```
commit 7d0a458e1963128ee5a85bf0584bea5e75149946
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Mon Sep 5 10:31:36 2022 +0200

    Linux 5.19.7

    Link: https://lore.kernel.org/r/20220902121404.772492078@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Rudi Heitbaum <rudi@heitbaum.com>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Tested-by: Sudip Mukherjee <sudip.mukherjee@codethink.co.uk>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a754ee1c66bd0a23e613f0bf865053b29cb90e16
Author: Sudeep Holla <sudeep.holla@arm.com>
Date:   Mon Aug 8 09:46:40 2022 +0100

    arm64: cacheinfo: Fix incorrect assignment of signed error value to unsigned fw_level

    commit e75d18cecbb3805895d8ed64da4f78575ec96043 upstream.

    Though acpi_find_last_cache_level() always returned signed value and the
    document states it will return any errors caused by lack of a PPTT table,
    it never returned negative values before.

    Commit 0c80f9e165f8 ("ACPI: PPTT: Leave the table mapped for the runtime usage")
    however changed it by returning -ENOENT if no PPTT was found. The value
    returned from acpi_find_last_cache_level() is then assigned to unsigned
    fw_level.

    It will result in the number of cache leaves calculated incorrectly as
    a huge value which will then cause the following warning from __alloc_pages
    as the order would be great than MAX_ORDER because of incorrect and huge
    cache leaves value.

      |  WARNING: CPU: 0 PID: 1 at mm/page_alloc.c:5407 __alloc_pages+0x74/0x314
      |  Modules linked in:
      |  CPU: 0 PID: 1 Comm: swapper/0 Not tainted 5.19.0-10393-g7c2a8d3ac4c0 #73
      |  pstate: 20000005 (nzCv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
      |  pc : __alloc_pages+0x74/0x314
      |  lr : alloc_pages+0xe8/0x318
      |  Call trace:
      |   __alloc_pages+0x74/0x314
      |   alloc_pages+0xe8/0x318
      |   kmalloc_order_trace+0x68/0x1dc
      |   __kmalloc+0x240/0x338
      |   detect_cache_attributes+0xe0/0x56c
      |   update_siblings_masks+0x38/0x284
      |   store_cpu_topology+0x78/0x84
      |   smp_prepare_cpus+0x48/0x134
      |   kernel_init_freeable+0xc4/0x14c
      |   kernel_init+0x2c/0x1b4
      |   ret_from_fork+0x10/0x20

    Fix the same by changing fw_level to be signed integer and return the
    error from init_cache_level() early in case of error.

    Reported-and-Tested-by: Bruno Goncalves <bgoncalv@redhat.com>
```

```
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Link: https://lore.kernel.org/r/20220808084640.3165368-1-sudeep.holla@arm.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 123bf15c1a15f8dc7b5aceb5186bed37f2e80af5
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Mon Aug 22 10:53:46 2022 +0800

    net: neigh: don't call kfree_skb() under spin_lock_irqsave()

    commit d5485d9dd24e1d04e5509916515260186eb1455c upstream.

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with interrupts being disabled. So add all skb to
    a tmp list, then free them after spin_unlock_irqrestore() at
    once.

    Fixes: 66ba215cb513 ("neigh: fix possible DoS due to net iface start/stop loop")
    Suggested-by: Denis V. Lunev <den@openvz.org>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Reviewed-by: Nikolay Aleksandrov <razor@blackwall.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec274d8f1e649b12530bbba3f01478e20bf8bc23
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Wed Jul 27 17:33:12 2022 +0800

    net/af_packet: check len when min_header_len equals to 0

    commit dc633700f00f726e027846a318c5ffeb8deaaeda upstream.

    User can use AF_PACKET socket to send packets with the length of 0.
    When min_header_len equals to 0, packet_snd will call __dev_queue_xmit
    to send packets, and sock->type can be any type.

    Reported-by: syzbot+5ea725c25d06fb9114c4@syzkaller.appspotmail.com
    Fixes: fd1894224407 ("bpf: Don't redirect packets with invalid pkt_len")
    Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d25d6744c46b8289d61180bafb30f3140fef24a4
Author: Liam Howlett <liam.howlett@oracle.com>
Date:   Mon Jun 27 15:18:59 2022 +0000

    android: binder: fix lockdep check on clearing vma

    commit b0cab80ecd54ae3b2356bb081af0bffd538c8265 upstream.

    When munmapping a vma, the mmap_lock can be degraded to a write before
    calling close() on the file handle.  The binder close() function calls
    binder_alloc_set_vma() to clear the vma address, which now has a lock dep
    check for writing on the mmap_lock.  Change the lockdep check to ensure
    the reading lock is held while clearing and keep the write check while
    writing.

    Link: https://lkml.kernel.org/r/20220627151857.2316964-1-Liam.Howlett@oracle.com
    Fixes: 472a68df605b ("android: binder: stop saving a pointer to the VMA")
    Signed-off-by: Liam R. Howlett <Liam.Howlett@oracle.com>
    Reported-by: syzbot+da54fa8d793ca89c741f@syzkaller.appspotmail.com
    Acked-by: Todd Kjos <tkjos@google.com>
    Cc: "Arve Hjønnevåg" <arve@android.com>
    Cc: Christian Brauner (Microsoft) <brauner@kernel.org>
```

```
    Cc: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Cc: Hridya Valsaraju <hridya@google.com>
    Cc: Joel Fernandes <joel@joelfernandes.org>
    Cc: Martijn Coenen <maco@android.com>
    Cc: Suren Baghdasaryan <surenb@google.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8a239dd886e25d03bb2581c8ab30e8a1a12446b1
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Wed Aug 3 14:28:47 2022 -0400

    btrfs: tree-checker: check for overlapping extent items

    [ Upstream commit 899b7f69f244e539ea5df1b4d756046337de44a5 ]

    We're seeing a weird problem in production where we have overlapping
    extent items in the extent tree.  It's unclear where these are coming
    from, and in debugging we realized there's no check in the tree checker
    for this sort of problem.  Add a check to the tree-checker to make sure
    that the extents do not overlap each other.

    Reviewed-by: Qu Wenruo <wqu@suse.com>
    Signed-off-by: Josef Bacik <josef@toxicpanda.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 b83e1567af28055b9a0bb7b10fcdd58b02d00954
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Tue Jul 26 16:24:04 2022 -0400

    btrfs: fix lockdep splat with reloc root extent buffers

    [ Upstream commit b40130b23ca4a08c5785d5a3559805916bddba3c ]

    We have been hitting the following lockdep splat with btrfs/187 recently

      WARNING: possible circular locking dependency detected
      5.19.0-rc8+ #775 Not tainted
      ------------------------------------------------------
      btrfs/752500 is trying to acquire lock:
      ffff97e1875a97b8 (btrfs-treloc-02#2){+.+.}-{3:3}, at: __btrfs_tree_lock+0x24/0x110

      but task is already holding lock:
      ffff97e1875a9278 (btrfs-tree-01/1){+.+.}-{3:3}, at: __btrfs_tree_lock+0x24/0x110

      which lock already depends on the new lock.

      the existing dependency chain (in reverse order) is:

      -> #2 (btrfs-tree-01/1){+.+.}-{3:3}:
             down_write_nested+0x41/0x80
             __btrfs_tree_lock+0x24/0x110
             btrfs_init_new_buffer+0x7d/0x2c0
             btrfs_alloc_tree_block+0x120/0x3b0
             __btrfs_cow_block+0x136/0x600
             btrfs_cow_block+0x10b/0x230
             btrfs_search_slot+0x53b/0xb70
             btrfs_lookup_inode+0x2a/0xa0
             __btrfs_update_delayed_inode+0x5f/0x280
             btrfs_async_run_delayed_root+0x24c/0x290
             btrfs_work_helper+0xf2/0x3e0
             process_one_work+0x271/0x590
             worker_thread+0x52/0x3b0
```

```
               kthread+0xf0/0x120
               ret_from_fork+0x1f/0x30

       -> #1 (btrfs-tree-01){++++}-{3:3}:
               down_write_nested+0x41/0x80
               __btrfs_tree_lock+0x24/0x110
               btrfs_search_slot+0x3c3/0xb70
               do_relocation+0x10c/0x6b0
               relocate_tree_blocks+0x317/0x6d0
               relocate_block_group+0x1f1/0x560
               btrfs_relocate_block_group+0x23e/0x400
               btrfs_relocate_chunk+0x4c/0x140
               btrfs_balance+0x755/0xe40
               btrfs_ioctl+0x1ea2/0x2c90
               __x64_sys_ioctl+0x88/0xc0
               do_syscall_64+0x38/0x90
               entry_SYSCALL_64_after_hwframe+0x63/0xcd

       -> #0 (btrfs-treloc-02#2){+.+.}-{3:3}:
               __lock_acquire+0x1122/0x1e10
               lock_acquire+0xc2/0x2d0
               down_write_nested+0x41/0x80
               __btrfs_tree_lock+0x24/0x110
               btrfs_lock_root_node+0x31/0x50
               btrfs_search_slot+0x1cb/0xb70
               replace_path+0x541/0x9f0
               merge_reloc_root+0x1d6/0x610
               merge_reloc_roots+0xe2/0x260
               relocate_block_group+0x2c8/0x560
               btrfs_relocate_block_group+0x23e/0x400
               btrfs_relocate_chunk+0x4c/0x140
               btrfs_balance+0x755/0xe40
               btrfs_ioctl+0x1ea2/0x2c90
               __x64_sys_ioctl+0x88/0xc0
               do_syscall_64+0x38/0x90
               entry_SYSCALL_64_after_hwframe+0x63/0xcd

       other info that might help us debug this:

       Chain exists of:
         btrfs-treloc-02#2 --> btrfs-tree-01 --> btrfs-tree-01/1

        Possible unsafe locking scenario:

               CPU0                    CPU1
               ----                    ----
         lock(btrfs-tree-01/1);
                                       lock(btrfs-tree-01);
                                       lock(btrfs-tree-01/1);
         lock(btrfs-treloc-02#2);

        *** DEADLOCK ***

       7 locks held by btrfs/752500:
        #0: ffff97e292fdf460 (sb_writers#12){.+.+}-{0:0}, at: btrfs_ioctl+0x208/0x2c90
        #1: ffff97e284c02050 (&fs_info->reclaim_bgs_lock){+.+.}-{3:3}, at:
 btrfs_balance+0x55f/0xe40
        #2: ffff97e284c00878 (&fs_info->cleaner_mutex){+.+.}-{3:3}, at:
 btrfs_relocate_block_group+0x236/0x400
        #3: ffff97e292fdf650 (sb_internal#2){.+.+}-{0:0}, at: merge_reloc_root+0xef/0x610
        #4: ffff97e284c02378 (btrfs_trans_num_writers){++++}-{0:0}, at:
 join_transaction+0x1a8/0x5a0
        #5: ffff97e284c023a0 (btrfs_trans_num_extwriters){++++}-{0:0}, at:
 join_transaction+0x1a8/0x5a0
        #6: ffff97e1875a9278 (btrfs-tree-01/1){+.+.}-{3:3}, at: __btrfs_tree_lock+0x24/0x110
```

```
  stack backtrace:
  CPU: 1 PID: 752500 Comm: btrfs Not tainted 5.19.0-rc8+ #775
  Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.13.0-2.fc32 04/01/2014
  Call Trace:

   dump_stack_lvl+0x56/0x73
   check_noncircular+0xd6/0x100
   ? lock_is_held_type+0xe2/0x140
   __lock_acquire+0x1122/0x1e10
   lock_acquire+0xc2/0x2d0
   ? __btrfs_tree_lock+0x24/0x110
   down_write_nested+0x41/0x80
   ? __btrfs_tree_lock+0x24/0x110
   __btrfs_tree_lock+0x24/0x110
   btrfs_lock_root_node+0x31/0x50
   btrfs_search_slot+0x1cb/0xb70
   ? lock_release+0x137/0x2d0
   ? _raw_spin_unlock+0x29/0x50
   ? release_extent_buffer+0x128/0x180
   replace_path+0x541/0x9f0
   merge_reloc_root+0x1d6/0x610
   merge_reloc_roots+0xe2/0x260
   relocate_block_group+0x2c8/0x560
   btrfs_relocate_block_group+0x23e/0x400
   btrfs_relocate_chunk+0x4c/0x140
   btrfs_balance+0x755/0xe40
   btrfs_ioctl+0x1ea2/0x2c90
   ? lock_is_held_type+0xe2/0x140
   ? lock_is_held_type+0xe2/0x140
   ? __x64_sys_ioctl+0x88/0xc0
   __x64_sys_ioctl+0x88/0xc0
   do_syscall_64+0x38/0x90
   entry_SYSCALL_64_after_hwframe+0x63/0xcd
```

This isn't necessarily new, it's just tricky to hit in practice. There are two competing things going on here. With relocation we create a snapshot of every fs tree with a reloc tree. Any extent buffers that get initialized here are initialized with the reloc root lockdep key. However since it is a snapshot, any blocks that are currently in cache that originally belonged to the fs tree will have the normal tree lockdep key set. This creates the lock dependency of

```
  reloc tree -> normal tree
```

for the extent buffer locking during the first phase of the relocation as we walk down the reloc root to relocate blocks.

However this is problematic because the final phase of the relocation is merging the reloc root into the original fs root. This involves searching down to any keys that exist in the original fs root and then swapping the relocated block and the original fs root block. We have to search down to the fs root first, and then go search the reloc root for the block we need to replace. This creates the dependency of

```
  normal tree -> reloc tree
```

which is why lockdep complains.

Additionally even if we were to fix this particular mismatch with a different nesting for the merge case, we're still slotting in a block that has a owner of the reloc root objectid into a normal tree, so that block will have its lockdep key set to the tree reloc root, and create a lockdep splat later on when we wander into that block from the fs root.

```
    Unfortunately the only solution here is to make sure we do not set the
    lockdep key to the reloc tree lockdep key normally, and then reset any
    blocks we wander into from the reloc root when we're doing the merged.

    This solves the problem of having mixed tree reloc keys intermixed with
    normal tree keys, and then allows us to make sure in the merge case we
    maintain the lock order of

      normal tree -> reloc tree

    We handle this by setting a bit on the reloc root when we do the search
    for the block we want to relocate, and any block we search into or COW
    at that point gets set to the reloc tree key.  This works correctly
    because we only ever COW down to the parent node, so we aren't resetting
    the key for the block we're linking into the fs root.

    With this patch we no longer have the lockdep splat in btrfs/187.

    Signed-off-by: Josef Bacik <josef@toxicpanda.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 98f803acf8ac18a1cc159473137d563db4df0aa8
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Tue Jul 26 16:24:03 2022 -0400

    btrfs: move lockdep class helpers to locking.c

    [ Upstream commit 0a27a0474d146eb79e09ec88bf0d4229f4cfc1b8 ]

    These definitions exist in disk-io.c, which is not related to the
    locking.  Move this over to locking.h/c where it makes more sense.

    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Signed-off-by: Josef Bacik <josef@toxicpanda.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 82af833b99dc51c8795d171dfd376b451506544a
Author: Stefan Binding <sbinding@opensource.cirrus.com>
Date:   Tue Aug 16 16:19:01 2022 +0100

    ALSA: hda/cs8409: Support new Dolphin Variants

    [ Upstream commit 1ff954f9ab054675b9eb02dd14add8f7aa376d71 ]

    Add 4 new Dolphin Systems, same configuration as older systems.

    Signed-off-by: Stefan Binding <sbinding@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220816151901.1398007-1-sbinding@opensource.cirrus.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f31a350fd9a7c750bcde02f92eec3293b6bb667
Author: Lucas Tanure <tanureal@opensource.cirrus.com>
Date:   Wed Jul 27 10:59:24 2022 +0100

    platform/x86: serial-multi-instantiate: Add CLSA0101 Laptop

    [ Upstream commit 88392a0dd0ab263edb4ca416ebdecabd8289158a ]

    The device CLSA0101 has two instances of CS35L41
    connected by I2C.
```

```
    Signed-off-by: Lucas Tanure <tanureal@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220727095924.80884-5-tanureal@opensource.cirrus.com
    Link: https://lore.kernel.org/r/20220816194639.13870-1-cam@neo-zeon.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 503728838bdf9b8fd50ff1f89d47668e922880aa
Author: Florian Westphal <fw@strlen.de>
Date:   Tue Aug 16 14:15:21 2022 +0200

    testing: selftests: nft_flowtable.sh: use random netns names

    [ Upstream commit b71b7bfeac38c7a21c423ddafb29aa6258949df8 ]

    "ns1" is a too generic name, use a random suffix to avoid
    errors when such a netns exists.  Also allows to run multiple
    instances of the script in parallel.

    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d4bfe34b9cbb0395cb9508fa64324d4a1379e00
Author: Geert Uytterhoeven <geert@linux-m68k.org>
Date:   Mon Aug 15 12:39:20 2022 +0200

    netfilter: conntrack: NF_CONNTRACK_PROCFS should no longer default to y

    [ Upstream commit aa5762c34213aba7a72dc58e70601370805fa794 ]

    NF_CONNTRACK_PROCFS was marked obsolete in commit 54b07dca68557b09
    ("netfilter: provide config option to disable ancient procfs parts") in
    v3.3.

    Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 933020e6d0572f9997e8faaed8c082fda0825718
Author: Mukul Joshi <mukul.joshi@amd.com>
Date:   Fri Aug 12 15:23:51 2022 -0400

    drm/amdgpu: Fix interrupt handling on ih_soft ring

    [ Upstream commit de8341ee3ce7316883e836a2c4e9bf01ab651e0f ]

    There are no backing hardware registers for ih_soft ring.
    As a result, don't try to access hardware registers for read
    and write pointers when processing interrupts on the IH soft
    ring.

    Signed-off-by: Mukul Joshi <mukul.joshi@amd.com>
    Acked-by: Christian König <christian.koenig@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 09186eeffd3949a7e6ed39085bc151c1396a99a8
Author: Shane Xiao <shane.xiao@amd.com>
Date:   Mon Aug 15 16:32:15 2022 +0800

    drm/amdgpu: Add secure display TA load for Renoir

    [ Upstream commit e42dfa66d59240afbdd8d4b47b87486db39504aa ]
```

```
    Add secure display TA load for Renoir

    Signed-off-by: Shane Xiao <shane.xiao@amd.com>
    Reviewed-by: Aaron Liu <aaron.liu@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d278f332d105d3340179965dc14d6f7157fb9450
Author: Harish Kasiviswanathan <Harish.Kasiviswanathan@amd.com>
Date:   Fri Aug 5 18:04:08 2022 -0400

    drm/amdgpu: Add decode_iv_ts helper for ih_v6 block

    [ Upstream commit 1af9add1f1512b10d9ce44ec7137612bc81ff069 ]

    Was missing.  Add it.

    Signed-off-by: Harish Kasiviswanathan <Harish.Kasiviswanathan@amd.com>
    Reviewed-by: Hawking Zhang <Hawking.Zhang@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4eb2be1d0827627b5878b26dfa21d68ffe677fd0
Author: Charlene Liu <Charlene.Liu@amd.com>
Date:   Fri Aug 5 12:59:47 2022 -0400

    drm/amd/display: avoid doing vm_init multiple time

    [ Upstream commit 5544a7b5a07480192eb5fd3536462faed2c21528 ]

    [why]
    this is to ensure that driver will not reprogram hvm_prefetch_req again if
    it is done.

    Reviewed-by: Martin Leung <Martin.Leung@amd.com>
    Acked-by: Brian Chang <Brian.Chang@amd.com>
    Signed-off-by: Charlene Liu <Charlene.Liu@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 bd71ef2c30244c25243d84b593ea0ba88164ebb6
Author: Tom Chung <chiahsuan.chung@amd.com>
Date:   Wed Aug 3 15:15:17 2022 +0800

    drm/amd/display: Fix plug/unplug external monitor will hang while playback MPO video

    [ Upstream commit e98459c06e3d45c2229b097f7b8cdd412357fa2f ]

    [Why]
    Pipes for MPO primary and overlay will be power down and power up during
    plug/unplug external monitor while MPO video playback.
    But the pipes were the same after plug/unplug and should not need to be
    power down and power up or it will make page flip interrupt disabled and
    cause hang issue.

    [How]
    Add pipe split change condition that not only check the top pipe pointer
    but also check the index of top pipe if both top pipes are available.

    Reviewed-by: Sun peng Li <Sunpeng.Li@amd.com>
    Acked-by: Brian Chang <Brian.Chang@amd.com>
    Signed-off-by: Tom Chung <chiahsuan.chung@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 db1d9deb84d7870486688a5d13602962443a5500
Author: Dusica Milinkovic <Dusica.Milinkovic@amd.com>
Date:   Wed Aug 10 09:43:15 2022 +0200

    drm/amdgpu: Increase tlb flush timeout for sriov

    [ Upstream commit 373008bfc9cdb0f050258947fa5a095f0657e1bc ]

    [Why]
    During multi-vf executing benchmark (Luxmark) observed kiq error timeout.
    It happenes because all of VFs do the tlb invalidation at the same time.
    Although each VF has the invalidate register set, from hardware side
    the invalidate requests are queue to execute.

    [How]
    In case of 12 VF increase timeout on 12*100ms

    Signed-off-by: Dusica Milinkovic <Dusica.Milinkovic@amd.com>
    Acked-by: Shaoyun Liu <shaoyun.liu@amd.com>
    Acked-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 db0669c9731b849dee3dcaa6a0b4a891a0055338
Author: Ilya Bakoulin <Ilya.Bakoulin@amd.com>
Date:   Tue Jul 26 16:19:38 2022 -0400

    drm/amd/display: Fix pixel clock programming

    [ Upstream commit 04fb918bf421b299feaee1006e82921d7d381f18 ]

    [Why]
    Some pixel clock values could cause HDMI TMDS SSCPs to be misaligned
    between different HDMI lanes when using YCbCr420 10-bit pixel format.

    BIOS functions for transmitter/encoder control take pixel clock in kHz
    increments, whereas the function for setting the pixel clock is in 100Hz
    increments. Setting pixel clock to a value that is not on a kHz boundary
    will cause the issue.

    [How]
    Round pixel clock down to nearest kHz in 10/12-bpc cases.

    Reviewed-by: Aric Cyr <Aric.Cyr@amd.com>
    Acked-by: Brian Chang <Brian.Chang@amd.com>
    Signed-off-by: Ilya Bakoulin <Ilya.Bakoulin@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 22a75c616f1971c23838506b14971a4ef4a66bd7
Author: Evan Quan <evan.quan@amd.com>
Date:   Wed Aug 10 11:11:16 2022 +0800

    drm/amd/pm: add missing ->fini_xxxx interfaces for some SMU13 asics

    [ Upstream commit 4bac1c846eff8042dd59ddecd0a43f3b9de5fd23 ]

    Without these, potential memory leak may be induced.

    Signed-off-by: Evan Quan <evan.quan@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 4d21584ac6392aa66171b7efd647ecd1a447556b
Author: Evan Quan <evan.quan@amd.com>
Date:   Wed Aug 10 11:08:31 2022 +0800

    drm/amd/pm: add missing ->fini_microcode interface for Sienna Cichlid

    [ Upstream commit 0a2d922a5618377cdf8fa476351362733ef55342 ]

    To avoid any potential memory leak.

    Signed-off-by: Evan Quan <evan.quan@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 d9d70547cba1d88fe52bfb37655d09a12926a0f3
Author: Evan Quan <evan.quan@amd.com>
Date:   Wed Aug 3 16:13:56 2022 +0800

    drm/amdgpu: disable 3DCGCG/CGLS temporarily due to stability issue

    [ Upstream commit 1b586595df6d04c27088ef348b8202204ce26d45 ]

    Some stability issues were reported with these features.

    Signed-off-by: Evan Quan <evan.quan@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 dc441eda061720bab0117c45e30a2eb065b858ec
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Sun Aug 14 22:40:25 2022 +0900

    ksmbd: don't remove dos attribute xattr on O_TRUNC open

    [ Upstream commit 17661ecf6a64eb11ae7f1108fe88686388b2acd5 ]

    When smb client open file in ksmbd share with O_TRUNC, dos attribute
    xattr is removed as well as data in file. This cause the FSCTL_SET_SPARSE
    request from the client fails because ksmbd can't update the dos attribute
    after setting ATTR_SPARSE_FILE. And this patch fix xfstests generic/469
    test also.

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

commit 895096784ab9eb17126e3099a6ef2649e96f8d84
Author: Juergen Gross <jgross@suse.com>
Date:   Mon Jun 20 11:45:34 2022 +0200

    s390/hypfs: avoid error message under KVM

    [ Upstream commit 7b6670b03641ac308aaa6fa2e6f964ac993b5ea3 ]

    When booting under KVM the following error messages are issued:

    hypfs.7f5705: The hardware system does not support hypfs
    hypfs.7a79f0: Initialization of hypfs failed with rc=-61

    Demote the severity of first message from "error" to "info" and issue
```

```
    the second message only in other error cases.

    Signed-off-by: Juergen Gross <jgross@suse.com>
    Acked-by: Heiko Carstens <hca@linux.ibm.com>
    Acked-by: Christian Borntraeger <borntraeger@linux.ibm.com>
    Link: https://lore.kernel.org/r/20220620094534.18967-1-jgross@suse.com
    [arch/s390/hypfs/hypfs_diag.c changed description]
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2432f152c175c149b6365b07d110701ca1df988e
Author: Stefan Binding <sbinding@opensource.cirrus.com>
Date:   Mon Aug 15 15:19:53 2022 +0100

    ALSA: hda/realtek: Add quirks for ASUS Zenbooks using CS35L41

    [ Upstream commit 461122b999bda2ebef2086a35d8990f9ccac5ab8 ]

    These Asus Zenbook laptop use Realtek HDA codec combined with
    2xCS35L41 Amplifiers using SPI.

    Signed-off-by: Stefan Binding <sbinding@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20220815141953.25197-1-sbinding@opensource.cirrus.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2dd5ed474115150d8175825bc3b56c6385c3a83b
Author: Denis V. Lunev <den@openvz.org>
Date:   Thu Aug 11 18:20:11 2022 +0300

    neigh: fix possible DoS due to net iface start/stop loop

    [ Upstream commit 66ba215cb51323e4e55e38fd5f250e0fae0cbc94 ]

    Normal processing of ARP request (usually this is Ethernet broadcast
    packet) coming to the host is looking like the following:
    * the packet comes to arp_process() call and is passed through routing
      procedure
    * the request is put into the queue using pneigh_enqueue() if
      corresponding ARP record is not local (common case for container
      records on the host)
    * the request is processed by timer (within 80 jiffies by default) and
      ARP reply is sent from the same arp_process() using
      NEIGH_CB(skb)->flags & LOCALLY_ENQUEUED condition (flag is set inside
      pneigh_enqueue())

    And here the problem comes. Linux kernel calls pneigh_queue_purge()
    which destroys the whole queue of ARP requests on ANY network interface
    start/stop event through __neigh_ifdown().

    This is actually not a problem within the original world as network
    interface start/stop was accessible to the host 'root' only, which
    could do more destructive things. But the world is changed and there
    are Linux containers available. Here container 'root' has an access
    to this API and could be considered as untrusted user in the hosting
    (container's) world.

    Thus there is an attack vector to other containers on node when
    container's root will endlessly start/stop interfaces. We have observed
    similar situation on a real production node when docker container was
    doing such activity and thus other containers on the node become not
    accessible.

    The patch proposed doing very simple thing. It drops only packets from
    the same namespace in the pneigh_queue_purge() where network interface
```

```
    state change is detected. This is enough to prevent the problem for the
    whole node preserving original semantics of the code.

    v2:
            - do del_timer_sync() if queue is empty after pneigh_queue_purge()
    v3:
            - rebase to net tree

    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: Daniel Borkmann <daniel@iogearbox.net>
    Cc: David Ahern <dsahern@kernel.org>
    Cc: Yajun Deng <yajun.deng@linux.dev>
    Cc: Roopa Prabhu <roopa@nvidia.com>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: netdev@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Cc: Alexey Kuznetsov <kuznet@ms2.inr.ac.ru>
    Cc: Alexander Mikhalitsyn <alexander.mikhalitsyn@virtuozzo.com>
    Cc: Konstantin Khorenko <khorenko@virtuozzo.com>
    Cc: kernel@openvz.org
    Cc: devel@openvz.org
    Investigated-by: Alexander Mikhalitsyn <alexander.mikhalitsyn@virtuozzo.com>
    Signed-off-by: Denis V. Lunev <den@openvz.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d462a681d7d98ecb1b48605f6625802acd85d31
Author: Li Qiong <liqiong@nfschina.com>
Date:   Fri Aug 12 11:09:54 2022 +0800

    net: lan966x: fix checking for return value of platform_get_irq_byname()

    [ Upstream commit 40b4ac880e21d917da7f3752332fa57564a4c202 ]

    The platform_get_irq_byname() returns non-zero IRQ number
    or negative error number. "if (irq)" always true, chang it
    to "if (irq > 0)"

    Signed-off-by: Li Qiong <liqiong@nfschina.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 40f424dc1ee328832ec86df29b8939e2e1618153
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Mon Aug 8 21:56:48 2022 +0900

    ksmbd: return STATUS_BAD_NETWORK_NAME error status if share is not configured

    [ Upstream commit fe54833dc8d97ef387e86f7c80537d51c503ca75 ]

    If share is not configured in smb.conf, smb2 tree connect should return
    STATUS_BAD_NETWORK_NAME instead of STATUS_BAD_NETWORK_PATH.

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

commit 4b25bdb54578f3b96ff055e5d27bc1cb82950e51
Author: Zhen Ni <nizhen@uniontech.com>
Date:   Wed Aug 3 17:19:58 2022 +0800
```

```
    drm/amd/pm: Fix a potential gpu_metrics_table memory leak

    [ Upstream commit 5afb76522a0af0513b6dc01f84128a73206b051b ]

    Memory is allocated for gpu_metrics_table in
    smu_v13_0_4_init_smc_tables(), but not freed in
    smu_v13_0_4_fini_smc_tables(). This may cause memory leaks, fix it.

    Reviewed-by: Evan Quan <evan.quan@amd.com>
    Signed-off-by: Zhen Ni <nizhen@uniontech.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 21e3891abc83f2d5482c67cca2fb7f228635daaf
Author: Felix Kuehling <Felix.Kuehling@amd.com>
Date:   Thu Aug 4 18:19:38 2022 -0400

    drm/amdkfd: Handle restart of kfd_ioctl_wait_events

    [ Upstream commit bea9a56afbc4b5a41ea579b8b0dc5e189b439504 ]

    When kfd_ioctl_wait_events needs to restart due to a signal, we need to
    update the timeout to account for the time already elapsed. We also need
    to undo auto_reset of events that have signaled already, so that the
    restarted ioctl will be able to count those signals again.

    This fixes infinite hangs when kfd_ioctl_wait_events is interrupted by a
    signal.

    Signed-off-by: Felix Kuehling <Felix.Kuehling@amd.com>
    Reviewed-and-tested-by: Xiaogang Chen <Xiaogang.Chen@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 952d4c72d9032eec5c38249ca353d31c43be06e0
Author: Kenneth Feng <kenneth.feng@amd.com>
Date:   Tue Aug 9 10:13:54 2022 +0800

    drm/amd/pm: skip pptable override for smu_v13_0_7

    [ Upstream commit 4e64b529c5b04e7944b41de554ee686ecab00744 ]

    skip pptable override for smu_v13_0_7 secure boards only.

    Signed-off-by: Kenneth Feng <kenneth.feng@amd.com>
    Reviewed-by: Feifei Xu <Feifei.Xu@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 909dae582c8f3a9a231c7bc41594787b1620a5f1
Author: Meenakshikumar Somasundaram <meenakshikumar.somasundaram@amd.com>
Date:   Mon Jul 11 18:37:41 2022 -0400

    drm/amd/display: Fix TDR eDP and USB4 display light up issue

    [ Upstream commit 30456ffa65469d1d2e5e1da05017e6728d24c11c ]

    [Why]
    After TDR recovery, eDP and USB4 display does not light up. Because
    dmub outbox notifications are not enabled after dmub reload and link
    encoder assignments for the streams are not cleared before dc state
    reset.

    [How]
    - Dmub outbox notification is enabled after tdr recovery by issuing
```

```
       inbox command to dmub.
     - Link encoders for the streams are unassigned before dc state reset.

     Reviewed-by: Jimmy Kizito <Jimmy.Kizito@amd.com>
     Reviewed-by: Jun Lei <Jun.Lei@amd.com>
     Acked-by: Tom Chung <chiahsuan.chung@amd.com>
     Signed-off-by: Meenakshikumar Somasundaram <meenakshikumar.somasundaram@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 3101839b080137c367f3f88c2a040f791de880aa
Author: Fudong Wang <Fudong.Wang@amd.com>
Date:   Wed Jul 27 12:01:29 2022 +0800

    drm/amd/display: clear optc underflow before turn off odm clock

    [ Upstream commit b2a93490201300a749ad261b5c5d05cb50179c44 ]

    [Why]
    After ODM clock off, optc underflow bit will be kept there always and clear not work.
    We need to clear that before clock off.

    [How]
    Clear that if have when clock off.

    Reviewed-by: Alvin Lee <alvin.lee2@amd.com>
    Acked-by: Tom Chung <chiahsuan.chung@amd.com>
    Signed-off-by: Fudong Wang <Fudong.Wang@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 c6f7ba837efc8f12c51aecaf9570d86adb80ab90
Author: Alvin Lee <alvin.lee2@amd.com>
Date:   Thu Jul 28 09:51:05 2022 -0400

    drm/amd/display: For stereo keep "FLIP_ANY_FRAME"

    [ Upstream commit 84ef99c728079dfd21d6bc70b4c3e4af20602b3c ]

    [Description]
    Observed in stereomode that programming FLIP_LEFT_EYE
    can cause hangs. Keep FLIP_ANY_FRAME in stereo mode so
    the surface flip can take place before left or right eye

    Reviewed-by: Martin Leung <Martin.Leung@amd.com>
    Acked-by: Tom Chung <chiahsuan.chung@amd.com>
    Signed-off-by: Alvin Lee <alvin.lee2@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 71f2cb739a3cbcda4612889ce660fd8c380e3e1d
Author: Leo Ma <hanghong.ma@amd.com>
Date:   Fri Jul 22 13:42:58 2022 -0400

    drm/amd/display: Fix HDMI VSIF V3 incorrect issue

    [ Upstream commit 0591183699fceeafb4c4141072d47775de83ecfb ]

    [Why]
    Reported from customer the checksum in AMD VSIF V3 is incorrect and
    causing blank screen issue.
```

```
    [How]
    Fix the packet length issue on AMD HDMI VSIF V3.

    Reviewed-by: Anthony Koo <Anthony.Koo@amd.com>
    Acked-by: Tom Chung <chiahsuan.chung@amd.com>
    Signed-off-by: Leo Ma <hanghong.ma@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 b26218b2349ac2b2141fd07430ded3786bee6910
Author: Josip Pavic <Josip.Pavic@amd.com>
Date:   Thu Jul 21 15:33:00 2022 -0400

    drm/amd/display: Avoid MPC infinite loop

    [ Upstream commit 8de297dc046c180651c0500f8611663ae1c3828a ]

    [why]
    In some cases MPC tree bottom pipe ends up point to itself.  This causes
    iterating from top to bottom to hang the system in an infinite loop.

    [how]
    When looping to next MPC bottom pipe, check that the pointer is not same
    as current to avoid infinite loop.

    Reviewed-by: Josip Pavic <Josip.Pavic@amd.com>
    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Alex Hung <alex.hung@amd.com>
    Signed-off-by: Aric Cyr <aric.cyr@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 ab01ac7fa985755e29663283ec13179e99bf8fcf
Author: Chiawen Huang <chiawen.huang@amd.com>
Date:   Thu Jul 21 21:57:05 2022 +0800

    drm/amd/display: Device flash garbage before get in OS

    [ Upstream commit 9c580e8f6cd6524d4e2c3490c440110526f7ddd6 ]

    [Why]
    Enabling stream with tg lock makes config settings
    pending causing the garbage until tg unlock.

    [How]
    Keep the original lock mechanism
    The driver doesn't lock tg if plane_state is null.

    Reviewed-by: Anthony Koo <Anthony.Koo@amd.com>
    Acked-by: Tom Chung <chiahsuan.chung@amd.com>
    Signed-off-by: Chiawen Huang <chiawen.huang@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 ae29c0c466812726347813ab44bc144f0b4644b6
Author: Aurabindo Pillai <aurabindo.pillai@amd.com>
Date:   Fri Jul 29 13:36:31 2022 -0400

    drm/amd/display: Add a missing register field for HPO DP stream encoder

    [ Upstream commit 37bc31f0e7da4fbad4664e64d906ae7b9009e550 ]
```

```
    [Why&How]
    Add the missing definition to set the register field
    HBLANK_MINIMUM_SYMBOL_WIDTH

    Signed-off-by: Aurabindo Pillai <aurabindo.pillai@amd.com>
    Acked-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 de3f7eb0b8642ded71fae5bbf4db584933da6c7d
Author: Alexandre Vicenzi <alexandre.vicenzi@suse.com>
Date:   Mon Aug 8 20:03:43 2022 +0200

    rtla: Fix tracer name

    [ Upstream commit f1432cd24c240cedf78c0d026631e3b10052c8e1 ]

    The correct tracer name is timerlat and not timelat.

    Link: https://lore.kernel.org/linux-trace-devel/20220808180343.22262-1-
alexandre.vicenzi@suse.com

    Signed-off-by: Alexandre Vicenzi <alexandre.vicenzi@suse.com>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33f15a70fe9d26554d6b8f306f088a5e3ae31e49
Author: Oder Chiou <oder_chiou@realtek.com>
Date:   Mon Aug 8 13:28:36 2022 +0800

    ASoC: rt5640: Fix the JD voltage dropping issue

    [ Upstream commit afb176d45870048eea540991b082208270824037 ]

    The patch fixes the JD voltage dropping issue in the HDA JD using.

    Signed-off-by: Oder Chiou <oder_chiou@realtek.com>
    Reported-by: Mohan Kumar D <mkumard@nvidia.com>
    Link: https://lore.kernel.org/r/20220808052836.25791-1-oder_chiou@realtek.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc4351487a55899ad73a286391c1b58604754956
Author: Biju Das <biju.das.jz@bp.renesas.com>
Date:   Thu Jul 28 10:26:12 2022 +0100

    ASoC: sh: rz-ssi: Improve error handling in rz_ssi_probe() error path

    [ Upstream commit c75ed9f54ce8d349fee557f2b471a4d637ed2a6b ]

    We usually do cleanup in reverse order of init. Currently in case of
    error rz_ssi_release_dma_channels() done in the reverse order. This
    patch improves error handling in rz_ssi_probe() error path.

    While at it, use "goto cleanup" style to reduce code duplication.

    Reported-by: Pavel Machek <pavel@denx.de>
    Signed-off-by: Biju Das <biju.das.jz@bp.renesas.com>
    Link: https://lore.kernel.org/r/20220728092612.38858-1-biju.das.jz@bp.renesas.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d875be57a21586f2ebdacb04afa134932b4d4d91
Author: Konstantin Komarov <almaz.alexandrovich@paragon-software.com>
Date:   Fri May 13 19:54:23 2022 +0300
```

```
    fs/ntfs3: Fix work with fragmented xattr

    [ Upstream commit 42f86b1226a42bfc79a7125af435432ad4680a32 ]

    In some cases xattr is too fragmented,
    so we need to load it before writing.

    Signed-off-by: Konstantin Komarov <almaz.alexandrovich@paragon-software.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9215e9018bd17ac5b4dfb1a032268e27f45a819
Author: Liming Sun <limings@nvidia.com>
Date:   Tue Aug 9 13:37:42 2022 -0400

    mmc: sdhci-of-dwcmshc: Re-enable support for the BlueField-3 SoC

    [ Upstream commit a0753ef66c34c1739580219dca664eda648164b7 ]

    The commit 08f3dff799d4 (mmc: sdhci-of-dwcmshc: add rockchip platform
    support") introduces the use of_device_get_match_data() to check for some
    chips. Unfortunately, it also breaks the BlueField-3 FW, which uses ACPI.

    To fix the problem, let's add the ACPI match data and the corresponding
    quirks to re-enable the support for the BlueField-3 SoC.

    Reviewed-by: David Woods <davwoods@nvidia.com>
    Signed-off-by: Liming Sun <limings@nvidia.com>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Fixes: 08f3dff799d4 ("mmc: sdhci-of-dwcmshc: add rockchip platform support")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20220809173742.178440-1-limings@nvidia.com
    [Ulf: Clarified the commit message a bit]
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 30841ba786ee2e55e05fa98a7c255b83313694f4
Author: Sebastian Reichel <sebastian.reichel@collabora.com>
Date:   Wed May 4 23:32:40 2022 +0200

    mmc: sdhci-of-dwcmshc: rename rk3568 to rk35xx

    [ Upstream commit 86e1a8e1f9b555af342c53ae06284eeeab9a4263 ]

    Prepare driver for rk3588 support by renaming the internal data
    structures.

    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Link: https://lore.kernel.org/r/20220504213251.264819-11-sebastian.reichel@collabora.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1541cdf5aecb0b85c26624478b154cce920281b
Author: Yifeng Zhao <yifeng.zhao@rock-chips.com>
Date:   Wed May 4 23:32:39 2022 +0200

    mmc: sdhci-of-dwcmshc: add reset call back for rockchip Socs

    [ Upstream commit 70f832206fe72e9998b46363e8e59e89b0b757bc ]

    The reset function build in the SDHCI will not reset the logic
    circuit related to the tuning function, which may cause data
    reading errors. Resetting the complete SDHCI controller through
    the reset controller fixes the issue.
```

```
    Signed-off-by: Yifeng Zhao <yifeng.zhao@rock-chips.com>
    [rebase, use optional variant of reset getter]
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Link: https://lore.kernel.org/r/20220504213251.264819-10-sebastian.reichel@collabora.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1c137c46d16e33b71ef92daf1882d5443fa2ec93
Author: Wenbin Mei <wenbin.mei@mediatek.com>
Date:   Thu Jul 28 16:00:48 2022 +0800

    mmc: mtk-sd: Clear interrupts when cqe off/disable

    [ Upstream commit cc5d1692600613e72f32af60e27330fe0c79f4fe ]

    Currently we don't clear MSDC interrupts when cqe off/disable, which led
    to the data complete interrupt will be reserved for the next command.
    If the next command with data transfer after cqe off/disable, we process
    the CMD ready interrupt and trigger DMA start for data, but the data
    complete interrupt is already exists, then SW assume that the data transfer
    is complete, SW will trigger DMA stop, but the data may not be transmitted
    yet or is transmitting, so we may encounter the following error:
    mtk-msdc 11230000.mmc: CMD bus busy detected.

    Signed-off-by: Wenbin Mei <wenbin.mei@mediatek.com>
    Fixes: 88bd652b3c74 ("mmc: mediatek: command queue support")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20220728080048.21336-1-wenbin.mei@mediatek.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1d7df13d223b4bdf52ba51eec15d7161235bdde4
Author: Even Xu <even.xu@intel.com>
Date:   Tue Aug 23 09:10:59 2022 +0800

    HID: intel-ish-hid: ipc: Add Meteor Lake PCI device ID

    commit 467249a7dff68451868ca79696aef69764193a8a upstream.

    Add device ID of Meteor Lake P into ishtp support list.

    Signed-off-by: Even Xu <even.xu@intel.com>
    Acked-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a32046896fe1ad944fd61aec3127485066c4eff4
Author: Michael Hübner <michaelh.95@t-online.de>
Date:   Fri Aug 5 10:05:23 2022 +0200

    HID: thrustmaster: Add sparco wheel and fix array length

    commit d9a17651f3749e69890db57ca66e677dfee70829 upstream.

    Add device id for the Sparco R383 Mod wheel.

    Fix wheel info array length to match actual wheel count present in the array.

    Signed-off-by: Michael Hübner <michaelh.95@t-online.de>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7c6e6c334154be16740b44dcd7638fb510b9bd91
```

```
Author: Daniel J. Ogorchock <djogorchock@gmail.com>
Date:   Wed Jul 13 16:20:59 2022 -0400

    HID: nintendo: fix rumble worker null pointer deref

    commit 1ff89e06c2e5fab30274e4b02360d4241d6e605e upstream.

    We can dereference a null pointer trying to queue work to a destroyed
    workqueue.

    If the device is disconnected, nintendo_hid_remove is called, in which
    the rumble_queue is destroyed. Avoid using that queue to defer rumble
    work once the controller state is set to JOYCON_CTLR_STATE_REMOVED.

    This eliminates the null pointer dereference.

    Signed-off-by: Daniel J. Ogorchock <djogorchock@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2b32e820ccf5a0385e46d1c038b321aba9a5ec2d
Author: Josh Kilmer <srjek2@gmail.com>
Date:   Thu Jul 28 12:51:11 2022 -0500

    HID: asus: ROG NKey: Ignore portion of 0x5a report

    commit 1c0cc9d11c665020cbeb80e660fb8929164407f4 upstream.

    On an Asus G513QY, of the 5 bytes in a 0x5a report, only the first byte
    is a meaningful keycode. The other bytes are zeroed out or hold garbage
    from the last packet sent to the keyboard.

    This patch fixes up the report descriptor for this event so that the
    general hid code will only process 1 byte for keycodes, avoiding
    spurious key events and unmapped Asus vendor usagepage code warnings.

    Signed-off-by: Josh Kilmer <srjek2@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5b3063adff7da583e61b284beb3308f7588a4c8a
Author: Aditya Garg <gargaditya08@live.com>
Date:   Sun Aug 21 08:04:45 2022 +0000

    HID: Add Apple Touchbar on T2 Macs in hid_have_special_driver list

    commit 750ec977288d96e9a11424e3507ede097af732c4 upstream.

    The touchbar on Apple T2 Macs has 2 modes, one that shows the function
    keys and other that shows the media controls. The user can use the fn
    key on his keyboard to switch between the 2 modes.

    On Linux, if people were using an external keyboard or mouse, the
    touchbar failed to change modes on pressing the fn key with the following
    in dmesg :-

    [   10.661445] apple-ib-als 0003:05AC:8262.0001: : USB HID v1.01 Device [Apple Inc.
Ambient Light Sensor] on usb-bce-vhci-3/input0
    [   11.830992] apple-ib-touchbar 0003:05AC:8302.0007: input: USB HID v1.01 Keyboard
[Apple Inc. Touch Bar Display] on usb-bce-vhci-6/input0
    [   12.139407] apple-ib-touchbar 0003:05AC:8102.0008: : USB HID v1.01 Device [Apple Inc.
Touch Bar Backlight] on usb-bce-vhci-7/input0
    [   12.211824] apple-ib-touchbar 0003:05AC:8102.0009: : USB HID v1.01 Device [Apple Inc.
Touch Bar Backlight] on usb-bce-vhci-7/input1
    [   14.219759] apple-ib-touchbar 0003:05AC:8302.0007: tb: Failed to set touch bar mode
```

```
to 2 (-110)
    [   24.395670] apple-ib-touchbar 0003:05AC:8302.0007: tb: Failed to set touch bar mode
to 2 (-110)
    [   34.635791] apple-ib-touchbar 0003:05AC:8302.0007: tb: Failed to set touch bar mode
to 2 (-110)
    [  269.579233] apple-ib-touchbar 0003:05AC:8302.0007: tb: Failed to set touch bar mode
to 1 (-110)

    Add the USB IDs of the touchbar found in T2 Macs to HID have special
    driver list to fix the issue.

    Signed-off-by: Aditya Garg <gargaditya08@live.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2803225ac89bc761ade8673105db8a8d84335054
Author: Akihiko Odaki <akihiko.odaki@gmail.com>
Date:   Tue Aug 16 19:21:20 2022 +0900

    HID: AMD_SFH: Add a DMI quirk entry for Chromebooks

    commit adada3f4930ac084740ea340bd8e94028eba4f22 upstream.

    Google Chromebooks use Chrome OS Embedded Controller Sensor Hub instead
    of Sensor Hub Fusion and leaves MP2 uninitialized, which disables all
    functionalities, even including the registers necessary for feature
    detections.

    The behavior was observed with Lenovo ThinkPad C13 Yoga.

    Signed-off-by: Akihiko Odaki <akihiko.odaki@gmail.com>
    Suggested-by: Mario Limonciello <mario.limonciello@amd.com>
    Acked-by: Basavaraj Natikar <Basavaraj.Natikar@amd.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4af0f12367b2cdf687e58044a71916edb60b0f4d
Author: Steev Klimaszewski <steev@kali.org>
Date:   Thu Aug 18 21:39:24 2022 -0500

    HID: add Lenovo Yoga C630 battery quirk

    commit 3a47fa7b14c7d9613909a844aba27f99d3c58634 upstream.

    Similar to the Surface Go devices, the Elantech touchscreen/digitizer in
    the Lenovo Yoga C630 mistakenly reports the battery of the stylus, and
    always reports an empty battery.

    Apply the HID_BATTERY_QUIRK_IGNORE quirk to ignore this battery and
    prevent the erroneous low battery warnings.

    Signed-off-by: Steev Klimaszewski <steev@kali.org>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b61400c4e784065c9783442e8ef096fefa811901
Author: Benjamin Tissoires <benjamin.tissoires@redhat.com>
Date:   Mon Aug 22 08:22:47 2022 +0200

    HID: input: fix uclogic tablets

    commit 8db8be9cfc89935c97d791c7e6264e710a7e8a56 upstream.

    commit 87562fcd1342 ("HID: input: remove the need for HID_QUIRK_INVERT")
    made the assumption that it was the only one handling tablets and thus
```

```
    kept an internal state regarding the tool.

    Turns out that the uclogic driver has a timer to release the in range
    bit, effectively making hid-input ignoring all in range information
    after the very first one.

    Fix that by having a more rationale approach which consists in forwarding
    every event and let the input stack filter out the duplicates.

    Reported-by: Stefan Hansson <newbie13xd@gmail.com>
    Fixes: 87562fcd1342 ("HID: input: remove the need for HID_QUIRK_INVERT")
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec45c77833192fa8da4681b1a4c6847fe7e4bf4c
Author: Takashi Iwai <tiwai@suse.de>
Date:   Sun Aug 28 09:41:43 2022 +0200

    ALSA: usb-audio: Add quirk for LH Labs Geek Out HD Audio 1V5

    commit 5f3d9e8161bb8cb23ab3b4678cd13f6e90a06186 upstream.

    The USB DAC from LH Labs (2522:0007) seems requiring the same quirk as
    Sony Walkman to set up the interface like UAC1; otherwise it gets the
    constant errors "usb_set_interface failed (-71)".  This patch adds a
    quirk entry for addressing the buggy behavior.

    Reported-by: Lennert Van Alboom <lennert@vanalboom.org>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/T3VPXtCc4uFws9Gfh2RjX6OdwM1RqfC6VqQr--
_LMDyB2x5N3p9_q6AtPna17IXhHwBtcJVdXuS80ZZSCMjh_BafIbnzJPhbrkmhmWS6DlI=@vanalboom.org
    Link: https://lore.kernel.org/r/20220828074143.14736-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7877eaa1131147b4d6a063962f3aac0ab1b8ea1c
Author: Jann Horn <jannh@google.com>
Date:   Wed Aug 31 19:06:00 2022 +0200

    mm/rmap: Fix anon_vma->degree ambiguity leading to double-reuse

    commit 2555283eb40df89945557273121e9393ef9b542b upstream.

    anon_vma->degree tracks the combined number of child anon_vmas and VMAs
    that use the anon_vma as their ->anon_vma.

    anon_vma_clone() then assumes that for any anon_vma attached to
    src->anon_vma_chain other than src->anon_vma, it is impossible for it to
    be a leaf node of the VMA tree, meaning that for such VMAs ->degree is
    elevated by 1 because of a child anon_vma, meaning that if ->degree
    equals 1 there are no VMAs that use the anon_vma as their ->anon_vma.

    This assumption is wrong because the ->degree optimization leads to leaf
    nodes being abandoned on anon_vma_clone() - an existing anon_vma is
    reused and no new parent-child relationship is created.  So it is
    possible to reuse an anon_vma for one VMA while it is still tied to
    another VMA.

    This is an issue because is_mergeable_anon_vma() and its callers assume
    that if two VMAs have the same ->anon_vma, the list of anon_vmas
    attached to the VMAs is guaranteed to be the same.  When this assumption
    is violated, vma_merge() can merge pages into a VMA that is not attached
    to the corresponding anon_vma, leading to dangling page->mapping
    pointers that will be dereferenced during rmap walks.
```

```
    Fix it by separately tracking the number of child anon_vmas and the
    number of VMAs using the anon_vma as their ->anon_vma.

    Fixes: 7a3ef208e662 ("mm: prevent endless growth of anon_vma hierarchy")
    Cc: stable@kernel.org
    Acked-by: Michal Hocko <mhocko@suse.com>
    Acked-by: Vlastimil Babka <vbabka@suse.cz>
    Signed-off-by: Jann Horn <jannh@google.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 72f2dc8993f10262092745a88cb2dd0fef094f23
Author: Zhengchao Shao <shaozhengchao@huawei.com>
Date:   Fri Jul 15 19:55:59 2022 +0800

    bpf: Don't redirect packets with invalid pkt_len

    commit fd1894224407c484f652ad456e1ce423e89bb3eb upstream.

    Syzbot found an issue [1]: fq_codel_drop() try to drop a flow whitout any
    skbs, that is, the flow->head is null.
    The root cause, as the [2] says, is because that bpf_prog_test_run_skb()
    run a bpf prog which redirects empty skbs.
    So we should determine whether the length of the packet modified by bpf
    prog or others like bpf_prog_test is valid before forwarding it directly.

    LINK: [1] https://syzkaller.appspot.com/bug?id=0b84da80c2917757915afa89f7738a9d16ec96c5
    LINK: [2] https://www.spinics.net/lists/netdev/msg777503.html

    Reported-by: syzbot+7a12909485b94426aceb@syzkaller.appspotmail.com
    Signed-off-by: Zhengchao Shao <shaozhengchao@huawei.com>
    Reviewed-by: Stanislav Fomichev <sdf@google.com>
    Link: https://lore.kernel.org/r/20220715115559.139691-1-shaozhengchao@huawei.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d81bd6671f45fde4c3ac7fd7733c6e3082ae9d8e
Author: Yang Jihong <yangjihong1@huawei.com>
Date:   Thu Aug 18 11:26:59 2022 +0800

    ftrace: Fix NULL pointer dereference in is_ftrace_trampoline when ftrace is dead

    commit c3b0f72e805f0801f05fa2aa52011c4bfc694c44 upstream.

    ftrace_startup does not remove ops from ftrace_ops_list when
    ftrace_startup_enable fails:

    register_ftrace_function
      ftrace_startup
        __register_ftrace_function
          ...
          add_ftrace_ops(&ftrace_ops_list, ops)
          ...
        ...
        ftrace_startup_enable // if ftrace failed to modify, ftrace_disabled is set to 1
        ...
      return 0 // ops is in the ftrace_ops_list.

    When ftrace_disabled = 1, unregister_ftrace_function simply returns without doing
anything:
    unregister_ftrace_function
      ftrace_shutdown
        if (unlikely(ftrace_disabled))
                return -ENODEV;  // return here, __unregister_ftrace_function is not
```

```
executed,
                                // as a result, ops is still in the ftrace_ops_list
        __unregister_ftrace_function
        ...

    If ops is dynamically allocated, it will be free later, in this case,
    is_ftrace_trampoline accesses NULL pointer:

    is_ftrace_trampoline
      ftrace_ops_trampoline
        do_for_each_ftrace_op(op, ftrace_ops_list) // OOPS! op may be NULL!

    Syzkaller reports as follows:
    [ 1203.506103] BUG: kernel NULL pointer dereference, address: 000000000000010b
    [ 1203.508039] #PF: supervisor read access in kernel mode
    [ 1203.508798] #PF: error_code(0x0000) - not-present page
    [ 1203.509558] PGD 800000011660b067 P4D 800000011660b067 PUD 130fb8067 PMD 0
    [ 1203.510560] Oops: 0000 [#1] SMP KASAN PTI
    [ 1203.511189] CPU: 6 PID: 29532 Comm: syz-executor.2 Tainted: G    B   W         5.10.0
#8
    [ 1203.512324] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-
g155821a1990b-prebuilt.qemu.org 04/01/2014
    [ 1203.513895] RIP: 0010:is_ftrace_trampoline+0x26/0xb0
    [ 1203.514644] Code: ff eb d3 90 41 55 41 54 49 89 fc 55 53 e8 f2 00 fd ff 48 8b 1d 3b
35 5d 03 e8 e6 00 fd ff 48 8d bb 90 00 00 00 e8 2a 81 26 00 <48> 8b ab 90 00 00 00 48 85 ed
74 1d e8 c9 00 fd ff 48 8d bb 98 00
    [ 1203.518838] RSP: 0018:ffffc900012cf960 EFLAGS: 00010246
    [ 1203.520092] RAX: 0000000000000000 RBX: 000000000000007b RCX: ffffffff8a331866
    [ 1203.521469] RDX: 0000000000000000 RSI: 0000000000000008 RDI: 000000000000010b
    [ 1203.522583] RBP: 0000000000000000 R08: 0000000000000000 R09: ffffffff8df18b07
    [ 1203.523550] R10: fffffbfff1be3160 R11: 0000000000000001 R12: 0000000000478399
    [ 1203.524596] R13: 0000000000000000 R14: ffff888145088000 R15: 0000000000000008
    [ 1203.525634] FS:  00007f429f5f4700(0000) GS:ffff8881daf00000(0000)
knlGS:0000000000000000
    [ 1203.526801] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [ 1203.527626] CR2: 000000000000010b CR3: 0000000170e1e001 CR4: 0000000003706e0
    [ 1203.528611] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [ 1203.529605] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400

    Therefore, when ftrace_startup_enable fails, we need to rollback registration
    process and remove ops from ftrace_ops_list.

    Link: https://lkml.kernel.org/r/20220818032659.56209-1-yangjihong1@huawei.com

    Suggested-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Yang Jihong <yangjihong1@huawei.com>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cb4bb011a683532841344ca7f281b5e04389b4f8
Author: Letu Ren <fantasquex@gmail.com>
Date:   Thu Aug 18 18:44:24 2022 +0800

    fbdev: fb_pm2fb: Avoid potential divide by zero error

    commit 19f953e7435644b81332dd632ba1b2d80b1e37af upstream.

    In `do_fb_ioctl()` of fbmem.c, if cmd is FBIOPUT_VSCREENINFO, var will be
    copied from user, then go through `fb_set_var()` and
    `info->fbops->fb_check_var()` which could may be `pm2fb_check_var()`.
    Along the path, `var->pixclock` won't be modified. This function checks
    whether reciprocal of `var->pixclock` is too high. If `var->pixclock` is
    zero, there will be a divide by zero error. So, it is necessary to check
    whether denominator is zero to avoid crash. As this bug is found by
    Syzkaller, logs are listed below.
```

```
     divide error in pm2fb_check_var
     Call Trace:
      <TASK>
      fb_set_var+0x367/0xeb0 drivers/video/fbdev/core/fbmem.c:1015
      do_fb_ioctl+0x234/0x670 drivers/video/fbdev/core/fbmem.c:1110
      fb_ioctl+0xdd/0x130 drivers/video/fbdev/core/fbmem.c:1189

     Reported-by: Zheyu Ma <zheyuma97@gmail.com>
     Signed-off-by: Letu Ren <fantasquex@gmail.com>
     Signed-off-by: Helge Deller <deller@gmx.de>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 86026be8535c16fcc5e4f960286faf04d7f77815
Author: Hawkins Jiawei <yin31149@gmail.com>
Date:   Fri Aug 5 15:48:34 2022 +0800

    net: fix refcount bug in sk_psock_get (2)

    commit 2a0133723f9ebeb751cfce19f74ec07e108bef1f upstream.

    Syzkaller reports refcount bug as follows:
    ------------[ cut here ]------------
    refcount_t: saturated; leaking memory.
    WARNING: CPU: 1 PID: 3605 at lib/refcount.c:19 refcount_warn_saturate+0xf4/0x1e0
lib/refcount.c:19
    Modules linked in:
    CPU: 1 PID: 3605 Comm: syz-executor208 Not tainted 5.18.0-syzkaller-03023-g7e062cda7d90
#0
     <TASK>
     __refcount_add_not_zero include/linux/refcount.h:163 [inline]
     __refcount_inc_not_zero include/linux/refcount.h:227 [inline]
     refcount_inc_not_zero include/linux/refcount.h:245 [inline]
     sk_psock_get+0x3bc/0x410 include/linux/skmsg.h:439
     tls_data_ready+0x6d/0x1b0 net/tls/tls_sw.c:2091
     tcp_data_ready+0x106/0x520 net/ipv4/tcp_input.c:4983
     tcp_data_queue+0x25f2/0x4c90 net/ipv4/tcp_input.c:5057
     tcp_rcv_state_process+0x1774/0x4e80 net/ipv4/tcp_input.c:6659
     tcp_v4_do_rcv+0x339/0x980 net/ipv4/tcp_ipv4.c:1682
     sk_backlog_rcv include/net/sock.h:1061 [inline]
     __release_sock+0x134/0x3b0 net/core/sock.c:2849
     release_sock+0x54/0x1b0 net/core/sock.c:3404
     inet_shutdown+0x1e0/0x430 net/ipv4/af_inet.c:909
     __sys_shutdown_sock net/socket.c:2331 [inline]
     __sys_shutdown_sock net/socket.c:2325 [inline]
     __sys_shutdown+0xf1/0x1b0 net/socket.c:2343
     __do_sys_shutdown net/socket.c:2351 [inline]
     __se_sys_shutdown net/socket.c:2349 [inline]
     __x64_sys_shutdown+0x50/0x70 net/socket.c:2349
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80
     entry_SYSCALL_64_after_hwframe+0x46/0xb0
     </TASK>

    During SMC fallback process in connect syscall, kernel will
    replaces TCP with SMC. In order to forward wakeup
    smc socket waitqueue after fallback, kernel will sets
    clcsk->sk_user_data to origin smc socket in
    smc_fback_replace_callbacks().

    Later, in shutdown syscall, kernel will calls
    sk_psock_get(), which treats the clcsk->sk_user_data
    as psock type, triggering the refcnt warning.

    So, the root cause is that smc and psock, both will use
```

```
    sk_user_data field. So they will mismatch this field
    easily.

    This patch solves it by using another bit(defined as
    SK_USER_DATA_PSOCK) in PTRMASK, to mark whether
    sk_user_data points to a psock object or not.
    This patch depends on a PTRMASK introduced in commit f1ff5ce2cd5e
    ("net, sk_msg: Clear sk_user_data pointer on clone if tagged").

    For there will possibly be more flags in the sk_user_data field,
    this patch also refactor sk_user_data flags code to be more generic
    to improve its maintainability.

    Reported-and-tested-by: syzbot+5f26f85569bd179c18ce@syzkaller.appspotmail.com
    Suggested-by: Jakub Kicinski <kuba@kernel.org>
    Acked-by: Wen Gu <guwen@linux.alibaba.com>
    Signed-off-by: Hawkins Jiawei <yin31149@gmail.com>
    Reviewed-by: Jakub Sitnicki <jakub@cloudflare.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 52a3c62a815161c2dcf38ac421f6c41d8679462b
Author: Karthik Alapati <mail@karthek.com>
Date:   Thu Jul 28 21:13:17 2022 +0530

    HID: hidraw: fix memory leak in hidraw_release()

    commit a5623a203cffe2d2b84d2f6c989d9017db1856af upstream.

    Free the buffered reports before deleting the list entry.

    BUG: memory leak
    unreferenced object 0xffff88810e72f180 (size 32):
      comm "softirq", pid 0, jiffies 4294945143 (age 16.080s)
      hex dump (first 32 bytes):
        64 f3 c6 6a d1 88 07 04 00 00 00 00 00 00 00 00  d..j............
        00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
      backtrace:
        [<ffffffff814ac6c3>] kmemdup+0x23/0x50 mm/util.c:128
        [<ffffffff8357c1d2>] kmemdup include/linux/fortify-string.h:440 [inline]
        [<ffffffff8357c1d2>] hidraw_report_event+0xa2/0x150 drivers/hid/hidraw.c:521
        [<ffffffff8356ddad>] hid_report_raw_event+0x27d/0x740 drivers/hid/hid-core.c:1992
        [<ffffffff8356e41e>] hid_input_report+0x1ae/0x270 drivers/hid/hid-core.c:2065
        [<ffffffff835f0d3f>] hid_irq_in+0x1ff/0x250 drivers/hid/usbhid/hid-core.c:284
        [<ffffffff82d3c7f9>] __usb_hcd_giveback_urb+0xf9/0x230 drivers/usb/core/hcd.c:1670
        [<ffffffff82d3cc26>] usb_hcd_giveback_urb+0x1b6/0x1d0 drivers/usb/core/hcd.c:1747
        [<ffffffff82ef1e14>] dummy_timer+0x8e4/0x14c0
  drivers/usb/gadget/udc/dummy_hcd.c:1988
        [<ffffffff812f50a8>] call_timer_fn+0x38/0x200 kernel/time/timer.c:1474
        [<ffffffff812f5586>] expire_timers kernel/time/timer.c:1519 [inline]
        [<ffffffff812f5586>] __run_timers.part.0+0x316/0x430 kernel/time/timer.c:1790
        [<ffffffff812f56e4>] __run_timers kernel/time/timer.c:1768 [inline]
        [<ffffffff812f56e4>] run_timer_softirq+0x44/0x90 kernel/time/timer.c:1803
        [<ffffffff848000e6>] __do_softirq+0xe6/0x2ea kernel/softirq.c:571
        [<ffffffff81246db0>] invoke_softirq kernel/softirq.c:445 [inline]
        [<ffffffff81246db0>] __irq_exit_rcu kernel/softirq.c:650 [inline]
        [<ffffffff81246db0>] irq_exit_rcu+0xc0/0x110 kernel/softirq.c:662
        [<ffffffff84574f02>] sysvec_apic_timer_interrupt+0xa2/0xd0
  arch/x86/kernel/apic/apic.c:1106
        [<ffffffff84600c8b>] asm_sysvec_apic_timer_interrupt+0x1b/0x20
  arch/x86/include/asm/idtentry.h:649
        [<ffffffff8458a070>] native_safe_halt arch/x86/include/asm/irqflags.h:51 [inline]
        [<ffffffff8458a070>] arch_safe_halt arch/x86/include/asm/irqflags.h:89 [inline]
        [<ffffffff8458a070>] acpi_safe_halt drivers/acpi/processor_idle.c:111 [inline]
        [<ffffffff8458a070>] acpi_idle_do_entry+0xc0/0xd0 drivers/acpi/processor_idle.c:554
```

```
    Link: https://syzkaller.appspot.com/bug?id=19a04b43c75ed1092021010419b5e560a8172c4f
    Reported-by: syzbot+f59100a0428e6ded9443@syzkaller.appspotmail.com
    Signed-off-by: Karthik Alapati <mail@karthek.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f44b0b95d50fffeca036e1ba36770390e0b519dd
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Thu Jul 21 11:07:10 2022 -0400

    USB: gadget: Fix use-after-free Read in usb_udc_uevent()

    commit 2191c00855b03aa59c20e698be713d952d51fc18 upstream.

    The syzbot fuzzer found a race between uevent callbacks and gadget
    driver unregistration that can cause a use-after-free bug:

    ------------------------------------------------------------
    BUG: KASAN: use-after-free in usb_udc_uevent+0x11f/0x130
    drivers/usb/gadget/udc/core.c:1732
    Read of size 8 at addr ffff888078ce2050 by task udevd/2968

    CPU: 1 PID: 2968 Comm: udevd Not tainted 5.19.0-rc4-next-20220628-syzkaller #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
    06/29/2022
    Call Trace:
     <TASK>
     __dump_stack lib/dump_stack.c:88 [inline]
     dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106
     print_address_description mm/kasan/report.c:317 [inline]
     print_report.cold+0x2ba/0x719 mm/kasan/report.c:433
     kasan_report+0xbe/0x1f0 mm/kasan/report.c:495
     usb_udc_uevent+0x11f/0x130 drivers/usb/gadget/udc/core.c:1732
     dev_uevent+0x290/0x770 drivers/base/core.c:2424
    ------------------------------------------------------------

    The bug occurs because usb_udc_uevent() dereferences udc->driver but
    does so without acquiring the udc_lock mutex, which protects this
    field.  If the gadget driver is unbound from the udc concurrently with
    uevent processing, the driver structure may be accessed after it has
    been deallocated.

    To prevent the race, we make sure that the routine holds the mutex
    around the racing accesses.

    Link: <https://lore.kernel.org/all/0000000000004de90405a719c951@google.com>
    CC: stable@vger.kernel.org # fc274c1e9973
    Reported-and-tested-by: syzbot+b0de012ceb1e2a97891b@syzkaller.appspotmail.com
    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Link: https://lore.kernel.org/r/YtlrnhHyrHsSky9m@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c02d2a91a85c4c4d05826cd1ea74a9b8d42e4280
Author: Dongliang Mu <mudongliangabcd@gmail.com>
Date:   Thu Jun 9 08:35:28 2022 +0100

    media: pvrusb2: fix memory leak in pvr_probe

    commit 945a9a8e448b65bec055d37eba58f711b39f66f0 upstream.

    The error handling code in pvr2_hdw_create forgets to unregister the
    v4l2 device. When pvr2_hdw_create returns back to pvr2_context_create,
    it calls pvr2_context_destroy to destroy context, but mp->hdw is NULL,
    which leads to that pvr2_hdw_destroy directly returns.
```

```
    Fix this by adding v4l2_device_unregister to decrease the refcount of
    usb interface.

    Reported-by: syzbot+77b432d57c4791183ed4@syzkaller.appspotmail.com
    Signed-off-by: Dongliang Mu <mudongliangabcd@gmail.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f2f6ea1a8da1317430a84701fc0170449ee88315
Author: Vivek Kasireddy <vivek.kasireddy@intel.com>
Date:   Fri May 20 13:52:35 2022 -0700

    udmabuf: Set the DMA mask for the udmabuf device (v2)

    commit 9e9fa6a9198b767b00f48160800128e83a038f9f upstream.

    If the DMA mask is not set explicitly, the following warning occurs
    when the userspace tries to access the dma-buf via the CPU as
    reported by syzbot here:

    WARNING: CPU: 1 PID: 3595 at kernel/dma/mapping.c:188
    __dma_map_sg_attrs+0x181/0x1f0 kernel/dma/mapping.c:188
    Modules linked in:
    CPU: 0 PID: 3595 Comm: syz-executor249 Not tainted
    5.17.0-rc2-syzkaller-00316-g0457e5153e0e #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS
    Google 01/01/2011
    RIP: 0010:__dma_map_sg_attrs+0x181/0x1f0 kernel/dma/mapping.c:188
    Code: 00 00 00 00 fc ff df 48 c1 e8 03 80 3c 10 00 75 71 4c 8b 3d c0
    83 b5 0d e9 db fe ff ff e8 b6 0f 13 00 0f 0b e8 af 0f 13 00 <0f> 0b 45
       31 e4 e9 54 ff ff ff e8 a0 0f 13 00 49 8d 7f 50 48 b8 00
    RSP: 0018:ffffc90002a07d68 EFLAGS: 00010293
    RAX: 0000000000000000 RBX: 0000000000000000 RCX: 0000000000000000
    RDX: ffff88807e25e2c0 RSI: ffffffff81649e91 RDI: ffff88801b848408
    RBP: ffff88801b848000 R08: 0000000000000002 R09: ffff88801d86c74f
    R10: ffffffff81649d72 R11: 0000000000000001 R12: 0000000000000002
    R13: ffff88801d86c680 R14: 0000000000000001 R15: 0000000000000000
    FS:  0000555556e30300(0000) GS:ffff8880b9d00000(0000)
    knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00000000200000cc CR3: 000000001d74a000 CR4: 00000000003506e0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     <TASK>
     dma_map_sgtable+0x70/0xf0 kernel/dma/mapping.c:264
     get_sg_table.isra.0+0xe0/0x160 drivers/dma-buf/udmabuf.c:72
     begin_cpu_udmabuf+0x130/0x1d0 drivers/dma-buf/udmabuf.c:126
     dma_buf_begin_cpu_access+0xfd/0x1d0 drivers/dma-buf/dma-buf.c:1164
     dma_buf_ioctl+0x259/0x2b0 drivers/dma-buf/dma-buf.c:363
     vfs_ioctl fs/ioctl.c:51 [inline]
     __do_sys_ioctl fs/ioctl.c:874 [inline]
     __se_sys_ioctl fs/ioctl.c:860 [inline]
     __x64_sys_ioctl+0x193/0x200 fs/ioctl.c:860
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80
     entry_SYSCALL_64_after_hwframe+0x44/0xae
    RIP: 0033:0x7f62fcf530f9
    Code: 28 c3 e8 2a 14 00 00 66 2e 0f 1f 84 00 00 00 00 00 48 89 f8 48 89
    f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01
    f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d8 64 89 01 48
    RSP: 002b:00007ffe3edab9b8 EFLAGS: 00000246 ORIG_RAX: 0000000000000010
    RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 00007f62fcf530f9
```

```
     RDX: 0000000020000200 RSI: 0000000040086200 RDI: 0000000000000006
     RBP: 00007f62fcf170e0 R08: 0000000000000000 R09: 0000000000000000
     R10: 0000000000000000 R11: 0000000000000246 R12: 00007f62fcf17170
     R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000000
      </TASK>

     v2: Dont't forget to deregister if DMA mask setup fails.

     Reported-by: syzbot+10e27961f4da37c443b2@syzkaller.appspotmail.com
     Cc: Gerd Hoffmann <kraxel@redhat.com>
     Signed-off-by: Vivek Kasireddy <vivek.kasireddy@intel.com>
     Link: http://patchwork.freedesktop.org/patch/msgid/20220520205235.3687336-1-
 vivek.kasireddy@intel.com
     Signed-off-by: Gerd Hoffmann <kraxel@redhat.com>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit dee1e51b54794e90763e70a3c78f27ba4fa930ec
 Author: Lee Jones <lee.jones@linaro.org>
 Date:   Fri Jul 8 08:40:09 2022 +0100

     HID: steam: Prevent NULL pointer dereference in steam_{recv,send}_report

     commit cd11d1a6114bd4bc6450ae59f6e110ec47362126 upstream.

     It is possible for a malicious device to forgo submitting a Feature
     Report.  The HID Steam driver presently makes no prevision for this
     and de-references the 'struct hid_report' pointer obtained from the
     HID devices without first checking its validity.  Let's change that.

     Cc: Jiri Kosina <jikos@kernel.org>
     Cc: Benjamin Tissoires <benjamin.tissoires@redhat.com>
     Cc: linux-input@vger.kernel.org
     Fixes: c164d6abf3841 ("HID: add driver for Valve Steam Controller")
     Signed-off-by: Lee Jones <lee.jones@linaro.org>
     Signed-off-by: Jiri Kosina <jkosina@suse.cz>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit adb9caabcd16a4a7535f3fcce79d98eab185d434
 Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
 Date:   Thu Sep 1 13:01:03 2022 +0200

     Revert "PCI/portdrv: Don't disable AER reporting in get_port_device_capability()"

     This reverts commit 65e393fddc5379b2c41ca7e73cd4bb9572c4d90e which is
     commit 8795e182b02dc87e343c79e73af6b8b7f9c5e635 upstream.

     It is reported to cause problems, so drop it from the stable trees for
     now until it gets sorted out.

     Link: https://lore.kernel.org/r/47b775c5-57fa-5edf-b59e-8a9041ffbee7@candelatech.com
     Reported-by: Ben Greear <greearb@candelatech.com>
     Cc: Stefan Roese <sr@denx.de>
     Cc: Bjorn Helgaas <bhelgaas@google.com>
     Cc: Pali Rohár <pali@kernel.org>
     Cc: Rafael J. Wysocki <rjw@rjwysocki.net>
     Cc: Bharat Kumar Gogada <bharat.kumar.gogada@xilinx.com>
     Cc: Michal Simek <michal.simek@xilinx.com>
     Cc: Yao Hongbo <yaohongbo@linux.alibaba.com>
     Cc: Naveen Naidu <naveennaidu479@gmail.com>
     Cc: Sasha Levin <sashal@kernel.org>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit e8fcbce5424966ddbe9e426c163b2e63ff91b789
 Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
 Date:   Fri Aug 12 15:33:57 2022 -0700
```

```
    Bluetooth: L2CAP: Fix build errors in some archs

    commit b840304fb46cdf7012722f456bce06f151b3e81b upstream.

    This attempts to fix the follow errors:

    In function 'memcmp',
        inlined from 'bacmp' at ./include/net/bluetooth/bluetooth.h:347:9,
        inlined from 'l2cap_global_chan_by_psm' at
        net/bluetooth/l2cap_core.c:2003:15:
    ./include/linux/fortify-string.h:44:33: error: '__builtin_memcmp'
    specified bound 6 exceeds source size 0 [-Werror=stringop-overread]
       44 | #define __underlying_memcmp     __builtin_memcmp
          |                                 ^
    ./include/linux/fortify-string.h:420:16: note: in expansion of macro
    '__underlying_memcmp'
      420 |         return __underlying_memcmp(p, q, size);
          |                ^~~~~~~~~~~~~~~~~~~
    In function 'memcmp',
        inlined from 'bacmp' at ./include/net/bluetooth/bluetooth.h:347:9,
        inlined from 'l2cap_global_chan_by_psm' at
        net/bluetooth/l2cap_core.c:2004:15:
    ./include/linux/fortify-string.h:44:33: error: '__builtin_memcmp'
    specified bound 6 exceeds source size 0 [-Werror=stringop-overread]
       44 | #define __underlying_memcmp     __builtin_memcmp
          |                                 ^
    ./include/linux/fortify-string.h:420:16: note: in expansion of macro
    '__underlying_memcmp'
      420 |         return __underlying_memcmp(p, q, size);
          |                ^~~~~~~~~~~~~~~~~~~

    Fixes: 332f1795ca20 ("Bluetooth: L2CAP: Fix l2cap_global_chan_by_psm regression")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Cc: Sudip Mukherjee <sudipm.mukherjee@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d3916da4bc6c022c9c41c576f6538fa78e5ce746
Author: James Morse <james.morse@arm.com>
Date:   Mon Jul 4 16:57:32 2022 +0100

    arm64: errata: Add Cortex-A510 to the repeat tlbi list

    commit 39fdb65f52e9a53d32a6ba719f96669fd300ae78 upstream.

    Cortex-A510 is affected by an erratum where in rare circumstances the
    CPUs may not handle a race between a break-before-make sequence on one
    CPU, and another CPU accessing the same page. This could allow a store
    to a page that has been unmapped.

    Work around this by adding the affected CPUs to the list that needs
    TLB sequences to be done twice.

    Signed-off-by: James Morse <james.morse@arm.com>
    Link: https://lore.kernel.org/r/20220704155732.21216-1-james.morse@arm.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Cc: Lucas Wei <lucaswei@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a8ba16e9f171e9638401ff39daebdd2b84553e2d
Author: Akira Yokosawa <akiyks@gmail.com>
Date:   Mon Aug 8 17:53:57 2022 +0900

    docs: kerneldoc-preamble: Test xeCJK.sty before loading
```

```
    commit cee7db1b0239468b22c295cf04a8c40c34ecd35a upstream.

    On distros whose texlive packaging is fine-grained, texlive-xecjk
    can be installed/removed independently of other texlive packages.
    Conditionally loading xeCJK depending only on the existence of the
    "Noto Sans CJK SC" font might end up in xelatex error of
    "xeCJK.sty not found!".

    Improve the situation by testing existence of xeCJK.sty before
    loading it.

    This is useful on RHEL 9 and its clone distros where texlive-xecjk
    doesn't work at the moment due to a missing dependency [1].
    "make pdfdocs" for non-CJK contents should work after removing
    texlive-xecjk.

    Link: [1] https://bugzilla.redhat.com/show_bug.cgi?id=2086254
    Fixes: 398f7abdcb7e ("docs: pdfdocs: Pull LaTeX preamble part out of conf.py")
    Cc: stable@vger.kernel.org # v5.18+
    Signed-off-by: Akira Yokosawa <akiyks@gmail.com>
    Acked-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Link: https://lore.kernel.org/r/c24c2a87-70b2-5342-bcc9-de467940466e@gmail.com
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e120a63532946bad71dc6255579a3219a9caf66c
Author: Eric Biggers <ebiggers@google.com>
Date:   Thu Aug 25 22:04:56 2022 -0700

    crypto: lib - remove unneeded selection of XOR_BLOCKS

    commit 874b301985ef2f89b8b592ad255e03fb6fbfe605 upstream.

    CRYPTO_LIB_CHACHA_GENERIC doesn't need to select XOR_BLOCKS.  It perhaps
    was thought that it's needed for __crypto_xor, but that's not the case.

    Enabling XOR_BLOCKS is problematic because the XOR_BLOCKS code runs a
    benchmark when it is initialized.  That causes a boot time regression on
    systems that didn't have it enabled before.

    Therefore, remove this unnecessary and problematic selection.

    Fixes: e56e18985596 ("lib/crypto: add prompts back to crypto libraries")
    Cc: stable@vger.kernel.org
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 515af71adf3d5e05587c7bb27c1725c20cd259a0
Author: Timo Alho <talho@nvidia.com>
Date:   Wed Jun 22 16:22:59 2022 +0300

    firmware: tegra: bpmp: Do only aligned access to IPC memory area

    commit a4740b148a04dc60e14fe6a1dfe216d3bae214fd upstream.

    Use memcpy_toio and memcpy_fromio variants of memcpy to guarantee no
    unaligned access to IPC memory area. This is to allow the IPC memory to
    be mapped as Device memory to further suppress speculative reads from
    happening within the 64 kB memory area above the IPC memory when 64 kB
    memory pages are used.

    Signed-off-by: Timo Alho <talho@nvidia.com>
    Signed-off-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
```

```
    Cc: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1cf82931822dbf3b24945147b5330d8fdd86cdff
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Jun 29 14:34:36 2022 +0200

    drm/vc4: hdmi: Depends on CONFIG_PM

    commit 72e2329e7c9bbe15e7a813670497ec9c6f919af3 upstream.

    We already depend on runtime PM to get the power domains and clocks for
    most of the devices supported by the vc4 driver, so let's just select it
    to make sure it's there.

    Link: https://lore.kernel.org/r/20220629123510.1915022-38-maxime@cerno.tech
    Acked-by: Thomas Zimmermann <tzimmermann@suse.de>
    Tested-by: Stefan Wahren <stefan.wahren@i2se.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    (cherry picked from commit f1bc386b319e93e56453ae27e9e83817bb1f6f95)
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Cc: "Sudip Mukherjee (Codethink)" <sudipm.mukherjee@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1a289a16857062e704da52683a49aaf16e826cbf
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Jun 29 14:34:37 2022 +0200

    drm/vc4: hdmi: Rework power up

    commit 258e483a4d5e97a6a8caa74381ddc1f395ac1c71 upstream.

    The current code tries to handle the case where CONFIG_PM isn't selected
    by first calling our runtime_resume implementation and then properly
    report the power state to the runtime_pm core.

    This allows to have a functionning device even if pm_runtime_get_*
    functions are nops.

    However, the device power state if CONFIG_PM is enabled is
    RPM_SUSPENDED, and thus our vc4_hdmi_write() and vc4_hdmi_read() calls
    in the runtime_pm hooks will now report a warning since the device might
    not be properly powered.

    Even more so, we need CONFIG_PM enabled since the previous RaspberryPi
    have a power domain that needs to be powered up for the HDMI controller
    to be usable.

    The previous patch has created a dependency on CONFIG_PM, now we can
    just assume it's there and only call pm_runtime_resume_and_get() to make
    sure our device is powered in bind.

    Link: https://lore.kernel.org/r/20220629123510.1915022-39-maxime@cerno.tech
    Acked-by: Thomas Zimmermann <tzimmermann@suse.de>
    Tested-by: Stefan Wahren <stefan.wahren@i2se.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    (cherry picked from commit 53565c28e6af2cef6bbf438c34250135e3564459)
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Cc: "Sudip Mukherjee (Codethink)" <sudipm.mukherjee@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```