. Fix that by adding a dummy callback.

    Looks like the release_device() callback is not mandatory to be
    implemented as IOMMU framework makes sure that callback is initialized
    before dereferencing.

    Reported-by: Viresh Kumar <viresh.kumar@linaro.org>
    Fixes: 57365a04c921 ("iommu: Move bus setup to IOMMU device registration")
    Signed-off-by: Oleksandr Tyshchenko <oleksandr_tyshchenko@epam.com>
    Tested-by: Viresh Kumar <viresh.kumar@linaro.org>
    Reviewed-by: Stefano Stabellini <sstabellini@kernel.org>
    Link: https://lore.kernel.org/r/20230208153649.3604857-1-olekstysh@gmail.com
    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9903313fb8b43372c4338693b5da4f72590b1b1f
Author: Ilya Leoshkevich <iii@linux.ibm.com>
```

```
Date:   Fri Feb 10 01:12:01 2023 +0100

    libbpf: Fix alen calculation in libbpf_nla_dump_errormsg()

    [ Upstream commit 17bcd27a08a21397698edf143084d7c87ce17946 ]

    The code assumes that everything that comes after nlmsgerr are nlattrs.
    When calculating their size, it does not account for the initial
    nlmsghdr. This may lead to accessing uninitialized memory.

    Fixes: bbf48c18ee0c ("libbpf: add error reporting in XDP")
    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20230210001210.395194-8-iii@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70a6dadf36496f54d8f1ee3800e2e8186b509f0c
Author: Halil Pasic <pasic@linux.ibm.com>
Date:   Thu Feb 9 00:00:24 2023 +0100

    s390/ap: fix status returned by ap_qact()

    [ Upstream commit a2522c80f074c35254974fec39fffe8b8d75befe ]

    Since commit 159491f3b509 ("s390/ap: rework assembler functions to use
    unions for in/out register variables") the  function ap_qact() tries to
    grab the status from the wrong part of the register. Thus we always end
    up with zeros. Which is wrong, among others, because we detect failures
    via status.response_code.

    Signed-off-by: Halil Pasic <pasic@linux.ibm.com>
    Reported-by: Harald Freudenberger <freude@linux.ibm.com>
    Fixes: 159491f3b509 ("s390/ap: rework assembler functions to use unions for in/out
register variables")
    Reviewed-by: Harald Freudenberger <freude@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c670b0e412d9c087c425b59032c8dd3de25232ae
Author: Halil Pasic <pasic@linux.ibm.com>
Date:   Thu Feb 9 00:00:23 2023 +0100

    s390/ap: fix status returned by ap_aqic()

    [ Upstream commit 394740d7645ea767795074287769dd26dbd4d782 ]

    There function ap_aqic() tries to grab the status from the
    wrong part of the register. Thus we always end up with
    zeros. Which is wrong, among others, because we detect
    failures via status.response_code.

    Signed-off-by: Halil Pasic <pasic@linux.ibm.com>
    Reported-by: Janosch Frank <frankja@linux.ibm.com>
    Fixes: 159491f3b509 ("s390/ap: rework assembler functions to use unions for in/out
register variables")
    Reviewed-by: Harald Freudenberger <freude@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fcff42d7df5574c4a679b8928a0bb4b6b2b7f240
Author: Halil Pasic <pasic@linux.ibm.com>
Date:   Thu Feb 9 00:00:22 2023 +0100

    s390: vfio-ap: tighten the NIB validity check
```

```
    [ Upstream commit a64a6d23874c574d30a9816124b2dc37467f3811 ]

    The NIB is architecturally invalid if the address designates a
    storage location that is not installed or if it is zero.

    Signed-off-by: Halil Pasic <pasic@linux.ibm.com>
    Reported-by: Janosch Frank <frankja@linux.ibm.com>
    Fixes: ec89b55e3bce ("s390: ap: implement PAPQ AQIC interception in kernel")
    Reviewed-by: Tony Krowiak <akrowiak@linux.ibm.com>
    Reviewed-by: Pierre Morel <pmorel@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01fc68c867041072ff36a02dfd263d1a9a710721
Author: Alex Elder <elder@linaro.org>
Date:   Wed Feb 8 14:56:45 2023 -0600

    net: ipa: generic command param fix

    [ Upstream commit 2df181f09c961377a55510a864216d48d787fe49 ]

    Starting at IPA v4.11, the GSI_GENERIC_COMMAND GSI register got a
    new PARAMS field.  The code that encodes a value into that field
    sets it unconditionally, which is wrong.

    We currently only provide 0 as the field's value, so this error has
    no real effect.  Still, it's a bug, so let's fix it.

    Fix an (unrelated) incorrect comment as well.  Fields in the
    ERROR_LOG GSI register actually *are* defined for IPA versions
    prior to v3.5.1.

    Fixes: fe68c43ce388 ("net: ipa: support enhanced channel flow control")
    Signed-off-by: Alex Elder <elder@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c4717782aabee1934602802f07b73c72c3e5f4c
Author: Zhengping Jiang <jiangzp@google.com>
Date:   Thu Feb 2 09:47:01 2023 -0800

    Bluetooth: hci_qca: get wakeup status from serdev device handle

    [ Upstream commit 03b0093f7b310493bc944a20f725228cfe0d3fea ]

    Bluetooth controller attached via the UART is handled by the serdev driver.
    Get the wakeup status from the device handle through serdev, instead of the
    parent path.

    Fixes: c1a74160eaf1 ("Bluetooth: hci_qca: Add device_may_wakeup support")
    Signed-off-by: Zhengping Jiang <jiangzp@google.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 994e3e18908f5c4a12d07b44018e6aa85f071048
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Feb 1 14:01:11 2023 -0800

    Bluetooth: L2CAP: Fix potential user-after-free

    [ Upstream commit df5703348813235874d851934e957c3723d71644 ]

    This fixes all instances of which requires to allocate a buffer calling
    alloc_skb which may release the chan lock and reacquire later which
    makes it possible that the chan is disconnected in the meantime.
```

```
    Fixes: a6a5568c03c4 ("Bluetooth: Lock the L2CAP channel when sending")
    Reported-by: Alexander Coffin <alex.coffin@matician.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6ef26a77e5b2ad93b9452555a1453c1783cda0a
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Jan 27 16:51:54 2023 -0800

    Bluetooth: hci_conn: Refactor hci_bind_bis() since it always succeeds

    [ Upstream commit a00a29b0eeea6caaaf9edc3dd284f81b072ee343 ]

    The compiler thinks "conn" might be NULL after a call to hci_bind_bis(),
    which cannot happen. Avoid any confusion by just making it not return a
    value since it cannot fail. Fixes the warnings seen with GCC 13:

    In function 'arch_atomic_dec_and_test',
        inlined from 'atomic_dec_and_test' at ../include/linux/atomic/atomic-
instrumented.h:576:9,
        inlined from 'hci_conn_drop' at ../include/net/bluetooth/hci_core.h:1391:6,
        inlined from 'hci_connect_bis' at ../net/bluetooth/hci_conn.c:2124:3:
    ../arch/x86/include/asm/rmwcc.h:37:9: warning: array subscript 0 is outside array bounds
of 'atomic_t[0]' [-Warray-bounds=]
       37 |         asm volatile (fullop CC_SET(cc) \
          |         ^~~
    ...
    In function 'hci_connect_bis':
    cc1: note: source object is likely at address zero

    Fixes: eca0ae4aea66 ("Bluetooth: Add initial implementation of BIS connections")
    Cc: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Cc: Marcel Holtmann <marcel@holtmann.org>
    Cc: Johan Hedberg <johan.hedberg@gmail.com>
    Cc: "David S. Miller" <davem@davemloft.net>
    Cc: Eric Dumazet <edumazet@google.com>
    Cc: Jakub Kicinski <kuba@kernel.org>
    Cc: Paolo Abeni <pabeni@redhat.com>
    Cc: linux-bluetooth@vger.kernel.org
    Cc: netdev@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab05ae4ab831f64bbc427592c86f599ed9c4324f
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Wed Feb 8 10:26:54 2023 +0100

    cpufreq: davinci: Fix clk use after free

    [ Upstream commit 5d8f384a9b4fc50f6a18405f1c08e5a87a77b5b3 ]

    The remove function first frees the clks and only then calls
    cpufreq_unregister_driver(). If one of the cpufreq callbacks is called
    just before cpufreq_unregister_driver() is run, the freed clks might be
    used.

    Fixes: 6601b8030de3 ("davinci: add generic CPUFreq driver for DaVinci")
    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Acked-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 902d709f0db661acbb93e194faada5f180ed1334
Author: Qi Zheng <zhengqi.arch@bytedance.com>
Date:   Wed Feb 8 12:00:37 2023 +0800

    OPP: fix error checking in opp_migrate_dentry()

    [ Upstream commit eca4c0eea53432ec4b711b2a8ad282cbad231b4f ]

    Since commit ff9fb72bc077 ("debugfs: return error values,
    not NULL") changed return value of debugfs_rename() in
    error cases from %NULL to %ERR_PTR(-ERROR), we should
    also check error values instead of NULL.

    Fixes: ff9fb72bc077 ("debugfs: return error values, not NULL")
    Signed-off-by: Qi Zheng <zhengqi.arch@bytedance.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 002189a1e23faec05e68e00290d651938dac4bc8
Author: David Howells <dhowells@redhat.com>
Date:   Tue Nov 29 12:37:37 2022 +0000

    rxrpc: Fix overwaking on call poking

    [ Upstream commit a33395ab85b9b9cff83948a03a1d6d96347935d8 ]

    If an rxrpc call is given a poke, it will get woken up unconditionally,
    even if there's already a poke pending (for which there will have been a
    wake) or if the call refcount has gone to 0.

    Fix this by only waking the call if it is still referenced and if it
    doesn't already have a poke pending.

    Fixes: 15f661dc95da ("rxrpc: Implement a mechanism to send an event notification to a
call")
    Signed-off-by: David Howells <dhowells@redhat.com>
    cc: Marc Dionne <marc.dionne@auristor.com>
    cc: linux-afs@lists.infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fce60a29cc0cf888687e2686538a23d1a0db0468
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sat Feb 4 17:39:22 2023 +0000

    tap: tap_open(): correctly initialize socket uid

    [ Upstream commit 66b2c338adce580dfce2199591e65e2bab889cff ]

    sock_init_data() assumes that the `struct socket` passed in input is
    contained in a `struct socket_alloc` allocated with sock_alloc().
    However, tap_open() passes a `struct socket` embedded in a `struct
    tap_queue` allocated with sk_alloc().
    This causes a type confusion when issuing a container_of() with
    SOCK_INODE() in sock_init_data() which results in assigning a wrong
    sk_uid to the `struct sock` in input.
    On default configuration, the type confused field overlaps with
    padding bytes between `int vnet_hdr_sz` and `struct tap_dev __rcu
    *tap` in `struct tap_queue`, which makes the uid of all tap sockets 0,
    i.e., the root one.
    Fix the assignment by using sock_init_data_uid().

    Fixes: 86741ec25462 ("net: core: Add a UID field to struct sock.")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
```

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

commit 4aa4b4b3b3e9551c4de2bf2987247c28805fb8f6
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sat Feb 4 17:39:21 2023 +0000

    tun: tun_chr_open(): correctly initialize socket uid

    [ Upstream commit a096ccca6e503a5c575717ff8a36ace27510ab0a ]

    sock_init_data() assumes that the `struct socket` passed in input is
    contained in a `struct socket_alloc` allocated with sock_alloc().
    However, tun_chr_open() passes a `struct socket` embedded in a `struct
    tun_file` allocated with sk_alloc().
    This causes a type confusion when issuing a container_of() with
    SOCK_INODE() in sock_init_data() which results in assigning a wrong
    sk_uid to the `struct sock` in input.
    On default configuration, the type confused field overlaps with the
    high 4 bytes of `struct tun_struct __rcu *tun` of `struct tun_file`,
    NULL at the time of call, which makes the uid of all tun sockets 0,
    i.e., the root one.
    Fix the assignment by using sock_init_data_uid().

    Fixes: 86741ec25462 ("net: core: Add a UID field to struct sock.")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a0968828a8b7234c1be0b599489de14fa117cdf
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sat Feb 4 17:39:20 2023 +0000

    net: add sock_init_data_uid()

    [ Upstream commit 584f3742890e966d2f0a1f3c418c9ead70b2d99e ]

    Add sock_init_data_uid() to explicitly initialize the socket uid.
    To initialise the socket uid, sock_init_data() assumes a the struct
    socket* sock is always embedded in a struct socket_alloc, used to
    access the corresponding inode uid. This may not be true.
    Examples are sockets created in tun_chr_open() and tap_open().

    Fixes: 86741ec25462 ("net: core: Add a UID field to struct sock.")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b0543d6e2d229f814e207e1325663892fe73209
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Mon Jan 23 12:49:47 2023 +0100

    s390/boot: fix mem_detect extended area allocation

    [ Upstream commit 22476f47b6b7fb7d066c71f67ebc11892adb0849 ]

    Allocation of mem_detect extended area was not considered neither
    in commit 9641b8cc733f ("s390/ipl: read IPL report at early boot")
    nor in commit b2d24b97b2a9 ("s390/kernel: add support for kernel address
    space layout randomization (KASLR)"). As a result mem_detect extended
    theoretically may overlap with ipl report or randomized kernel image
    position. But as mem_detect code will allocate extended area only
    upon exceeding 255 online regions (which should alternate with offline
    memory regions) it is not seen in practice.
```

```
    To make sure mem_detect extended area does not overlap with ipl report
    or randomized kernel position extend usage of "safe_addr". Make initrd
    handling and mem_detect extended area allocation code move it further
    right and make KASLR takes in into consideration as well.

    Fixes: 9641b8cc733f ("s390/ipl: read IPL report at early boot")
    Fixes: b2d24b97b2a9 ("s390/kernel: add support for kernel address space layout
randomization (KASLR)")
    Reviewed-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 520cefabe936892de1abb79a46d0ded3765bc91e
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Fri Jan 27 14:57:43 2023 +0100

    s390/mem_detect: rely on diag260() if sclp_early_get_memsize() fails

    [ Upstream commit eb33f9eb304a4c18beb5ba6362eaa5c4beaf40d8 ]

    In case sclp_early_get_memsize() fails but diag260() succeeds make sure
    some sane value is returned. This error scenario is highly unlikely,
    but this change makes system able to boot in such case.

    Suggested-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Reviewed-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Stable-dep-of: 22476f47b6b7 ("s390/boot: fix mem_detect extended area allocation")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 769058254fa219f1d35aca24c88354d00ff8f564
Author: Alexander Gordeev <agordeev@linux.ibm.com>
Date:   Thu May 5 16:54:54 2022 +0200

    s390/boot: cleanup decompressor header files

    [ Upstream commit 9c3205b2b062420c26b33924b910880889acf832 ]

    Move declarations to appropriate header files. Instead of cryptic
    casting directly assign struct vmlinux_info type to _vmlinux_info
    linker script variable - wich it actually is.

    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Stable-dep-of: 22476f47b6b7 ("s390/boot: fix mem_detect extended area allocation")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3c47c31293f7540edab266657980a00b5c199a0
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Sat Jan 28 17:35:12 2023 +0100

    s390/vmem: fix empty page tables cleanup under KASAN

    [ Upstream commit 108303b0a2d27cb14eed565e33e64ad9eefe5d7e ]

    Commit b9ff81003cf1 ("s390/vmem: cleanup empty page tables") introduced
    empty page tables cleanup in vmem code, but when the kernel is built
    with KASAN enabled the code has no effect due to wrong KASAN shadow
    memory intersection condition, which effectively ignores any memory
    range below KASAN shadow. Fix intersection condition to make code
    work as anticipated.
```

```
    Fixes: b9ff81003cf1 ("s390/vmem: cleanup empty page tables")
    Reviewed-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4bca757d397bfa2ab805af529f871a2eefb23138
Author: Vasily Gorbik <gor@linux.ibm.com>
Date:   Fri Jan 27 14:03:07 2023 +0100

    s390/mem_detect: fix detect_memory() error handling

    [ Upstream commit 3400c35a4090704e6c465449616ab7e67a9209e7 ]

    Currently if for some reason sclp_early_read_info() fails,
    sclp_early_get_memsize() will not set max_physmem_end and it
    will stay uninitialized. Any garbage value other than 0 will lead
    to detect_memory() taking wrong path or returning a garbage value
    as max_physmem_end. To avoid that simply initialize max_physmem_end.

    Fixes: 73045a08cf55 ("s390: unify identity mapping limits handling")
    Reported-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Reviewed-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ae40c20f1519e1767ba01609abc7e8d6485fc0c
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 12:56:10 2023 +0400

    irqchip/ti-sci: Fix refcount leak in ti_sci_intr_irq_domain_probe

    [ Upstream commit 02298b7bae12936ca313975b02e7f98b06670d37 ]

    of_irq_find_parent() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: cd844b0715ce ("irqchip/ti-sci-intr: Add support for Interrupt Router driver")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230102085611.3955984-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 91e149b201bdb68c77011d50d011e47fadbcc8bd
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 12:42:08 2023 +0400

    irqchip/irq-mvebu-gicp: Fix refcount leak in mvebu_gicp_probe

    [ Upstream commit 9419e700021a393f67be36abd0c4f3acc6139041 ]

    of_irq_find_parent() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: a68a63cb4dfc ("irqchip/irq-mvebu-gicp: Add new driver for Marvell GICP")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230102084208.3951758-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5fbf2cc39b62a4afe44f3d42ee3dcf8f012c1926
```

```
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 12:28:10 2023 +0400

    irqchip/alpine-msi: Fix refcount leak in alpine_msix_init_domains

    [ Upstream commit 071d068b89e95d1b078aa6bbcb9d0961b77d6aa1 ]

    of_irq_find_parent() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: e6b78f2c3e14 ("irqchip: Add the Alpine MSIX interrupt controller")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230102082811.3947760-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c32fb16331f612e66a7fa8930164e0dc15725b72
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 16:13:18 2023 +0400

    irqchip: Fix refcount leak in platform_irqchip_probe

    [ Upstream commit 6caa5a2b78f5f53c433d3a3781e53325da22f0ac ]

    of_irq_find_parent() returns a node pointer with refcount incremented,
    We should use of_node_put() on it when not needed anymore.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: f8410e626569 ("irqchip: Add IRQCHIP_PLATFORM_DRIVER_BEGIN/END and IRQCHIP_MATCH
helper macros")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230102121318.3990586-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7471c6a987c1ddeb5e3f4f40bb1a441a6ff5f612
Author: Jack Morgenstein <jackm@nvidia.com>
Date:   Wed Jan 18 19:57:04 2023 +0200

    net/mlx5: Enhance debug print in page allocation failure

    [ Upstream commit 7eef93003e5d20e1a6a6e59e12d914b5431cbda2 ]

    Provide more details to aid debugging.

    Fixes: bf0bf77f6519 ("mlx5: Support communicating arbitrary host page size to firmware")
    Signed-off-by: Eran Ben Elisha <eranbe@nvidia.com>
    Signed-off-by: Majd Dibbiny <majd@nvidia.com>
    Signed-off-by: Jack Morgenstein <jackm@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0765b5b4719f0435bb019370b317d2fb8138eb34
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Tue Jan 17 01:16:07 2023 +0100

    wifi: mt76: mt7996: rely on mt76_connac2_mac_tx_rate_val

    [ Upstream commit 0b8e2d69467f78a7c9d87b452220e87012435e33 ]

    In order to fix a possible NULL pointer dereference in
    mt7996_mac_write_txwi() of vif pointer, export
    mt76_connac2_mac_tx_rate_val utility routine and reuse it
    in mt7996 driver.
```

```
    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59c2b0aec7a5aa350aa3edfb4363b19348c2269f
Author: Aaron Ma <aaron.ma@canonical.com>
Date:   Mon Jan 16 23:22:35 2023 +0800

    wifi: mt76: mt7921: fix error code of return in mt7921_acpi_read

    [ Upstream commit 888d89034f9eaeab9b5b75f13dbe35376c7dd471 ]

    Kernel NULL pointer dereference when ACPI SAR table isn't implemented well.
    Fix the error code of return to mark the ACPI SAR table as invalid.

    [    5.077128] mt7921e 0000:06:00.0: sar cnt = 0
    [    5.077381] BUG: kernel NULL pointer dereference, address:
    0000000000000004
    [    5.077630] #PF: supervisor read access in kernel mode
    [    5.077883] #PF: error_code(0x0000) - not-present page
    [    5.078138] PGD 0 P4D 0
    [    5.078398] Oops: 0000 [#1] PREEMPT SMP NOPTI
    [    5.079202] RIP: 0010:mt7921_init_acpi_sar+0x106/0x220
    [mt7921_common]
    ...
    [    5.080786] Call Trace:
    [    5.080786]  <TASK>
    [    5.080786]  mt7921_register_device+0x37d/0x490 [mt7921_common]
    [    5.080786]  mt7921_pci_probe.part.0+0x2ee/0x310 [mt7921e]
    [    5.080786]  mt7921_pci_probe+0x52/0x70 [mt7921e]
    [    5.080786]  local_pci_probe+0x47/0x90
    [    5.080786]  pci_call_probe+0x55/0x190
    [    5.080786]  pci_device_probe+0x84/0x120

    Fixes: f965333e491e ("mt76: mt7921: introduce ACPI SAR support")
    Signed-off-by: Aaron Ma <aaron.ma@canonical.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e6fffb4ed9e42f0f95a69598ff72b9c0c099b1a
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Sun Jan 15 12:44:57 2023 +0800

    wifi: mt76: add memory barrier to SDIO queue kick

    [ Upstream commit 5f54237ad798f41cb6a503271aa9ca47188cfb9b ]

    Ensure the entry has been fully updated before SDIO bus worker access
    it. This patch would fix potential memory risk in both mt7663s and
    mt7921s.

    Fixes: 764dee47e2c1 ("mt76: sdio: move common code in mt76_sdio module")
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c6485d749fb5514da344cd7a4f322ce2e4788dd7
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Jan 7 05:40:42 2023 +0800

    wifi: mt76: mt7915: fix WED TxS reporting
```

```
    [ Upstream commit 0d7084e209a9e2c924cb0d6e7f1f978db2a54127 ]

    The previous commit forgot to remove a leftover check.

    Fixes: 43eaa3689507 ("wifi: mt76: add PPDU based TxS support for WED device")
    Reported-By: Sujuan Chen <sujuan.chen@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c3f4a3feb31a0155b0693b44e1cfc37e42df3d5
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Tue Jan 3 14:56:19 2023 +0100

    wifi: mt76: fix switch default case in mt7996_reverse_frag0_hdr_trans

    [ Upstream commit aed8d9b506c22848add261ffad8fdabc57f21964 ]

    The switch default case for frame_contorl is invalid. Return -EINVAL
    error for it.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 604990fee0a6d608a6cca179ae474f2a1c6add8a
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun Jan 1 12:50:05 2023 +0100

    wifi: mt76: dma: fix memory leak running mt76_dma_tx_cleanup

    [ Upstream commit 3f7dda36e0b6dfa2cd26191f754ba061ab8191f2 ]

    Fix device unregister memory leak and alway cleanup all configured
    rx queues in mt76_dma_tx_cleanup routine.

    Fixes: 52546e27787e ("wifi: mt76: add WED RX support to dma queue alloc")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b539d35e13e5d6b3dca76271261106b2356aa64c
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun Jan 1 12:43:52 2023 +0100

    wifi: mt76: mt7996: fix memory leak in mt7996_mcu_exit

    [ Upstream commit da5b4d93e141b52c5a71d0c41a042d1bcaf70d2e ]

    Always purge mcu skb queues in mt7996_mcu_exit routine even if
    mt7996_firmware_state fails.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3095fe1e1d3198e62a3c7116c4cf7c734871475f
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun Jan 1 12:42:04 2023 +0100
```

```
    wifi: mt76: mt7915: fix memory leak in mt7915_mcu_exit

    [ Upstream commit 49bd78282e79ad177d14f37f4049f0605bf92dad ]

    Always purge mcu skb queues in mt7915_mcu_exit routine even if
    mt7915_firmware_state fails.

    Fixes: e57b7901469f ("mt76: add mac80211 driver for MT7915 PCIe-based chipsets")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a0559fc28ba70c7a0ecdb3e785a0ac1b843339f8
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Sat Dec 31 17:01:10 2022 +0800

    wifi: mt76: mt7921: fix invalid remain_on_channel duration

    [ Upstream commit c36457a8f34d1e0fc55fbbd7b8b2d716af3f6289 ]

    round_jiffies_up() may change the duration reported from chip. We should
    take the real timeout for current channel privilege.

    Fixes: 034ae28b56f1 ("wifi: mt76: mt7921: introduce remain_on_channel support")
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba9dfddcf198eaeffc1e8bae0f739a69dff451a2
Author: Shayne Chen <shayne.chen@mediatek.com>
Date:   Fri Dec 23 18:13:48 2022 +0800

    wifi: mt76: connac: fix POWER_CTRL command name typo

    [ Upstream commit 0d82fc956edb67b5755cc64ac6b9aee79cfbbff0 ]

    Fix typo MCU_UNI_CMD_POWER_CREL to MCU_UNI_CMD_POWER_CTRL.

    Fixes: 779d34de055e ("wifi: mt76: connac: add more unified command IDs")
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0add6bb83c7fe946888180d7930ba8141113b6af
Author: Shayne Chen <shayne.chen@mediatek.com>
Date:   Fri Dec 23 18:13:46 2022 +0800

    wifi: mt76: mt7996: update register for CFEND_RATE

    [ Upstream commit 793445cf812506375cbe4c59d0fb9f648f716e88 ]

    In newer chipsets, CFEND_RATE setting has been moved to different hw
    module.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec34f5f7745ce82406d959b67914bf3bd7a7e196
Author: Shayne Chen <shayne.chen@mediatek.com>
Date:   Fri Dec 23 18:13:45 2022 +0800
```

```
    wifi: mt76: mt7996: fix chainmask calculation in mt7996_set_antenna()

    [ Upstream commit eb1fdb9f5a2280de6820624cd02e0863babab683 ]

    Fix per-band chainmask when restoring from the dev chainmask.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fbe5e644556d666298cef04ef221956f206e2eb3
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Fri Dec 16 19:04:38 2022 +0800

    wifi: mt76: mt7921: fix channel switch fail in monitor mode

    [ Upstream commit 914189af23b83c9a83a0137a3a40f17de7d2c618 ]

    When the chanctx enabled, IEEE80211_CONF_CHANGE_CHANNEL in .config()
    would not be called anymore. mt76 cannot control RF channel for monitor.
    Add monitor type set-channel flow in .change_chanctx().

    Fixes: 41ac53c899bd ("wifi: mt76: mt7921: introduce chanctx support")
    Reviewed-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06c153e5f91ccf0af5af85d1dfa5310a42487653
Author: Howard Hsu <howard-yh.hsu@mediatek.com>
Date:   Wed Dec 14 10:28:08 2022 +0800

    wifi: mt76: mt7915: rework mt7915_thermal_temp_store()

    [ Upstream commit ecaccdae7a7e8f76eeb6544295ca0593c2f65a33 ]

    Call mt7915_mcu_set_thermal_protect() through
    mt7915_thermal_temp_store() to update firmware trigger/restore temp
    directly.

    Fixes: 02ee68b95d81 ("mt76: mt7915: add control knobs for thermal throttling")
    Reviewed-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Howard Hsu <howard-yh.hsu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bebfd3d514b479943e16114fc6efabf62f9619a3
Author: Howard Hsu <howard-yh.hsu@mediatek.com>
Date:   Wed Dec 14 10:28:07 2022 +0800

    wifi: mt76: mt7915: rework mt7915_mcu_set_thermal_throttling

    [ Upstream commit 5ad42d19f6596e54b091827c397fdb7c091d45f7 ]

    This patch includes 2 changes:
    1. Firmware expects to disable thermal protect first before
    reconfiguring.
    2. Separate setting thermal_protect and setting thermal_tx_duty into
    different functions. These two firmware commands do not need to send
    together.

    Fixes: 34b877d972be ("mt76: mt7915: add thermal cooling device support")
    Reviewed-by: Ryder Lee <ryder.lee@mediatek.com>
```

```
    Signed-off-by: Howard Hsu <howard-yh.hsu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df822d9a700cc69b9150dc610ec6e477d9c80968
Author: Howard Hsu <howard-yh.hsu@mediatek.com>
Date:   Wed Dec 14 10:28:06 2022 +0800

    wifi: mt76: mt7915: call mt7915_mcu_set_thermal_throttling() only after init_work

    [ Upstream commit 7d12b38ab6f6b77198cd3a66db19587bbdd3308c ]

    Enable thermal management by default shall not be executed during mcu
    init. This causes thermal configuration being reset to the firmware
    default settings.

    Fixes: 0063b86c9120 ("mt76: mt7915e: Enable thermal management by default")
    Reviewed-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Howard Hsu <howard-yh.hsu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edf45a6228f840fe06da7af4e69fbed1f6ba75c0
Author: Felix Fietkau <nbd@nbd.name>
Date:   Wed Dec 21 11:47:47 2022 +0100

    wifi: mt76: mt7921: fix deadlock in mt7921_abort_roc

    [ Upstream commit de19b9c83654e323d83f839a550ca4af37fea15b ]

    When mt7921_abort_roc is called with dev->mutex held, it can deadlock while
    calling cancel_work_sync(&phy->roc_work), because the work function could
    be waiting to acquire the mutex.

    Fix this by flushing the work before taking the mutex

    Reported-by: Mikhail Gavrilov <mikhail.v.gavrilov@gmail.com>
    Fixes: 034ae28b56f1 ("wifi: mt76: mt7921: introduce remain_on_channel support")
    Fixes: 41ac53c899bd ("wifi: mt76: mt7921: introduce chanctx support")
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 997c7474b2d11c8894fb1156cc52c6c0c21e1e8f
Author: Tonghao Zhang <tong@infragraf.org>
Date:   Thu Feb 2 21:17:01 2023 +0800

    bpftool: profile online CPUs instead of possible

    [ Upstream commit 377c16fa3f3c60d21e4b05314c8be034ce37f2eb ]

    The number of online cpu may be not equal to possible cpu.
    "bpftool prog profile" can not create pmu event on possible
    but on online cpu.

    $ dmidecode -s system-product-name
    PowerEdge R620
    $ cat /sys/devices/system/cpu/possible
    0-47
    $ cat /sys/devices/system/cpu/online
    0-31

    Disable cpu dynamically:
    $ echo 0 > /sys/devices/system/cpu/cpuX/online

    If one cpu is offline, perf_event_open will return ENODEV.
```

```
    To fix this issue:
    * check value returned and skip offline cpu.
    * close pmu_fd immediately on error path, avoid fd leaking.

    Fixes: 47c09d6a9f67 ("bpftool: Introduce "prog profile" command")
    Signed-off-by: Tonghao Zhang <tong@infragraf.org>
    Cc: Quentin Monnet <quentin@isovalent.com>
    Cc: Alexei Starovoitov <ast@kernel.org>
    Cc: Daniel Borkmann <daniel@iogearbox.net>
    Cc: Andrii Nakryiko <andrii@kernel.org>
    Cc: Martin KaFai Lau <martin.lau@linux.dev>
    Cc: Song Liu <song@kernel.org>
    Cc: Yonghong Song <yhs@fb.com>
    Cc: John Fastabend <john.fastabend@gmail.com>
    Cc: KP Singh <kpsingh@kernel.org>
    Cc: Stanislav Fomichev <sdf@google.com>
    Cc: Hao Luo <haoluo@google.com>
    Cc: Jiri Olsa <jolsa@kernel.org>
    Acked-by: John Fastabend <john.fastabend@gmail.com>
    Link: https://lore.kernel.org/r/20230202131701.29519-1-tong@infragraf.org
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e48ea40782634a82026d592651d4c2e0edb92bae
Author: Tom Lendacky <thomas.lendacky@amd.com>
Date:   Mon Jan 23 16:53:08 2023 -0600

    crypto: ccp - Flush the SEV-ES TMR memory before giving it to firmware

    [ Upstream commit 46a334a98f585ef78d51d8f5736596887bdd7f54 ]

    Perform a cache flush on the SEV-ES TMR memory after allocation to prevent
    any possibility of the firmware encountering an error should dirty cache
    lines be present. Use clflush_cache_range() to flush the SEV-ES TMR memory.

    Fixes: 97f9ac3db661 ("crypto: ccp - Add support for SEV-ES to the PSP driver")
    Signed-off-by: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7c11ddbca1c37d28e804282a06f76ea16ed8ebc
Author: Ilya Leoshkevich <iii@linux.ibm.com>
Date:   Fri Feb 3 00:53:35 2023 +0100

    selftests/bpf: Initialize tc in xdp_synproxy

    [ Upstream commit 354bb4a0e0b6be8f55bacbe7f08c94b4741f5658 ]

    xdp_synproxy/xdp fails in CI with:

        Error: bpf_tc_hook_create: File exists

    The XDP version of the test should not be calling bpf_tc_hook_create();
    the reason it's happening anyway is that if we don't specify --tc on the
    command line, tc variable remains uninitialized.

    Fixes: 784d5dc0efc2 ("selftests/bpf: Add selftests for raw syncookie helpers in TC
mode")
    Reported-by: Alexei Starovoitov <ast@kernel.org>
    Reported-by: Joanne Koong <joannelkoong@gmail.com>
    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230202235335.3403781-1-iii@linux.ibm.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 3a40c5ca783d97aff4e5693e710e8d5f229a1fed
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Jan 23 19:56:07 2023 +0100

    can: rcar_canfd: Fix R-Car V3U GAFLCFG field accesses

    [ Upstream commit 9be8c5583581244d8a77e41afa16b8b0a5ddabc0 ]

    Each Global Acceptance Filter List Configuration Register (GAFLCFG)
    contains two fields, and stores the number of channel rules for one
    channel pair.

    As R-Car V3U and later can have more than 2 channels, the field
    selection should be based on the LSB (even or odd) of the channel
    number, instead of on the full channel number.

    Fixes: 45721c406dcf50d4 ("can: rcar_canfd: Add support for r8a779a0 SoC")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/all/36bcf0ffb96d6aaed970751f9546b901af638bcf.1674499048.git.geert+re
nesas@glider.be
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e753635ad25201c342ac0e990a3f6947d8c03e45
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Jan 23 19:56:06 2023 +0100

    can: rcar_canfd: Fix R-Car V3U CAN mode selection

    [ Upstream commit 0a016639ef92b28eb74ba4fed21bee1f1328513a ]

    When adding support for R-Car V3U, the Global FD Configuration register
    (CFDGFDCFG) and the Channel-specific CAN-FD Configuration Registers
    (CFDCmFDCFG) were mixed up.  Use the correct register, and apply the
    selected CAN mode to all available channels.

    Annotate the corresponding register bits, to make it clear they do
    not exist on older variants.

    Fixes: 45721c406dcf50d4 ("can: rcar_canfd: Add support for r8a779a0 SoC")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/all/388ddf312917eb9f6cc460a481f68402a876f9b5.1674499048.git.geert+re
nesas@glider.be
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 200a6ed20b431e485d650adb357e57769e743517
Author: Mark Brown <broonie@kernel.org>
Date:   Tue Jan 31 22:56:34 2023 +0000

    kselftest/arm64: Fix enumeration of systems without 128 bit SME

    [ Upstream commit 5f389238534ac8ca4ee3ab12eeb89d3984d303a1 ]

    The current signal handling tests for SME do not account for the fact that
    unlike SVE all SME vector lengths are optional so we can't guarantee that
    we will encounter the minimum possible VL, they will hang enumerating VLs
    on such systems. Abort enumeration when we find the lowest VL.

    Fixes: 4963aeb35a9e ("kselftest/arm64: signal: Add SME signal handling tests")
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20230131-arm64-kselftest-sig-sme-no-128-v1-1-
d47c13dc8e1e@kernel.org
```

```
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 25d931e7dd7097f71a4e6da791c0af59c493e4f8
Author: Gregory Greenman <gregory.greenman@intel.com>
Date:   Fri Jan 27 00:28:21 2023 +0200

    wifi: iwlwifi: mei: fix compilation errors in rfkill()

    [ Upstream commit 9cbd5a8abca904441e36861e3a92961bec41d13f ]

    The rfkill() callback was invoked with wrong parameters.
    It was missed since MEI is defined now as depending on BROKEN.
    Fix that.

    Fixes: d288067ede4b ("wifi: iwlwifi: mei: avoid blocking sap messages handling due to
rtnl lock")
    Fixes: 5aa7ce31bd84 ("wifi: iwlwifi: mei: make sure ownership confirmed message is
sent")
    Fixes: 95170a46b7dd ("wifi: iwlwifi: mei: don't send SAP commands if AMT is disabled")
    Link: https://lore.kernel.org/r/20230126222821.305122-2-gregory.greenman@intel.com
    Signed-off-by: Gregory Greenman <gregory.greenman@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d253003b34bb683380b46a802336c1e95a0479d
Author: Ilya Leoshkevich <iii@linux.ibm.com>
Date:   Sun Jan 29 20:04:55 2023 +0100

    s390/bpf: Add expoline to tail calls

    [ Upstream commit bb4ef8fc3d193ed8d5583fb47cbeff5d8fb8302f ]

    All the indirect jumps in the eBPF JIT already use expolines, except
    for the tail call one.

    Fixes: de5cb6eb514e ("s390: use expoline thunks in the BPF JIT")
    Signed-off-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230129190501.1624747-3-iii@linux.ibm.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b9ea8780077b32ac7180325a92418ba458fba3e
Author: Kees Cook <keescook@chromium.org>
Date:   Sun Nov 27 10:30:41 2022 -0800

    drm/nouveau/disp: Fix nvif_outp_acquire_dp() argument size

    [ Upstream commit 4076ea2419cf15bc1e1580f8b24ddf675fbdb02c ]

    Both Coverity and GCC with -Wstringop-overflow noticed that
    nvif_outp_acquire_dp() accidentally defined its second argument with 1
    additional element:

    drivers/gpu/drm/nouveau/dispnv50/disp.c: In function 'nv50_pior_atomic_enable':
    drivers/gpu/drm/nouveau/dispnv50/disp.c:1813:17: error: 'nvif_outp_acquire_dp' accessing
16 bytes in a region of size 15 [-Werror=stringop-overflow=]
     1813 |                 nvif_outp_acquire_dp(&nv_encoder->outp, nv_encoder->dp.dpcd, 0,
0, false, false);
          |
^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    drivers/gpu/drm/nouveau/dispnv50/disp.c:1813:17: note: referencing argument 2 of type
'u8[16]' {aka 'unsigned char[16]'}
    drivers/gpu/drm/nouveau/include/nvif/outp.h:24:5: note: in a call to function
'nvif_outp_acquire_dp'
       24 | int nvif_outp_acquire_dp(struct nvif_outp *, u8 dpcd[16],
```

```
            |     ^~~~~~~~~~~~~~~~~~~
```

    Avoid these warnings by defining the argument size using the matching
    define (DP_RECEIVER_CAP_SIZE, 15) instead of having it be a literal
    (and incorrect) value (16).

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527269 ("Memory - corruptions")
    Addresses-Coverity-ID: 1527268 ("Memory - corruptions")
    Link: https://lore.kernel.org/lkml/202211100848.FFBA2432@keescook/
    Link: https://lore.kernel.org/lkml/202211100848.F4C2819BB@keescook/
    Fixes: 813443721331 ("drm/nouveau/disp: move DP link config into acquire")
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Cc: Ben Skeggs <bskeggs@redhat.com>
    Cc: Karol Herbst <kherbst@redhat.com>
    Cc: David Airlie <airlied@gmail.com>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Cc: Dave Airlie <airlied@redhat.com>
    Cc: "Gustavo A. R. Silva" <gustavo@embeddedor.com>
    Cc: dri-devel@lists.freedesktop.org
    Cc: nouveau@lists.freedesktop.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20221127183036.never.139-kees@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ab47c887b46dcda08f84d50d293f4912a069d2c
Author: Hans de Goede <hdegoede@redhat.com>
Date: Fri Jan 20 12:45:14 2023 +0100

    leds: led-class: Add missing put_device() to led_put()

    [Upstream commit 445110941eb94709216363f9d807d2508e64abd7]

    led_put() is used to "undo" a successful of_led_get() call,
    of_led_get() uses class_find_device_by_of_node() which returns
    a reference to the device which must be free-ed with put_device()
    when the caller is done with it.

    Add a put_device() call to led_put() to free the reference returned
    by class_find_device_by_of_node().

    And also add a put_device() in the error-exit case of try_module_get()
    failing.

    Fixes: 699a8c7c4bd3 ("leds: Add of_led_get() and led_put()")
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Lee Jones <lee@kernel.org>
    Link: https://lore.kernel.org/r/20230120114524.408368-2-hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d5870848879291700fe6c5257dcb48aadd10425c
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date: Sun Jan 22 16:07:37 2023 +0800

    crypto: xts - Handle EBUSY correctly

    [Upstream commit 51c082514c2dedf2711c99d93c196cc4eedceb40]

    As it is xts only handles the special return value of EINPROGRESS,
    which means that in all other cases it will free data related to the
    request.

    However, as the caller of xts may specify MAY_BACKLOG, we also need

```
    to expect EBUSY and treat it in the same way.  Otherwise backlogged
    requests will trigger a use-after-free.

    Fixes: 8083b1bf8163 ("crypto: xts - add support for ciphertext stealing")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Acked-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a4350bd6ebc64d92041afb91f238651009e2883
Author: Daniel T. Lee <danieltimlee@gmail.com>
Date:   Wed Jan 25 19:04:40 2023 +0900

    selftests/bpf: Fix vmtest static compilation error

    [ Upstream commit 2514a31241e1e9067d379e0fbdb60e4bc2bf4659 ]

    As stated in README.rst, in order to resolve errors with linker errors,
    'LDLIBS=-static' should be used. Most problems will be solved by this
    option, but in the case of urandom_read, this won't fix the problem. So
    the Makefile is currently implemented to strip the 'static' option when
    compiling the urandom_read. However, stripping this static option isn't
    configured properly on $(LDLIBS) correctly, which is now causing errors
    on static compilation.

        # LDLIBS=-static ./vmtest.sh
        ld.lld: error: attempted static link of dynamic object liburandom_read.so
        clang: error: linker command failed with exit code 1 (use -v to see invocation)
        make: *** [Makefile:190: /linux/tools/testing/selftests/bpf/urandom_read] Error 1
        make: *** Waiting for unfinished jobs....

    This commit fixes this problem by configuring the strip with $(LDLIBS).

    Fixes: 68084a136420 ("selftests/bpf: Fix building bpf selftests statically")
    Signed-off-by: Daniel T. Lee <danieltimlee@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230125100440.21734-1-danieltimlee@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f3f2164f534920050652cfb39e2e8e30057a73b
Author: Siddharth Vadapalli <s-vadapalli@ti.com>
Date:   Fri Jan 20 12:37:31 2023 +0530

    net: ethernet: ti: am65-cpsw/cpts: Fix CPTS release action

    [ Upstream commit 4ad8766cd3982744e53f107f378d2c65b76ff9a8 ]

    The am65_cpts_release() function is registered as a devm_action in the
    am65_cpts_create() function in am65-cpts driver. When the am65-cpsw driver
    invokes am65_cpts_create(), am65_cpts_release() is added in the set of devm
    actions associated with the am65-cpsw driver's device.

    In the event of probe failure or probe deferral, the platform_drv_probe()
    function invokes dev_pm_domain_detach() which powers off the CPSW and the
    CPSW's CPTS hardware, both of which share the same power domain. Since the
    am65_cpts_disable() function invoked by the am65_cpts_release() function
    attempts to reset the CPTS hardware by writing to its registers, the CPTS
    hardware is assumed to be powered on at this point. However, the hardware
    is powered off before the devm actions are executed.

    Fix this by getting rid of the devm action for am65_cpts_release() and
    invoking it directly on the cleanup and exit paths.

    Fixes: f6bd59526ca5 ("net: ethernet: ti: introduce am654 common platform time sync
driver")
```

```
    Signed-off-by: Siddharth Vadapalli <s-vadapalli@ti.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Reviewed-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Reviewed-by: Roger Quadros <rogerq@kernel.org>
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e5f2a79b8181844420058a9d4acfa027031c668f
Author: Ashok Raj <ashok.raj@intel.com>
Date:   Mon Jan 9 07:35:52 2023 -0800

    x86/microcode: Adjust late loading result reporting message

    [ Upstream commit 6eab3abac7043226e5375e9ead0c7607ced6767b ]

    During late microcode loading, the "Reload completed" message is issued
    unconditionally, regardless of success or failure.

    Adjust the message to report the result of the update.

      [ bp: Massage. ]

    Fixes: 9bd681251b7c ("x86/microcode: Announce reload operation's completion")
    Suggested-by: Thomas Gleixner <tglx@linutronix.de>
    Signed-off-by: Ashok Raj <ashok.raj@intel.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Reviewed-by: Tony Luck <tony.luck@intel.com>
    Link: https://lore.kernel.org/lkml/874judpqqd.ffs@tglx/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70dd912e0cd0b0292e1d6a5d2e1afcaaba3c9489
Author: Ashok Raj <ashok.raj@intel.com>
Date:   Mon Jan 9 07:35:51 2023 -0800

    x86/microcode: Check CPU capabilities after late microcode update correctly

    [ Upstream commit c0dd9245aa9e25a697181f6085692272c9ec61bc ]

    The kernel caches each CPU's feature bits at boot in an x86_capability[]
    structure. However, the capabilities in the BSP's copy can be turned off
    as a result of certain command line parameters or configuration
    restrictions, for example the SGX bit. This can cause a mismatch when
    comparing the values before and after the microcode update.

    Another example is X86_FEATURE_SRBDS_CTRL which gets added only after
    microcode update:

    #  --- cpuid.before     2023-01-21 14:54:15.652000747 +0100
    #  +++ cpuid.after      2023-01-21 14:54:26.632001024 +0100
    #  @@ -10,7 +10,7 @@ CPU:
    #       0x00000004 0x04: eax=0x00000000 ebx=0x00000000 ecx=0x00000000 edx=0x00000000
    #       0x00000005 0x00: eax=0x00000040 ebx=0x00000040 ecx=0x00000003 edx=0x11142120
    #       0x00000006 0x00: eax=0x000027f7 ebx=0x00000002 ecx=0x00000001 edx=0x00000000
    #  -    0x00000007 0x00: eax=0x00000000 ebx=0x029c6fbf ecx=0x40000000 edx=0xbc002400
    #  +    0x00000007 0x00: eax=0x00000000 ebx=0x029c6fbf ecx=0x40000000 edx=0xbc002e00
                                                                                ^^^

    and which proves for a gazillionth time that late loading is a bad bad
    idea.

    microcode_check() is called after an update to report any previously
    cached CPUID bits which might have changed due to the update.

    Therefore, store the cached CPU caps before the update and compare them
    with the CPU caps after the microcode update has succeeded.
```

```
    Thus, the comparison is done between the CPUID *hardware* bits before
    and after the upgrade instead of using the cached, possibly runtime
    modified values in BSP's boot_cpu_data copy.

    As a result, false warnings about CPUID bits changes are avoided.

      [ bp:
            - Massage.
            - Add SRBDS_CTRL example.
            - Add kernel-doc.
            - Incorporate forgotten review feedback from dhansen.
            ]

    Fixes: 1008c52c09dc ("x86/CPU: Add a microcode loader callback")
    Signed-off-by: Ashok Raj <ashok.raj@intel.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230109153555.4986-3-ashok.raj@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc0215538d0a31132358f2f0fd1fd1af2496c834
Author: Ashok Raj <ashok.raj@intel.com>
Date:   Mon Jan 9 07:35:50 2023 -0800

    x86/microcode: Add a parameter to microcode_check() to store CPU capabilities

    [ Upstream commit ab31c74455c64e69342ddab21fd9426fcbfefde7 ]

    Add a parameter to store CPU capabilities before performing a microcode
    update so that CPU capabilities can be compared before and after update.

      [ bp: Massage. ]

    Signed-off-by: Ashok Raj <ashok.raj@intel.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230109153555.4986-2-ashok.raj@intel.com
    Stable-dep-of: c0dd9245aa9e ("x86/microcode: Check CPU capabilities after late microcode
update correctly")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c33007812a55612d9b2a7b85c8d04cefeeaf0d21
Author: Kumar Kartikeya Dwivedi <memxor@gmail.com>
Date:   Sat Jan 21 05:52:32 2023 +0530

    bpf: Fix partial dynptr stack slot reads/writes

    [ Upstream commit ef8fc7a07c0e161841779d6fe3f6acd5a05c547c ]

    Currently, while reads are disallowed for dynptr stack slots, writes are
    not. Reads don't work from both direct access and helpers, while writes
    do work in both cases, but have the effect of overwriting the slot_type.

    While this is fine, handling for a few edge cases is missing. Firstly,
    a user can overwrite the stack slots of dynptr partially.

    Consider the following layout:
    spi: [d][d][?]
          2  1  0

    First slot is at spi 2, second at spi 1.
    Now, do a write of 1 to 8 bytes for spi 1.

    This will essentially either write STACK_MISC for all slot_types or
    STACK_MISC and STACK_ZERO (in case of size < BPF_REG_SIZE partial write
    of zeroes). The end result is that slot is scrubbed.
```

```
    Now, the layout is:
    spi: [d][m][?]
          2  1  0

    Suppose if user initializes spi = 1 as dynptr.
    We get:
    spi: [d][d][d]
          2  1  0

    But this time, both spi 2 and spi 1 have first_slot = true.

    Now, when passing spi 2 to dynptr helper, it will consider it as
    initialized as it does not check whether second slot has first_slot ==
    false. And spi 1 should already work as normal.

    This effectively replaced size + offset of first dynptr, hence allowing
    invalid OOB reads and writes.

    Make a few changes to protect against this:
    When writing to PTR_TO_STACK using BPF insns, when we touch spi of a
    STACK_DYNPTR type, mark both first and second slot (regardless of which
    slot we touch) as STACK_INVALID. Reads are already prevented.

    Second, prevent writing to stack memory from helpers if the range may
    contain any STACK_DYNPTR slots. Reads are already prevented.

    For helpers, we cannot allow it to destroy dynptrs from the writes as
    depending on arguments, helper may take uninit_mem and dynptr both at
    the same time. This would mean that helper may write to uninit_mem
    before it reads the dynptr, which would be bad.

    PTR_TO_MEM: [?????dd]

    Depending on the code inside the helper, it may end up overwriting the
    dynptr contents first and then read those as the dynptr argument.

    Verifier would only simulate destruction when it does byte by byte
    access simulation in check_helper_call for meta.access_size, and
    fail to catch this case, as it happens after argument checks.

    The same would need to be done for any other non-trivial objects created
    on the stack in the future, such as bpf_list_head on stack, or
    bpf_rb_root on stack.

    A common misunderstanding in the current code is that MEM_UNINIT means
    writes, but note that writes may also be performed even without
    MEM_UNINIT in case of helpers, in that case the code after handling meta
    && meta->raw_mode will complain when it sees STACK_DYNPTR. So that
    invalid read case also covers writes to potential STACK_DYNPTR slots.
    The only loophole was in case of meta->raw_mode which simulated writes
    through instructions which could overwrite them.

    A future series sequenced after this will focus on the clean up of
    helper access checks and bugs around that.

    Fixes: 97e03f521050 ("bpf: Add verifier support for dynptrs")
    Signed-off-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230121002241.2113993-4-memxor@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 489b67f268ae0270a6c3f2b49144aeeb8eee301a
Author: Kumar Kartikeya Dwivedi <memxor@gmail.com>
Date:   Sat Jan 21 05:52:31 2023 +0530
```

```
    bpf: Fix missing var_off check for ARG_PTR_TO_DYNPTR

    [ Upstream commit 79168a669d8125453c8a271115f1ffd4294e61f6 ]

    Currently, the dynptr function is not checking the variable offset part
    of PTR_TO_STACK that it needs to check. The fixed offset is considered
    when computing the stack pointer index, but if the variable offset was
    not a constant (such that it could not be accumulated in reg->off), we
    will end up a discrepency where runtime pointer does not point to the
    actual stack slot we mark as STACK_DYNPTR.

    It is impossible to precisely track dynptr state when variable offset is
    not constant, hence, just like bpf_timer, kptr, bpf_spin_lock, etc.
    simply reject the case where reg->var_off is not constant. Then,
    consider both reg->off and reg->var_off.value when computing the stack
    pointer index.

    A new helper dynptr_get_spi is introduced to hide over these details
    since the dynptr needs to be located in multiple places outside the
    process_dynptr_func checks, hence once we know it's a PTR_TO_STACK, we
    need to enforce these checks in all places.

    Note that it is disallowed for unprivileged users to have a non-constant
    var_off, so this problem should only be possible to trigger from
    programs having CAP_PERFMON. However, its effects can vary.

    Without the fix, it is possible to replace the contents of the dynptr
    arbitrarily by making verifier mark different stack slots than actual
    location and then doing writes to the actual stack address of dynptr at
    runtime.

    Fixes: 97e03f521050 ("bpf: Add verifier support for dynptrs")
    Acked-by: Joanne Koong <joannelkoong@gmail.com>
    Signed-off-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230121002241.2113993-3-memxor@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 720d2504791a93becde81c335abcea2f42d066a7
Author: Kumar Kartikeya Dwivedi <memxor@gmail.com>
Date:   Sat Jan 21 05:52:30 2023 +0530

    bpf: Fix state pruning for STACK_DYNPTR stack slots

    [ Upstream commit d6fefa1105dacc8a742cdcf2f4bfb501c9e61349 ]

    The root of the problem is missing liveness marking for STACK_DYNPTR
    slots. This leads to all kinds of problems inside stacksafe.

    The verifier by default inside stacksafe ignores spilled_ptr in stack
    slots which do not have REG_LIVE_READ marks. Since this is being checked
    in the 'old' explored state, it must have already done clean_live_states
    for this old bpf_func_state. Hence, it won't be receiving any more
    liveness marks from to be explored insns (it has received REG_LIVE_DONE
    marking from liveness point of view).

    What this means is that verifier considers that it's safe to not compare
    the stack slot if was never read by children states. While liveness
    marks are usually propagated correctly following the parentage chain for
    spilled registers (SCALAR_VALUE and PTR_* types), the same is not the
    case for STACK_DYNPTR.

    clean_live_states hence simply rewrites these stack slots to the type
    STACK_INVALID since it sees no REG_LIVE_READ marks.
```

```
    The end result is that we will never see STACK_DYNPTR slots in explored
    state. Even if verifier was conservatively matching !REG_LIVE_READ
    slots, very next check continuing the stacksafe loop on seeing
    STACK_INVALID would again prevent further checks.

    Now as long as verifier stores an explored state which we can compare to
    when reaching a pruning point, we can abuse this bug to make verifier
    prune search for obviously unsafe paths using STACK_DYNPTR slots
    thinking they are never used hence safe.

    Doing this in unprivileged mode is a bit challenging. add_new_state is
    only set when seeing BPF_F_TEST_STATE_FREQ (which requires privileges)
    or when jmps_processed difference is >= 2 and insn_processed difference
    is >= 8. So coming up with the unprivileged case requires a little more
    work, but it is still totally possible. The test case being discussed
    below triggers the heuristic even in unprivileged mode.

    However, it no longer works since commit
    8addbfc7b308 ("bpf: Gate dynptr API behind CAP_BPF").

    Let's try to study the test step by step.

    Consider the following program (C style BPF ASM):

    0  r0 = 0;
    1  r6 = &ringbuf_map;
    3  r1 = r6;
    4  r2 = 8;
    5  r3 = 0;
    6  r4 = r10;
    7  r4 -= -16;
    8  call bpf_ringbuf_reserve_dynptr;
    9  if r0 == 0 goto pc+1;
    10 goto pc+1;
    11 *(r10 - 16) = 0xeB9F;
    12 r1 = r10;
    13 r1 -= -16;
    14 r2 = 0;
    15 call bpf_ringbuf_discard_dynptr;
    16 r0 = 0;
    17 exit;

    We know that insn 12 will be a pruning point, hence if we force
    add_new_state for it, it will first verify the following path as
    safe in straight line exploration:
    0 1 3 4 5 6 7 8 9 -> 10 -> (12) 13 14 15 16 17

    Then, when we arrive at insn 12 from the following path:
    0 1 3 4 5 6 7 8 9 -> 11 (12)

    We will find a state that has been verified as safe already at insn 12.
    Since register state is same at this point, regsafe will pass. Next, in
    stacksafe, for spi = 0 and spi = 1 (location of our dynptr) is skipped
    seeing !REG_LIVE_READ. The rest matches, so stacksafe returns true.
    Next, refsafe is also true as reference state is unchanged in both
    states.

    The states are considered equivalent and search is pruned.

    Hence, we are able to construct a dynptr with arbitrary contents and use
    the dynptr API to operate on this arbitrary pointer and arbitrary size +
    offset.

    To fix this, first define a mark_dynptr_read function that propagates
```

```
    liveness marks whenever a valid initialized dynptr is accessed by dynptr
    helpers. REG_LIVE_WRITTEN is marked whenever we initialize an
    uninitialized dynptr. This is done in mark_stack_slots_dynptr. It allows
    screening off mark_reg_read and not propagating marks upwards from that
    point.

    This ensures that we either set REG_LIVE_READ64 on both dynptr slots, or
    none, so clean_live_states either sets both slots to STACK_INVALID or
    none of them. This is the invariant the checks inside stacksafe rely on.

    Next, do a complete comparison of both stack slots whenever they have
    STACK_DYNPTR. Compare the dynptr type stored in the spilled_ptr, and
    also whether both form the same first_slot. Only then is the later path
    safe.

    Fixes: 97e03f521050 ("bpf: Add verifier support for dynptrs")
    Acked-by: Eduard Zingerman <eddyz87@gmail.com>
    Signed-off-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230121002241.2113993-2-memxor@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a029332fcfde366dceb78c79e0ff80af73ab00e
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Jan 3 20:57:26 2023 +0800

    powercap: fix possible name leak in powercap_register_zone()

    [ Upstream commit 1b6599f741a4525ca761ecde46e5885ff1e6ba58 ]

    In the error path after calling dev_set_name(), the device
    name is leaked. To fix this, calling dev_set_name() before
    device_register(), and call put_device() if it returns error.

    All the resources is released in powercap_release(), so it
    can return from powercap_register_zone() directly.

    Fixes: 75d2364ea0ca ("PowerCap: Add class driver")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9477db935eb690f697d9bcc4f608927841bc8b36
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Fri Jan 13 18:27:51 2023 +0800

    crypto: seqiv - Handle EBUSY correctly

    [ Upstream commit 32e62025e5e52fbe4812ef044759de7010b15dbc ]

    As it is seqiv only handles the special return value of EINPROGERSS,
    which means that in all other cases it will free data related to the
    request.

    However, as the caller of seqiv may specify MAY_BACKLOG, we also need
    to expect EBUSY and treat it in the same way.  Otherwise backlogged
    requests will trigger a use-after-free.

    Fixes: 0a270321dbf9 ("[CRYPTO] seqiv: Add Sequence Number IV Generator")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69c67d451fc19d88e54f7d97e8e7c093e08357e1
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Fri Jan 13 18:24:09 2023 +0800
```

```
    crypto: essiv - Handle EBUSY correctly

    [ Upstream commit b5a772adf45a32c68bef28e60621f12617161556 ]

    As it is essiv only handles the special return value of EINPROGERSS,
    which means that in all other cases it will free data related to the
    request.

    However, as the caller of essiv may specify MAY_BACKLOG, we also need
    to expect EBUSY and treat it in the same way.  Otherwise backlogged
    requests will trigger a use-after-free.

    Fixes: be1eb7f78aa8 ("crypto: essiv - create wrapper template...")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Acked-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3fea4acff6583886cde1b87f8aaa14e2bccb8bbf
Author: Koba Ko <koba.taiwan@gmail.com>
Date:   Mon Jan 9 10:15:02 2023 +0800

    crypto: ccp - Failure on re-initialization due to duplicate sysfs filename

    [ Upstream commit 299bf602b3f92f1456aef59c6413591fb02e762a ]

    The following warning appears during the CCP module re-initialization:

    [  140.965403] sysfs: cannot create duplicate filename
    '/devices/pci0000:00/0000:00:07.1/0000:03:00.2/dma/dma0chan0'
    [  140.975736] CPU: 0 PID: 388 Comm: kworker/0:2 Kdump: loaded Not
    tainted 6.2.0-0.rc2.18.eln124.x86_64 #1
    [  140.985185] Hardware name: HPE ProLiant DL325 Gen10/ProLiant DL325
    Gen10, BIOS A41 07/17/2020
    [  140.993761] Workqueue: events work_for_cpu_fn
    [  140.998151] Call Trace:
    [  141.000613]  <TASK>
    [  141.002726]  dump_stack_lvl+0x33/0x46
    [  141.006415]  sysfs_warn_dup.cold+0x17/0x23
    [  141.010542]  sysfs_create_dir_ns+0xba/0xd0
    [  141.014670]  kobject_add_internal+0xba/0x260
    [  141.018970]  kobject_add+0x81/0xb0
    [  141.022395]  device_add+0xdc/0x7e0
    [  141.025822]  ? complete_all+0x20/0x90
    [  141.029510]  __dma_async_device_channel_register+0xc9/0x130
    [  141.035119]  dma_async_device_register+0x19e/0x3b0
    [  141.039943]  ccp_dmaengine_register+0x334/0x3f0 [ccp]
    [  141.045042]  ccp5_init+0x662/0x6a0 [ccp]
    [  141.049000]  ? devm_kmalloc+0x40/0xd0
    [  141.052688]  ccp_dev_init+0xbb/0xf0 [ccp]
    [  141.056732]  ? __pci_set_master+0x56/0xd0
    [  141.060768]  sp_init+0x70/0x90 [ccp]
    [  141.064377]  sp_pci_probe+0x186/0x1b0 [ccp]
    [  141.068596]  local_pci_probe+0x41/0x80
    [  141.072374]  work_for_cpu_fn+0x16/0x20
    [  141.076145]  process_one_work+0x1c8/0x380
    [  141.080181]  worker_thread+0x1ab/0x380
    [  141.083953]  ? __pfx_worker_thread+0x10/0x10
    [  141.088250]  kthread+0xda/0x100
    [  141.091413]  ? __pfx_kthread+0x10/0x10
    [  141.095185]  ret_from_fork+0x2c/0x50
    [  141.098788]  </TASK>
    [  141.100996] kobject_add_internal failed for dma0chan0 with -EEXIST,
    don't try to register things with the same name in the same directory.
```

```
    [  141.113703] ccp 0000:03:00.2: ccp initialization failed

    The /dma/dma0chan0 sysfs file is not removed since dma_chan object
    has been released in ccp_dma_release() before releasing dma device.
    A correct procedure would be: release dma channels first => unregister
    dma device => release ccp dma object.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216888
    Fixes: 68dbe80f5b51 ("crypto: ccp - Release dma channels before dmaengine unrgister")
    Tested-by: Vladis Dronov <vdronov@redhat.com>
    Signed-off-by: Koba Ko <koba.ko@canonical.com>
    Reviewed-by: Vladis Dronov <vdronov@redhat.com>
    Acked-by: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f716bd5150547407e34a9b870be20b8a484d13b
Author: Tiezhu Yang <yangtiezhu@loongson.cn>
Date:   Wed Jan 18 15:56:44 2023 +0800

    selftests/bpf: Fix build errors if CONFIG_NF_CONNTRACK=m

    [ Upstream commit 92afc5329a5b23d876b215b783d200352d5aaea6 ]

    If CONFIG_NF_CONNTRACK=m, there are no definitions of NF_NAT_MANIP_SRC
    and NF_NAT_MANIP_DST in vmlinux.h, build test_bpf_nf.c failed.

    $ make -C tools/testing/selftests/bpf/

      CLNG-BPF [test_maps] test_bpf_nf.bpf.o
    progs/test_bpf_nf.c:160:42: error: use of undeclared identifier 'NF_NAT_MANIP_SRC'
                    bpf_ct_set_nat_info(ct, &saddr, sport, NF_NAT_MANIP_SRC);
                                                           ^
    progs/test_bpf_nf.c:163:42: error: use of undeclared identifier 'NF_NAT_MANIP_DST'
                    bpf_ct_set_nat_info(ct, &daddr, dport, NF_NAT_MANIP_DST);
                                                           ^
    2 errors generated.

    Copy the definitions in include/net/netfilter/nf_nat.h to test_bpf_nf.c,
    in order to avoid redefinitions if CONFIG_NF_CONNTRACK=y, rename them with
    ___local suffix. This is similar with commit 1058b6a78db2 ("selftests/bpf:
    Do not fail build if CONFIG_NF_CONNTRACK=m/n").

    Fixes: b06b45e82b59 ("selftests/bpf: add tests for bpf_ct_set_nat_info kfunc")
    Signed-off-by: Tiezhu Yang <yangtiezhu@loongson.cn>
    Acked-by: Jiri Olsa <jolsa@kernel.org>
    Tested-by: Jiri Olsa <jolsa@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Link: https://lore.kernel.org/r/1674028604-7113-1-git-send-email-yangtiezhu@loongson.cn
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8970be2c488b403941c9ca47bf89f4ebcb8389eb
Author: Armin Wolf <W_Armin@gmx.de>
Date:   Sat Jan 14 09:50:50 2023 +0100

    ACPI: battery: Fix missing NUL-termination with large strings

    [ Upstream commit f2ac14b5f197e4a2dec51e5ceaa56682ff1592bc ]

    When encountering a string bigger than the destination buffer (32 bytes),
    the string is not properly NUL-terminated, causing buffer overreads later.

    This for example happens on the Inspiron 3505, where the battery
    model name is larger than 32 bytes, which leads to sysfs showing
```

```
    the model name together with the serial number string (which is
    NUL-terminated and thus prevents worse).

    Fix this by using strscpy() which ensures that the result is
    always NUL-terminated.

    Fixes: 106449e870b3 ("ACPI: Battery: Allow extract string from integer")
    Signed-off-by: Armin Wolf <W_Armin@gmx.de>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07be332c87b2a70fcd9f6d5b394f0e9bd6e787b7
Author: Shivani Baranwal <quic_shivbara@quicinc.com>
Date:   Tue Dec 6 20:07:14 2022 +0530

    wifi: cfg80211: Fix extended KCK key length check in nl80211_set_rekey_data()

    [ Upstream commit df4969ca135b9b3b2c38c07514aaa775112ac835 ]

    The extended KCK key length check wrongly using the KEK key attribute
    for validation. Due to this GTK rekey offload is failing when the KCK
    key length is 24 bytes even though the driver advertising
    WIPHY_FLAG_SUPPORTS_EXT_KEK_KCK flag. Use correct attribute to fix the
    same.

    Fixes: 093a48d2aa4b ("cfg80211: support bigger kek/kck key length")
    Signed-off-by: Shivani Baranwal <quic_shivbara@quicinc.com>
    Signed-off-by: Veerendranath Jakkam <quic_vjakkam@quicinc.com>
    Link: https://lore.kernel.org/r/20221206143715.1802987-2-quic_vjakkam@quicinc.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64a78ec4f4579798d8e885aca9bdd707bca6b16b
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Mon Jan 2 12:11:42 2023 +0400

    wifi: ath11k: Fix memory leak in ath11k_peer_rx_frag_setup

    [ Upstream commit ed3f83b3459a67a3ab9d806490ac304b567b1c2d ]

    crypto_alloc_shash() allocates resources, which should be released by
    crypto_free_shash(). When ath11k_peer_find() fails, there has memory
    leak. Add missing crypto_free_shash() to fix this.

    Fixes: 243874c64c81 ("ath11k: handle RX fragments")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230102081142.3937570-1-linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 554048a72d7ecfdd58cc1bfb56e0a1864e64e82c
Author: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
Date:   Wed Jan 4 21:41:30 2023 +0900

    wifi: ath9k: Fix potential stack-out-of-bounds write in ath9k_wmi_rsp_callback()

    [ Upstream commit 8a2f35b9830692f7a616f2f627f943bc748af13a ]

    Fix a stack-out-of-bounds write that occurs in a WMI response callback
    function that is called after a timeout occurs in ath9k_wmi_cmd().
    The callback writes to wmi->cmd_rsp_buf, a stack-allocated buffer that
    could no longer be valid when a timeout occurs. Set wmi->last_seq_id to
    0 when a timeout occurred.
```

```
    Found by a modified version of syzkaller.

    BUG: KASAN: stack-out-of-bounds in ath9k_wmi_ctrl_rx
    Write of size 4
    Call Trace:
     memcpy
     ath9k_wmi_ctrl_rx
     ath9k_htc_rx_msg
     ath9k_hif_usb_reg_in_cb
     __usb_hcd_giveback_urb
     usb_hcd_giveback_urb
     dummy_timer
     call_timer_fn
     run_timer_softirq
     __do_softirq
     irq_exit_rcu
     sysvec_apic_timer_interrupt

    Fixes: fb9987d0f748 ("ath9k_htc: Support for AR9271 chipset.")
    Signed-off-by: Minsuk Kang <linuxlovemin@yonsei.ac.kr>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230104124130.10996-1-linuxlovemin@yonsei.ac.kr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c766e37fccd5a5c5059be7efcd9618bf8a2c17c3
Author: Fedor Pchelkin <pchelkin@ispras.ru>
Date:   Wed Jan 4 15:36:15 2023 +0300

    wifi: ath9k: hif_usb: clean up skbs if ath9k_hif_usb_rx_stream() fails

    [ Upstream commit 0af54343a76263a12dbae7fafb64eb47c4a6ad38 ]

    Syzkaller detected a memory leak of skbs in ath9k_hif_usb_rx_stream().
    While processing skbs in ath9k_hif_usb_rx_stream(), the already allocated
    skbs in skb_pool are not freed if ath9k_hif_usb_rx_stream() fails. If we
    have an incorrect pkt_len or pkt_tag, the input skb is considered invalid
    and dropped. All the associated packets already in skb_pool should be
    dropped and freed. Added a comment describing this issue.

    The patch also makes remain_skb NULL after being processed so that it
    cannot be referenced after potential free. The initialization of hif_dev
    fields which are associated with remain_skb (rx_remain_len,
    rx_transfer_len and rx_pad_len) is moved after a new remain_skb is
    allocated.

    Found by Linux Verification Center (linuxtesting.org) with Syzkaller.

    Fixes: 6ce708f54cc8 ("ath9k: Fix out-of-bound memcpy in ath9k_hif_usb_rx_stream")
    Fixes: 44b23b488d44 ("ath9k: hif_usb: Reduce indent 1 column")
    Reported-by: syzbot+e9632e3eb038d93d6bc6@syzkaller.appspotmail.com
    Signed-off-by: Fedor Pchelkin <pchelkin@ispras.ru>
    Signed-off-by: Alexey Khoroshilov <khoroshilov@ispras.ru>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230104123615.51511-1-pchelkin@ispras.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbfababb4f899fe1556eac195f9774b6fe675fb6
Author: Fedor Pchelkin <pchelkin@ispras.ru>
Date:   Wed Jan 4 15:35:46 2023 +0300

    wifi: ath9k: htc_hst: free skb in ath9k_htc_rx_msg() if there is no callback function

    [ Upstream commit 9b25e3985477ac3f02eca5fc1e0cc6850a3f7e69 ]
```

```
    It is stated that ath9k_htc_rx_msg() either frees the provided skb or
    passes its management to another callback function. However, the skb is
    not freed in case there is no another callback function, and Syzkaller was
    able to cause a memory leak. Also minor comment fix.

    Found by Linux Verification Center (linuxtesting.org) with Syzkaller.

    Fixes: fb9987d0f748 ("ath9k_htc: Support for AR9271 chipset.")
    Reported-by: syzbot+e008dccab31bd3647609@syzkaller.appspotmail.com
    Reported-by: syzbot+6692c72009680f7c4eb2@syzkaller.appspotmail.com
    Signed-off-by: Fedor Pchelkin <pchelkin@ispras.ru>
    Signed-off-by: Alexey Khoroshilov <khoroshilov@ispras.ru>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230104123546.51427-1-pchelkin@ispras.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb2debface6782c92dab14c718dbaf2c8c85b2ff
Author: Viorel Suman <viorel.suman@nxp.com>
Date:   Tue Jan 17 11:19:55 2023 +0200

    thermal/drivers/imx_sc_thermal: Fix the loop condition

    [ Upstream commit 4b26b7c9cdefdcb478047c30901d2379ef9e50fc ]

    The minimal resource ID is 0: IMX_SC_R_AP_0=0, so fix
    the loop condition. Aside of this - constify the array.

    Fixes: 31fd4b9db13b ("thermal/drivers/imx_sc: Rely on the platform data to get the
resource id")
    Signed-off-by: Viorel Suman <viorel.suman@nxp.com>
    Reviewed-by: Dong Aisheng <Aisheng.dong@nxp.com>
    Link: https://lore.kernel.org/r/20230117091956.61729-1-viorel.suman@oss.nxp.com
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef946133140382c6cdd01ec2fb47f767ac4f4618
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Sun Jan 8 22:13:24 2023 +0100

    wifi: rtw88: Use non-atomic sta iterator in rtw_ra_mask_info_update()

    [ Upstream commit 2931978cd74f4a643deeea5b211523001d95aa44 ]

    USB and (upcoming) SDIO support may sleep in the read/write handlers.
    Use non-atomic rtw_iterate_stas() in rtw_ra_mask_info_update() because
    the iterator function rtw_ra_mask_info_update_iter() needs to read and
    write registers from within rtw_update_sta_info(). Using the non-atomic
    iterator ensures that we can sleep during USB and SDIO register reads
    and writes. This fixes "scheduling while atomic" or "Voluntary context
    switch within RCU read-side critical section!" warnings as seen by SDIO
    card users (but it also affects USB cards).

    Fixes: 78d5bf925f30 ("wifi: rtw88: iterate over vif/sta list non-atomically")
    Suggested-by: Ping-Ke Shih <pkshih@realtek.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Tested-by: Sascha Hauer <s.hauer@pengutronix.de>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230108211324.442823-4-
martin.blumenstingl@googlemail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 82a4774e6a0a7333c37127f0622f5cba4fecdf0e
```

```
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Sun Jan 8 22:13:23 2023 +0100

    wifi: rtw88: Use rtw_iterate_vifs() for rtw_vif_watch_dog_iter()

    [ Upstream commit 313f6dc7c5ed723d0c5691553eff4c0090f16bb8 ]

    USB and (upcoming) SDIO support may sleep in the read/write handlers.
    Make rtw_watch_dog_work() use rtw_iterate_vifs() to prevent "scheduling
    while atomic" or "Voluntary context switch within RCU read-side
    critical section!" warnings when accessing the registers using an SDIO
    card (which is where this issue has been spotted in the real world but
    it also affects USB cards).

    Fixes: 78d5bf925f30 ("wifi: rtw88: iterate over vif/sta list non-atomically")
    Suggested-by: Ping-Ke Shih <pkshih@realtek.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Tested-by: Sascha Hauer <s.hauer@pengutronix.de>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230108211324.442823-3-
martin.blumenstingl@googlemail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0eabd7b9d7fd454eb112f9978b3525d2029a2594
Author: Alexey Kodanev <aleksei.kodanev@bell-sw.com>
Date:   Tue Dec 27 16:33:06 2022 +0300

    wifi: orinoco: check return value of hermes_write_wordrec()

    [ Upstream commit 1e346cbb096a5351a637ec1992beffbf330547f0 ]

    There is currently no return check for writing an authentication
    type (HERMES_AUTH_SHARED_KEY or HERMES_AUTH_OPEN). It looks like
    it was accidentally skipped.

    This patch adds a return check similar to the other checks in
    __orinoco_hw_setup_enc() for hermes_write_wordrec().

    Detected using the static analysis tool - Svace.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Alexey Kodanev <aleksei.kodanev@bell-sw.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221227133306.201356-1-aleksei.kodanev@bell-sw.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f39a86b4efd270947ee252cc32a30b0aef492d65
Author: Bitterblue Smith <rtl8821cerfe2@gmail.com>
Date:   Thu Dec 22 13:48:04 2022 +0200

    wifi: rtl8xxxu: Fix memory leaks with RTL8723BU, RTL8192EU

    [ Upstream commit b39f662ce1648db0b9de32e6a849b098480793cb ]

    The wifi + bluetooth combo chip RTL8723BU can leak memory (especially?)
    when it's connected to a bluetooth audio device. The busy bluetooth
    traffic generates lots of C2H (card to host) messages, which are not
    freed correctly.

    To fix this, move the dev_kfree_skb() call in rtl8xxxu_c2hcmd_callback()
    inside the loop where skb_dequeue() is called.

    The RTL8192EU leaks memory because the C2H messages are added to the
    queue and left there forever. (This was fine in the past because it
```

```
    probably wasn't sending any C2H messages until commit e542e66b7c2e
    ("wifi: rtl8xxxu: gen2: Turn on the rate control"). Since that commit
    it sends a C2H message when the TX rate changes.)

    To fix this, delete the check for rf_paths > 1 and the goto. Let the
    function process the C2H messages from RTL8192EU like the ones from
    the other chips.

    Theoretically the RTL8188FU could also leak like RTL8723BU, but it
    most likely doesn't send C2H messages frequently enough.

    This change was tested with RTL8723BU by Erhard F. I tested it with
    RTL8188FU and RTL8192EU.

    Reported-by: Erhard F. <erhard_f@mailbox.org>
    Tested-by: Erhard F. <erhard_f@mailbox.org>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=215197
    Fixes: e542e66b7c2e ("rtl8xxxu: add bluetooth co-existence support for single antenna")
    Signed-off-by: Bitterblue Smith <rtl8821cerfe2@gmail.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/03b099c1-c671-d252-36f4-57b70d721f9d@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48ffa7dbe88ce12ffa9eb5bcd8b79c5c48821717
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Jan 4 22:29:01 2023 +0800

    wifi: rtw89: Add missing check for alloc_workqueue

    [ Upstream commit ed9e6166eb0984b718facb7ca59296098cc3aa64 ]

    Add check for the return value of alloc_workqueue since it may return
    NULL pointer.
    Moreover, add destroy_workqueue when rtw89_load_firmware fails.

    Fixes: e3ec7017f6a2 ("rtw89: add Realtek 802.11ax driver")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230104142901.1611-1-jiasheng@iscas.ac.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a10b6d73123bd480751d916575835abb615072fd
Author: Zong-Zhe Yang <kevin_yang@realtek.com>
Date:   Tue Jan 3 22:10:54 2023 +0800

    wifi: rtw89: fix potential leak in rtw89_append_probe_req_ie()

    [ Upstream commit 4a0e218cc9c42d1903ade8b5a371dcf48cf918c5 ]

    Do `kfree_skb(new)` before `goto out` to prevent potential leak.

    Fixes: 895907779752 ("rtw89: 8852a: add ieee80211_ops::hw_scan")
    Signed-off-by: Zong-Zhe Yang <kevin_yang@realtek.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230103141054.17372-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d3955094a2ce1d1fb256b88f5e57d9838b1c33ac
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sun Jan 1 21:40:22 2023 +0200
```

```
    thermal/drivers/tsens: limit num_sensors to 9 for msm8939

    [ Upstream commit 903238a33c116edf5f64f7a3fd246e6169cccfa6 ]

    On msm8939 last (hwid=10) sensor was added in the hw revision 3.0.
    Calibration data for it was placed outside of the main calibration data
    blob, so it is not accessible by the current blob-parsing code.

    Moreover data for the sensor's p2 is not contiguous in the fuses. This
    makes it hard to use nvmem_cell API to parse calibration data in a
    generic way.

    Since the sensor doesn't seem to be actually used by the existing
    hardware, disable the sensor for now.

    Fixes: 332bc8ebab2c ("thermal: qcom: tsens-v0_1: Add support for MSM8939")
    Cc: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Cc: Shawn Guo <shawn.guo@linaro.org>
    Acked-by: Shawn Guo <shawn.guo@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Link: https://lore.kernel.org/r/20230101194034.831222-9-dmitry.baryshkov@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2e90a3f7f48c9c990052a8072ac653f1fbc046f9
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sun Jan 1 21:40:21 2023 +0200

    thermal/drivers/tsens: fix slope values for msm8939

    [ Upstream commit 5aec3b035e0cbf3f042c2a03d654e5ad6748feb7 ]

    According to the vendor kernels (msm-3.10, 3.14 and 3.18), msm8939
    uses non-standard slope values for calibrating the sensors. Fill them
    accordingly.

    Fixes: 332bc8ebab2c ("thermal: qcom: tsens-v0_1: Add support for MSM8939")
    Cc: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Cc: Shawn Guo <shawn.guo@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Acked-by: Shawn Guo <shawn.guo@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Link: https://lore.kernel.org/r/20230101194034.831222-8-dmitry.baryshkov@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b96888ed77dfb83be068b308ab2531f98f18a39c
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sun Jan 1 21:40:20 2023 +0200

    thermal/drivers/tsens: Sort out msm8976 vs msm8956 data

    [ Upstream commit a7d3006be5ca7b04e4b84b5ceaae55a700e511bd ]

    Tsens driver mentions that msm8976 data should be used for both msm8976
    and msm8956 SoCs. This is not quite correct, as according to the
    vendor kernels, msm8976 should use standard slope values (3200), while
    msm8956 really uses the slope values found in the driver.

    Add separate compatibility string for msm8956, move slope value
    overrides to the corresponding init function and use the standard
    compute_intercept_slope() function for both platforms.
```

```
    Fixes: 0e580290170d ("thermal: qcom: tsens-v1: Add support for MSM8956 and MSM8976")
    Cc: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Link: https://lore.kernel.org/r/20230101194034.831222-7-dmitry.baryshkov@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 140eb1a410c8db3ed7d92996d7c672e8da8c61a3
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sun Jan 1 21:40:19 2023 +0200

    thermal/drivers/tsens: Drop msm8976-specific defines

    [ Upstream commit 3bf0ea99e2e32b0335106b86d84404cc85bcd113 ]

    Drop msm8976-specific defines, which duplicate generic ones.

    Fixes: 0e580290170d ("thermal: qcom: tsens-v1: Add support for MSM8956 and MSM8976")
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Link: https://lore.kernel.org/r/20230101194034.831222-6-dmitry.baryshkov@linaro.org
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b86f2b5d3a84fc0c095617d691482e494aad3819
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sat Jan 14 18:33:09 2023 +0100

    x86/signal: Fix the value returned by strict_sas_size()

    [ Upstream commit ef6dfc4b238a435ab552207ec09e4a82b6426d31 ]

    Functions used with __setup() return 1 when the argument has been
    successfully parsed.

    Reverse the returned value so that 1 is returned when kstrtobool() is
    successful (i.e. returns 0).

    My understanding of these __setup() functions is that returning 1 or 0
    does not change much anyway - so this is more of a cleanup than a
    functional fix.

    I spot it and found it spurious while looking at something else.
    Even if the output is not perfect, you'll get the idea with:

       $ git grep -B2 -A10 retu.*kstrtobool | grep __setup -B10

    Fixes: 3aac3ebea08f ("x86/signal: Implement sigaltstack size validation")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link:
https://lore.kernel.org/r/73882d43ebe420c9d8fb82d0560021722b243000.1673717552.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 423f406b84a092cc0875307e72ce7a0362686403
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sun Jul 31 18:09:14 2022 +0200

    s390/vfio-ap: fix an error handling path in vfio_ap_mdev_probe_queue()

    [ Upstream commit 08866d34c7099e981918d34aab5d6a437436628f ]
```

```
    The commit in Fixes: has switch the order of a sysfs_create_group() and a
    kzalloc().

    It correctly removed the now useless kfree() but forgot to add a
    sysfs_remove_group() in case of (unlikely) memory allocation failure.

    Add it now.

    Fixes: 260f3ea14138 ("s390/vfio-ap: move probe and remove callbacks to vfio_ap_ops.c")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Link:
https://lore.kernel.org/r/d0c0a35eec4fa87cb7f3910d8ac4dc0f7dc9008a.1659283738.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2686e112b03085d19313a12e5842b0c74354c2e
Author: Alexander Gordeev <agordeev@linux.ibm.com>
Date:   Thu Dec 15 08:00:34 2022 +0100

    s390/early: fix sclp_early_sccb variable lifetime

    [ Upstream commit 639886b71ddef085a0e7bb1f225b8ae3eda5c06f ]

    Commit ada1da31ce34 ("s390/sclp: sort out physical vs
    virtual pointers usage") fixed the notion of virtual
    address for sclp_early_sccb pointer. However, it did
    not take into account that kasan_early_init() can also
    output messages and sclp_early_sccb should be adjusted
    by the time kasan_early_init() is called.

    Currently it is not a problem, since virtual and physical
    addresses on s390 are the same. Nevertheless, should they
    ever differ, this would cause an invalid pointer access.

    Fixes: ada1da31ce34 ("s390/sclp: sort out physical vs virtual pointers usage")
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68b2e5d84aede94634f8b6f8a59882f4cd3b1e6c
Author: Lai Jiangshan <jiangshan.ljs@antgroup.com>
Date:   Thu Jan 12 16:14:27 2023 +0000

    workqueue: Protects wq_unbound_cpumask with wq_pool_attach_mutex

    [ Upstream commit 99c621ef243bda726fb8d982a274ded96570b410 ]

    When unbind_workers() reads wq_unbound_cpumask to set the affinity of
    freshly-unbound kworkers, it only holds wq_pool_attach_mutex. This isn't
    sufficient as wq_unbound_cpumask is only protected by wq_pool_mutex.

    Make wq_unbound_cpumask protected with wq_pool_attach_mutex and also
    remove the need of temporary saved_cpumask.

    Fixes: 10a5a651e3af ("workqueue: Restrict kworker in the offline CPU pool running on
housekeeping CPUs")
    Reported-by: Valentin Schneider <vschneid@redhat.com>
    Signed-off-by: Lai Jiangshan <jiangshan.ljs@antgroup.com>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 807e3596ef20257fcc61462b10a6f6573c395537
Author: Mark Brown <broonie@kernel.org>
Date:   Tue Dec 27 13:06:35 2022 +0000

    kselftest/arm64: Fix syscall-abi for systems without 128 bit SME

    [ Upstream commit 97ec597b26df774a257e3f8e97353fd1b4471615 ]

    SME does not mandate any specific VL so we may not have 128 bit SME but
    the algorithm used for enumerating VLs assumes that we will. Add the
    required check to ensure that the algorithm terminates.

    Fixes: 43e3f85523e4 ("kselftest/arm64: Add SME support to syscall ABI test")
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20221223-arm64-syscall-abi-sme-only-v1-1-
4fabfbd62087@kernel.org
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f4bf701d6897c5de4110ff7180370db3c30b7696
Author: Mark Brown <broonie@kernel.org>
Date:   Tue Dec 27 12:55:55 2022 +0000

    arm64/sysreg: Fix errors in 32 bit enumeration values

    [ Upstream commit 8c6e105558628b63292ae770198f11a1d5535660 ]

    The recently converted 32 bit ID registers have errors in MVFR0_EL1.FPSP,
    MVFR0_EL1.SIMDReg and MVFR1_EL1.SIMDHP where enumeration values which
    should be 0b0010 are specified as 0b0001. Correct these.

    Fixes: e79c94a2a487 ("arm64/sysreg: Convert MVFR0_EL1 to automatic generation")
    Fixes: c9b718eda706 ("arm64/sysreg: Convert MVFR1_EL1 to automatic generation")
    Reviewed-by: Mark Rutland <mark.rutland@arm.com>
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20221207-arm64-sysreg-helpers-v3-2-
0d71a7b174a8@kernel.org
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 26efe8b6f0f228d85b733f652fcfb38e8acdc4dc
Author: Mark Brown <broonie@kernel.org>
Date:   Tue Dec 27 12:55:54 2022 +0000

    arm64/cpufeature: Fix field sign for DIT hwcap detection

    [ Upstream commit 50daf5b7c4ec4efcaf49a4128930f872bec7dbc0 ]

    Since it was added our hwcap for DIT has specified that DIT is a signed
    field but this appears to be incorrect, the two values for the enumeration
    are:

            0b0000  NI
            0b0001  IMP

    which look like a normal unsigned enumeration and the in-kernel DIT usage
    added by 01ab991fc0ee ("arm64: Enable data independent timing (DIT) in the
    kernel") detects the feature with an unsigned enum. Fix the hwcap to specify
    the field as unsigned.

    Fixes: 7206dc93a58f ("arm64: Expose Arm v8.4 features")
    Reviewed-by: Mark Rutland <mark.rutland@arm.com>
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20221207-arm64-sysreg-helpers-v3-1-
0d71a7b174a8@kernel.org
```

```
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b016d8a231761487475f654053fdbe0b7ea6ef68
Author: Magnus Karlsson <magnus.karlsson@intel.com>
Date:   Wed Jan 11 10:35:15 2023 +0100

    selftests/xsk: print correct error codes when exiting

    [ Upstream commit 085dcccfb7d3dc52ed708fc588587f319541bc83 ]

    Print the correct error codes when exiting the test suite due to some
    terminal error. Some of these had a switched sign and some of them
    printed zero instead of errno.

    Fixes: facb7cb2e909 ("selftests/bpf: Xsk selftests - SKB POLL, NOPOLL")
    Signed-off-by: Magnus Karlsson <magnus.karlsson@intel.com>
    Acked-by: Maciej Fijalkowski <maciej.fijalkowski@intel.com>
    Link: https://lore.kernel.org/r/20230111093526.11682-5-magnus.karlsson@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a20b489a36bc7f5c4556d725e3e3b20e09a7c0b
Author: Magnus Karlsson <magnus.karlsson@intel.com>
Date:   Wed Jan 11 10:35:12 2023 +0100

    selftests/xsk: print correct payload for packet dump

    [ Upstream commit 2d0b2ae2871ae6d42a9f0a4280e0fb5bff8d38b8 ]

    Print the correct payload when the packet dump option is selected. The
    network to host conversion was forgotten and the payload was
    erronously declared to be an int instead of an unsigned int.

    Fixes: facb7cb2e909 ("selftests/bpf: Xsk selftests - SKB POLL, NOPOLL")
    Signed-off-by: Magnus Karlsson <magnus.karlsson@intel.com>
    Acked-by: Maciej Fijalkowski <maciej.fijalkowski@intel.com>
    Link: https://lore.kernel.org/r/20230111093526.11682-2-magnus.karlsson@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cc9a2d3b03af48ba88d7dac6ab671d2dffcea25
Author: Michal Suchanek <msuchanek@suse.de>
Date:   Mon Jan 9 12:34:42 2023 +0100

    bpf_doc: Fix build error with older python versions

    [ Upstream commit 5fbea42387eba1c7517fcad79099df706def7054 ]

    The ability to subscript match result as an array is only available
    since python 3.6. Existing code in bpf_doc uses the older group()
    interface but commit 8a76145a2ec2 adds code using the new interface.

    Use the old interface consistently to avoid build error on older
    distributions like the below:

    + make -j48 -s -C /dev/shm/kbuild/linux.33946/current ARCH=powerpc HOSTCC=gcc
CROSS_COMPILE=powerpc64-suse-linux- clean
    TypeError: '_sre.SRE_Match' object is not subscriptable

    Fixes: 8a76145a2ec2 ("bpf: explicitly define BPF_FUNC_xxx integer values")
    Signed-off-by: Michal Suchanek <msuchanek@suse.de>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Quentin Monnet <quentin@isovalent.com>
    Link: https://lore.kernel.org/bpf/20230109113442.20946-1-msuchanek@suse.de
```

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

commit 35812dcac9081e51927bf3410be56192236819c7
Author: Ludovic L'Hours <ludovic.lhours@gmail.com>
Date:   Sun Jan 8 19:20:18 2023 +0100

    libbpf: Fix map creation flags sanitization

    [ Upstream commit 6920b08661e3ad829206078b5c9879b24aea8dfc ]

    As BPF_F_MMAPABLE flag is now conditionnaly set (by map_is_mmapable),
    it should not be toggled but disabled if not supported by kernel.

    Fixes: 4fcac46c7e10 ("libbpf: only add BPF_F_MMAPABLE flag for data maps with global
vars")
    Signed-off-by: Ludovic L'Hours <ludovic.lhours@gmail.com>
    Acked-by: Jiri Olsa <jolsa@kernel.org>
    Link: https://lore.kernel.org/r/20230108182018.24433-1-ludovic.lhours@gmail.com
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bab09d99cbbe559da3c7dc417f4752ae9bfd163
Author: Daniil Tatianin <d-tatianin@yandex-team.ru>
Date:   Sat Jan 7 02:53:08 2023 +0300

    ACPICA: nsrepair: handle cases without a return value correctly

    [ Upstream commit ca843a4c79486e99a19b859ef0b9887854afe146 ]

    Previously acpi_ns_simple_repair() would crash if expected_btypes
    contained any combination of ACPI_RTYPE_NONE with a different type,
    e.g | ACPI_RTYPE_INTEGER because of slightly incorrect logic in the
    !return_object branch, which wouldn't return AE_AML_NO_RETURN_VALUE
    for such cases.

    Found by Linux Verification Center (linuxtesting.org) with the SVACE
    static analysis tool.

    Link: https://github.com/acpica/acpica/pull/811
    Fixes: 61db45ca2163 ("ACPICA: Restore code that repairs NULL package elements in return
values.")
    Signed-off-by: Daniil Tatianin <d-tatianin@yandex-team.ru>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 08fe0860477b49a82ca9823f3d2e3ce506c2b11a
Author: Prashant Malani <pmalani@chromium.org>
Date:   Wed Dec 28 00:45:08 2022 +0000

    platform/chrome: cros_ec_typec: Update port DP VDO

    [ Upstream commit 8d2b28df6c3dc1581d856f52d9f78059ef2a568f ]

    The port advertising DP support is a Type-C receptacle. Fix the port's
    DisplayPort VDO to reflect this.

    Fixes: 1903adae0464 ("platform/chrome: cros_ec_typec: Add bit offset for DP VDO")
    Signed-off-by: Prashant Malani <pmalani@chromium.org>
    Reviewed-by: Benson Leung <bleung@chromium.org>
    Acked-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Link: https://lore.kernel.org/r/20221228004648.793339-6-pmalani@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b2519474352ac93fa9911ff4ee7d32e54a1b331
Author: David Rientjes <rientjes@google.com>
```

```
Date:   Fri Dec 30 14:18:46 2022 -0800

    crypto: ccp - Avoid page allocation failure warning for SEV_GET_ID2

    [ Upstream commit 91dfd98216d817ec5f1c55890bacb7b4fe9b068a ]

    For SEV_GET_ID2, the user provided length does not have a specified
    limitation because the length of the ID may change in the future.  The
    kernel memory allocation, however, is implicitly limited to 4MB on x86 by
    the page allocator, otherwise the kzalloc() will fail.

    When this happens, it is best not to spam the kernel log with the warning.
    Simply fail the allocation and return ENOMEM to the user.

    Fixes: d6112ea0cb34 ("crypto: ccp - introduce SEV_GET_ID2 command")
    Reported-by: Andy Nguyen <theflow@google.com>
    Reported-by: Peter Gonda <pgonda@google.com>
    Suggested-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: David Rientjes <rientjes@google.com>
    Acked-by: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e83bf66998bf7c71bbe37755dd007949cc2629a6
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Tue Dec 27 15:27:39 2022 +0100

    lib/mpi: Fix buffer overrun when SG is too long

    [ Upstream commit 7361d1bc307b926cbca214ab67b641123c2d6357 ]

    The helper mpi_read_raw_from_sgl sets the number of entries in
    the SG list according to nbytes.  However, if the last entry
    in the SG list contains more data than nbytes, then it may overrun
    the buffer because it only allocates enough memory for nbytes.

    Fixes: 2d4d1eea540b ("lib/mpi: Add mpi sgl helpers")
    Reported-by: Roberto Sassu <roberto.sassu@huaweicloud.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Reviewed-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8901f928f65213c02a48567b3000b13a8f04e5ce
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Fri Nov 25 14:55:00 2022 +0100

    rcu-tasks: Fix synchronize_rcu_tasks() VS zap_pid_ns_processes()

    [ Upstream commit 28319d6dc5e2ffefa452c2377dd0f71621b5bff0 ]

    RCU Tasks and PID-namespace unshare can interact in do_exit() in a
    complicated circular dependency:

    1) TASK A calls unshare(CLONE_NEWPID), this creates a new PID namespace
       that every subsequent child of TASK A will belong to. But TASK A
       doesn't itself belong to that new PID namespace.

    2) TASK A forks() and creates TASK B. TASK A stays attached to its PID
       namespace (let's say PID_NS1) and TASK B is the first task belonging
       to the new PID namespace created by unshare()  (let's call it PID_NS2).

    3) Since TASK B is the first task attached to PID_NS2, it becomes the
       PID_NS2 child reaper.
```

```
     4) TASK A forks() again and creates TASK C which get attached to PID_NS2.
        Note how TASK C has TASK A as a parent (belonging to PID_NS1) but has
        TASK B (belonging to PID_NS2) as a pid_namespace child_reaper.

     5) TASK B exits and since it is the child reaper for PID_NS2, it has to
        kill all other tasks attached to PID_NS2, and wait for all of them to
        die before getting reaped itself (zap_pid_ns_process()).

     6) TASK A calls synchronize_rcu_tasks() which leads to
        synchronize_srcu(&tasks_rcu_exit_srcu).

     7) TASK B is waiting for TASK C to get reaped. But TASK B is under a
        tasks_rcu_exit_srcu SRCU critical section (exit_notify() is between
        exit_tasks_rcu_start() and exit_tasks_rcu_finish()), blocking TASK A.

     8) TASK C exits and since TASK A is its parent, it waits for it to reap
        TASK C, but it can't because TASK A waits for TASK B that waits for
        TASK C.

     Pid_namespace semantics can hardly be changed at this point. But the
     coverage of tasks_rcu_exit_srcu can be reduced instead.

     The current task is assumed not to be concurrently reapable at this
     stage of exit_notify() and therefore tasks_rcu_exit_srcu can be
     temporarily relaxed without breaking its constraints, providing a way
     out of the deadlock scenario.

     [ paulmck: Fix build failure by adding additional declaration. ]

     Fixes: 3f95aa81d265 ("rcu: Make TASKS_RCU handle tasks that are almost done exiting")
     Reported-by: Pengfei Xu <pengfei.xu@intel.com>
     Suggested-by: Boqun Feng <boqun.feng@gmail.com>
     Suggested-by: Neeraj Upadhyay <quic_neeraju@quicinc.com>
     Suggested-by: Paul E. McKenney <paulmck@kernel.org>
     Cc: Oleg Nesterov <oleg@redhat.com>
     Cc: Lai Jiangshan <jiangshanlai@gmail.com>
     Cc: Eric W . Biederman <ebiederm@xmission.com>
     Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
     Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1125f99a5c35a3d4af9a77a27c93024132aef58
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Fri Nov 25 14:54:59 2022 +0100

     rcu-tasks: Remove preemption disablement around srcu_read_[un]lock() calls

     [ Upstream commit 44757092958bdd749775022f915b7ac974384c2a ]

     Ever since the following commit:

             5a41344a3d83 ("srcu: Simplify __srcu_read_unlock() via this_cpu_dec()")

     SRCU doesn't rely anymore on preemption to be disabled in order to
     modify the per-CPU counter. And even then it used to be done from the API
     itself.

     Therefore and after checking further, it appears to be safe to remove
     the preemption disablement around __srcu_read_[un]lock() in
     exit_tasks_rcu_start() and exit_tasks_rcu_finish()

     Suggested-by: Boqun Feng <boqun.feng@gmail.com>
     Suggested-by: Paul E. McKenney <paulmck@kernel.org>
     Suggested-by: Neeraj Upadhyay <quic_neeraju@quicinc.com>
     Cc: Lai Jiangshan <jiangshanlai@gmail.com>
```

```
    Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Stable-dep-of: 28319d6dc5e2 ("rcu-tasks: Fix synchronize_rcu_tasks() VS
zap_pid_ns_processes()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e08fa8ffeaf30ed627c87fd27f921e6a5c33931
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Fri Nov 25 14:54:58 2022 +0100

    rcu-tasks: Improve comments explaining tasks_rcu_exit_srcu purpose

    [ Upstream commit e4e1e8089c5fd948da12cb9f4adc93821036945f ]

    Make sure we don't need to look again into the depths of git blame in
    order not to miss a subtle part about how rcu-tasks is dealing with
    exiting tasks.

    Suggested-by: Boqun Feng <boqun.feng@gmail.com>
    Suggested-by: Neeraj Upadhyay <quic_neeraju@quicinc.com>
    Suggested-by: Paul E. McKenney <paulmck@kernel.org>
    Cc: Oleg Nesterov <oleg@redhat.com>
    Cc: Lai Jiangshan <jiangshanlai@gmail.com>
    Cc: Eric W. Biederman <ebiederm@xmission.com>
    Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Stable-dep-of: 28319d6dc5e2 ("rcu-tasks: Fix synchronize_rcu_tasks() VS
zap_pid_ns_processes()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9116e48f933b2f9c82dccf276acb33e3a9986282
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Sat Nov 19 17:25:03 2022 +0800

    genirq: Fix the return type of kstat_cpu_irqs_sum()

    [ Upstream commit 47904aed898a08f028572b9b5a5cc101ddfb2d82 ]

    The type of member ->irqs_sum is unsigned long, but kstat_cpu_irqs_sum()
    returns int, which can result in truncation.  Therefore, change the
    kstat_cpu_irqs_sum() function's return value to unsigned long to avoid
    truncation.

    Fixes: f2c66cd8eedd ("/proc/stat: scalability of irq num per cpu")
    Reported-by: Elliott, Robert (Servers) <elliott@hpe.com>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Cc: Tejun Heo <tj@kernel.org>
    Cc: "Peter Zijlstra (Intel)" <peterz@infradead.org>
    Cc: Josh Don <joshdon@google.com>
    Cc: Andrew Morton <akpm@linux-foundation.org>
    Reviewed-by: Frederic Weisbecker <frederic@kernel.org>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d028f266127c07777f9256f5212e370d6869b32
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Thu Dec 15 09:51:20 2022 -0600

    ACPICA: Drop port I/O validation for some regions

    [ Upstream commit e1d9148582ab2c3dada5c5cf8ca7531ca269fee5 ]

    Microsoft introduced support in Windows XP for blocking port I/O
    to various regions.  For Windows compatibility ACPICA has adopted
    the same protections and will disallow writes to those
```

```
    (presumably) the same regions.

    On some systems the AML included with the firmware will issue 4 byte
    long writes to 0x80.  These writes aren't making it over because of this
    blockage. The first 4 byte write attempt is rejected, and then
    subsequently 1 byte at a time each offset is tried. The first at 0x80
    works, but then the next 3 bytes are rejected.

    This manifests in bizarre failures for devices that expected the AML to
    write all 4 bytes.  Trying the same AML on Windows 10 or 11 doesn't hit
    this failure and all 4 bytes are written.

    Either some of these regions were wrong or some point after Windows XP
    some of these regions blocks have been lifted.

    In the last 15 years there doesn't seem to be any reports popping up of
    this error in the Windows event viewer anymore.  There is no documentation
    at Microsoft's developer site indicating that Windows ACPI interpreter
    blocks these regions. Between the lack of documentation and the fact that
    the writes actually do work in Windows 10 and 11, it's quite likely
    Windows doesn't actually enforce this anymore.

    So to help the issue, only enforce Windows XP specific entries if the
    latest _OSI supported is Windows XP. Continue to enforce the
    ALWAYS_ILLEGAL entries.

    Link: https://github.com/acpica/acpica/pull/817
    Fixes: 7f0719039085 ("ACPICA: New: I/O port protection")
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c42ac108753b42620cb965daba509575f11f64b8
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Tue Dec 20 09:05:36 2022 +0100

    crypto: ux500 - update debug config after ux500 cryp driver removal

    [ Upstream commit 49bc6a7786b7d03eab6912a88d09a7991a32174e ]

    Commit 453de3eb08c4 ("crypto: ux500/cryp - delete driver") removes the
    config CRYPTO_DEV_UX500_CRYP, but leaves an obsolete reference in the
    dependencies of config CRYPTO_DEV_UX500_DEBUG.

    Remove that obsolete reference, and adjust the description while at it.

    Fixes: 453de3eb08c4 ("crypto: ux500/cryp - delete driver")
    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Acked-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 56255bf83850a5d627840b23c2e1faca54cb4235
Author: Eric Biggers <ebiggers@google.com>
Date:   Mon Dec 19 21:40:40 2022 -0800

    crypto: x86/ghash - fix unaligned access in ghash_setkey()

    [ Upstream commit 116db2704c193fff6d73ea6c2219625f0c9bdfc8 ]

    The key can be unaligned, so use the unaligned memory access helpers.

    Fixes: 8ceee72808d1 ("crypto: ghash-clmulni-intel - use C implementation for setkey()")
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
```

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

commit 6e042f650bfaaf367afff4bdd3d75de8d4bb1598
Author: Daniel T. Lee <danieltimlee@gmail.com>
Date:   Sat Dec 24 16:15:27 2022 +0900

    libbpf: Fix invalid return address register in s390

    [ Upstream commit 7244eb669397f309c3d014264823cdc9cb3f8e6b ]

    There is currently an invalid register mapping in the s390 return
    address register. As the manual[1] states, the return address can be
    found at r14. In bpf_tracing.h, the s390 registers were named
    gprs(general purpose registers). This commit fixes the problem by
    correcting the mistyped mapping.

    [1]: https://uclibc.org/docs/psABI-s390x.pdf#page=14

    Fixes: 3cc31d794097 ("libbpf: Normalize PT_REGS_xxx() macro definitions")
    Signed-off-by: Daniel T. Lee <danieltimlee@gmail.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20221224071527.2292-7-danieltimlee@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc81aba97718447101dbb2920f5b1931a0a9cc23
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 23:04:53 2022 +0800

    wifi: wl3501_cs: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 44bacbdf9066c590423259dbd6d520baac99c1a8 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So replace kfree_skb()
    with dev_kfree_skb_irq() under spin_lock_irqsave(). Compile
    tested only.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207150453.114742-1-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b467882341149d61ddd86e542b8378a3d66f97d5
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 23:00:08 2022 +0800

    wifi: libertas: cmdresp: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 708a49a64237f19bd404852f297aaadbc9e7fee0 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So replace kfree_skb()
    with dev_kfree_skb_irq() under spin_lock_irqsave(). Compile
    tested only.

    Fixes: f52b041aed77 ("libertas: Add spinlock to avoid race condition")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207150008.111743-5-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 08e4921ed386da07033d05f7a3006c560eeb6979
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 23:00:07 2022 +0800
```

```
    wifi: libertas: main: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit f393df151540bf858effbd29ff572ab94e76a4c4 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So replace kfree_skb()
    with dev_kfree_skb_irq() under spin_lock_irqsave(). Compile
    tested only.

    Fixes: d2e7b3425c47 ("libertas: disable functionality when interface is down")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207150008.111743-4-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9e820b5084e0d116e4ab17b11a761236f459ba22
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 23:00:06 2022 +0800

    wifi: libertas: if_usb: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 3968e81ba644f10a7d45bae2539560db9edac501 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So replace kfree_skb()
    with dev_kfree_skb_irq() under spin_lock_irqsave(). Compile
    tested only.

    Fixes: a3128feef6d5 ("libertas: use irqsave() in USB's complete callback")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207150008.111743-3-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e03d5afddc81d5901850ef2640e1f11370b939b
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 23:00:05 2022 +0800

    wifi: libertas_tf: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 9388ce97b98216833c969191ee6df61a7201d797 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So replace kfree_skb()
    with dev_kfree_skb_irq() under spin_lock_irqsave(). Compile
    tested only.

    Fixes: fc75122fabb5 ("libertas_tf: use irqsave() in USB's complete callback")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207150008.111743-2-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 55b81a6c138629de4318d680ed80abcb81eec606
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Wed Dec 7 09:31:14 2022 +0800

    wifi: brcmfmac: unmap dma buffer in brcmf_msgbuf_alloc_pktid()

    [ Upstream commit b9f420032f2ba1e634b22ca7b433e5c40ea663af ]

    After the DMA buffer is mapped to a physical address, address is stored
    in pktids in brcmf_msgbuf_alloc_pktid(). Then, pktids is parsed in
    brcmf_msgbuf_get_pktid()/brcmf_msgbuf_release_array() to obtain physaddr
```

```
    and later unmap the DMA buffer. But when count is always equal to
    pktids->array_size, physaddr isn't stored in pktids and the DMA buffer
    will not be unmapped anyway.

    Fixes: 9a1bb60250d2 ("brcmfmac: Adding msgbuf protocol.")
    Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
    Reviewed-by: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207013114.1748936-1-shaozhengchao@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a4d18318f473e97d628f410215b3fac32d07aed
Author: Zhang Changzhong <zhangchangzhong@huawei.com>
Date:   Thu Nov 17 19:33:01 2022 +0800

    wifi: brcmfmac: fix potential memory leak in brcmf_netdev_start_xmit()

    [ Upstream commit 212fde3fe76e962598ce1d47b97cc78afdfc71b3 ]

    The brcmf_netdev_start_xmit() returns NETDEV_TX_OK without freeing skb
    in case of pskb_expand_head() fails, add dev_kfree_skb() to fix it.
    Compile tested only.

    Fixes: 270a6c1f65fe ("brcmfmac: rework headroom check in .start_xmit()")
    Signed-off-by: Zhang Changzhong <zhangchangzhong@huawei.com>
    Reviewed-by: Arend van Spriel <arend.vanspriel@broadcom.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/1668684782-47422-1-git-send-email-
zhangchangzhong@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6da6ce086221803ed6c3b1db11096cecd3e58ec8
Author: Wang Yufen <wangyufen@huawei.com>
Date:   Thu Nov 24 19:38:22 2022 +0800

    wifi: wilc1000: add missing unregister_netdev() in wilc_netdev_ifc_init()

    [ Upstream commit 2b88974ecb358990e1c33fabcd0b9e142bab7f21 ]

    Fault injection test reports this issue:

    kernel BUG at net/core/dev.c:10731!
    invalid opcode: 0000 [#1] PREEMPT SMP KASAN PTI
    Call Trace:
      <TASK>
      wilc_netdev_ifc_init+0x19f/0x220 [wilc1000 884bf126e9e98af6a708f266a8dffd53f99e4bf5]
      wilc_cfg80211_init+0x30c/0x380 [wilc1000 884bf126e9e98af6a708f266a8dffd53f99e4bf5]
      wilc_bus_probe+0xad/0x2b0 [wilc1000_spi 1520a7539b6589cc6cde2ae826a523a33f8bacff]
      spi_probe+0xe4/0x140
      really_probe+0x17e/0x3f0
      __driver_probe_device+0xe3/0x170
      driver_probe_device+0x49/0x120

    The root case here is alloc_ordered_workqueue() fails, but
    cfg80211_unregister_netdevice() or unregister_netdev() not be called in
    error handling path. To fix add unregister_netdev goto lable to add the
    unregister operation in error handling path.

    Fixes: 09ed8bfc5215 ("wilc1000: Rename workqueue from "WILC_wq" to "NETDEV-wq"")
    Signed-off-by: Wang Yufen <wangyufen@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/1669289902-23639-1-git-send-email-wangyufen@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit baef42df7de7c35ba60b75a5f96d1eb039f4d782
```

```
Author: Zhang Changzhong <zhangchangzhong@huawei.com>
Date:   Thu Nov 17 19:36:03 2022 +0800

    wifi: wilc1000: fix potential memory leak in wilc_mac_xmit()

    [ Upstream commit deb962ec9e1c9a81babd3d37542ad4bd6ac3396e ]

    The wilc_mac_xmit() returns NETDEV_TX_OK without freeing skb, add
    dev_kfree_skb() to fix it. Compile tested only.

    Fixes: c5c77ba18ea6 ("staging: wilc1000: Add SDIO/SPI 802.11 driver")
    Signed-off-by: Zhang Changzhong <zhangchangzhong@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/1668684964-48622-1-git-send-email-
zhangchangzhong@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f590fb3d14e5db3a9e06ee141b1685c429278ce
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Fri Dec 9 09:24:22 2022 +0800

    wifi: ipw2200: fix memory leak in ipw_wdev_init()

    [ Upstream commit 9fe21dc626117fb44a8eb393713a86a620128ce3 ]

    In the error path of ipw_wdev_init(), exception value is returned, and
    the memory applied for in the function is not released. Also the memory
    is not released in ipw_pci_probe(). As a result, memory leakage occurs.
    So memory release needs to be added to the error path of ipw_wdev_init().

    Fixes: a3caa99e6c68 ("libipw: initiate cfg80211 API conversion (v2)")
    Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221209012422.182669-1-shaozhengchao@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3193cea2ae80c5b38cff5cd3b89e58ac65d05a52
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Dec 8 22:38:26 2022 +0800

    wifi: ipw2x00: don't call dev_kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 45fc6d7461f18df2f238caf0cbc5acc4163203d1 ]

    It is not allowed to call kfree_skb() or consume_skb() from hardware
    interrupt context or with hardware interrupts being disabled.

    It should use dev_kfree_skb_irq() or dev_consume_skb_irq() instead.
    The difference between them is free reason, dev_kfree_skb_irq() means
    the SKB is dropped in error and dev_consume_skb_irq() means the SKB
    is consumed in normal.

    In this case, dev_kfree_skb() is called to free and drop the SKB when
    it's reset, so replace it with dev_kfree_skb_irq(). Compile tested
    only.

    Fixes: 43f66a6ce8da ("Add ipw2200 wireless driver.")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221208143826.2385218-1-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59de322c4903eb4244b44b6869bba14870bb37e9
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Mon Dec 12 13:15:03 2022 -0800
```

```
    libbpf: Fix btf__align_of() by taking into account field offsets

    [ Upstream commit 25a4481b4136af7794e1df2d6c90ed2f354d60ce ]

    btf__align_of() is supposed to be return alignment requirement of
    a requested BTF type. For STRUCT/UNION it doesn't always return correct
    value, because it calculates alignment only based on field types. But
    for packed structs this is not enough, we need to also check field
    offsets and struct size. If field offset isn't aligned according to
    field type's natural alignment, then struct must be packed. Similarly,
    if struct size is not a multiple of struct's natural alignment, then
    struct must be packed as well.

    This patch fixes this issue precisely by additionally checking these
    conditions.

    Fixes: 3d208f4ca111 ("libbpf: Expose btf__align_of() API")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20221212211505.558851-5-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bcc2c3e0e44e4cc5dd31019ebb7d57956c001b87
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Mon Dec 12 13:15:00 2022 -0800

    libbpf: Fix single-line struct definition output in btf_dump

    [ Upstream commit 872aec4b5f635d94111d48ec3c57fbe078d64e7d ]

    btf_dump APIs emit unnecessary tabs when emitting struct/union
    definition that fits on the single line. Before this patch we'd get:

    struct blah {<tab>};

    This patch fixes this and makes sure that we get more natural:

    struct blah {};

    Fixes: 44a726c3f23c ("bpftool: Print newline before '}' for struct with padding only
fields")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20221212211505.558851-2-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 057b52461dc005ecd85a3e4998913b1492ec0f72
Author: Li Zetao <lizetao1@huawei.com>
Date:   Mon Dec 12 10:58:12 2022 +0800

    wifi: rtlwifi: Fix global-out-of-bounds bug in _rtl8812ae_phy_set_txpower_limit()

    [ Upstream commit 117dbeda22ec5ea0918254d03b540ef8b8a64d53 ]

    There is a global-out-of-bounds reported by KASAN:

      BUG: KASAN: global-out-of-bounds in
      _rtl8812ae_eq_n_byte.part.0+0x3d/0x84 [rtl8821ae]
      Read of size 1 at addr ffffffffa0773c43 by task NetworkManager/411

      CPU: 6 PID: 411 Comm: NetworkManager Tainted: G      D
      6.1.0-rc8+ #144 e1558508517267d37
      Hardware name: QEMU Standard PC (Q35 + ICH9, 2009),
      Call Trace:
```

```
        <TASK>
        ...
        kasan_report+0xbb/0x1c0
        _rtl8812ae_eq_n_byte.part.0+0x3d/0x84 [rtl8821ae]
        rtl8821ae_phy_bb_config.cold+0x346/0x641 [rtl8821ae]
        rtl8821ae_hw_init+0x1f5e/0x79b0 [rtl8821ae]
        ...
        </TASK>

    The root cause of the problem is that the comparison order of
    "prate_section" in _rtl8812ae_phy_set_txpower_limit() is wrong. The
    _rtl8812ae_eq_n_byte() is used to compare the first n bytes of the two
    strings from tail to head, which causes the problem. In the
    _rtl8812ae_phy_set_txpower_limit(), it was originally intended to meet
    this requirement by carefully designing the comparison order.
    For example, "pregulation" and "pbandwidth" are compared in order of
    length from small to large, first is 3 and last is 4. However, the
    comparison order of "prate_section" dose not obey such order requirement,
    therefore when "prate_section" is "HT", when comparing from tail to head,
    it will lead to access out of bounds in _rtl8812ae_eq_n_byte(). As
    mentioned above, the _rtl8812ae_eq_n_byte() has the same function as
    strcmp(), so just strcmp() is enough.

    Fix it by removing _rtl8812ae_eq_n_byte() and use strcmp() barely.
    Although it can be fixed by adjusting the comparison order of
    "prate_section", this may cause the value of "rate_section" to not be
    from 0 to 5. In addition, commit "21e4b0726dc6" not only moved driver
    from staging to regular tree, but also added setting txpower limit
    function during the driver config phase, so the problem was introduced
    by this commit.

    Fixes: 21e4b0726dc6 ("rtlwifi: rtl8821ae: Move driver from staging to regular tree")
    Signed-off-by: Li Zetao <lizetao1@huawei.com>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221212025812.1541311-1-lizetao1@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab5c38262ff91edfa74e7c2ee4e95270605ac691
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Fri Dec 9 10:09:39 2022 +0800

    wifi: rtw89: 8852c: rfk: correct DPK settings

    [ Upstream commit 21b5f159a2ee47d30f418559f6ece0088c80199f ]

    Some DPK settings are wrong, and causes bad TX performance occasionally.
    So, fix them by internal suggestions.

    Fixes: da4cea16cb13 ("rtw89: 8852c: rfk: add DPK")
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221209020940.9573-3-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d212c322a03078bdfb7229bf145701b480118c3
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Fri Dec 9 10:09:38 2022 +0800

    wifi: rtw89: 8852c: rfk: correct DACK setting

    [ Upstream commit b2bab7b14098dcf5d405fa8c76b2c3f6ce9184f9 ]

    After filling calibration parameters, set BIT(0) to enable the hardware
    circuit, but original set incorrect bit that affects a little TX
```

```
    performance.

    Fixes: 76599a8d0b7d ("rtw89: 8852c: rfk: add DACK")
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221209020940.9573-2-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8037542e5c45c70d315dfc0e5e778f85a54ccd1
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Dec 8 22:35:17 2022 +0800

    wifi: rtl8xxxu: don't call dev_kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 4c2005ac87685907b3719b4f40215b578efd27c4 ]

    It is not allowed to call kfree_skb() or consume_skb() from hardware
    interrupt context or with hardware interrupts being disabled.

    It should use dev_kfree_skb_irq() or dev_consume_skb_irq() instead.
    The difference between them is free reason, dev_kfree_skb_irq() means
    the SKB is dropped in error and dev_consume_skb_irq() means the SKB
    is consumed in normal.

    In this case, dev_kfree_skb() is called to free and drop the SKB when
    it's shutdown, so replace it with dev_kfree_skb_irq(). Compile tested
    only.

    Fixes: 26f1fad29ad9 ("New driver: rtl8xxxu (mac80211)")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221208143517.2383424-1-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84bd9ae0d054e9def57b8ae153d39e6137d2cc3c
Author: Bitterblue Smith <rtl8821cerfe2@gmail.com>
Date:   Thu Dec 8 21:34:02 2022 +0200

    wifi: rtl8xxxu: Fix assignment to bit field priv->cck_agc_report_type

    [ Upstream commit 639c26faf9b15922bc620af341545d6c5d6aab2d ]

    Just because priv->cck_agc_report_type is only one bit doesn't mean
    it works like a bool. The value assigned to it loses all bits except
    bit 0, so only assign 0 or 1 to it.

    This affects the RTL8192EU, but rtl8xxxu already can't connect to any
    networks with this chip, so it probably didn't bother anyone.

    Fixes: 2ad2a813b803 ("wifi: rtl8xxxu: Fix the CCK RSSI calculation")
    Signed-off-by: Bitterblue Smith <rtl8821cerfe2@gmail.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/7bb4858c-5cef-9cae-5e08-7e8444e8ba89@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1b3dead323986278a5791f527f889f8faddeb82
Author: Bitterblue Smith <rtl8821cerfe2@gmail.com>
Date:   Thu Dec 8 21:32:18 2022 +0200

    wifi: rtl8xxxu: Fix assignment to bit field priv->pi_enabled

    [ Upstream commit 9e32b4a709f0f7b7adf5d9939c3bd47c78c4f003 ]
```

```
    Just because priv->pi_enabled is only one bit doesn't mean it works
    like a bool. The value assigned to it loses all bits except bit 0,
    so only assign 0 or 1 to it.

    This affects the RTL8188FU, but fixing the assignment didn't make
    a difference for my device.

    Fixes: c888183b21f3 ("wifi: rtl8xxxu: Support new chip RTL8188FU")
    Signed-off-by: Bitterblue Smith <rtl8821cerfe2@gmail.com>
    Reviewed-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/4368d585-11ec-d3c7-ec12-7f0afdcedfda@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d46c33f667b05c22bc5c5b69aa730349c4b6fe31
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Thu Dec 8 20:14:48 2022 +0800

    wifi: libertas: fix memory leak in lbs_init_adapter()

    [ Upstream commit 16a03958618fb91bb1bc7077cf3211055162cc2f ]

    When kfifo_alloc() failed in lbs_init_adapter(), cmd buffer is not
    released. Add free memory to processing error path.

    Fixes: 7919b89c8276 ("libertas: convert libertas driver to use an event/cmdresp queue")
    Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
    Reviewed-by: Jiri Pirko <jiri@nvidia.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221208121448.2845986-1-shaozhengchao@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c486bf5ec56636e48dbe4bd99c7d61767c020c30
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 22:40:13 2022 +0800

    wifi: iwlegacy: common: don't call dev_kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 0c1528675d7a9787cb516b64d8f6c0f6f8efcb48 ]

    It is not allowed to call consume_skb() from hardware interrupt context
    or with interrupts being disabled. So replace dev_kfree_skb() with
    dev_consume_skb_irq() under spin_lock_irqsave(). Compile tested only.

    Fixes: 4bc85c1324aa ("Revert "iwlwifi: split the drivers for agn and legacy devices
3945/4965"")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Stanislaw Gruszka <stf_xl@wp.pl>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207144013.70210-1-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41d8fedf7e552fea52b62384ebc3422a34dc5fbf
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 22:14:11 2022 +0800

    wifi: rtlwifi: rtl8723be: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 313950c2114e7051c4e3020fd82495fa1fb526a8 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. All the SKBs have
    been dequeued from the old queue, so it's safe to enqueue these
    SKBs to a free queue, then free them after spin_unlock_irqrestore()
    at once. Compile tested only.
```

```
    Fixes: 5c99f04fec93 ("rtlwifi: rtl8723be: Update driver to match Realtek release of
06/28/14")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207141411.46098-4-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b85f2d6fd10b777df79a047e2b8bcb9f0dd9453
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 22:14:10 2022 +0800

    wifi: rtlwifi: rtl8188ee: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 2611687fa7ffc84190f92292de0b80468de17220 ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. All the SKBs have
    been dequeued from the old queue, so it's safe to enqueue these
    SKBs to a free queue, then free them after spin_unlock_irqrestore()
    at once. Compile tested only.

    Fixes: 7fe3b3abb5da ("rtlwifi: rtl8188ee: rtl8821ae: Fix a queue locking problem")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207141411.46098-3-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50bb23dbd85321fcc503addc887d809ff9a3e921
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Dec 7 22:14:09 2022 +0800

    wifi: rtlwifi: rtl8821ae: don't call kfree_skb() under spin_lock_irqsave()

    [ Upstream commit 106031c1f4a850915190d7ec1026696282f9359b ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. All the SKBs have
    been dequeued from the old queue, so it's safe to enqueue these
    SKBs to a free queue, then free them after spin_unlock_irqrestore()
    at once. Compile tested only.

    Fixes: 5c99f04fec93 ("rtlwifi: rtl8723be: Update driver to match Realtek release of
06/28/14")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221207141411.46098-2-yangyingliang@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ace789b1d465fae104cd37e49f6e1bcd1c8ff417
Author: Yuan Can <yuancan@huawei.com>
Date:   Mon Dec 5 06:14:41 2022 +0000

    wifi: rsi: Fix memory leak in rsi_coex_attach()

    [ Upstream commit 956fb851a6e19da5ab491e19c1bc323bb2c2cf6f ]

    The coex_cb needs to be freed when rsi_create_kthread() failed in
    rsi_coex_attach().

    Fixes: 2108df3c4b18 ("rsi: add coex support")
    Signed-off-by: Yuan Can <yuancan@huawei.com>
```

```
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20221205061441.114632-1-yuancan@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ead3cffd7510dc635d84cd4ea9dd1974fcb69a35
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Fri Dec 9 07:01:57 2022 +0800

    wifi: mt76: mt7921: resource leaks at mt7921_check_offload_capability()

    [ Upstream commit 47180ecf4541146836c5307c1d5526f8ac6a5a6d ]

    Fixed coverity issue with resource leaks at variable "fw" going out of
    scope leaks the storage it points to mt7921_check_offload_capability().

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527806 ("Resource leaks")
    Fixes: 034ae28b56f1 ("wifi: mt76: mt7921: introduce remain_on_channel support")
    Signed-off-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2e95d9922792b248f9bd049d89f3a601610adf82
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Thu Dec 8 00:03:10 2022 +0800

    wifi: mt76: fix coverity uninit_use_in_call in mt76_connac2_reverse_frag0_hdr_trans()

    [ Upstream commit 0ffcb2a68b15bd63d5555a923ae7dfe8bfdb14a7 ]

    The default case for frame_contorl is invalid. We should always
    assign addr3 of this frame properly.

    Coverity error message:
    if (ieee80211_has_a4(hdr.frame_control))
    (19) Event uninit_use_in_call:  Using uninitialized value "hdr".
    Field "hdr.addr3" is uninitialized when calling "memcpy".
            memcpy(skb_push(skb, sizeof(hdr)), &hdr, sizeof(hdr));
    else
            memcpy(skb_push(skb, sizeof(hdr) - 6), &hdr, sizeof(hdr) - 6);

    Fixes: 0880d40871d1 ("mt76: connac: move mt76_connac2_reverse_frag0_hdr_trans in mt76-
connac module")
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70f20a8e63e08691bc236073a0a4b576c8cf70ca
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Wed Dec 7 15:30:05 2022 +0800

    wifi: mt76: mt7915: fix unintended sign extension of mt7915_hw_queue_read()

    [ Upstream commit edb0406bda4629ef496f52eb11cbea7e92ed301b ]

    In the expression "map[i].qid << 24" starts as u8, but is promoted to
    "signed int", then sign-extended to type "unsigned long", which is not
    intended. Cast to u32 to avoid the sign extension.

    Fixes: 776ec4e77aa6 ("mt76: mt7915: rework debugfs queue info")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 946e01151b8264a5f4a147f6e752a36a172d38a3
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Wed Dec 7 15:30:04 2022 +0800

    wifi: mt76: mt7996: fix unintended sign extension of mt7996_hw_queue_read()

    [ Upstream commit 063cca0252b46970e7e9ca423d5e930be3179aa1 ]

    In the expression "map[i].qid << 24" starts as u8, but is promoted to
    "signed int", then sign-extended to type "unsigned long", which is not
    intended. Cast to u32 to avoid the sign extension.

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527813 ("Integer handling issues")
    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf425c5d7ef6fb4083c1e0d46440f886127b5ee5
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Mon Dec 5 17:59:04 2022 +0100

    wifi: mt76: mt76x0: fix oob access in mt76x0_phy_get_target_power

    [ Upstream commit 6e1abc51c945663bddebfa1beb9590ff5b250eb7 ]

    After 'commit ba45841ca5eb ("wifi: mt76: mt76x02: simplify struct
    mt76x02_rate_power")', mt76x02 relies on ht[0-7] rate_power data for
    vht mcs{0,7}, while it uses vth[0-1] rate_power for vht mcs {8,9}.
    Fix a possible out-of-bound access in mt76x0_phy_get_target_power routine.

    Fixes: ba45841ca5eb ("wifi: mt76: mt76x02: simplify struct mt76x02_rate_power")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e62f80ede601509d884add26f56e0530c8a1d75
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun Dec 4 23:01:37 2022 +0100

    wifi: mt76: mt7996: fix endianness warning in mt7996_mcu_sta_he_tlv

    [ Upstream commit 54ccb836ffb28eacba51b674cbe94cb5613f8441 ]

    Fix the following sparse warnings in mt7996_mcu_sta_he_tlv routine:

    warning: incorrect type in assignment (different base types)
       expected unsigned char
       got restricted __le16 [usertype]
    warning: incorrect type in assignment (different base types)
       expected unsigned char
       got restricted __le16 [usertype]

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4dbd32660640df1004a8b8b8e405a3bdb980e6cb
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sun Dec 4 15:18:15 2022 +0800
```

```
      wifi: mt76: mt7996: drop always true condition of __mt7996_reg_addr()

      [ Upstream commit ef1ea24cb0ecfd42c1ff266d92613163792aec77 ]

      addr <= MT_CBTOP2_PHY_END(0xffffffff) is always true (<= u32max),
      so drop it.

      Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
  devices")
      Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
      Signed-off-by: Felix Fietkau <nbd@nbd.name>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 3d51e056ca8d4b1208292d355c42c6d6f9b61c58
  Author: Ryder Lee <ryder.lee@mediatek.com>
  Date:   Sun Dec 4 15:18:14 2022 +0800

      wifi: mt76: mt7915: drop always true condition of __mt7915_reg_addr()

      [ Upstream commit b0f7b9563358493dfe70d3e4c3ebeffc92d4b494 ]

      smatch warnings:
      addr <= MT_CBTOP2_PHY_END(0xffffffff) is always true (<= u32max),
      so drop it.

      Fixes: cd4c314a65d3 ("mt76: mt7915: refine register definition")
      Reported-by: kernel test robot <lkp@intel.com>
      Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
      Signed-off-by: Felix Fietkau <nbd@nbd.name>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 278f9581f4fdf04c1c63ccc2f23610a1ec6e4017
  Author: Ryder Lee <ryder.lee@mediatek.com>
  Date:   Sun Dec 4 05:33:18 2022 +0800

      wifi: mt76: mt7996: check return value before accessing free_block_num

      [ Upstream commit 783ef7da7a53c7ab98471f47fbabab6cf6f45c8f ]

      Check return value of mt7996_mcu_get_eeprom_free_block() first before
      accessing free_block_num.

      Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
  devices")
      Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
      Signed-off-by: Felix Fietkau <nbd@nbd.name>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 5911840188891b2675ec981eae33547272b991e7
  Author: Ryder Lee <ryder.lee@mediatek.com>
  Date:   Sun Dec 4 05:33:17 2022 +0800

      wifi: mt76: mt7915: check return value before accessing free_block_num

      [ Upstream commit 59b27a7d472f100ac8998e15a63c47a03cced12a ]

      Check return value of mt7915_mcu_get_eeprom_free_block() first before
      accessing free_block_num.

      Fixes: bbc1d4154ec1 ("mt76: mt7915: add default calibrated data support")
      Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
      Signed-off-by: Felix Fietkau <nbd@nbd.name>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b1029145cca11ad02a4df261399195a57b6767e4
```

```
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Dec 3 11:22:25 2022 +0800

    wifi: mt76: mt7996: fix integer handling issue of mt7996_rf_regval_set()

    [ Upstream commit aab169ad3a7aa8678aed90d7fdbc243e3d4b32a6 ]

    This code is supposed to set a u32 value, but casting will not work on
    big endian systems.

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527816 ("Integer handling issues")
    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2861e7b1fa277cb78a302caf202bed2257fd4a12
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Dec 3 11:22:24 2022 +0800

    wifi: mt76: mt7996: fix insecure data handling of mt7996_mcu_rx_radar_detected()

    [ Upstream commit f37c6e5c75029443bc72c45acf92b2f2de2945be ]

    Coverity message:
    using tainted "r->band_idx" variable as an index into an array "(*dev).mt76.phys".

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527812 ("Insecure data handling")
    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 39f572e95818c1dec3a47d8f5dd927a6e206fd1b
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Dec 3 11:22:23 2022 +0800

    wifi: mt76: mt7996: fix insecure data handling of mt7996_mcu_ie_countdown()

    [ Upstream commit 5202b983f9894d31110e49c4ec6b57955b5eaa1a ]

    Coverity message:
    using tainted "hdr->band" variable as an index into an array "(*dev).mt76.phys".

    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527797 ("Insecure data handling")
    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ecff00802f60fa5d08c133d209ed984a455a41f8
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Dec 3 11:22:22 2022 +0800

    wifi: mt76: mt7915: fix mt7915_rate_txpower_get() resource leaks

    [ Upstream commit 8b25301af01566f4b5a301fc1ad7c5d2b1788d7f ]

    Coverity message: variable "buf" going out of scope leaks the storage.
```

```
    Reported-by: coverity-bot <keescook+coverity-bot@chromium.org>
    Addresses-Coverity-ID: 1527799 ("Resource leaks")
    Fixes: e3296759f347 ("wifi: mt76: mt7915: enable per bandwidth power limit support")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b358e36433d2c46a65488a146bf8b4623fc5bbb
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Thu Dec 1 23:53:37 2022 +0800

    wifi: mt76: mt7921s: fix slab-out-of-bounds access in sdio host

    [ Upstream commit aec4cf2ea0797e28f18f8dbe01943a56d987fe56 ]

    SDIO may need addtional 511 bytes to align bus operation. If the tailroom
    of this skb is not big enough, we would access invalid memory region.
    For low level operation, increase skb size to keep valid memory access in
    SDIO host.

    Error message:
    [69.951] BUG: KASAN: slab-out-of-bounds in sg_copy_buffer+0xe9/0x1a0
    [69.951] Read of size 64 at addr ffff88811c9cf000 by task kworker/u16:7/451
    [69.951] CPU: 4 PID: 451 Comm: kworker/u16:7 Tainted: G W  OE  6.1.0-rc5 #1
    [69.951] Workqueue: kvub300c vub300_cmndwork_thread [vub300]
    [69.951] Call Trace:
    [69.951]  <TASK>
    [69.952]  dump_stack_lvl+0x49/0x63
    [69.952]  print_report+0x171/0x4a8
    [69.952]  kasan_report+0xb4/0x130
    [69.952]  kasan_check_range+0x149/0x1e0
    [69.952]  memcpy+0x24/0x70
    [69.952]  sg_copy_buffer+0xe9/0x1a0
    [69.952]  sg_copy_to_buffer+0x12/0x20
    [69.952]  __command_write_data.isra.0+0x23c/0xbf0 [vub300]
    [69.952]  vub300_cmndwork_thread+0x17f3/0x58b0 [vub300]
    [69.952]  process_one_work+0x7ee/0x1320
    [69.952]  worker_thread+0x53c/0x1240
    [69.952]  kthread+0x2b8/0x370
    [69.952]  ret_from_fork+0x1f/0x30
    [69.952]  </TASK>

    [69.952] Allocated by task 854:
    [69.952]  kasan_save_stack+0x26/0x50
    [69.952]  kasan_set_track+0x25/0x30
    [69.952]  kasan_save_alloc_info+0x1b/0x30
    [69.952]  __kasan_kmalloc+0x87/0xa0
    [69.952]  __kmalloc_node_track_caller+0x63/0x150
    [69.952]  kmalloc_reserve+0x31/0xd0
    [69.952]  __alloc_skb+0xfc/0x2b0
    [69.952]  __mt76_mcu_msg_alloc+0xbf/0x230 [mt76]
    [69.952]  mt76_mcu_send_and_get_msg+0xab/0x110 [mt76]
    [69.952]  __mt76_mcu_send_firmware.cold+0x94/0x15d [mt76]
    [69.952]  mt76_connac_mcu_send_ram_firmware+0x415/0x54d [mt76_connac_lib]
    [69.952]  mt76_connac2_load_ram.cold+0x118/0x4bc [mt76_connac_lib]
    [69.952]  mt7921_run_firmware.cold+0x2e9/0x405 [mt7921_common]
    [69.952]  mt7921s_mcu_init+0x45/0x80 [mt7921s]
    [69.953]  mt7921_init_work+0xe1/0x2a0 [mt7921_common]
    [69.953]  process_one_work+0x7ee/0x1320
    [69.953]  worker_thread+0x53c/0x1240
    [69.953]  kthread+0x2b8/0x370
    [69.953]  ret_from_fork+0x1f/0x30
    [69.953] The buggy address belongs to the object at ffff88811c9ce800
                 which belongs to the cache kmalloc-2k of size 2048
```

```
    [69.953] The buggy address is located 0 bytes to the right of
                2048-byte region [ffff88811c9ce800, ffff88811c9cf000)

    [69.953] Memory state around the buggy address:
    [69.953]  ffff88811c9cef00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    [69.953]  ffff88811c9cef80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    [69.953] >ffff88811c9cf000: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
    [69.953]                    ^
    [69.953]  ffff88811c9cf080: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
    [69.953]  ffff88811c9cf100: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc

    Fixes: 764dee47e2c1 ("mt76: sdio: move common code in mt76_sdio module")
    Suggested-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Tested-by: YN Chen <YN.Chen@mediatek.com>
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d5528c89e7f6689713fdcedab4ff5f590a2b65fd
Author: Wang Yufen <wangyufen@huawei.com>
Date:   Fri Nov 25 17:06:07 2022 +0800

    wifi: mt76: mt7915: add missing of_node_put()

    [ Upstream commit 18425d7d74c5be88b13b970a21e52e2498abf4ba ]

    Add missing of_node_put() after of_reserved_mem_lookup()

    Fixes: 99ad32a4ca3a ("mt76: mt7915: add support for MT7986")
    Signed-off-by: Wang Yufen <wangyufen@huawei.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00219d96c48b9d1fcdcf46438fdb55c1a31d2296
Author: Jens Axboe <axboe@kernel.dk>
Date:   Thu Feb 16 19:39:15 2023 -0700

    block: use proper return value from bio_failfast()

    [ Upstream commit f3ca73862453ac1e64fc6968a14bf66d839cd2d8 ]

    kernel test robot complains about a type mismatch:

       block/blk-merge.c:984:42: sparse:     expected restricted blk_opf_t const [usertype]
ff
       block/blk-merge.c:984:42: sparse:     got unsigned int
       block/blk-merge.c:1010:42: sparse: sparse: incorrect type in initializer (different
base types) @@     expected restricted blk_opf_t const [usertype] ff @@     got unsigned int
@@
       block/blk-merge.c:1010:42: sparse:     expected restricted blk_opf_t const [usertype]
ff
       block/blk-merge.c:1010:42: sparse:     got unsigned int

    because bio_failfast() is return an unsigned int rather than the
    appropriate blk_opt_f type. Fix it up.

    Fixes: 3ce6a115980c ("block: sync mixed merged request's failfast with 1st bio's")
    Reported-by: kernel test robot <lkp@intel.com>
    Link: https://lore.kernel.org/oe-kbuild-all/202302170743.GXypM9Rt-lkp@intel.com/
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5d44ebbd588782f49a0323d9b838efe4f4e95a19
Author: Martin K. Petersen <martin.petersen@oracle.com>
Date:   Wed Feb 15 12:18:01 2023 -0500
```

```
    block: bio-integrity: Copy flags when bio_integrity_payload is cloned

    [ Upstream commit b6a4bdcda430e3ca43bbb9cb1d4d4d34ebe15c40 ]

    Make sure to copy the flags when a bio_integrity_payload is cloned.
    Otherwise per-I/O properties such as IP checksum flag will not be
    passed down to the HBA driver. Since the integrity buffer is owned by
    the original bio, the BIP_BLOCK_INTEGRITY flag needs to be masked off
    to avoid a double free in the completion path.

    Fixes: aae7df50190a ("block: Integrity checksum flag")
    Fixes: b1f01388574c ("block: Relocate bio integrity flags")
    Reported-by: Saurav Kashyap <skashyap@marvell.com>
    Tested-by: Saurav Kashyap <skashyap@marvell.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Chaitanya Kulkarni <kch@nvidia.com>
    Link: https://lore.kernel.org/r/20230215171801.21062-1-martin.petersen@oracle.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0187c0abe22cf40f6fa07cde6a8c7ecdcb1a672e
Author: Jinke Han <hanjinke.666@bytedance.com>
Date:   Thu Feb 16 11:22:50 2023 +0800

    block: Fix io statistics for cgroup in throttle path

    [ Upstream commit 0f7c8f0f7934c389b0f9fa1f151e753d8de6348f ]

    In the current code, io statistics are missing for cgroup when bio
    was throttled by blk-throttle. Fix it by moving the unreaching code
    to submit_bio_noacct_nocheck.

    Fixes: 3f98c753717c ("block: don't check bio in blk_throtl_dispatch_work_fn")
    Signed-off-by: Jinke Han <hanjinke.666@bytedance.com>
    Reviewed-by: Ming Lei <ming.lei@redhat.com>
    Acked-by: Muchun Song <songmuchun@bytedance.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230216032250.74230-1-hanjinke.666@bytedance.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f4a3b8597970b044e596485854261e1264a14a42
Author: Ming Lei <ming.lei@redhat.com>
Date:   Thu Feb 9 20:55:27 2023 +0800

    block: sync mixed merged request's failfast with 1st bio's

    [ Upstream commit 3ce6a115980c019928fcd06e01f64003886af79c ]

    We support mixed merge for requests/bios with different fastfail
    settings. When request fails, each time we only handle the portion
    with same failfast setting, then bios with failfast can be failed
    immediately, and bios without failfast can be retried.

    The idea is pretty good, but the current implementation has several
    defects:

    1) initially RA bio doesn't set failfast, however bio merge code
    doesn't consider this point, and just check its failfast setting for
    deciding if mixed merge is required. Fix this issue by adding helper
    of bio_failfast().

    2) when merging bio to request front, if this request is mixed
```

```
    merged, we have to sync request's faifast setting with 1st bio's
    failfast. Fix it by calling blk_update_mixed_merge().

    3) when merging bio to request back, if this request is mixed
    merged, we have to mark the bio as failfast, because blk_update_request
    simply updates request failfast with 1st bio's failfast. Fix
    it by calling blk_update_mixed_merge().

    Fixes one normal EXT4 READ IO failure issue, because it is observed
    that the normal READ IO is merged with RA IO, and the mixed merged
    request has different failfast setting with 1st bio's, so finally
    the normal READ IO doesn't get retried.

    Cc: Tejun Heo <tj@kernel.org>
    Fixes: 80a761fd33cf ("block: implement mixed merge of different failfast requests")
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Link: https://lore.kernel.org/r/20230209125527.667004-1-ming.lei@redhat.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0841d71ca71dafa72afd6d0f57d0be5c553b3d38
Author: Jingbo Xu <jefflexu@linux.alibaba.com>
Date:   Thu Feb 9 14:39:12 2023 +0800

    erofs: relinquish volume with mutex held

    [ Upstream commit 7032809a44d752b9e2275833787e0aa88a7540af ]

    Relinquish fscache volume with mutex held.  Otherwise if a new domain is
    registered when the old domain with the same name gets removed from the
    list but not relinquished yet, fscache may complain the collision.

    Fixes: 8b7adf1dff3d ("erofs: introduce fscache-based domain")
    Signed-off-by: Jingbo Xu <jefflexu@linux.alibaba.com>
    Reviewed-by: Jia Zhu <zhujia.zj@bytedance.com>
    Link: https://lore.kernel.org/r/20230209063913.46341-4-jefflexu@linux.alibaba.com
    Signed-off-by: Gao Xiang <hsiangkao@linux.alibaba.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 13e52f7afdc6458b12bde23078c154d9b5146d11
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Feb 13 22:29:30 2023 +0100

    arm64: dts: qcom: pmk8350: Use the correct PON compatible

    [ Upstream commit c0ee8e0ba5cc17623e63349a168b41e407b1eef0 ]

    A special compatible was introduced for PMK8350 both in the driver and
    the bindings to facilitate for 2 base registers (PBS & HLOS). Use it.

    Fixes: b2de43136058 ("arm64: dts: qcom: pmk8350: Add peripherals for pmk8350")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230213212930.2115182-1-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7caeed286338b7e71f0601c6a741f2700334ddc2
Author: Liu Xiaodong <xiaodong.liu@intel.com>
Date:   Fri Feb 10 09:13:56 2023 -0500

    block: ublk: check IO buffer based on flag need_get_data

    [ Upstream commit 2f1e07dda1e1310873647abc40bbc49eaf3b10e3 ]

    Currently, uring_cmd with UBLK_IO_FETCH_REQ or
```

```
    UBLK_IO_COMMIT_AND_FETCH_REQ is always checked whether
    userspace server has provided IO buffer even flag
    UBLK_F_NEED_GET_DATA is configured.

    This is a excessive check. If UBLK_F_NEED_GET_DATA is
    configured, FETCH_RQ doesn't need to provide IO buffer;
    COMMIT_AND_FETCH_REQ also doesn't need to do that if
    the IO type is not READ.

    Check ub_cmd->addr together with ublk_need_get_data()
    and IO type in ublk_ch_uring_cmd().

    With this fix, userspace server doesn't need to preserve
    buffers for every ublk_io when flag UBLK_F_NEED_GET_DATA
    is configured, in order to save memory.

    Signed-off-by: Liu Xiaodong <xiaodong.liu@intel.com>
    Fixes: c86019ff75c1 ("ublk_drv: add support for UBLK_IO_NEED_GET_DATA")
    Reviewed-by: Ming Lei <ming.lei@redhat.com>
    Link: https://lore.kernel.org/r/20230210141356.112321-1-xiaodong.liu@intel.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f4221b5e85dbd1d99cb0b2537add50b49c4f4e1d
Author: Denis Kenzior <denkenz@gmail.com>
Date:   Fri Aug 26 09:51:19 2022 -0500

    KEYS: asymmetric: Fix ECDSA use via keyctl uapi

    [ Upstream commit 10de7b54293995368c52d9aa153f3e7a359f04a1 ]

    When support for ECDSA keys was added, constraints for data & signature
    sizes were never updated.  This makes it impossible to use such keys via
    keyctl API from userspace.

    Update constraint on max_data_size to 64 bytes in order to support
    SHA512-based signatures. Also update the signature length constraints
    per ECDSA signature encoding described in RFC 5480.

    Fixes: 299f561a6693 ("x509: Add support for parsing x509 certs with ECDSA keys")
    Signed-off-by: Denis Kenzior <denkenz@gmail.com>
    Reviewed-by: Stefan Berger <stefanb@linux.ibm.com>
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9cdc137ea08cd24ca26a7260a73cf54e6856786
Author: silviazhao <silviazhao-oc@zhaoxin.com>
Date:   Wed Feb 8 16:27:22 2023 +0800

    x86/perf/zhaoxin: Add stepping check for ZXC

    [ Upstream commit fd636b6a9bc6034f2e5bb869658898a2b472c037 ]

    Some of Nano series processors will lead GP when accessing
    PMC fixed counter. Meanwhile, their hardware support for PMC
    has not announced externally. So exclude Nano CPUs from ZXC
    by checking stepping information. This is an unambiguous way
    to differentiate between ZXC and Nano CPUs.

    Following are Nano and ZXC FMS information:
    Nano FMS: Family=6, Model=F, Stepping=[0-A][C-D]
    ZXC FMS:  Family=6, Model=F, Stepping=E-F OR
              Family=6, Model=0x19, Stepping=0-3
```

```
    Fixes: 3a4ac121c2ca ("x86/perf: Add hardware performance events support for Zhaoxin
CPU.")

    Reported-by: Arjan <8vvbbqzo567a@nospam.xutrox.com>
    Reported-by: Kevin Brace <kevinbrace@gmx.com>
    Signed-off-by: silviazhao <silviazhao-oc@zhaoxin.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=212389
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 60892ab8bee94ca81ab97de40858a576a5d821fb
Author: Kan Liang <kan.liang@linux.intel.com>
Date:   Wed Jan 25 12:49:25 2023 -0800

    perf/x86/intel/ds: Fix the conversion from TSC to perf time

    [ Upstream commit 89e97eb8cec0f1af5ebf2380308913256ca7915a ]

    The time order is incorrect when the TSC in a PEBS record is used.

     $perf record -e cycles:upp dd if=/dev/zero of=/dev/null
      count=10000
     $ perf script --show-task-events
           perf-exec     0     0.000000: PERF_RECORD_COMM: perf-exec:915/915
                  dd   915   106.479872: PERF_RECORD_COMM exec: dd:915/915
                  dd   915   106.483270: PERF_RECORD_EXIT(915:915):(914:914)
                  dd   915   106.512429:          1 cycles:upp:
     ffffffff96c011b7 [unknown] ([unknown])
     ... ...

    The perf time is from sched_clock_cpu(). The current PEBS code
    unconditionally convert the TSC to native_sched_clock(). There is a
    shift between the two clocks. If the TSC is stable, the shift is
    consistent, __sched_clock_offset. If the TSC is unstable, the shift has
    to be calculated at runtime.

    This patch doesn't support the conversion when the TSC is unstable. The
    TSC unstable case is a corner case and very unlikely to happen. If it
    happens, the TSC in a PEBS record will be dropped and fall back to
    perf_event_clock().

    Fixes: 47a3aeb39e8d ("perf/x86/intel/pebs: Fix PEBS timestamps overwritten")
    Reported-by: Namhyung Kim <namhyung@kernel.org>
    Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lore.kernel.org/all/CAM9d7cgWDVAq8-11RbJ2uGfwkKD6fA-
OMwOKDrNUrU_=8MgEjg@mail.gmail.com/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1099004ae1664703ec573fc4c61ffb24144bcb63
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Mon Feb 6 22:33:54 2023 +0000

    sched/rt: pick_next_rt_entity(): check list_entry

    [ Upstream commit 7c4a5b89a0b5a57a64b601775b296abf77a9fe97 ]

    Commit 326587b84078 ("sched: fix goto retry in pick_next_task_rt()")
    removed any path which could make pick_next_rt_entity() return NULL.
    However, BUG_ON(!rt_se) in _pick_next_task_rt() (the only caller of
    pick_next_rt_entity()) still checks the error condition, which can
    never happen, since list_entry() never returns NULL.
    Remove the BUG_ON check, and instead emit a warning in the only
    possible error condition here: the queue being empty which should
    never happen.
```

```
    Fixes: 326587b84078 ("sched: fix goto retry in pick_next_task_rt()")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Phil Auld <pauld@redhat.com>
    Reviewed-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Link: https://lore.kernel.org/r/20230128-list-entry-null-check-sched-v3-1-
b1a71bd1ac6b@diag.uniroma1.it
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7df3d6b77e5193e7ca6fab7c14c4d601b52e35b8
Author: Richard Guy Briggs <rgb@redhat.com>
Date:   Wed Feb 1 15:33:33 2023 -0500

    io_uring,audit: don't log IORING_OP_MADVISE

    [ Upstream commit fbe870a72fd1ddc5e08c23764e23e5766f54aa87 ]

    fadvise and madvise both provide hints for caching or access pattern for
    file and memory respectively.  Skip them.

    Fixes: 5bd2182d58e9 ("audit,io_uring,io-wq: add some basic audit support to io_uring")
    Signed-off-by: Richard Guy Briggs <rgb@redhat.com>
    Link:
https://lore.kernel.org/r/b5dfdcd541115c86dbc774aa9dd502c964849c5f.1675282642.git.rgb@redhat
.com
    Acked-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aede5230d154b6b237985ec9df7ebbd1dce96810
Author: Qiheng Lin <linqiheng@huawei.com>
Date:   Fri Feb 10 01:02:53 2023 +0100

    s390/dasd: Fix potential memleak in dasd_eckd_init()

    [ Upstream commit 460e9bed82e49db1b823dcb4e421783854d86c40 ]

    `dasd_reserve_req` is allocated before `dasd_vol_info_req`, and it
    also needs to be freed before the error returns, just like the other
    cases in this function.

    Fixes: 9e12e54c7a8f ("s390/dasd: Handle out-of-space constraint")
    Signed-off-by: Qiheng Lin <linqiheng@huawei.com>
    Link: https://lore.kernel.org/r/20221208133809.16796-1-linqiheng@huawei.com
    Signed-off-by: Stefan Haberland <sth@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230210000253.1644903-3-sth@linux.ibm.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6025cb6ef365ada997c31177f90be7cb712c63d
Author: Petr Vorel <pvorel@suse.cz>
Date:   Fri Feb 3 11:09:52 2023 +0100

    arm64: dts: qcom: msm8992-lg-bullhead: Enable regulators

    [ Upstream commit 2866527093ddbc6356bb31f560f0b4b4decf3e2e ]

    Enable pm8994_s1, pm8994_l{26,29,30,32} regulators.
    Use values from downstream kernel on bullhead rev 1.01.

    NOTE: downstream kernel on angler rev 1.01 differences:
    * pm8994_l29: regulator-min-microvolt = <2700000>
    * pm8994_l{20,28,31}: use regulator-boot-on
```

```
    Verification:
    [    1.832460] s1: Bringing 0uV into 1025000-1025000uV
    ...
    [    2.057667] l26: Bringing 0uV into 987500-987500uV
    ...
    [    2.075722] l29: Bringing 0uV into 2800000-2800000uV
    [    2.076604] l30: Bringing 0uV into 1800000-1800000uV
    [    2.082431] l31: Bringing 0uV into 1262500-1262500uV
    [    2.095767] l32: Bringing 0uV into 1800000-1800000uV

    Fixes: f3b2c99e73be ("arm64: dts: Enable onboard SDHCI on msm8992")
    Signed-off-by: Petr Vorel <pvorel@suse.cz>
    Tested-by: Jamie Douglass <jamiemdouglass@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230203100952.13857-1-pvorel@suse.cz
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bcf342727b7ad7bae15c9f8c51d02f7571934920
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 2 11:44:52 2023 +0100

    arm64: dts: qcom: sm6115: correct TLMM gpio-ranges

    [ Upstream commit 272fc52431611d779de132c4e03de6e0c1199731 ]

    Correct the number of GPIOs in TLMM pin controller.

    Fixes: 97e563bf5ba1 ("arm64: dts: qcom: sm6115: Add basic soc dtsi")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Iskren Chernev <me@iskren.info>
    Acked-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230202104452.299048-11-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f61329e724ef7997a902fdecab0e38c8d6f7c0bf
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Feb 2 11:44:51 2023 +0100

    arm64: dts: qcom: msm8953: correct TLMM gpio-ranges

    [ Upstream commit a4fb71497df23cb0d02d70fa2b8f8786328e325d ]

    Correct the number of GPIOs in TLMM pin controller.

    Fixes: 9fb08c801923 ("arm64: dts: qcom: Add MSM8953 device tree")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Luca Weiss <luca@z3ntu.xyz>
    Acked-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230202104452.299048-10-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0dc611034e644d35e14708a3d4aa24717d0d87d8
Author: Jamie Douglass <jamiemdouglass@gmail.com>
Date:   Thu Feb 2 16:48:19 2023 +1100

    arm64: dts: qcom: msm8992-lg-bullhead: Correct memory overlaps with the SMEM and MPSS
memory regions

    [ Upstream commit d44106883d74992343710f18c4aaae937c7cefab ]

    The memory region reserved by a previous commit (see fixes tag below)
    overlaps with the SMEM and MPSS memory regions, causing error messages in
    dmesg:
```

```
              OF: reserved mem: OVERLAP DETECTED!
              reserved@5000000 (0x0000000005000000--0x0000000007200000)
              overlaps with smem_region@6a00000
              (0x0000000006a00000--0x0000000006c00000)

              OF: reserved mem: OVERLAP DETECTED!
              reserved@6c00000 (0x0000000006c00000--0x0000000007200000)
              overlaps with memory@7000000
              (0x0000000007000000--0x000000000ca00000)

      This patch resolves both of these by splitting the previously reserved
      memory region into two sections either side of the SMEM region and by
      cutting off the second memory region to 0x7000000.

      Fixes: 22c7e1a0fa45 ("arm64: dts: msm8992-bullhead: add memory hole region")
      Signed-off-by: Jamie Douglass <jamiemdouglass@gmail.com>
      Reviewed-by: Petr Vorel <pvorel@suse.cz>
      Tested-by: Petr Vorel <pvorel@suse.cz>
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20230202054819.16079-1-jamiemdouglass@gmail.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit dd35d6207cbd1f4653a13ef8623f391e9f3033b4
  Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
  Date:   Tue Jan 24 09:49:50 2023 +0100

      arm64: dts: qcom: sm8450: drop incorrect cells from serial

      [ Upstream commit 60d2da2c916956535cf37b7bf1ae8fefbf432e55 ]

      The serial/UART device node does not have children with unit addresses,
      so address/size cells are not correct.

      Fixes: f5837418479a ("arm64: dts: qcom: sm8450: add uart20 node")
      Fixes: 5188049c9b36 ("arm64: dts: qcom: Add base SM8450 DTSI")
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20230124084951.38195-3-krzysztof.kozlowski@linaro.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 6fca050624c8a4d85a575cbb4f85f8df50a41b48
  Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
  Date:   Tue Jan 24 09:49:49 2023 +0100

      arm64: dts: qcom: sm8350: drop incorrect cells from serial

      [ Upstream commit 6027331e6eae9eb957d1b73a7e3255f4151d6163 ]

      The serial/UART device node does not have children with unit addresses,
      so address/size cells are not correct.

      Fixes: cf03cd7e12bd ("arm64: dts: qcom: sm8350: Set up WRAP0 QUPs")
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20230124084951.38195-2-krzysztof.kozlowski@linaro.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit bac86c41081751587ab34b06853f6a1583c84566
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
  Date:   Fri Jan 20 08:14:15 2023 +0200

      arm64: dts: qcom: msm8996 switch from RPM_SMD_BB_CLK1 to RPM_SMD_XO_CLK_SRC
```

```
    [ Upstream commit 8ae72166c2b73b0f2ce498ea15d4feceb9fef50e ]

    The vendor kernel uses RPM_SMD_XO_CLK_SRC clock as an CXO clock rather
    than using the RPM_SMD_BB_CLK1 directly. Follow this example and switch
    msm8996.dtsi to use RPM_SMD_XO_CLK_SRC clock instead of RPM_SMB_BB_CLK1.

    Fixes: 2b8c9c77c268 ("arm64: dts: qcom: msm8996: convert xo_board to RPM_SMD_BB_CLK1")
    Suggested-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230120061417.2623751-7-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4c9cec2ca78dc383ecc33dbe06e6fb6e4234e540
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Fri Jan 13 14:05:44 2023 +0200

    arm64: dts: qcom: msm8996: support using GPLL0 as kryocc input

    [ Upstream commit ac0d84d4556cecf81ba0b1631d25d9a395235a5c ]

    In some cases the driver might need using GPLL0 to drive CPU clocks.
    Bring it in through the sys_apcs_aux clock.

    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113120544.59320-15-dmitry.baryshkov@linaro.org
    Stable-dep-of: 8ae72166c2b7 ("arm64: dts: qcom: msm8996 switch from RPM_SMD_BB_CLK1 to
RPM_SMD_XO_CLK_SRC")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97d9df3a6056c803d78bcf2550974575d208c812
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Wed Jan 18 17:37:26 2023 +0800

    blk-mq: correct stale comment of .get_budget

    [ Upstream commit 01542f651a9f58a9b176c3d3dc3eefbacee53b78 ]

    Commit 88022d7201e96 ("blk-mq: don't handle failure in .get_budget")
    remove BLK_STS_RESOURCE return value and we only check if we can get
    the budget from .get_budget() now.
    Correct stale comment that ".get_budget() returns BLK_STS_NO_RESOURCE"
    to ".get_budget() fails to get the budget".

    Fixes: 88022d7201e9 ("blk-mq: don't handle failure in .get_budget")
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d7c3d676434f7d5feeba9592a23687ef4297f50
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Wed Jan 18 17:37:16 2023 +0800

    blk-mq: Fix potential io hung for shared sbitmap per tagset

    [ Upstream commit 47df9ce95cd568d3f84218c4f65e9fbd4dfeda55 ]

    Commit f906a6a0f4268 ("blk-mq: improve tag waiting setup for non-shared
    tags") mark restart for unshared tags for improvement. At that time,
    tags is only shared betweens queues and we can check if tags is shared
    by test BLK_MQ_F_TAG_SHARED.
```

```
    Afterwards, commit 32bc15afed04b ("blk-mq: Facilitate a shared sbitmap per
    tagset") enabled tags share betweens hctxs inside a queue. We only
    mark restart for shared hctxs inside a queue and may cause io hung if
    there is no tag currently allocated by hctxs going to be marked restart.
    Wait on sbitmap_queue instead of mark restart for shared hctxs case to
    fix this.

    Fixes: 32bc15afed04 ("blk-mq: Facilitate a shared sbitmap per tagset")
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 559c1d69c1858cf7403c723da67181e134dcfd2f
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Wed Jan 18 17:37:15 2023 +0800

    blk-mq: wait on correct sbitmap_queue in blk_mq_mark_tag_wait

    [ Upstream commit 98b99e9412d0cde8c7b442bf5efb09528a2ede8b ]

    For shared queues case, we will only wait on bitmap_tags if we fail to get
    driver tag. However, rq could be from breserved_tags, then two problems
    will occur:
    1. io hung if no tag is currently allocated from bitmap_tags.
    2. unnecessary wakeup when tag is freed to bitmap_tags while no tag is
    freed to breserved_tags.
    Wait on the bitmap which rq from to fix this.

    Fixes: f906a6a0f426 ("blk-mq: improve tag waiting setup for non-shared tags")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec64867f8d8cb781d01e0e61c61245eb920d8761
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Wed Jan 18 17:37:14 2023 +0800

    blk-mq: remove stale comment for blk_mq_sched_mark_restart_hctx

    [ Upstream commit c31e76bcc379182fe67a82c618493b7b8868c672 ]

    Commit 97889f9ac24f8 ("blk-mq: remove synchronize_rcu() from
    blk_mq_del_queue_tag_set()") remove handle of TAG_SHARED in restart,
    then shared_hctx_restart counted for how many hardware queues are marked
    for restart is removed too.
    Remove the stale comment that we still count hardware queues need restart.

    Fixes: 97889f9ac24f ("blk-mq: remove synchronize_rcu() from blk_mq_del_queue_tag_set()")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6c2577612bddb5a616e55b99f039f2d97b1158b2
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Wed Jan 18 17:37:13 2023 +0800

    blk-mq: avoid sleep in blk_mq_alloc_request_hctx

    [ Upstream commit 6ee858a3d3270a68902d66bb47c151a83622535c ]

    Commit 1f5bd336b9150 ("blk-mq: add blk_mq_alloc_request_hctx") add
    blk_mq_alloc_request_hctx to send commands to a specific queue. If
```

```
    BLK_MQ_REQ_NOWAIT is not set in tag allocation, we may change to different
    hctx after sleep and get tag from unexpected hctx. So BLK_MQ_REQ_NOWAIT
    must be set in flags for blk_mq_alloc_request_hctx.
    After commit 600c3b0cea784 ("blk-mq: open code __blk_mq_alloc_request in
    blk_mq_alloc_request_hctx"), blk_mq_alloc_request_hctx return -EINVAL
    if both BLK_MQ_REQ_NOWAIT and BLK_MQ_REQ_RESERVED are not set instead of
    if BLK_MQ_REQ_NOWAIT is not set. So if BLK_MQ_REQ_NOWAIT is not set and
    BLK_MQ_REQ_RESERVED is set, blk_mq_alloc_request_hctx could alloc tag
    from unexpected hctx. I guess what we need here is that return -EINVAL
    if either BLK_MQ_REQ_NOWAIT or BLK_MQ_REQ_RESERVED is not set.

    Currently both BLK_MQ_REQ_NOWAIT and BLK_MQ_REQ_RESERVED will be set if
    specific hctx is needed in nvme_auth_submit, nvmf_connect_io_queue
    and nvmf_connect_admin_queue. Fix the potential BLK_MQ_REQ_NOWAIT missed
    case in future.

    Fixes: 600c3b0cea78 ("blk-mq: open code __blk_mq_alloc_request in
blk_mq_alloc_request_hctx")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eae792c478dd1f89d9b12198e5c4d02a657eddb7
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Feb 3 15:23:09 2023 +0100

    arm64: dts: qcom: sm8450-nagara: Correct firmware paths

    [ Upstream commit e27f38e6255306527e32af85592d805f3360ff94 ]

    Nagara is definitely not SM8350, fix it!

    Fixes: c53532f7825c ("arm64: dts: qcom: pdx223: correct firmware paths")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230203142309.1106349-1-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 205cbe7ff26f4e5a8312858066bc36370dc28553
Author: Patrick Delaunay <patrick.delaunay@foss.st.com>
Date:   Wed Jan 18 13:49:51 2023 +0100

    ARM: dts: stm32: Update part number NVMEM description on stm32mp131

    [ Upstream commit 366384e495511bea8583e44173629a3012d62db0 ]

    The STM32MP13x Device Part Number (also named RPN in reference manual)
    only uses the first 12 bits in OTP4, all the other bit are reserved and
    they can be different of zero; they must be masked in NVMEM result, so
    the number of bits must be defined in the nvmem cell description.

    Fixes: 1da8779c0029 ("ARM: dts: stm32: add STM32MP13 SoCs support")
    Signed-off-by: Patrick Delaunay <patrick.delaunay@foss.st.com>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edc4610f35571636a57f13dd816be51cb42e27ce
Author: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
Date:   Tue Nov 8 11:32:05 2022 +0800

    arm64: dts: mediatek: mt7986: Fix watchdog compatible

    [ Upstream commit 70d24df30d06e5c822ba94751166ef55d0e28a89 ]
```

```
    MT7986's watchdog embeds a reset controller and needs only the
    mediatek,mt7986-wdt compatible string as the MT6589 one is there
    for watchdogs that don't have any reset controller capability.

    Fixes: 50137c150f5f ("arm64: dts: mediatek: add basic mt7986 support")
    Signed-off-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Link: https://lore.kernel.org/r/20221108033209.22751-4-allen-kh.cheng@mediatek.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9b720f30d271931619cb11e2e626612f4cb5bba
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Tue Nov 8 11:32:04 2022 +0800

    arm64: dts: mediatek: mt8195: Fix watchdog compatible

    [ Upstream commit 02938f460cde0d360dde48056c4d1c0a4bd49230 ]

    MT8195's watchdog embeds a reset controller and needs only the
    mediatek,mt8195-wdt compatible string as the MT6589 one is there
    for watchdogs that don't have any reset controller capability.

    Fixes: 37f2582883be ("arm64: dts: Add mediatek SoC mt8195 and evaluation board")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Co-developed-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Signed-off-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Reviewed-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Link: https://lore.kernel.org/r/20221108033209.22751-3-allen-kh.cheng@mediatek.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3876dcb6f894191f685f3c15060dc38428c8b231
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Tue Nov 8 11:32:03 2022 +0800

    arm64: dts: mediatek: mt8186: Fix watchdog compatible

    [ Upstream commit e5e961628d696237ddc3d53d9d5ac11f43e0bf67 ]

    MT8186's watchdog embeds a reset controller and needs only the
    mediatek,mt8186-wdt compatible string as the MT6589 one is there
    for watchdogs that don't have any reset controller capability.

    Fixes: 2e78620b1350 ("arm64: dts: Add MediaTek MT8186 dts and evaluation board and
Makefile")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Co-developed-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Signed-off-by: Allen-KH Cheng <allen-kh.cheng@mediatek.com>
    Reviewed-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Link: https://lore.kernel.org/r/20221108033209.22751-2-allen-kh.cheng@mediatek.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0feed1a4debc8992bf60dee49f7d8e6a82b685a1
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Mon Nov 28 12:20:27 2022 +0100

    arm64: dts: mediatek: mt7622: Add missing pwm-cells to pwm node

    [ Upstream commit 22925af785fa3470efdf566339616d801119d348 ]

    Specify #pwm-cells on pwm@11006000 to make it actually usable.
```

```
    Fixes: ae457b7679c4 ("arm64: dts: mt7622: add SoC and peripheral related device nodes")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221128112028.58021-2-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64fdd028c907ff392b88e01032a0d3bbf71a1cbc
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Thu Jan 26 11:35:23 2023 +0100

    arm64: dts: mt8186: Fix CPU map for single-cluster SoC

    [ Upstream commit 1c473804b0c8a68c6ef2cf519b38ec6725ca4aa5 ]

    MT8186 features the ARM DynamIQ technology and combines both two
    Cortex-A76 (big) and six Cortex-A55 (LITTLE) CPUs in one cluster:
    fix the CPU map to reflect that.

    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Fixes: 2e78620b1350 ("arm64: dts: Add MediaTek MT8186 dts and evaluation board and
Makefile")
    Link: https://lore.kernel.org/r/20230126103526.417039-4-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 75ce55f17f899f1c0560b645501e4ddf9979898b
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Thu Jan 26 11:35:22 2023 +0100

    arm64: dts: mt8192: Fix CPU map for single-cluster SoC

    [ Upstream commit 160ce54d635455ffb5e9b42c5ba9cb9aaa98cdb2 ]

    MT8192 features the ARM DynamIQ technology and combines both four
    Cortex-A76 (big) and four Cortex-A55 (LITTLE) CPUs in one cluster:
    fix the CPU map to reflect that.

    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Fixes: 48489980e27e ("arm64: dts: Add Mediatek SoC MT8192 and evaluation board dts and
Makefile")
    Link: https://lore.kernel.org/r/20230126103526.417039-3-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b4489039e39c7986322ec0aaa64691dc0bcc71a6
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Thu Jan 26 11:35:21 2023 +0100

    arm64: dts: mt8195: Fix CPU map for single-cluster SoC

    [ Upstream commit cc4f0b13a887b483faa45084616998a21b63889d ]

    MT8195 features the ARM DynamIQ technology and combines both four
    Cortex-A78 (big) and four Cortex-A55 (LITTLE) CPUs in one cluster:
    fix the CPU map to reflect that.

    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Fixes: 37f2582883be ("arm64: dts: Add mediatek SoC mt8195 and evaluation board")
    Link: https://lore.kernel.org/r/20230126103526.417039-2-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5d22d1b2b9a77d70b4866d3dd47cce56e5710d6f
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Tue Jan 17 04:50:59 2023 +0800

    sbitmap: correct wake_batch recalculation to avoid potential IO hung

    [ Upstream commit b5fcf7871acb7f9a3a8ed341a68bd86aba3e254a ]

    Commit 180dccb0dba4f ("blk-mq: fix tag_get wait task can't be awakened")
    mentioned that in case of shared tags, there could be just one real
    active hctx(queue) because of lazy detection of tag idle. Then driver tag
    allocation may wait forever on this real active hctx(queue) if wake_batch
    is > hctx_max_depth where hctx_max_depth is available tags depth for the
    actve hctx(queue). However, the condition wake_batch > hctx_max_depth is
    not strong enough to avoid IO hung as the sbitmap_queue_wake_up will only
    wake up one wait queue for each wake_batch even though there is only one
    waiter in the woken wait queue. After this, there is only one tag to free
    and wake_batch may not be reached anymore. Commit 180dccb0dba4f ("blk-mq:
    fix tag_get wait task can't be awakened") methioned that driver tag
    allocation may wait forever. Actually, the inactive hctx(queue) will be
    truely idle after at most 30 seconds and will call blk_mq_tag_wakeup_all
    to wake one waiter per wait queue to break the hung. But IO hung for 30
    seconds is also not acceptable. Set batch size to small enough that depth
    of the shared hctx(queue) is enough to wake up all of the queues like
    sbq_calc_wake_batch do to fix this potential IO hung.

    Although hctx_max_depth will be clamped to at least 4 while wake_batch
    recalculation does not do the clamp, the wake_batch will be always
    recalculated to 1 when hctx_max_depth <= 4.

    Fixes: 180dccb0dba4 ("blk-mq: fix tag_get wait task can't be awakened")
    Reviewed-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Link: https://lore.kernel.org/r/20230116205059.3821738-6-shikemeng@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c3f9a4b2ca361868bcbb7c5b10e7991285b0e5d
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Tue Jan 17 04:50:56 2023 +0800

    sbitmap: remove redundant check in __sbitmap_queue_get_batch

    [ Upstream commit 903e86f3a64d9573352bbab2f211fdbbaa5772b7 ]

    Commit fbb564a557809 ("lib/sbitmap: Fix invalid loop in
    __sbitmap_queue_get_batch()") mentioned that "Checking free bits when
    setting the target bits. Otherwise, it may reuse the busying bits."
    This commit add check to make sure all masked bits in word before
    cmpxchg is zero. Then the existing check after cmpxchg to check any
    zero bit is existing in masked bits in word is redundant.

    Actually, old value of word before cmpxchg is stored in val and we
    will filter out busy bits in val by "(get_mask & ~val)" after cmpxchg.
    So we will not reuse busy bits methioned in commit fbb564a557809
    ("lib/sbitmap: Fix invalid loop in __sbitmap_queue_get_batch()"). Revert
    new-added check to remove redundant check.

    Fixes: fbb564a55780 ("lib/sbitmap: Fix invalid loop in __sbitmap_queue_get_batch()")
    Reviewed-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Link: https://lore.kernel.org/r/20230116205059.3821738-3-shikemeng@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit ad781e6bc538779982b0b8616830bde71d57214c
Author: Peng Fan <peng.fan@nxp.com>
Date:   Fri Dec 23 11:04:33 2022 +0800

    ARM: dts: imx7s: correct iomuxc gpr mux controller cells

    [ Upstream commit 0e3e1946606a2919b1dda9967ab2e1c5af2fedd6 ]

    Per binding doc reg-mux.yaml, the #mux-control-cells should be 1

    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Marco Felsch <m.felsch@pengutronix.de>
    Fixes: 94a905a79f2c ("ARM: dts: imx7s: add multiplexer controls")
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17d0008b84a522bb4dd9e5006b6ea92159498449
Author: Ming Lei <ming.lei@redhat.com>
Date:   Fri Jan 6 12:17:07 2023 +0800

    ublk_drv: don't probe partitions if the ubq daemon isn't trusted

    [ Upstream commit 73a166d9749230d598320fdae3b687cdc0e2e205 ]

    If any ubq daemon is unprivileged, the ublk char device is allowed
    for unprivileged user actually, and we can't trust the current user,
    so not probe partitions.

    Fixes: 71f28f3136af ("ublk_drv: add io_uring based userspace block driver")
    Reviewed-by: ZiyangZhang <ZiyangZhang@linux.alibaba.com>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Link: https://lore.kernel.org/r/20230106041711.914434-3-ming.lei@redhat.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 431b82071d4225305a5406efcf0a3028f490f72d
Author: Ming Lei <ming.lei@redhat.com>
Date:   Fri Jan 6 12:17:06 2023 +0800

    ublk_drv: remove nr_aborted_queues from ublk_device

    [ Upstream commit ed878d1c1c641c4a6bd366658fc8e6bc842b80d1 ]

    No one uses 'nr_aborted_queues' any more, so remove it.

    Reviewed-by: ZiyangZhang <ZiyangZhang@linux.alibaba.com>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Link: https://lore.kernel.org/r/20230106041711.914434-2-ming.lei@redhat.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 73a166d97492 ("ublk_drv: don't probe partitions if the ubq daemon isn't
trusted")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f6be4faa4308ada8775ce9bbc57012d40d2fc19
Author: Samuel Holland <samuel@sholland.org>
Date:   Sat Dec 31 16:58:54 2022 -0600

    ARM: dts: sun8i: nanopi-duo2: Fix regulator GPIO reference

    [ Upstream commit 2177d4ae971f79b4a9a3c411f2fb8ae6113d1430 ]

    The property named in the schema is 'enable-gpios', not 'enable-gpio'.
    This makes no difference at runtime, because the regulator is marked as
    always-on, but it breaks validation.
```

```
    Fixes: 4701fc6e5dd9 ("ARM: dts: sun8i: add FriendlyARM NanoPi Duo2")
    Reviewed-by: Andre Przywara <andre.przywara@arm.com>
    Acked-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Signed-off-by: Samuel Holland <samuel@sholland.org>
    Link: https://lore.kernel.org/r/20221231225854.16320-2-samuel@sholland.org
    Signed-off-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 801138b36d016f3e66de9f010d022f1927e0e113
Author: Christian Hewitt <christianshewitt@gmail.com>
Date:   Fri Jan 27 14:22:17 2023 +0000

    arm64: dts: meson: bananapi-m5: switch VDDIO_C pin to OPEN_DRAIN

    [ Upstream commit 856968e066bd77b113965f1a355ec7401edff65f ]

    For proper warm (re)boot from SD card the BPI-M5 board requires TFLASH_VDD_EN
    and VDDIO_C pins to be switched to high impedance mode. This can be achieved
    using OPEN_DRAIN instead of ACTIVE_HIGH to leave the GPIO pins in input mode
    and retain high state (pin has the pull-up).

    This change is inspired by meson-sm1-odroid.dtsi where OPEN_DRAIN has been
    used to resolve similar problems with the Odroid C4 board (TF_IO in the C4
    dts is the equivalent regulator).

    Fixes: 976e920183e4 ("arm64: dts: meson-sm1: add Banana PI BPI-M5 board dts")
    Suggested-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Christian Hewitt <christianshewitt@gmail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230127142221.3718184-2-christianshewitt@gmail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3680b694baa72b6dd2dca18f2c079c1f8a4e0794
Author: Christian Hewitt <christianshewitt@gmail.com>
Date:   Fri Jan 27 10:39:13 2023 +0000

    arm64: dts: meson: radxa-zero: allow usb otg mode

    [ Upstream commit ce43ea00b927805c1fd0450ccc9b4b6069e292c5 ]

    Setting dr_mode to "host" prevents otg which can be useful on a board
    with limited connectivity options. So don't force host mode.

    Fixes: 26d1400f7457 ("arm64: dts: amlogic: add support for Radxa Zero")
    Signed-off-by: Christian Hewitt <christianshewitt@gmail.com>
    Acked-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230127103913.3386435-1-christianshewitt@gmail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 08f4f3bdc836a76c3b8a4e88b724e416ca04b381
Author: Adam Ford <aford173@gmail.com>
Date:   Sat Jan 14 16:56:45 2023 -0600

    arm64: dts: renesas: beacon-renesom: Fix gpio expander reference

    [ Upstream commit d7f9492dfc03153ac56ab59066a196558748f575 ]

    The board used to originally introduce the Beacon Embedded RZ/G2[M/N/H]
    boards had a GPIO expander with address 20, but this was changed when
    the final board went to production.

    The production boards changed both the part itself and the address.
```

```
    With the incorrect address, the LCD cannot come up.  If the LCD fails,
    the rcar-du driver fails to come up, and that also breaks HDMI.

    Pre-release board were not shipped to the general public, so it should
    be safe to push this as a fix.  Anyone with a production board would
    have video fail due to this GPIO expander change.

    Fixes: a1d8a344f1ca ("arm64: dts: renesas: Introduce r8a774a1-beacon-rzg2m-kit")
    Signed-off-by: Adam Ford <aford173@gmail.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20230114225647.227972-1-aford173@gmail.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17d0b6fd7e6d3e5d30ef4ce9d80294eaa15bd24f
Author: Mikko Perttunen <mperttunen@nvidia.com>
Date:   Thu Jan 19 15:38:58 2023 +0200

    arm64: tegra: Mark host1x as dma-coherent on Tegra194/234

    [ Upstream commit 361238cdc52523fd7b1f3aa447c0579f42448b00 ]

    Ensure appropriate configuration is done to make the host1x device
    and context devices DMA coherent by adding the dma-coherent flag.

    Fixes: b35f5b53a87b ("arm64: tegra: Add context isolation domains on Tegra234")
    Signed-off-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 933fc8f3d8831092042be85f32a4fee10d7fa241
Author: Waiman Long <longman@redhat.com>
Date:   Wed Jan 25 19:36:26 2023 -0500

    locking/rwsem: Disable preemption in all down_read*() and up_read() code paths

    [ Upstream commit 3f5245538a1964ae186ab7e1636020a41aa63143 ]

    Commit:

      91d2a812dfb9 ("locking/rwsem: Make handoff writer optimistically spin on owner")

    ... assumes that when the owner field is changed to NULL, the lock will
    become free soon. But commit:

      48dfb5d2560d ("locking/rwsem: Disable preemption while trying for rwsem lock")

    ... disabled preemption when acquiring rwsem for write.

    However, preemption has not yet been disabled when acquiring a read lock
    on a rwsem.  So a reader can add a RWSEM_READER_BIAS to count without
    setting owner to signal a reader, got preempted out by a RT task which
    then spins in the writer slowpath as owner remains NULL leading to live lock.

    One easy way to fix this problem is to disable preemption at all the
    down_read*() and up_read() code paths as implemented in this patch.

    Fixes: 91d2a812dfb9 ("locking/rwsem: Make handoff writer optimistically spin on owner")
    Reported-by: Mukesh Ojha <quic_mojha@quicinc.com>
    Suggested-by: Peter Zijlstra <peterz@infradead.org>
    Signed-off-by: Waiman Long <longman@redhat.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link: https://lore.kernel.org/r/20230126003628.365092-3-longman@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 280c8cfe5729ce5ede65273ecfdeb924c2606bad
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 12:39:08 2023 +0100

    arm64: dts: amlogic: meson-sm1-odroid-hc4: fix active fan thermal trip

    [ Upstream commit 1d2f14117aa7773efff50f832b85fc7779e586e0 ]

    Add an active trip tied to the on-board fan cooling device, which is better
    than describing it along the passive cooling maps.

    Fixes: 33b14f663df8 ("arm64: dts: meson: add initial device-tree for ODROID-HC4")
    Reported-by: Ricardo Pardini <ricardo@pardini.net>
    Link: https://lore.kernel.org/r/20230124-topic-odroid-hc4-upstream-fix-fan-trip-v1-1-
b0c6aa355d93@linaro.org
    Tested-by: Ricardo Pardini <ricardo@pardini.net>
    [narmstrong: added Ricardo's tested-by from off-list chat]
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9bc9f185e0bafdf848658d7ea418909c5f2680a4
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:35 2023 +0100

    arm64: dts: amlogic: meson-g12b-odroid-go-ultra: fix rk818 pmic properties

    [ Upstream commit ce9999722656f2433af9029571bc2b99448dda74 ]

    Fixes:
    pmic@1c: '#clock-cells' is a required property
    pmic@1c: 'switch-supply' does not match any of the regexes: 'pinctrl-[0-9]+'

    The switch supply is described by vcc9-supply per bindings documentation.

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-14-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e4c1d85675b6379a61839a8ffd0cc6fd58d70b3
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:34 2023 +0100

    arm64: dts: amlogic: meson-gxbb-kii-pro: fix led node name

    [ Upstream commit afdef3b188c934f79ad4b0a7bd8c692742f9b5af ]

    Fixes:
    leds: status: {...} is not of type 'array'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-13-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed1133865ef9435fb0d2e7f235429918029eb27f
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:33 2023 +0100

    arm64: dts: amlogic: meson-gxl-s905d-phicomm-n1: fix led node name

    [ Upstream commit eee64d8fbbdaab72bbab3e462f3a7b742d20c8c2 ]

    Fixes:
    leds: status: {...} is not of type 'array'
```

```
    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-12-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4590168a22d17561918e02772820d5c2786f5c6
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:32 2023 +0100

    arm64: dts: amlogic: meson-sm1-bananapi-m5: fix adc keys node names

    [ Upstream commit d519a73332b6c3d14e15f8fd20d7c6f29ed13d41 ]

    Fixes:
    adc_keys: 'key' does not match any of the regexes: '^button-', 'pinctrl-[0-9]+'

    Also fix the invalid "adc_keys" node name.

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-11-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 092728063954c50ccbcd211bafea3ca62587e760
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:31 2023 +0100

    arm64: dts: amlogic: meson-gx-libretech-pc: fix update button name

    [ Upstream commit 6bb506ed36968207a8832f0143ebc127f0770eef ]

    Fixes:
     adc-keys: 'update-button' does not match any of the regexes: '^button-', 'pinctrl-[0-
9]+'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-10-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff24f4cfe12cd16f91f98fdd94b965a55b51eb34
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:30 2023 +0100

    arm64: dts: amlogic: meson-gxl: add missing unit address to eth-phy-mux node name

    [ Upstream commit d19189f70ba596798ea49166d2d1ef36a8df5289 ]

    Fixes:
    bus@c8834000: eth-phy-mux: {...} should not be valid under {'type': 'object'}

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-9-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7bbd664c1dc697da4624b9cf7caea0781d619b68
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:29 2023 +0100

    arm64: dts: amlogic: meson-axg-jethome-jethub-j1xx: fix invalid rtc node name

    [ Upstream commit 956f52025c5dd92c80c12e31c99c854086a6fc55 ]
```

```
    Fixes:
    pcf8563@51: $nodename:0: 'pcf8563@51' does not match '^rtc(@.*|-[0-9a-f])*$'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-8-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9ce75ea6c5462ece9d076a78ab1a7f735f9d472
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:28 2023 +0100

    arm64: dts: amlogic: meson-gxl-s905w-jethome-jethub-j80: fix invalid rtc node name

    [ Upstream commit 11172a97c092eaeb0a65c6434df0fc73f886a495 ]

    Fixes:
    pcf8563@51: $nodename:0: 'pcf8563@51' does not match '^rtc(@.*|-[0-9a-f])*$'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-7-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63eae3e98bdc9a063c783dbcb64f250131003140
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:27 2023 +0100

    arm64: dts: amlogic: meson-gx: add missing unit address to rng node name

    [ Upstream commit 61ff70708b98a85516eccb3755084ac97b42cf48 ]

    Fixes:
    bus@c8834000: rng: {...} should not be valid under {'type': 'object'}

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-6-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b559693685afa94702181167dc865e476657043
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:26 2023 +0100

    arm64: dts: amlogic: meson-gxl-s905d-sml5442tw: drop invalid clock-names property

    [ Upstream commit e3bd275ccbacf5eb18eaa311cea39f8bf8655feb ]

    Fixes:
    bluetooth: 'clock-names' does not match any of the regexes: 'pinctrl-[0-9]+'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-5-
44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 49780ea4d75a0535cc3790c2d9bf15c1092ef14f
Author: Neil Armstrong <neil.armstrong@linaro.org>
Date:   Tue Jan 24 11:34:25 2023 +0100

    arm64: dts: amlogic: meson-axg-jethome-jethub-j1xx: fix supply name of USB controller
node

    [ Upstream commit a69cb1042cea840bc7b60fea1c26a6b259e68bf2 ]
```

```
    Fixes:
    usb@ffe09080: 'phy-supply' does not match any of the regexes: '^usb@[0-9a-f]+$',
 'pinctrl-[0-9]+'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-4-
 44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit de9ab0963f50762fdd8a0ba82df178f37c374e91
 Author: Neil Armstrong <neil.armstrong@linaro.org>
 Date:   Tue Jan 24 11:34:24 2023 +0100

    arm64: dts: amlogic: meson-gx: add missing SCPI sensors compatible

    [ Upstream commit 2ff650051493d5bdb6dd09d4c2850bb37db6be31 ]

    Fixes:
    scpi: sensors:compatible: 'oneOf' conditional failed, one must be fixed:
            ['amlogic,meson-gxbb-scpi-sensors'] is too short
            'arm,scpi-sensors' was expected

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-3-
 44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 8a03cf13f9b138e33ea9a2a09cfaadaf6f8ba039
 Author: Neil Armstrong <neil.armstrong@linaro.org>
 Date:   Tue Jan 24 11:34:23 2023 +0100

    arm64: dts: amlogic: meson-axg: fix SCPI clock dvfs node name

    [ Upstream commit 5b7069d72f03c92a0ab919725017394ebce03a81 ]

    Fixes:
    scpi: clocks: 'clock-controller' does not match any of the regexes: '^clocks-[0-9a-
 f]+$', 'pinctrl-[0-9]+'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-2-
 44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 8d833c25032179be461f49afad63c97962571a86
 Author: Neil Armstrong <neil.armstrong@linaro.org>
 Date:   Tue Jan 24 11:34:22 2023 +0100

    arm64: dts: amlogic: meson-gx: fix SCPI clock dvfs node name

    [ Upstream commit 127f79212b07c5d9a6657a87e3eafdd889335814 ]

    Fixes:
    scpi: clocks: 'clock-controller' does not match any of the regexes: '^clocks-[0-9a-
 f]+$', 'pinctrl-[0-9]+'

    Link: https://lore.kernel.org/r/20230124-b4-amlogic-bindings-fixups-v1-1-
 44351528957e@linaro.org
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 3e29428dc2f09abe8e34c2dc207fcb52fbbfe4a8
 Author: Angus Chen <angus.chen@jaguarmicro.com>
 Date:   Thu Jan 5 14:11:23 2023 +0800
```

```
      ARM: imx: Call ida_simple_remove() for ida_simple_get

      [ Upstream commit ebeb49f43c8952f12aa20f03f00d7009edc2d1c5 ]

      The function call ida_simple_get maybe fail,we should deal with it.
      And if ida_simple_get success ,it need to call ida_simple_remove also.
      BTW,devm_kasprintf can handle id is zero for consistency.

      Fixes: e76bdfd7403a ("ARM: imx: Added perf functionality to mmdc driver")
      Signed-off-by: Angus Chen <angus.chen@jaguarmicro.com>
      Signed-off-by: Shawn Guo <shawnguo@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c0ca74d0611757eed7a39d42449097416160736e
  Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
  Date:   Fri Jan 20 16:53:54 2023 +0100

      ARM: dts: exynos: correct wr-active property in Exynos3250 Rinato

      [ Upstream commit d15d2a617499882971ddb773a583015bf36fa492 ]

      The property is wr-active:

        exynos3250-rinato.dtb: fimd@11c00000: i80-if-timings: 'wr-act' does not match any of
  the regexes: 'pinctrl-[0-9]+'

      Fixes: b59b3afb94d4 ("ARM: dts: add fimd device support for exynos3250-rinato")
      Link: https://lore.kernel.org/r/20230120155404.323386-2-krzysztof.kozlowski@linaro.org
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 6186f6be6230d4372d8564258d94d9f283dbcf71
  Author: Vaishnav Achath <vaishnav.a@ti.com>
  Date:   Thu Jan 19 09:56:22 2023 +0530

      arm64: dts: ti: k3-j7200: Fix wakeup pinmux range

      [ Upstream commit 9ae21ac445e911e3541985c20052fc05d60f6879 ]

      The WKUP_PADCONFIG register region in J7200 has multiple non-addressable
      regions, split the existing wkup_pmx region as follows to avoid the
      non-addressable regions and include all valid WKUP_PADCONFIG registers.
      Also update references to old nodes with new ones.

      wkup_pmx0 -> 13 pins (WKUP_PADCONFIG 0 - 12)
      wkup_pmx1 -> 2 pins (WKUP_PADCONFIG 14 - 15)
      wkup_pmx2 -> 59 pins (WKUP_PADCONFIG 26 - 84)
      wkup_pmx3 -> 8 pins (WKUP_PADCONFIG 93 - 100)

      J7200 Datasheet (Table 6-106, Section 6.4 Pin Multiplexing) :
              https://www.ti.com/lit/ds/symlink/dra821u.pdf

      Fixes: d361ed88455f ("arm64: dts: ti: Add support for J7200 SoC")

      Signed-off-by: Vaishnav Achath <vaishnav.a@ti.com>
      Reviewed-by: Jayesh Choudhary <j-choudhary@ti.com>
      Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
      Link: https://lore.kernel.org/r/20230119042622.22310-1-vaishnav.a@ti.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b5cb36a0c14763cccf95598435dcc024f200725b
  Author: Arnd Bergmann <arnd@arndb.de>
  Date:   Wed Jan 18 10:02:12 2023 +0100

      ARM: s3c: fix s3c64xx_set_timer_source prototype
```

```
    [ Upstream commit 5bf52f5e4d12b8109f348cab60cb7d51092c4270 ]

    The prototype does not match the definition, as gcc-13 points
    out:

    arch/arm/mach-s3c/s3c64xx.c:169:13: error: conflicting types for
's3c64xx_set_timer_source' due to enum/integer mismatch; have 'void(unsigned int,  unsigned
int)' [-Werror=enum-int-mismatch]
      169 | void __init s3c64xx_set_timer_source(unsigned int event, unsigned int source)
          |             ^~~~~~~~~~~~~~~~~~~~~~~~
    In file included from arch/arm/mach-s3c/s3c64xx.c:50:
    arch/arm/mach-s3c/s3c64xx.h:62:20: note: previous declaration of
's3c64xx_set_timer_source' with type 'void(enum s3c64xx_timer_mode,  enum
s3c64xx_timer_mode)'
       62 | extern void __init s3c64xx_set_timer_source(enum s3c64xx_timer_mode event,
          |                    ^~~~~~~~~~~~~~~~~~~~~~~~

    Fixes: 4280506ac9bb ("ARM: SAMSUNG: Move all platforms to new clocksource driver")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20230118090224.2162863-1-arnd@kernel.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2a50ba19d3615f4b64a5b144d2cfab9475a95ac
Author: Stefan Wahren <stefan.wahren@i2se.com>
Date:   Fri Jan 13 21:58:41 2023 +0100

    ARM: bcm2835_defconfig: Enable the framebuffer

    [ Upstream commit afc8dd99840b7fb7190e769a893cda673bc3a907 ]

    Booting Linux on a Raspberry Pi based on bcm2835_defconfig there is
    no display activity.

    Enable CONFIG_FB which is nowadays required for CONFIG_FB_SIMPLE
    and CONFIG_FRAMEBUFFER_CONSOLE.

    Fixes: f611b1e7624c ("drm: Avoid circular dependencies for CONFIG_FB")
    Signed-off-by: Stefan Wahren <stefan.wahren@i2se.com>
    Link: https://lore.kernel.org/r/20230113205842.17051-1-stefan.wahren@i2se.com
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 160e14b4c4ab3c858d97c6b3b4bfe2b3dc8f6593
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Thu Dec 29 18:12:02 2022 +0800

    arm64: dts: mediatek: mt8192: Mark scp_adsp clock as broken

    [ Upstream commit 089cd717e6ef03cf9cf7865777d67775de41339b ]

    The scp_adsp clock controller is under the SCP_ADSP power domain. This
    power domain is currently not supported nor defined.

    Mark the clock controller as broken for now, to avoid the system from
    trying to access it, and causing the CPU or bus to stall.

    Fixes: 5d2b897bc6f5 ("arm64: dts: mediatek: Add mt8192 clock controllers")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Tested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Link: https://lore.kernel.org/r/20221229101202.1655924-1-wenst@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8b3d81a5bf51e9e12ebaad775bf951d2db7bc5fc
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Jan 19 11:57:54 2023 +0200

    ARM: OMAP1: call platform_device_put() in error case in omap1_dm_timer_init()

    [ Upstream commit 0414a100d6ab32721efa70ab55524540fdfe0ede ]

    If platform_device_add() is not called or failed, it should call
    platform_device_put() in error case.

    Fixes: 97933d6ced60 ("ARM: OMAP1: dmtimer: conversion to platform devices")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Message-Id: <20220701094602.2365099-1-yangyingliang@huawei.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33d89002381c32f4e0f010fda03b36b0d2b38544
Author: Christian Hewitt <christianshewitt@gmail.com>
Date:   Thu Jan 19 05:30:31 2023 +0000

    arm64: dts: meson: remove CPU opps below 1GHz for G12A boards

    [ Upstream commit 3cbd431c2b34d84605d358c8c57654193fd661fb ]

    Amlogic G12A devices experience CPU stalls and random board wedges when
    the system idles and CPU cores clock down to lower opp points. Recent
    vendor kernels include a change to remove 100-250MHz and other distro
    sources also remove the 500/667MHz points. Unless all 100-667Mhz opps
    are removed or the CPU governor forced to performance stalls are still
    observed, so let's remove them to improve stability and uptime.

    Fixes: b190056fa9ee ("arm64: dts: meson-g12a: add cpus OPP table")
    Signed-off-by: Christian Hewitt <christianshewitt@gmail.com>
    Link: https://lore.kernel.org/r/20230119053031.21400-1-christianshewitt@gmail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b87b8b4aad11bb73939c24158ebccad888744f2f
Author: Robert Marko <robimarko@gmail.com>
Date:   Fri Jan 13 17:44:49 2023 +0100

    arm64: dts: qcom: ipq8074: correct PCIe QMP PHY output clock names

    [ Upstream commit 0e8b90c0256cf9c9589e2cee517dedc987a34355 ]

    Current PCIe QMP PHY output name were changed in ("arm64: dts: qcom: Fix
    IPQ8074 PCIe PHY nodes") however it did not account for the fact that GCC
    driver is relying on the old names to match them as they are being used as
    the parent for the gcc_pcie0_pipe_clk and gcc_pcie1_pipe_clk.

    This broke parenting as GCC could not find the parent clock, so fix it by
    changing to the names that driver is expecting.

    Fixes: 942bcd33ed45 ("arm64: dts: qcom: Fix IPQ8074 PCIe PHY nodes")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113164449.906002-9-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6a45cbd92bfe5af1b21c5ad65e047446d28412f3
Author: Robert Marko <robimarko@gmail.com>
Date:   Fri Jan 13 17:44:48 2023 +0100
```

```
    arm64: dts: qcom: ipq8074: fix Gen3 PCIe node

    [ Upstream commit 3e83a9c41ab0244a45a4a2800b9adb8de0d15f82 ]

    IPQ8074 comes in 2 silicon versions:
    * v1 with 2x Gen2 PCIe ports and QMP PHY-s
    * v2 with 1x Gen3 and 1x Gen2 PCIe ports and QMP PHY-s

    v2 is the final and production version that is actually supported by the
    kernel, however it looks like PCIe related nodes were added for the v1 SoC.

    Finish the PCIe fixup by using the correct compatible, adding missing ATU
    register space, declaring max-link-speed, use correct ranges, add missing
    clocks and resets.

    Fixes: 33057e1672fe ("ARM: dts: ipq8074: Add pcie nodes")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113164449.906002-8-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89c25ddcdf24ebea5a4b8125f58030da459e266e
Author: Robert Marko <robimarko@gmail.com>
Date:   Fri Jan 13 17:44:43 2023 +0100

    arm64: dts: qcom: ipq8074: correct Gen2 PCIe ranges

    [ Upstream commit 2055cb7dccea16bafa3adf9c5e3216949512c34a ]

    Current ranges property set in Gen2 PCIe node is incorrect, replace it
    with the downstream 5.4 QCA kernel value.

    Fixes: 33057e1672fe ("ARM: dts: ipq8074: Add pcie nodes")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113164449.906002-3-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4af8ed2d6a842aae2e15aed328eb3061d8aa4f12
Author: Robert Marko <robimarko@gmail.com>
Date:   Fri Jan 13 17:44:42 2023 +0100

    arm64: dts: qcom: ipq8074: fix Gen3 PCIe QMP PHY

    [ Upstream commit 7ba33591b45f9d547a317e42f1c2acd19c925eb6 ]

    IPQ8074 comes in 2 silicon versions:
    * v1 with 2x Gen2 PCIe ports and QMP PHY-s
    * v2 with 1x Gen3 and 1x Gen2 PCIe ports and QMP PHY-s

    v2 is the final and production version that is actually supported by the
    kernel, however it looks like PCIe related nodes were added for the v1 SoC.

    Now that we have Gen3 QMP PHY support, we can start fixing the PCIe support
    by fixing the Gen3 QMP PHY node first.

    Change the compatible to the Gen3 QMP PHY, correct the register space start
    and size, add the missing misc PCS register space.

    Fixes: 33057e1672fe ("ARM: dts: ipq8074: Add pcie nodes")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113164449.906002-2-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5248b8c001db00f65802edcdd8cbb319bd44269d
Author: Robert Marko <robimarko@gmail.com>
Date:   Fri Jan 13 17:44:41 2023 +0100

    arm64: dts: qcom: ipq8074: fix Gen2 PCIe QMP PHY

    [ Upstream commit 100d9c94ccf15b02742c326cd04f422ab729153b ]

    Serdes register space sizes are incorrect, update them to match the
    actual sizes from downstream QCA 5.4 kernel.

    Fixes: 942bcd33ed45 ("arm64: dts: qcom: Fix IPQ8074 PCIe PHY nodes")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230113164449.906002-1-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b40421e41183382f27a5dbf98ccca4ba41a25281
Author: Robert Marko <robimarko@gmail.com>
Date:   Sun Jan 8 14:04:40 2023 +0100

    arm64: dts: qcom: ipq8074: correct USB3 QMP PHY-s clock output names

    [ Upstream commit 877cff3568c0f54511d77918ae16b2d6e9a0dfce ]

    It seems that clock-output-names for the USB3 QMP PHY-s where set without
    actually checking what is the GCC clock driver expecting, so clock core
    could never actually find the parents for usb0_pipe_clk_src and
    usb1_pipe_clk_src clocks in the GCC driver.

    So, correct the names to be what the driver expects so that parenting
    works.

    Before:
    gcc_usb0_pipe_clk_src                0        0        0   125000000          0     0
50000         Y
    gcc_usb1_pipe_clk_src                0        0        0   125000000          0     0
50000         Y

    After:
     usb3phy_0_cc_pipe_clk                1        1        0   125000000          0     0
50000         Y
        usb0_pipe_clk_src                 1        1        0   125000000          0     0
50000         Y
           gcc_usb0_pipe_clk              1        1        0   125000000          0     0
50000         Y
     usb3phy_1_cc_pipe_clk                1        1        0   125000000          0     0
50000         Y
        usb1_pipe_clk_src                 1        1        0   125000000          0     0
50000         Y
           gcc_usb1_pipe_clk              1        1        0   125000000          0     0
50000         Y

    Fixes: 5e09bc51d07b ("arm64: dts: ipq8074: enable USB support")
    Signed-off-by: Robert Marko <robimarko@gmail.com>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230108130440.670181-2-robimarko@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19d8eed19add0ec4bc65b251cb9535569d9a36c4
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Sun Jan 1 21:40:29 2023 +0200
```

```
    arm64: dts: qcom: msm8956: use SoC-specific compat for tsens

    [ Upstream commit 0b3aa9aa629c10928e86f6c9666a5e0b14655b2a ]

    The slope values used during tsens calibration differ between msm8976
    and msm8956 SoCs. Use SoC-specific compat value for the msm8956 SoC.

    Fixes: 0484d3ce0902 ("arm64: dts: qcom: Add DTS for MSM8976 and MSM8956 SoCs")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230101194034.831222-16-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3cdd762a60ae80d0c2281f3218782405d650db2c
Author: Petr Vorel <petr.vorel@gmail.com>
Date:   Mon Dec 26 19:54:39 2022 +0100

    arm64: dts: qcom: msm8992-bullhead: Disable dfps_data_mem

    [ Upstream commit 4dee5aa44b924036511a744ceb3abb1ceeb96bb6 ]

    It's disabled on downstream [1] thus not shown on downstream dmesg.

    Removing it fixes warnings on v6.1:

    [    0.000000] OF: reserved mem: OVERLAP DETECTED!
    [    0.000000] dfps_data_mem@3400000 (0x0000000003400000--0x0000000003401000) overlaps
with memory@3400000 (0x0000000003400000--0x0000000004600000)

    [1] https://android.googlesource.com/kernel/msm.git/+/android-
7.0.0_r0.17/arch/arm64/boot/dts/lge/msm8992-bullhead.dtsi#137

    Fixes: 976d321f32dc ("arm64: dts: qcom: msm8992: Make the DT an overlay on top of 8994")

    Signed-off-by: Petr Vorel <petr.vorel@gmail.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221226185440.440968-3-pevik@seznam.cz
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 076651247257e6490d29097451c017ea85b0e6a9
Author: Petr Vorel <petr.vorel@gmail.com>
Date:   Mon Dec 26 19:54:38 2022 +0100

    arm64: dts: qcom: msm8992-bullhead: Fix cont_splash_mem size

    [ Upstream commit 26a91359aea4d89e7d3646d806eed0f3755b74bd ]

    Original google firmware reports 12 MiB:
    [    0.000000] cma: Found cont_splash_mem@0, memory base 0x0000000003400000, size 12
MiB, limit 0xffffffffffffffff

    which is actually 12*1024*1024 = 0xc00000.

    This matches the aosp source [1]:
    &cont_splash_mem {
            reg = <0 0x03400000 0 0xc00000>;
    };

    Fixes: 3cb6a271f4b0 ("arm64: dts: qcom: msm8992-bullhead: Fix cont_splash_mem mapping")
    Fixes: 976d321f32dc ("arm64: dts: qcom: msm8992: Make the DT an overlay on top of 8994")

    [1] https://android.googlesource.com/kernel/msm.git/+/android-
7.0.0_r0.17/arch/arm64/boot/dts/lge/msm8992-bullhead.dtsi#141
```

```
    Signed-off-by: Petr Vorel <petr.vorel@gmail.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221226185440.440968-2-pevik@seznam.cz
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32c4a01d9b04a11e1a5abbfabf2cbb85bc50c7e6
Author: Thierry Reding <treding@nvidia.com>
Date:   Fri Nov 18 07:35:31 2022 +0100

    arm64: tegra: Fix duplicate regulator on Jetson TX1

    [ Upstream commit 29bcc1eaca315326d1cc883fbe9b451d1f9e3fa5 ]

    When the top-level regulators were renamed, the 1.2V camera regulator
    accidentally ended up with the same DT node name as the 1.8V camera
    regulator.

    Fixes: 097e01c61015 ("arm64: tegra: Rename top-level regulators")
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 80df6265c68033f9d1be231257d6dd26370ec025
Author: Dhruva Gole <d-gole@ti.com>
Date:   Tue Jan 3 11:18:40 2023 +0530

    arm64: dts: ti: k3-am62-main: Fix clocks for McSPI

    [ Upstream commit 6be5d8e5d1804eb4cec29cd8a85dc9cb18683b5d ]

    Fixes the clock Device ID's in the DT according to the tisci docs clock
    identifiers for AM62x

    Fixes: c37c58fdeb8a ("arm64: dts: ti: k3-am62: Add more peripheral nodes")
    Reviewed-by: Bryan Brattlof <bb@ti.com>
    Signed-off-by: Dhruva Gole <d-gole@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20230103054840.1133711-1-d-gole@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dca64f4bea7669f2056662e1f2776054d62f0153
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Thu Jan 12 20:43:32 2023 +0100

    cpuidle, intel_idle: Fix CPUIDLE_FLAG_IRQ_ENABLE *again*

    [ Upstream commit 6d9c7f51b1d9179bf7c3542267c656a934e8af23 ]

    So objtool found this bug:

      vmlinux.o: warning: objtool: intel_idle_irq+0x10c: call to trace_hardirqs_off() leaves
.noinstr.text section

    As per commit 32d4fd5751ea ("cpuidle,intel_idle: Fix CPUIDLE_FLAG_IRQ_ENABLE"):

      "must not have tracing in idle functions"

    Clearly people can't read and tinker along until splat dissapears.
    This straight up reverts commit d295ad34f236 ("intel_idle: Fix false
    positive RCU splats due to incorrect hardirqs state").

    It doesn't re-introduce the problem because preceding patches fixed it
    properly.
```

```
    Fixes: d295ad34f236 ("intel_idle: Fix false positive RCU splats due to incorrect
hardirqs state")
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Tested-by: Tony Lindgren <tony@atomide.com>
    Tested-by: Ulf Hansson <ulf.hansson@linaro.org>
    Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Acked-by: Frederic Weisbecker <frederic@kernel.org>
    Link: https://lore.kernel.org/r/20230112195540.434302128@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad1370df2f363845e93cb3be2f4e15d360f64095
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:50 2023 +0100

    arm64: dts: meson-gx: Fix the SCPI DVFS node name and unit address

    [ Upstream commit f189c869ad92787ddd753558bcbae89d75825bb6 ]

    Node names should be generic and use hyphens instead of underscores to
    not cause warnings. Also nodes without a reg property should not have a
    unit-address. Change the scpi_dvfs node to use clock-controller as node
    name without a unit address (since it does not have a reg property).

    Fixes: 70db166a2baa ("ARM64: dts: meson-gxbb: Add SCPI with cpufreq & sensors Nodes")
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-7-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cccc56de9d4fb5432b438478ff968d0d874d555c
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:49 2023 +0100

    arm64: dts: meson-g12a: Fix internal Ethernet PHY unit name

    [ Upstream commit e7303651bbc76c848007f1cfac1fbeaa65f600d1 ]

    Documentation/devicetree/bindings/net/ethernet-phy.yaml defines that the
    node name for Ethernet PHYs should match the following pattern:
      ^ethernet-phy(@[a-f0-9]+)?$

    Replace the underscore with a hyphen to adhere to this binding.

    Fixes: 280c17df8fbf ("arm64: dts: meson: g12a: add mdio multiplexer")
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-6-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20f281b8c8cb3cdd5af4a3d41f60534a32280ae6
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:48 2023 +0100

    arm64: dts: meson-gx: Fix Ethernet MAC address unit name

    [ Upstream commit 8ed5310356bfa47cc6bb4221ae6b21258c52e3d1 ]

    Unit names should use hyphens instead of underscores to not cause
    warnings.
```

```
    Fixes: bfe59f92d306 ("ARM64: dts: amlogic: gxbb: Enable NVMEM")
    Suggested-by: Vyacheslav Bocharov <adeep@lexina.in>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-5-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b969307e9786efe99e7867d64380e6f67aa31c4
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:47 2023 +0100

    arm64: dts: meson-axg: jethub-j1xx: Fix MAC address node names

    [ Upstream commit 2f66eeb06e3e8b1cac9e9093be3baadbac2709eb ]

    Node names should use hyphens instead of underscores to not cause
    warnings.

    Fixes: 59ec069d5055 ("arm64: dts: meson-axg: add support for JetHub D1p (j110)")
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-4-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69e890aa8dc1f2d27b433464c8a77221391f9cfe
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:46 2023 +0100

    arm64: dts: meson-gxl: jethub-j80: Fix Bluetooth MAC node name

    [ Upstream commit cb199de1d3aecb02556d8a6e26393015effa0a9f ]

    Node names should use hyphens instead of underscores to not cause
    warnings.

    Fixes: abfaae24ecf3 ("arm64: dts: meson-gxl: add support for JetHub H1")
    Suggested-by: Vyacheslav Bocharov <adeep@lexina.in>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-3-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 23965b460cb03ce0c5aeda5e780f95bde4a2a55d
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 11 22:13:45 2023 +0100

    arm64: dts: meson-gxl: jethub-j80: Fix WiFi MAC address node

    [ Upstream commit f95acdb2b4af21caae2c76a48e565158181386ca ]

    Unit addresses should be written using lower-case hex characters. Use
    wifi_mac@c to fix a yaml schema validation error once the eFuse
    dt-bindings have been converted to a yaml schema:
      efuse: Unevaluated properties are not allowed ('wifi_mac@C' was
      unexpected)

    Also node names should use hyphens instead of underscores as the latter
    can also cause warnings.
```

```
    Fixes: abfaae24ecf3 ("arm64: dts: meson-gxl: add support for JetHub H1")
    Acked-by: Vyacheslav Bocharov <adeep@lexina.in>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230111211350.1461860-2-
martin.blumenstingl@googlemail.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2566dda99eab388853537939e93d4fce9adf1b90
Author: Bjorn Andersson <quic_bjorande@quicinc.com>
Date:   Thu Jan 12 05:51:17 2023 -0800

    arm64: dts: qcom: sc8280xp: Vote for CX in USB controllers

    [ Upstream commit fe07640280cd29ac2997a617a1fb5487feef9387 ]

    Running GCC_USB30_*_MASTER_CLK at 200MHz requires CX at nominal level,
    not doing so results in occasional lockups. This was previously hidden
    by the fact that the display stack incorrectly voted for CX (instead of
    MMCX).

    Fixes: 152d1faf1e2f ("arm64: dts: qcom: add SC8280XP platform")
    Signed-off-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230112135117.3836655-1-quic_bjorande@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9f289a4930712e2e1a06cd943f4c14cd81722360
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Mon Jan 9 06:24:06 2023 +0200

    arm64: dts: qcom: msm8996-oneplus-common: drop vdda-supply from DSI PHY

    [ Upstream commit 1de4e112b97c77efb5cbee39db8541e33dd2b0d5 ]

    14nm DSI PHY has the only supply, vcca. Drop the extra vdda-supply.

    Fixes: 5a134c940cd3 ("arm64: dts: qcom: msm8996: add support for oneplus3(t)")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230109042406.312047-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfeeda0f655163be53082ef6141dee1878c9c678
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Tue Jan 10 06:21:26 2023 +0200

    arm64: dts: qcom: sdm845: make DP node follow the schema

    [ Upstream commit 19e509d5ea229244b9205f8671bb75b13738a1f9 ]

    Drop the #clock-cells (probably a leftover from the times before the DP
    PHY split)

    Fixes: eaac4e55a6f4 ("arm64: dts: qcom: sdm845: add displayport node")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230110042126.702147-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e7253fb93d246b9d8a15914dc9c96df9acafac37
```

```
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Fri Dec 23 14:21:21 2022 +0100

    arm64: dts: qcom: sm8450: correct Soundwire wakeup interrupt name

    [ Upstream commit 5eafe69af43d77cb117f27201076ee50f48363f1 ]

    The bindings expect second Soundwire interrupt to be "wakeup" (Linux
    driver takes by index):

      sm8450-hdk.dtb: soundwire-controller@33b0000: interrupt-names:1: 'wakeup' was expected

    Fixes: 14341e76dbc7 ("arm64: dts: qcom: sm8450: add Soundwire and LPASS")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221223132121.81130-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 621284ceedb4f4e3b361115313336d84b4da8613
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Tue Dec 13 11:19:19 2022 +0100

    arm64: dts: qcom: sc8280xp: correct SPMI bus address cells

    [ Upstream commit 76d9e8b4d54ae2cb91a68f0cb82624887de767a7 ]

    The SPMI bus uses two address cells and zero size cells (second reg
    entry - SPMI_USID - is not the size):

      spmi@c440000: #address-cells:0:0: 2 was expected

    Fixes: 152d1faf1e2f ("arm64: dts: qcom: add SC8280XP platform")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Reviewed-by: Johan Hovold <johan+linaro@kernel.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221213101921.47924-3-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 82dfa286f0b03093187a2dfb9070604f33d6a31e
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Tue Dec 13 11:19:18 2022 +0100

    arm64: dts: qcom: sc7280: correct SPMI bus address cells

    [ Upstream commit 8da3786a91e56fe0c4aeb2c2209744474af6e517 ]

    The SPMI bus uses two address cells and zero size cells (second reg
    entry - SPMI_USID - is not the size):

      spmi@c440000: #address-cells:0:0: 2 was expected

    Fixes: 14abf8dfe364 ("arm64: dts: qcom: sc7280: Add SPMI PMIC arbiter device for
SC7280")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221213101921.47924-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 464926d65c4a1a747638077ff302880c517273e0
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
```

```
Date:   Tue Dec 13 11:19:17 2022 +0100

    arm64: dts: qcom: sc7180: correct SPMI bus address cells

    [ Upstream commit 1f75745537222172f84783d369bbd1fb2d4b6414 ]

    The SPMI bus uses two address cells and zero size cells (second reg
    entry - SPMI_USID - is not the size):

      spmi@c440000: #address-cells:0:0: 2 was expected

    Fixes: 0f9dc5f09fbd ("arm64: dts: qcom: sc7180: Add SPMI PMIC arbiter device")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221213101921.47924-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0af3caccd2fcc4d275924cceeece8fb98c834ec
Author: Kishon Vijay Abraham I <kvijayab@amd.com>
Date:   Thu Jan 5 04:10:59 2023 +0000

    x86/acpi/boot: Do not register processors that cannot be onlined for x2APIC

    [ Upstream commit e2869bd7af608c343988429ceb1c2fe99644a01f ]

    Section 5.2.12.12 Processor Local x2APIC Structure in the ACPI v6.5
    spec mandates that both "enabled" and "online capable" Local APIC Flags
    should be used to determine if the processor is usable or not.

    However, Linux doesn't use the "online capable" flag for x2APIC to
    determine if the processor is usable. As a result, cpu_possible_mask has
    incorrect value and results in more memory getting allocated for per_cpu
    variables than it is going to be used.

    Make sure Linux parses both "enabled" and "online capable" flags for
    x2APIC to correctly determine if the processor is usable.

    Fixes: aa06e20f1be6 ("x86/ACPI: Don't add CPUs that are not online capable")
    Reported-by: Leo Duran <leo.duran@amd.com>
    Signed-off-by: Kishon Vijay Abraham I <kvijayab@amd.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Reviewed-by: Borislav Petkov (AMD) <bp@alien8.de>
    Reviewed-by: Zhang Rui <rui.zhang@intel.com>
    Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Link: https://lore.kernel.org/r/20230105041059.39366-1-kvijayab@amd.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a2ee25fc3abec40a4fe1440b55e0c394b680764e
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Dec 22 16:13:17 2022 +0100

    arm64: dts: qcom: sdm845-xiaomi-beryllium: fix audio codec interrupt pin name

    [ Upstream commit e5011447376e1b050847ccb2ef7933176ce4de41 ]

    The pin config entry should have a string, not number, for the GPIO used
    as WCD9340 audio codec interrupt.

    Fixes: dd6459a0890a ("arm64: dts: qcom: split beryllium dts into common dtsi and tianma
dts")
    Reported-by: Doug Anderson <dianders@chromium.org>
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
```

```
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221222151319.122398-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 862bc64e0279078fbfc89851423702e9540c0d7e
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Dec 22 16:13:16 2022 +0100

    arm64: dts: qcom: sdm845-db845c: fix audio codec interrupt pin name

    [ Upstream commit 740862bb5f59b93efb390a417995f88a64bdc323 ]

    The pin config entry should have a string, not number, for the GPIO used
    as WCD9340 audio codec interrupt.

    Fixes: 89a32a4e769c ("arm64: dts: qcom: db845c: add analog audio support")
    Reported-by: Doug Anderson <dianders@chromium.org>
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221222151319.122398-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01ba5cd9b31bd4ce31fb65f85091256a15696886
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Thu Dec 1 16:42:29 2022 +0800

    arm64: dts: mediatek: mt8186: Fix systimer 13 MHz clock description

    [ Upstream commit b391efba57ff085233d5ead5e01817bf4b71d999 ]

    The systimer block derives its 13 MHz clock by dividing the main 26 MHz
    oscillator clock by 2 internally. The 13 MHz clock is not a separate
    oscillator.

    Fix this by making the 13 MHz clock a divide-by-2 fixed factor clock,
    taking its input from the main 26 MHz oscillator.

    Fixes: 2e78620b1350 ("arm64: dts: Add MediaTek MT8186 dts and evaluation board and
Makefile")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221201084229.3464449-5-wenst@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fe54ab05284c5b9f00ac1d57489dad4f9e204f4b
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Thu Dec 1 16:42:28 2022 +0800

    arm64: dts: mediatek: mt8195: Fix systimer 13 MHz clock description

    [ Upstream commit 0f1c806b65d136a5fe0b88adad5ff1cb451fc401 ]

    The systimer block derives its 13 MHz clock by dividing the main 26 MHz
    oscillator clock by 2 internally, not through the TOPCKGEN clock
    controller.

    On the MT8195 this divider is set either by power-on-reset or by the
    bootloader. The bootloader may then make the divider unconfigurable to,
    but can be read out by, the operating system.

    Making the systimer block take the 26 MHz clock directly requires
```

```
    changing the implementations. As an ABI compatible fix, change the
    input clock of the systimer block a fixed factor divide-by-2 clock
    that takes the 26 MHz oscillator as its input.

    Fixes: 37f2582883be ("arm64: dts: Add mediatek SoC mt8195 and evaluation board")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221201084229.3464449-4-wenst@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae6110b0ae42cae71b517c2192a7f94eb47d09db
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Thu Dec 1 16:42:27 2022 +0800

    arm64: dts: mediatek: mt8192: Fix systimer 13 MHz clock description

    [ Upstream commit f19f68e56b0c6631984a9f5023035d4bd09612bb ]

    The systimer block derives its 13 MHz clock by dividing the main 26 MHz
    oscillator clock by 2 internally, not through the TOPCKGEN clock
    controller.

    On the MT8192 this divider is fixed to /2 and is not configurable.

    Making the systimer block take the 26 MHz clock directly requires
    changing the implementations. As an ABI compatible fix, change the
    input clock of the systimer block a fixed factor divide-by-2 clock
    that takes the 26 MHz oscillator as its input.

    Fixes: 48489980e27e ("arm64: dts: Add Mediatek SoC MT8192 and evaluation board dts and
Makefile")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221201084229.3464449-3-wenst@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69f15d8ba1088f130b601c4f76515b908aa8ed0a
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Thu Dec 1 16:42:26 2022 +0800

    arm64: dts: mediatek: mt8183: Fix systimer 13 MHz clock description

    [ Upstream commit ce8a06b5bac75ccce99c0cf91b96b767d64f28a7 ]

    The systimer block derives its 13 MHz clock by dividing the main 26 MHz
    oscillator clock by 2 internally, not through the TOPCKGEN clock
    controller.

    On the MT8183 this divider is set either by power-on-reset or by the
    bootloader. The bootloader may then make the divider unconfigurable to,
    but can be read out by, the operating system.

    Making the systimer block take the 26 MHz clock directly requires
    changing the implementations. As an ABI compatible fix, change the
    input clock of the systimer block a fixed factor divide-by-2 clock
    that takes the 26 MHz oscillator as its input.

    Fixes: 5bc8e2875ffb ("arm64: dts: mt8183: add systimer0 device node")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221201084229.3464449-2-wenst@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit a2bbaed6518a89ede39ca961744749f157a8a5a9
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Wed Dec 14 14:11:17 2022 +0100

    arm64: dts: mediatek: mt8195: Add power domain to U3PHY1 T-PHY

    [ Upstream commit a9f6721a3c92764582ed12296292fda4a7f2dd25 ]

    Assign power domain to the U3PHY1 T-PHY in otder to keep this PHY
    alive after unused PD shutdown and to be able to completely cut
    and restore power to it, for example, to save some power during
    system suspend/sleep.

    Fixes: 2b515194bf0c ("arm64: dts: mt8195: Add power domains controller")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20221214131117.108008-2-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a43a0987ac1d396b0696aa3ca77511256901ba8c
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Jan 3 19:37:49 2023 +0800

    fs: dlm: fix return value check in dlm_memory_init()

    [ Upstream commit 8113aa91360a013ebe00763bb0823b5a41b11c4d ]

    It should check 'cb_cache', after calling kmem_cache_create("dlm_cb").

    Fixes: 61bed0baa4db ("fs: dlm: use a non-static queue for callbacks")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4c22ee805202087c2553c9175968e9e922d75bc1
Author: Qiheng Lin <linqiheng@huawei.com>
Date:   Tue Nov 29 22:05:44 2022 +0800

    ARM: zynq: Fix refcount leak in zynq_early_slcr_init

    [ Upstream commit 9eedb910a3be0005b88c696a8552c0d4c9937cd4 ]

    of_find_compatible_node() returns a node pointer with refcount incremented,
    we should use of_node_put() on error path.
    Add missing of_node_put() to avoid refcount leak.

    Fixes: 3329659df030 ("ARM: zynq: Simplify SLCR initialization")
    Signed-off-by: Qiheng Lin <linqiheng@huawei.com>
    Link: https://lore.kernel.org/r/20221129140544.41293-1-linqiheng@huawei.com
    Signed-off-by: Michal Simek <michal.simek@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ae2deb463ee8007b1a9cacb3d19a9addfc54bec
Author: Marek Vasut <marex@denx.de>
Date:   Fri Dec 2 17:23:50 2022 +0100

    arm64: dts: imx8m: Align SoC unique ID node unit address

    [ Upstream commit ee0d68f219be8618f53d3f8808952e20525e3f30 ]

    Align the SoC unique ID DT node unit address with its reg property.

    Reviewed-by: Peng Fan <peng.fan@nxp.com>
```

```
    Fixes: cbff23797fa1 ("arm64: dts: imx8m: add NVMEM provider and consumer to read soc
unique ID")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b0bf08601dd9bf41b1b6329b7a65574d7105805
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Thu Dec 22 20:24:43 2022 +0100

    arm64: dts: qcom: sm6125-seine: Clean up gpio-keys (volume down)

    [ Upstream commit a9f6a13da473bb6c7406d2784d9e3792f6763cba ]

    - Remove autorepeat (leave key repetition to userspace);
    - Remove unneeded status = "okay" (this is the default);
    - Remove unneeded linux,input-type <EV_KEY> (this is the default for
      gpio-keys);
    - Allow the interrupt line for this button to be disabled;
    - Use a full, descriptive node name;
    - Set proper bias on the GPIO via pinctrl;
    - Sort properties;
    - Replace deprecated gpio-key,wakeup property with wakeup-source.

    Fixes: 82e1783890b7 ("arm64: dts: qcom: sm6125: Add support for Sony Xperia 10II")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221222192443.119103-1-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64b8e7dec189698aedbf8d0ec3796d8462274240
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Fri Dec 16 22:33:43 2022 +0100

    arm64: dts: qcom: sm6125: Reorder HSUSB PHY clocks to match bindings

    [ Upstream commit 8416262b0ea46d84767141b074748f4d4f37736a ]

    Reorder the clocks and corresponding names to match the QUSB2 phy
    schema, fixing the following CHECK_DTBS errors:

        arch/arm64/boot/dts/qcom/sm6125-sony-xperia-seine-pdx201.dtb: phy@1613000: clock-
names:0: 'cfg_ahb' was expected
                From schema: /newdata/aosp-
r/kernel/mainline/kernel/Documentation/devicetree/bindings/phy/qcom,qusb2-phy.yaml
        arch/arm64/boot/dts/qcom/sm6125-sony-xperia-seine-pdx201.dtb: phy@1613000: clock-
names:1: 'ref' was expected
                From schema: /newdata/aosp-
r/kernel/mainline/kernel/Documentation/devicetree/bindings/phy/qcom,qusb2-phy.yaml

    Fixes: cff4bbaf2a2d ("arm64: dts: qcom: Add support for SM6125")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Martin Botka <martin.botka@somainline.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221216213343.1140143-1-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6db7df875be872bd9e396959a3f7bba2e011ebdf
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Thu Dec 22 22:59:06 2022 +0100

    arm64: dts: qcom: sm6350-lena: Flatten gpio-keys pinctrl state
```

```
    [ Upstream commit a40f5ae1ea64ab9e981faf47c31817dc4d7923e4 ]

    Pinctrl states typically collate multiple related pins.  In the case of
    gpio-keys there's no hardware-defined relation at all except all pins
    representing a key; and especially on Sony's lena board there's only one
    pin regardless. Flatten it similar to other boards [1].

    As a drive-by fix, clean up the label string.

    [1]: https://lore.kernel.org/linux-arm-msm/11174eb6-0a9d-7df1-6f06-
da4010f76453@linaro.org/

    Fixes: 2b8bbe985659 ("arm64: dts: qcom: sm6350-lena: Include pm6350 and configure
buttons")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221222215906.324092-1-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 83da279cd23f7b418ee315a8cbdfd67abd39045a
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Dec 29 11:27:12 2022 +0100

    arm64: dts: qcom: sm8350-sagami: Rectify GPIO keys

    [ Upstream commit dcc7cd5c46ca5e7bb8e4910ed8259597439c7246 ]

    With enough pins set properly, the hardware buttons now also work
    like a charm.

    Fixes: c2721b0c23d9 ("arm64: dts: qcom: Add support for Xperia 1 III / 5 III")
    Tested-by: Marijn Suijten <marijn.suijten@somainline.org> # On Xperia 1 III and Xperia 5
III
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221229102712.983306-1-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit faabae0cdae04c4a336c2983d1dbdc8c6148227f
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Nov 18 16:20:28 2022 +0100

    arm64: dts: qcom: sm8350-sagami: Add GPIO line names for PMIC GPIOs

    [ Upstream commit 7c679f2a2af84edbec0c28171af8c42c6da9af14 ]

    Sony ever so graciously provides GPIO line names in their downstream
    kernel (though sometimes they are not 100% accurate and you can judge
    that by simply looking at them and with what drivers they are used).

    Add these to the PDX213&214 DTSIs to better document the hardware.

    Diff between 223 and 224:

    pm8350b
    <       gpio-line-names = "NC", /* GPIO_1 */
    >       gpio-line-names = "CAM_PWR_A_CS", /* GPIO_1 */
    <                         "NC",
    >                         "CAM_PWR_LD_EN",

    pm8350c
    <                         "NC",
    >                         "WLC_TXPWR_EN",
```

```
    Which is due to different camera power wiring on 213 and lack of an
    additional SLG51000 PMIC on 214.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221118152028.59312-3-konrad.dybcio@linaro.org
    Stable-dep-of: dcc7cd5c46ca ("arm64: dts: qcom: sm8350-sagami: Rectify GPIO keys")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1877117a132690733cb1f585e85234eaa8a87db8
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Nov 18 16:20:27 2022 +0100

    arm64: dts: qcom: sm8350-sagami: Configure SLG51000 PMIC on PDX215

    [ Upstream commit 8875b1d71f112b30e4c7e65ed337096bc0cc396b ]

    Remove the mention of this PMIC from the common DTSI, as it's not
    used on PDX214. Add the required nodes to support it on PDX215.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221118152028.59312-2-konrad.dybcio@linaro.org
    Stable-dep-of: dcc7cd5c46ca ("arm64: dts: qcom: sm8350-sagami: Rectify GPIO keys")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2eae55655f9d48d0d41f3516bb4d4207ca7ceb66
Author: Dzmitry Sankouski <dsankouski@gmail.com>
Date:   Wed Dec 28 14:52:43 2022 +0300

    arm64: dts: qcom: Re-enable resin on MSM8998 and SDM845 boards

    [ Upstream commit 4c881ab73a64cdbf8691e258ef17b740d27040a0 ]

    resin node declaration was moved to pm8998.dtsi file (in disabled state).
    MSM8998 and SDM845 boards defining resin node did not previously have
    status="okay" and ended up disabled.
    Re-enable it by using resin node link from pm8998.dtsi with status="okay".

    Fixes: f86ae6f23a9e ("arm64: dts: qcom: sagit: add initial device tree for sagit")
    Signed-off-by: Dzmitry Sankouski <dsankouski@gmail.com>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reported-by: Marijn Suijten <marijn.suijten@somainline.org>
    Link: https://lore.kernel.org/linux-arm-
msm/20221222115922.jlachctn4lxopp7a@SoMainline.org/
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221228115243.201038-1-dsankouski@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6040263f44f8ab73cd05d290aadb3d6a5fc8bf27
Author: Richard Acayan <mailingradian@gmail.com>
Date:   Mon Dec 5 17:52:37 2022 -0500

    arm64: dts: qcom: sdm670-google-sargo: keep pm660 ldo8 on

    [ Upstream commit 861b67fbdccd62a9319d7350b1924d95f597db09 ]

    According to the downstream device tree, the regulator that powers the
    I/O for eMMC should not be turned off. Keep it always on just in case
    the eMMC driver fails and doesn't enable it, or unloads and disables it.

    Fixes: 07c8ded6e373 ("arm64: dts: qcom: add sdm670 and pixel 3a device trees")
    Link:
https://android.googlesource.com/kernel/msm/+/9ed6ddbe955d3b84d1416a1cf77e83904d1e8421/arch/
```

```
  arm64/boot/dts/google/sdm670-bonito-common.dtsi#105
      Signed-off-by: Richard Acayan <mailingradian@gmail.com>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20221205225237.200564-1-mailingradian@gmail.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 55fe06e36876a2f66c42fc132cc7bf18199bd9d8
  Author: Konrad Dybcio <konrad.dybcio@linaro.org>
  Date:   Sat Dec 10 11:25:59 2022 +0100

      arm64: dts: qcom: sm6350: Fix up the ramoops node

      [ Upstream commit 3b2ff50da499178cc418f4b319e279d1b52958ed ]

      Fix up the ramoops node to make it match bindings and style:

      - remove "removed-dma-pool"
      - don't pad size to 8 hex digits
      - change cc-size to ecc-size so that it's used
      - increase ecc-size from to 16
      - remove the zeroed ftrace-size

      Fixes: 5f82b9cda61e ("arm64: dts: qcom: Add SM6350 device tree")
      Reported-by: Luca Weiss <luca.weiss@fairphone.com>
      Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20221210102600.589028-1-konrad.dybcio@linaro.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 724e4426f08842297e041e779bd5d49ec2b1b5b4
  Author: Marijn Suijten <marijn.suijten@somainline.org>
  Date:   Fri Dec 9 22:54:37 2022 +0100

      arm64: dts: qcom: pmi8950: Correct rev_1250v channel label to mv

      [ Upstream commit 3c3d2cb221b8647d1c547b4c44d2d6060cc742a9 ]

      This was pointed out in review but never followed up on thanks to
      sidetracked discussions about labels vs node names.

      Fixes: 0d97fdf380b4 ("arm64: dts: qcom: Add configuration for PMI8950 peripheral")
      Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Reviewed-by: Luca Weiss <luca@z3ntu.xyz>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20221209215437.1783067-1-marijn.suijten@somainline.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4b6eabd62b1d08593978f53b59e3b8f747bf8754
  Author: Marijn Suijten <marijn.suijten@somainline.org>
  Date:   Fri Dec 9 20:17:33 2022 +0100

      arm64: dts: qcom: sm8150-kumano: Panel framebuffer is 2.5k instead of 4k

      [ Upstream commit be8de06dc397c45cb0f3fe04084089c3f06c419f ]

      The framebuffer configuration for kumano griffin, written in kumano dtsi
      (which is overwritten in bahamut dts for its smaller panel) has to use a
      1096x2560 configuration as this is what the panel (and framebuffer area)
      has been initialized to.  Downstream userspace also has access to (and
      uses) this 2.5k mode by default, and only switches the panel to 4k when
      requested.

      Fixes: d0a6ce59ea4e ("arm64: dts: qcom: sm8150: Add support for SONY Xperia 1 / 5
  (Kumano platform)")
```

```
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221209191733.1458031-1-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0400d75106bb68da5db768e69b0f2066a489e4e
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Dec 8 21:13:58 2022 +0100

    arm64: dts: qcom: sm6115: Provide xo clk to rpmcc

    [ Upstream commit ad9514be8ddb9d3a8c262aa415c2f1c1f4cc97f9 ]

    rpmcc used to rely on global clock lookup (and still does so for
    backwards compat reasons) of "xo_board", which was common back
    when we did not care about things like underscores in node names.
    Nowadays it expects to be fed a reference to the fixed clock.
    Satisfy that requirement to make sure rpm clock rates are not all
    stuck at zero.

    Fixes: 97e563bf5ba1 ("arm64: dts: qcom: sm6115: Add basic soc dtsi")
    Reported-by: Adam Skladowski <a39.skl@gmail.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221208201401.530555-2-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b37aeb8c510ec6eed1de78c16e1ba22fc11c3f9
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Dec 8 21:13:57 2022 +0100

    arm64: dts: qcom: sm6115: Fix UFS node

    [ Upstream commit 01b6041454e8bc4f5feb76e6bcdc83a48cea21f2 ]

    In its current form, UFS did not even probe successfully - it failed
    when trying to set XO (ref_clk) to 300 MHz instead of doing so to
    the ICE clk. Moreover, the missing reg-names prevented ICE from
    working or being discovered at all. Fix both of these issues.

    As a sidenote, the log reveals that this SoC uses UFS ICE v3.1.0.

    Fixes: 97e563bf5ba1 ("arm64: dts: qcom: sm6115: Add basic soc dtsi")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Iskren Chernev <me@iskren.info>
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20221208201401.530555-1-konrad.dybcio@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ba8e8c28ebac831c0809abcd0d75f4131308a27
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Nov 24 23:01:47 2022 +0100

    arm64: dts: qcom: msm8996-tone: Fix USB taking 6 minutes to wake up

    [ Upstream commit 43069b9cd358aebc692e654de91ee06ff66e26af ]

    The hardware turns out to be pretty sluggish at assuming it can only
    do USB2 with just a USB2 phy assigned to it - before it needed about
    6 minutes to acknowledge that.

    Limit it to USB-HS explicitly to make USB come up about 720x faster.

    Fixes: 9da65e441d4d ("arm64: dts: qcom: Add support for SONY Xperia X Performance / XZ /
```

```
  XZs (msm8996, Tone platform)")
      Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20221124220147.102611-1-konrad.dybcio@linaro.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 664d30109a21f347d9df0f3862e1a55f6fe500a1
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
  Date:   Mon Dec 26 06:21:51 2022 +0200

      arm64: dts: qcom: qcs404: use symbol names for PCIe resets

      [ Upstream commit 41a37d157a613444c97e8f71a5fb2a21116b70d7 ]

      The commit e5bbbff5b7d7 ("clk: gcc-qcs404: Add PCIe resets") added names
      for PCIe resets, but it did not change the existing qcs404.dtsi to use
      these names. Do it now and use symbol names to make it easier to check
      and modify the dtsi in future.

      Fixes: e5bbbff5b7d7 ("clk: gcc-qcs404: Add PCIe resets")
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Link: https://lore.kernel.org/r/20221226042154.2666748-14-dmitry.baryshkov@linaro.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4f7ad1b08533247c4bf29217ba499ea4138cc2c1
  Author: Chen Hui <judy.chenhui@huawei.com>
  Date:   Tue Nov 8 22:19:17 2022 +0800

      ARM: OMAP2+: Fix memory leak in realtime_counter_init()

      [ Upstream commit ed8167cbf65c2b6ff6faeb0f96ded4d6d581e1ac ]

      The "sys_clk" resource is malloced by clk_get(),
      it is not released when the function return.

      Fixes: fa6d79d27614 ("ARM: OMAP: Add initialisation for the real-time counter.")
      Signed-off-by: Chen Hui <judy.chenhui@huawei.com>
      Message-Id: <20221108141917.46796-1-judy.chenhui@huawei.com>
      Signed-off-by: Tony Lindgren <tony@atomide.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit a2bb058f8600171b3f1549990f37203f33b50e31
  Author: Damien Le Moal <damien.lemoal@opensource.wdc.com>
  Date:   Fri Mar 3 18:29:07 2023 +0900

      ata: ahci: Revert "ata: ahci: Add Tiger Lake UP{3,4} AHCI controller"

      commit 6210038aeaf49c395c2da57572246d93ec67f6d4 upstream.

      Commit 104ff59af73a ("ata: ahci: Add Tiger Lake UP{3,4} AHCI
      controller") enabled low power mode for the Tiger Lake AHIC adapter in
      the author system but created regressions for others. Revert this patch
      for now until a better solution is found to make this adapter
      eco-friendly.

      Link: https://bugzilla.kernel.org/show_bug.cgi?id=217114
      CC: stable@vger.kernel.org
      Signed-off-by: Damien Le Moal <damien.lemoal@opensource.wdc.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit a4afda515608d087a71917c3ad39f40a58f92d95
  Author: Anders Roxell <anders.roxell@linaro.org>
```

```
Date:   Wed Aug 10 13:43:18 2022 +0200

    powerpc/mm: Rearrange if-else block to avoid clang warning

    commit d78c8e32890ef7eca79ffd67c96022c7f9d8cce4 upstream.

    Clang warns:

      arch/powerpc/mm/book3s64/radix_tlb.c:1191:23: error: variable 'hstart' is
uninitialized when used here
        __tlbiel_va_range(hstart, hend, pid,
                          ^~~~~~
      arch/powerpc/mm/book3s64/radix_tlb.c:1191:31: error: variable 'hend' is uninitialized
when used here
        __tlbiel_va_range(hstart, hend, pid,
                                  ^~~~

    Rework the 'if (IS_ENABLE(CONFIG_TRANSPARENT_HUGEPAGE))' so hstart/hend
    is always initialized to silence the warnings. That will also simplify
    the 'else' path. Clang is getting confused with these warnings, but the
    warnings is a false-positive.

    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Suggested-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Anders Roxell <anders.roxell@linaro.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20220810114318.3220630-1-anders.roxell@linaro.org
    Signed-off-by: Daniel Díaz <daniel.diaz@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c21cd34d902de861ea02364fbdc3964895bc09d7
Author: Vasant Hegde <vasant.hegde@amd.com>
Date:   Wed Feb 15 05:26:40 2023 +0000

    iommu: Attach device group to old domain in error path

    commit 2cc73c5712f97de98c38c2fafc1f288354a9f3c3 upstream.

    iommu_attach_group() attaches all devices in a group to domain and then
    sets group domain (group->domain). Current code (__iommu_attach_group())
    does not handle error path. This creates problem as devices to domain
    attachment is in inconsistent state.

    Flow:
      - During boot iommu attach devices to default domain
      - Later some device driver (like amd/iommu_v2 or vfio) tries to attach
        device to new domain.
      - In iommu_attach_group() path we detach device from current domain.
        Then it tries to attach devices to new domain.
      - If it fails to attach device to new domain then device to domain link
        is broken.
      - iommu_attach_group() returns error.
      - At this stage iommu_attach_group() caller thinks, attaching device to
        new domain failed and devices are still attached to old domain.
      - But in reality device to old domain link is broken. It will result
        in all sort of failures (like IO page fault) later.

    To recover from this situation, we need to attach all devices back to the
    old domain. Also log warning if it fails attach device back to old domain.

    Suggested-by: Lu Baolu <baolu.lu@linux.intel.com>
    Reported-by: Matt Fagnani <matt.fagnani@bell.net>
    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
```

```
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Tested-by: Matt Fagnani <matt.fagnani@bell.net>
    Link: https://lore.kernel.org/r/20230215052642.6016-1-vasant.hegde@amd.com
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216865
    Link: https://lore.kernel.org/lkml/15d0f9ff-2a56-b3e9-5b45-e6b23300ae3b@leemhuis.info/
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 446080b353f048b1fddaec1434cb3d27b5de7efe
Author: Vasant Hegde <vasant.hegde@amd.com>
Date:   Wed Feb 15 05:26:42 2023 +0000

    iommu/amd: Improve page fault error reporting

    commit 996d120b4de2b0d6b592bd9fbbe6e244b81ab3cc upstream.

    If IOMMU domain for device group is not setup properly then we may hit
    IOMMU page fault. Current page fault handler assumes that domain is
    always setup and it will hit NULL pointer derefence (see below sample log).

    Lets check whether domain is setup or not and log appropriate message.

    Sample log:
    ----------
     amdgpu 0000:00:01.0: amdgpu: SE 1, SH per SE 1, CU per SH 8, active_cu_number 6
     BUG: kernel NULL pointer dereference, address: 0000000000000058
     #PF: supervisor read access in kernel mode
     #PF: error_code(0x0000) - not-present page
     PGD 0 P4D 0
     Oops: 0000 [#1] PREEMPT SMP NOPTI
     CPU: 2 PID: 56 Comm: irq/24-AMD-Vi Not tainted 6.2.0-rc2+ #89
     Hardware name: xxx
     RIP: 0010:report_iommu_fault+0x11/0x90
     [...]
     Call Trace:
      <TASK>
      amd_iommu_int_thread+0x60c/0x760
      ? __pfx_irq_thread_fn+0x10/0x10
      irq_thread_fn+0x1f/0x60
      irq_thread+0xea/0x1a0
      ? preempt_count_add+0x6a/0xa0
      ? __pfx_irq_thread_dtor+0x10/0x10
      ? __pfx_irq_thread+0x10/0x10
      kthread+0xe9/0x110
      ? __pfx_kthread+0x10/0x10
      ret_from_fork+0x2c/0x50
      </TASK>

    Reported-by: Matt Fagnani <matt.fagnani@bell.net>
    Suggested-by: Joerg Roedel <joro@8bytes.org>
    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216865
    Link: https://lore.kernel.org/lkml/15d0f9ff-2a56-b3e9-5b45-e6b23300ae3b@leemhuis.info/
    Link: https://lore.kernel.org/r/20230215052642.6016-3-vasant.hegde@amd.com
    Cc: stable@vger.kernel.org
    [joro: Edit commit message]
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Cc: "Limonciello, Mario" <Mario.Limonciello@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 539258a204cb3b480c8dd276badd074c5dda4f25
Author: Vasant Hegde <vasant.hegde@amd.com>
Date:   Wed Feb 15 05:26:41 2023 +0000

    iommu/amd: Skip attach device domain is same as new domain
```

```
    commit f451c7a5a3b818ecfeba2ba258570769998baf3a upstream.

    If device->domain is same as new domain then we can skip the
    device attach process.

    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
    Link: https://lore.kernel.org/r/20230215052642.6016-2-vasant.hegde@amd.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Cc: "Limonciello, Mario" <Mario.Limonciello@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit df5d89e2a685d741c2e5f57ccd44b7d74d3e98a4
Author: Vasant Hegde <vasant.hegde@amd.com>
Date:   Wed Jan 11 12:15:03 2023 +0000

    iommu/amd: Fix error handling for pdev_pri_ats_enable()

    commit 080920e52148b4fbbf9360d5345fdcd7846e4841 upstream.

    Current code throws kernel warning if it fails to enable pasid/pri [1].
    Do not call pci_disable_[pasid/pri] if pci_enable_[pasid/pri] failed.

    [1] https://lore.kernel.org/linux-iommu/15d0f9ff-2a56-b3e9-5b45-
e6b23300ae3b@leemhuis.info/

    Reported-by: Matt Fagnani <matt.fagnani@bell.net>
    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
    Reviewed-by: Suravee Suthikulpanit <suravee.suthikulpanit@amd.com>
    Link: https://lore.kernel.org/r/20230111121503.5931-1-vasant.hegde@amd.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Cc: "Limonciello, Mario" <Mario.Limonciello@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b08bcfb4c97d7bd41b362cff44b2c537ce9e8540
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 19:00:03 2023 +0000

    HID: asus: use spinlock to safely schedule workers

    commit 4ab3a086d10eeec1424f2e8a968827a6336203df upstream.

    Use spinlocks to deal with workers introducing a wrapper
    asus_schedule_work(), and several spinlock checks.
    Otherwise, asus_kbd_backlight_set() may schedule led->work after the
    structure has been freed, causing a use-after-free.

    Fixes: af22a610bc38 ("HID: asus: support backlight on USB keyboards")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230125-hid-unregister-leds-v4-5-
7860c5763c38@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Stefan Ghinea <stefan.ghinea@windriver.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3cd664c22de74fc09881d0869034d9863acd7a43
Author: Pietro Borrello <borrello@diag.uniroma1.it>
Date:   Sun Feb 12 19:00:02 2023 +0000

    HID: asus: use spinlock to protect concurrent accesses

    commit 315c537068a13f0b5681d33dd045a912f4bece6f upstream.

    asus driver has a worker that may access data concurrently.
    Proct the accesses using a spinlock.
```

```
    Fixes: af22a610bc38 ("HID: asus: support backlight on USB keyboards")
    Signed-off-by: Pietro Borrello <borrello@diag.uniroma1.it>
    Link: https://lore.kernel.org/r/20230125-hid-unregister-leds-v4-4-
7860c5763c38@diag.uniroma1.it
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Stefan Ghinea <stefan.ghinea@windriver.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```